

IBM Information

>>> On Demand

2007



SQL Tuning for Everyone: Using DB2 Optimization Expert for z/OS

*Jay Bruce, DB2 for z/OS autonomic
optimization, IBM, jmbruce@us.ibm.com*

*1846 Data Servers - DB2 and Tools for Linux, UNIX,
Windows*



Act.Right.Now.

IBM INFORMATION ON DEMAND 2007

October 14 - 19, 2007

Mandalay Bay

Las Vegas, Nevada



Overview



Four Tiered Implementation

- DB2 9 for z/OS enhancements
 - Profile based Query Performance normal/exception monitoring
 - Virtual Index “What-if” Optimizer support
 - Server based Administrative Scheduler
 - Expanded Explain facility
- Built on Eclipse
 - New Outside In Design process
- Optimization Service Center (OSC)
 - Visual Explain functional replacement
 - Workload Control Center
- Optimization Expert (OE)
 - Includes all OSC functions
 - Index, Query and Access Path Advisors



OSC / OE Availability

- DB2 for z/OS Optimization Service Center
 - New no-charge offering
 - Included in DB2 Accessories Suite for z/OS V1.1
 - Available via web download
- DB2 Optimization Expert for z/OS V1.1
 - New One Time Charge (OTC) offering
 - Value Unit, Reference based pricing
- Simultaneous availability with DB2 9 for z/OS



Family Product Positioning

Functions	Visual Explain	Optimization Service Center	Optimization Expert
Queries from Cache, Catalog	Yes	Yes	Yes
Query Formatter, Annotation		Yes	Yes
Access Plan Graph	Yes	Yes	Yes
Visual Plan Hint		Yes	Yes
Query Statistics Advisor	Yes	Yes	Yes
Workload Statistics Advisor		Yes	Yes
Profile based Monitoring		Yes	Yes
Query Index Advisor			Yes
Workload Index Advisor			Yes
Query Advisor			Yes
Workload Query Advisor			Yes
Access Path Advisor			Yes





Introduction



Application Performance Problem

- Applications are designed and implemented very quickly
- Often there is insufficient skill and resource 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 in the workload, the frequency of execution, and cost of individual operations
- The overwhelming amount of resource required to perform the review often means the analysis is either not done, or done incompletely



DB2 Solution

- Problem Query Identification
 - ✓ Snapping queries from various sources
 - ✓ Monitoring performance exceptions

- Problem Query Resolution
 - ✓ Design Advisors for recommendation of stats, index, and query design
 - ✓ Turing tools for deep analysis of problem query
 - ✓ Query format to present a readable query
 - ✓ Annotation of optimizer rewritten query to embed critical information
 - ✓ Intelligent report to show the underlying physical design with critical information
 - ✓ Visual explain to show the access path choice
 - ✓ Visual optimization hints to implement emergency solution
 - ✓ Service SQL to send relevant doc to IBM for diagnosis

- Tuning Application Workload
 - ✓ Automatic collection of application queries
 - ✓ Design Advisors for regular health check
 - ✓ Monitoring application performance exceptions
 - ✓ Problem resolution with Design Advisors and tuning tools

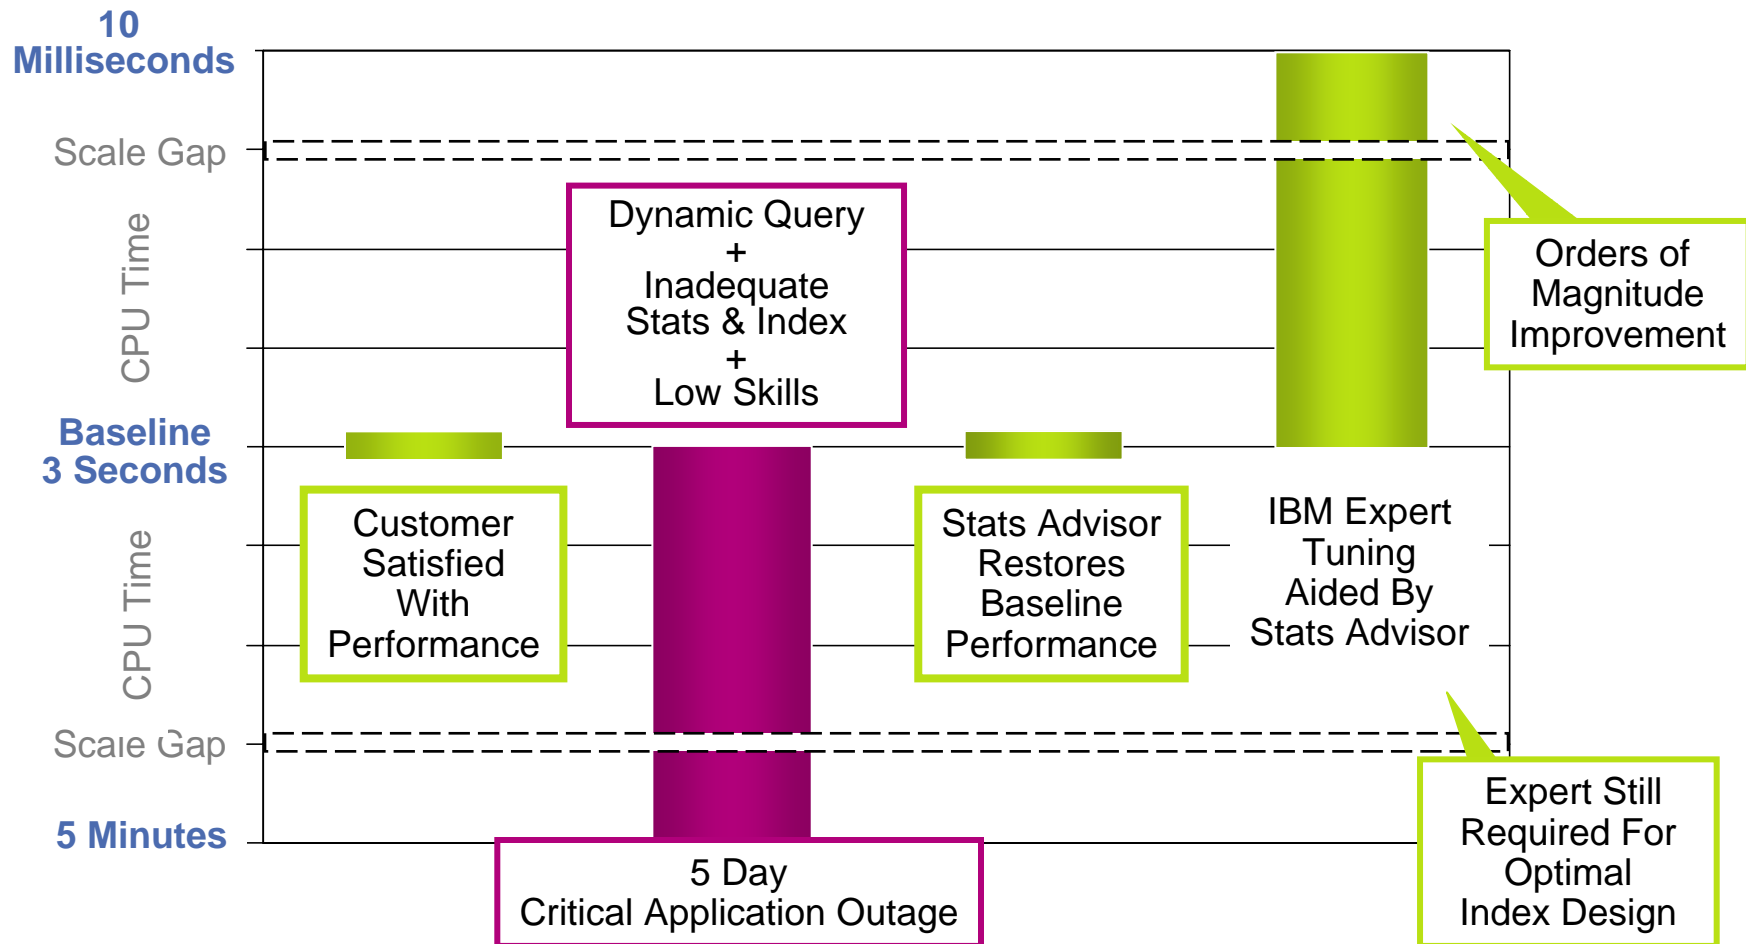




Identifying Problem Query



Scenario: Critical application outage



Scenario: Critical application outage

Post-mortem Analysis

- It sounds easy but actually not
 - ✓ Inadequate query tuning and physical database design skills
 - ✓ Took 3 days to identify the problem query
 - ✓ No idea how to investigate the performance problem

- Customer costs
 - ✓ Time and money lost due to application outage
 - ✓ Loss of confidence in IT team
 - ✓ Extensive and costly performance review performed for simple regression

- Lack of expert design rules and methodology
 - ✓ Original query performance was actual suboptimal
 - ✓ Good design rules and methodology would result in faster & more stable performance

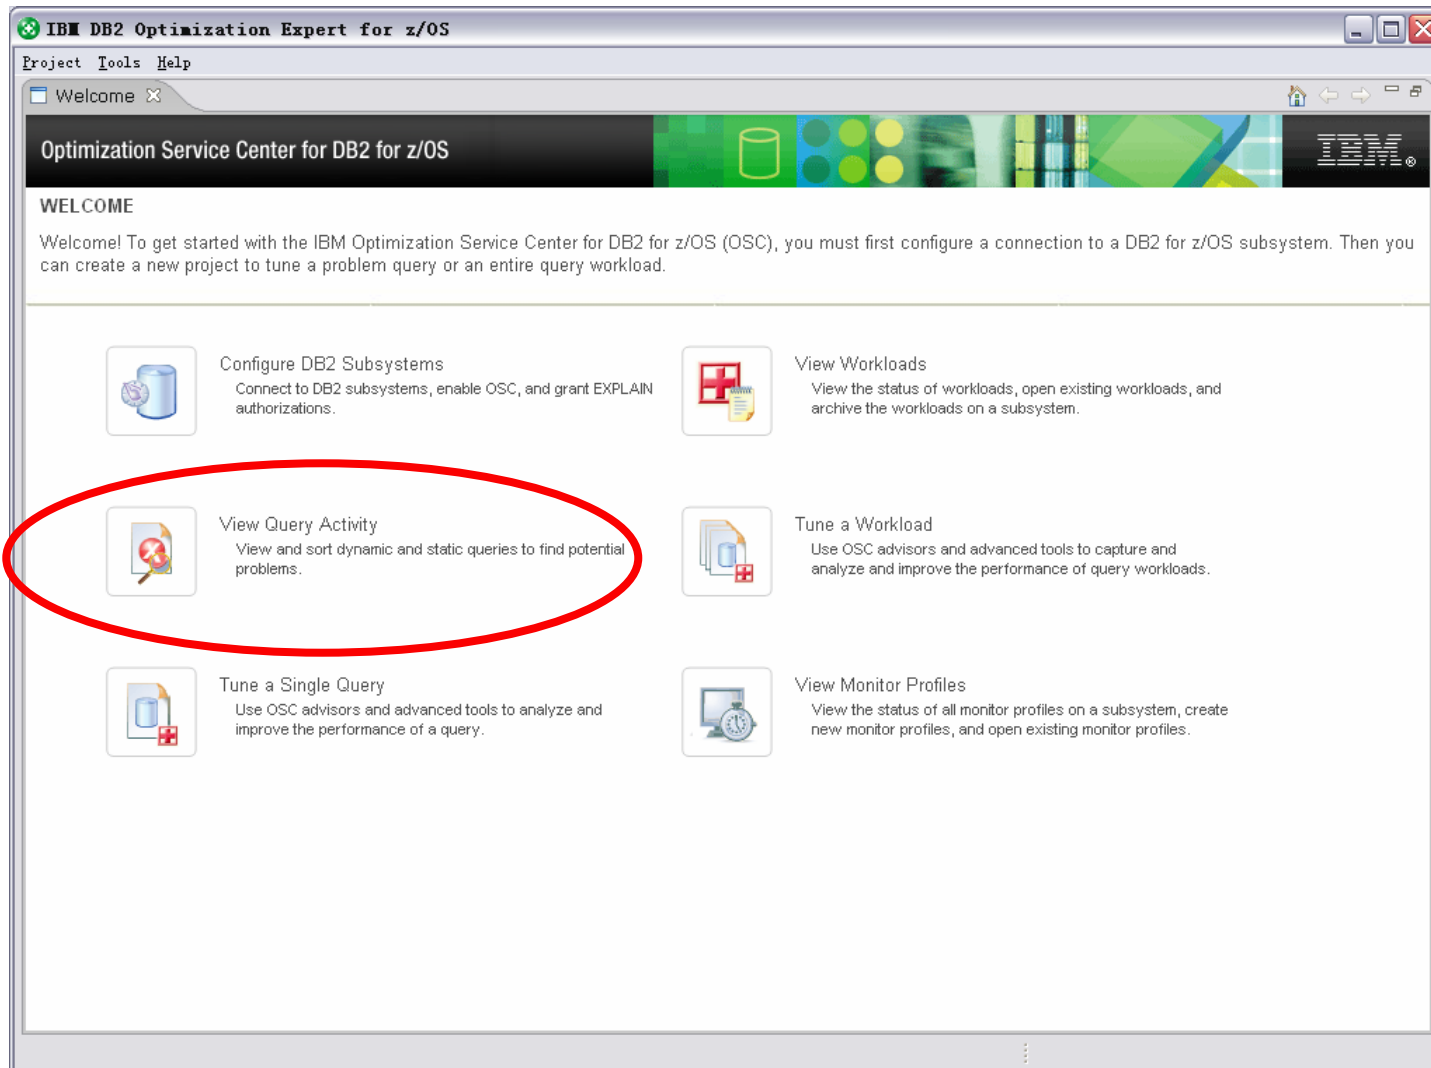


- Identifying Problem Query
 - From Dynamic Statement Cache
 - Many other query sources are supported

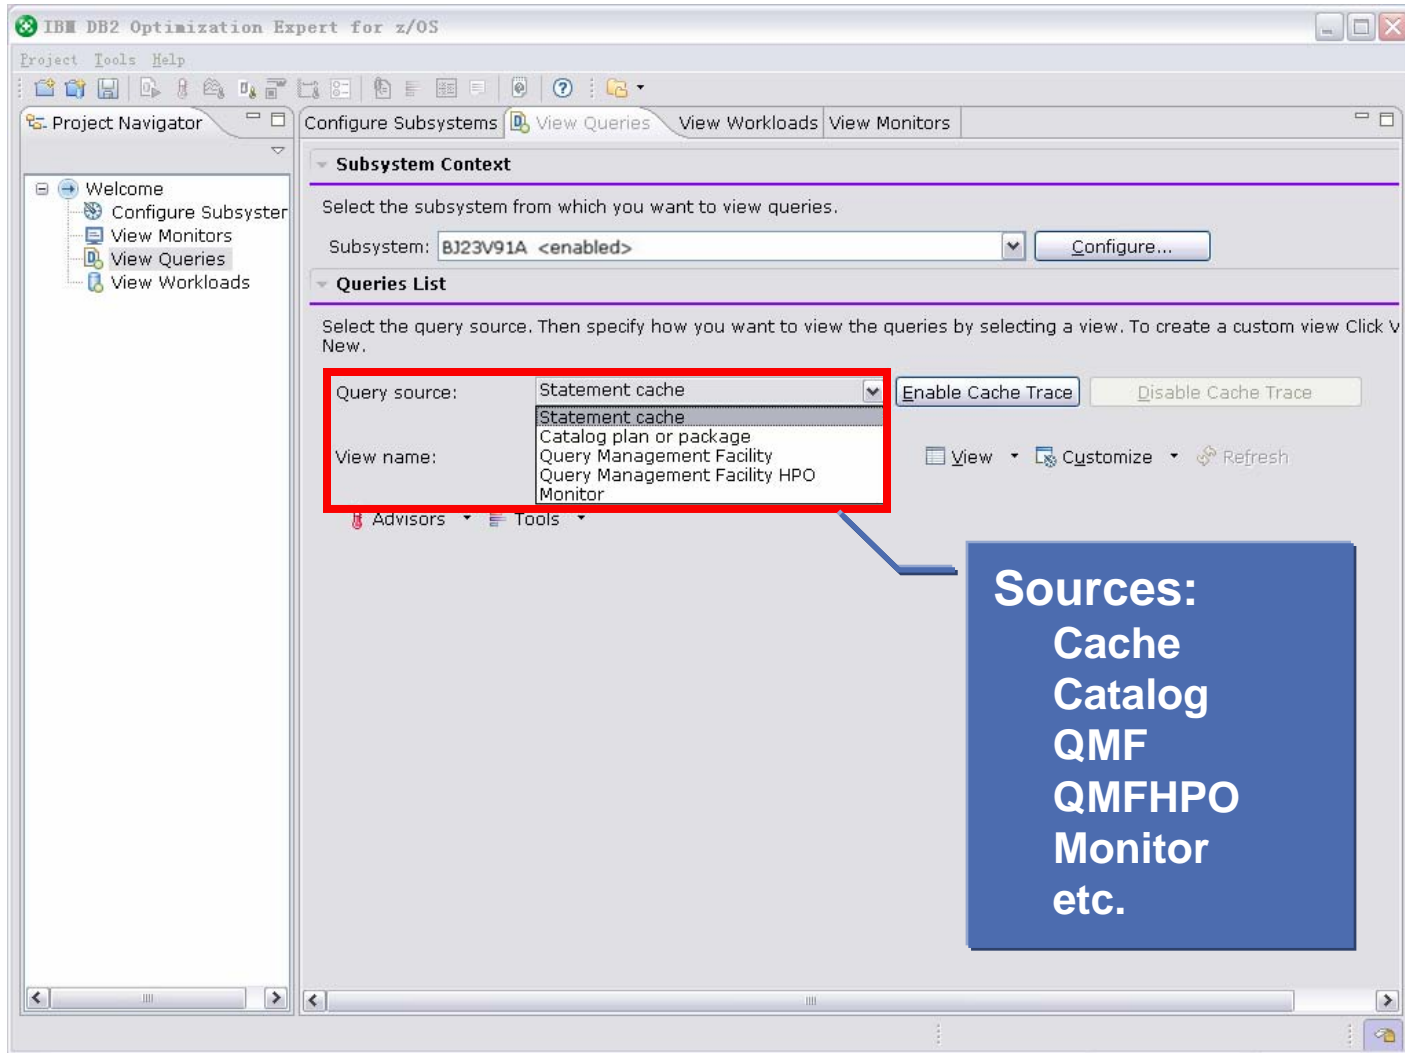
- Identifying Problem Query with Query Monitors
 - Automatic notification for performance exceptions



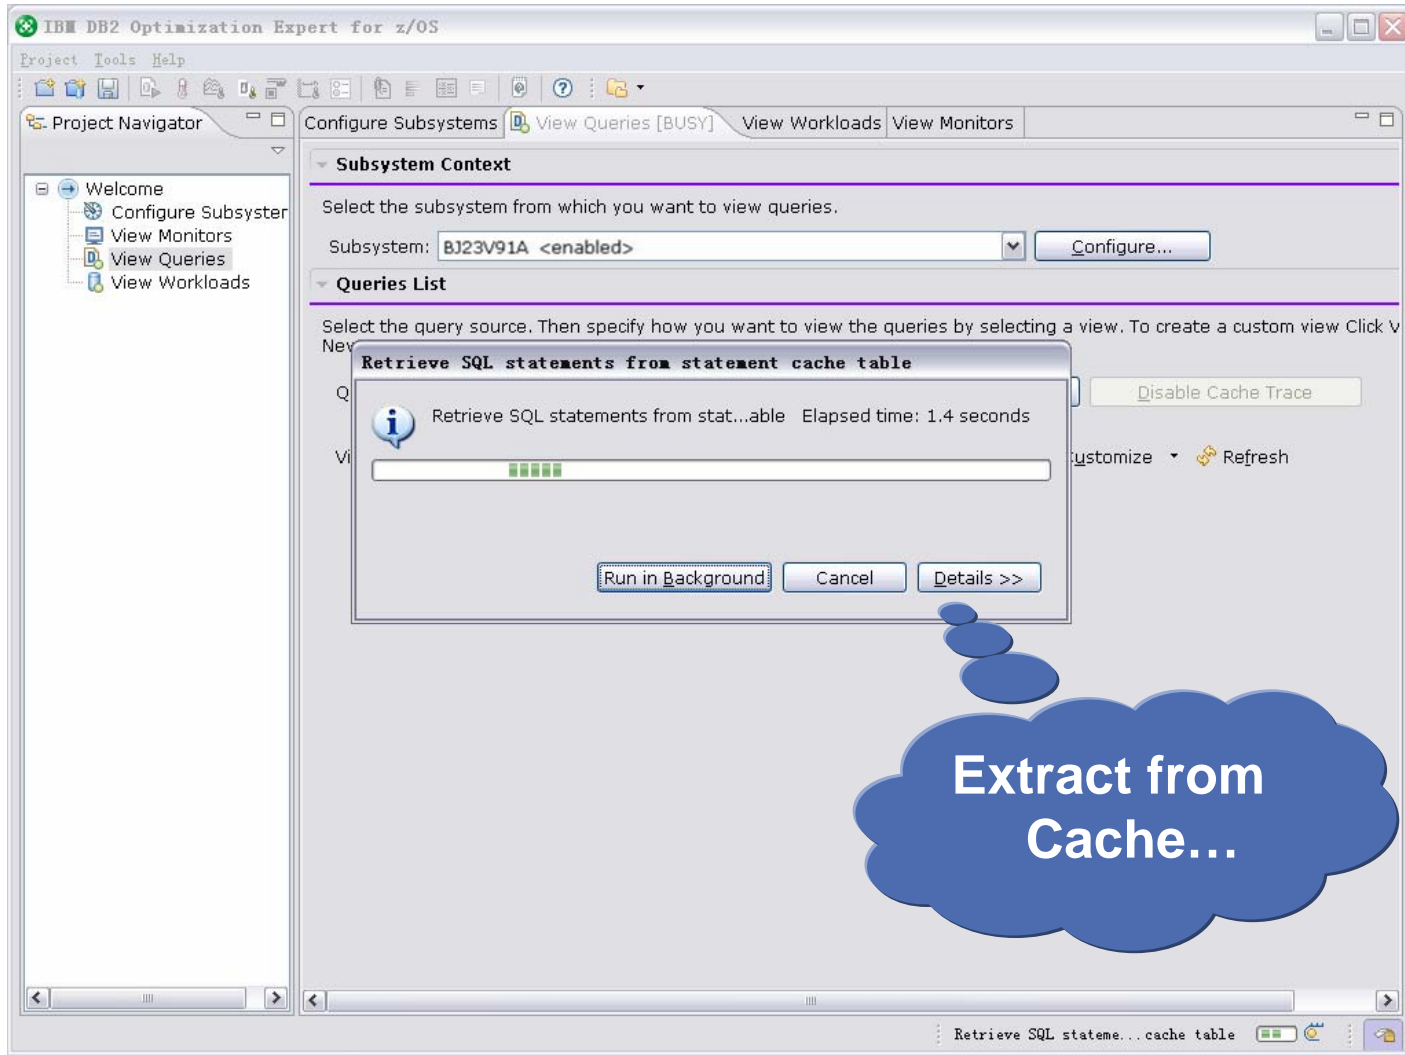
Identifying Problem Query – Welcome Panel



Identifying Problem Query – Dynamic Statement Cache



Identifying Problem Query – Dynamic Statement Cache



Identifying Problem Query – Dynamic Statement Cache

The screenshot shows the 'IBM DB2 Optimization Expert for z/OS' interface. The 'Subsystem Context' section shows 'Subsystem: BJ23V91A <enabled>'. The 'Queries List' section shows 'Query source: Statement cache' and 'View name: ACCUM_CPU_DESC'. A table of cached statements is displayed, sorted by CPU time. A red box highlights the table content.

STAT_EXEC	STAT_CPU	STAT_ELAP	STMT_TEXT
5	9.888036967...	189.97887...	SELECT L_ORDERKEY, SUM(L_EXTENDEDPRICE) AS REVENUE, O_ORDERDATE, O_SHIPPRIORITY FROM
5	8.001804591...	88.504505...	SELECT L_SUPPKEY, COUNT(*), MIN(L_TAX), MAX(L_TAX), SUM(L_EXTENDEDPRICE), AVG(L_EXT
5	2.412140131...	31.538043...	SELECT DISTINCT O_ORDERKEY FROM SYSADM.LINEITEM, SYSADM.ORDER WHERE L_ORDERKEY = O
5	1.040755511...	247.02081...	SELECT N_NAME, L_EXTENDEDPRICE AS REVENUE FROM CUSTOMER, ORDER, LINEITEM, SUPPLI
1	0.804309323...	1.3527300...	DELETE FROM SYSIBM.SYSCOLDISTSTATS WHERE TOWNER = 'SYSADM'
5	0.557750597...	33.846088...	SELECT S_SUPPKEY, S_NAME, SUM(L_EXTENDEDPRICE*(1-L_DISCOUNT)) AS REVENUE FROM SYSA
5	0.111229434...	10.281691...	SELECT L_ORDERKEY, L_SUPPKEY, L_SHIPDATE, L_RETURNFLAG FROM SYSADM.LINEITEM WHERE L_S
5	0.110829547...	33.650882...	SELECT * FROM SYSADM.LINEITEM L, SYSADM.ORDER O, SYSADM.SUPPLIER S WHERE L.L_RECEI
1	0.033946693...	0.1935956...	DELETE FROM SYSIBM.SYSCOLDIST WHERE TOWNER = 'SYSADM'
1	0.008425281...	0.0234471...	DELETE FROM DSN.STATEMENT_CACHE_TABLE
0	0.001826585...	0.0018725...	SELECT STMT_ID, STMT_TOKEN, COLLID, PROGRAM_NAME, INV_DROPALT, INV_REVOKE, INV_LRU,
0	4.801565308...	0.0019182...	SELECT CURRENT SQLID FROM SYSIBM.SYSDUMMY1
0	3.659374853...	3.6578119...	SELECT 1 FROM SYSIBM.SYSDUMMY1 WHERE 0 = 1;
0	0.0	0.0	SELECT CURRENT TIMESTAMP AS TIMESTAMP FROM SYSIBM.SYSDUMMY1

Cached Statements, sorted by descending accumulated CPU time



Identifying Problem Query – Dynamic Statement Cache

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Project Navigator

Welcome
Configure S
View Monit
View Quer
View Work

Create View

Filter Rows

Specify criteria to limit the query rows that are returned by typing values in the Value column and specifying operators in the Operator column. If you do not want to filter the queries, do not

Steps

- 1. Filter
- 2. Sort
- 3. Columns

View name: ACCUM_CPU_DESC

Maximum rows: 100

Column Name	Operator	Value	Comment
STAT_ELAP	=		The accumulated elapsed time that i
STAT_CPU	=		The accumulated CPU time that is us
STAT_SUS_SYN...	=		The accumulated wait time for synch
STAT_SUS_LOCK...	=		The accumulated wait time for lock a
STAT_SUS_SWIT...	=		The accumulated wait time for synch
STAT_SUS_GLCK...	=		The accumulated wait time for globa
STAT_SUS_OTHR...	=		The accumulated wait time for read
STAT_SUS_OT...	=		The accumulated wait time for write
AVG_STAT_CPU	=		Average accumulated CPU time.
AVG_STAT_ELAP	>	40	Average accumulated elapsed time used.
AVG_STAT_SUS...	=		Average accumulated wait time for a
AVG_STAT_SUS...	=		Average accumulated wait time for a
AVG_STAT_SUS...	=		Average accumulated wait time for a
AVG_STAT_SUS...	=		Average accumulated wait time for a
AVG_STAT_SUS...	=		Average accumulated wait time for a
AVG_STAT_SUS...	=		Average accumulated wait time for a
AVG_STAT_SUS...	=		Average accumulated wait time for a

< Back Next > Finish Cancel

Retrieve SQL stateme... table: (0%)

custom view Click V

SYSADM.LINEITEM WHO

User-defined view:
filters, ordering, attributes



Identifying Problem Query – Dynamic Statement Cache

Sort Rows

If you want to sort the queries, specify which columns are to be used to sort the result rows by moving columns from the list of available columns to the list of sorted columns list. Then select

Steps

1. Filter
2. Sort
3. Columns

Available columns

- STMT_ID
- STMT_TOKEN
- COLLID
- PROGRAM_NAME
- INV_DROPALT
- INV_REVOKE
- INV_LRU
- INV_RUNSTATS
- CACHED_TS
- USERS
- COPIES
- LINES
- PRIMAUTH
- CURSQLID
- BIND_QUALIFIER
- BIND_ISO
- BIND_CDATA
- BIND_DYNRL
- BIND_DEGRE
- BIND_SQLRL
- BIND_CHOLD
- STAT_TS

Sort columns

- STAT_CPU

Up
Down
 Ascending
 Descending

< Back Next > Finish Cancel

Retrieve SQL statements table: (0%)

User-defined view:
filters, **ordering**,
attributes



Identifying Problem Query – Dynamic Statement Cache

Steps

1. Filter
2. Sort
3. Columns

Available columns

- STMT_ID
- STMT_TOKEN
- COLLID
- PROGRAM_NAME
- INV_DROPALT
- INV_REVOKE
- INV_LRU
- INV_RUNSTATS
- USERS
- COPIES
- LINES
- PRIMAUTH
- CURSQLID
- BIND_QUALIFIER
- BIND_ISO
- BIND_CDATA
- BIND_DYNRL
- BIND_DEGRE
- BIND_SQLRL
- BIND_CHOLD
- STAT_GPAG
- STAT_SYNR

Display columns

- STAT_EXEC
- STAT_CPU
- STAT_ELAP
- STMT_TEXT

Back **Next >** **Finish** **Cancel**

1	0.008425281...	0.0234471...	DELETE FROM DSN_STATEMENT_CACHE_TABLE
0	0.001826585...	0.0018725...	SELECT STMT_ID, STMT_TOKEN, COLLID, PROGRAM_NAME, INV_DROPALT, INV_REVOKE, INV_LRU,
0	4.601565308...	0.0019182...	SELECT CURRENT SQLID FROM SYSIBM.SYSDUMMY1
0	3.659374853...	3.6578119...	SELECT 1 FROM SYSIBM.SYSDUMMY1 WHERE 0 = 1;
0	0.0	0.0	SELECT CURRENT TIMESTAMP AS TIMESTAMP FROM SYSIBM.SYSDUMMY1

User-defined view:
filters, ordering,
attributes



Identifying Problem Query – Dynamic Statement Cache

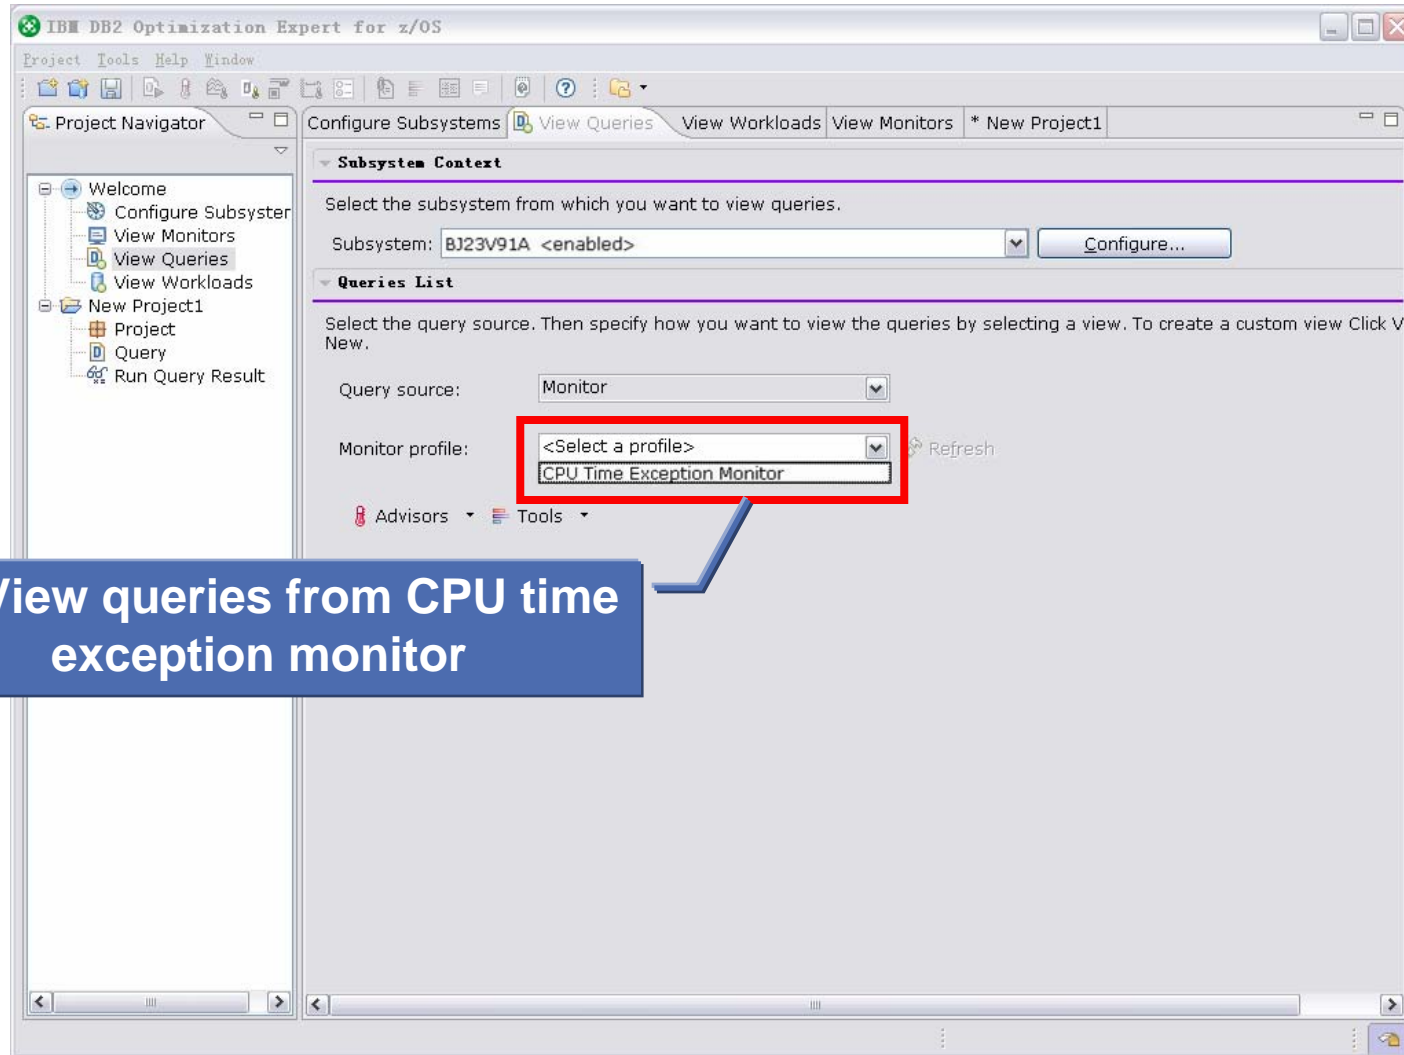
The screenshot shows the IBM DB2 Optimization Expert for z/OS interface. The 'View Queries' tab is active, displaying the 'View Queries List' section. The 'Query source' is set to 'Statement cache' and the 'View name' is 'ACCUM_CPU_DESC'. A table of results is shown below, with the first row highlighted in red. A blue callout box points to this row with the text 'Problem Query Identified!'.

STAT_EXEC	STAT_CPU	STAT_ELAP	STMT_TEXT
5	1.040755511285175	247.0208168179381	SELECT N_NAME, L_EXTENDEDPRIE AS REVENUE FROM CUSTOMER

Problem Query Identified!



Identifying Problem Query – Query Monitor



View queries from CPU time exception monitor



Identifying Problem Query – Query Monitor

IBM DB2 Optimization Expert for z/OS

Project Navigator

- Welcome
- Configure Subsystem
- View Monitors
- View Queries
- View Workloads
- New Project1
 - Project
 - Query
 - Run Query Result

Configure Subsystems | View Queries | View Workloads | View Monitors | * New Project1

Subsystem Context

Select the subsystem from which you want to view queries.

Subsystem: BJ23V91A <enabled>

Queries List

Select the query source. Then specify how you want to view the queries by selecting a view. To create a custom view Click V New.

Query source: Monitor

Monitor profile: CPU Time Exception Monitor

Advisors Tools

All of the rows are displayed. The number of rows is 7.

STMT_TEXT

```
SELECT * FROM SYSADM.LINEITEM L ,SYSADM.ORDER O ,SYSADM.SUPPLIER S WHERE L.L_RECEIPTDATE <= '1999-12-
SELECT DISTINCT O_ORDERKEY FROM SYSADM.LINEITEM , SYSADM.ORDER WHERE L_ORDERKEY = O_ORDERKEY AND
SELECT L_ORDERKEY, SUM(L_EXTENDEDPRICE) AS REVENUE , O_ORDERDATE, O_SHIPPRIORITY FROM SYSADM.CUSTO
SELECT L_ORDERKEY,L_SUPPKEY ,L_SHIPDATE,L_RETURNFLAG FROM SYSADM.LINEITEM WHERE L_SUPPKEY BETWEEN
SELECT L_SUPPKEY, COUNT(*) ,MIN(L_TAX) ,MAX(L_TAX) , SUM(L_EXTENDEDPRICE) , AVG(L_EXTENDEDPRICE) FROM SY
SELECT O_ORDERPRIORITY, COUNT(*) FROM SYSADM.ORDER , SYSADM.LINEITEM WHERE O_ORDERKEY = L_ORDERKE
SELECT S_SUPPKEY, S_NAME , SUM(L_EXTENDEDPRICE*(1-L_DISCOUNT)) AS REVENUE FROM SYSADM.ORDER , SYSAD
```

Problem Query Identified!



Ready for Problem Resolution – All Users

The screenshot shows the IBM DB2 Optimization Expert for z/OS interface. The 'View Queries' tab is active, displaying the 'Subsystem Context' and 'View Queries List' sections. The 'View Queries List' section includes a 'Query source' dropdown set to 'Statement cache', 'Enable Cache Trace' and 'Disable Cache Trace' buttons, and a 'View name' dropdown set to 'ACCUM_CPU_DESC'. Below these, there are 'Advisors...' and 'Tools...' dropdown menus. A red box highlights the 'Run All Advisors' menu option, which is expanded to show a list of advisor options: 'Run All Advisors', 'Run Statistics Advisor', 'Run Query Advisor', 'Run Access Path Advisor', 'Run Index Advisor', and 'Show Advisor Options'. A blue callout box points to this menu with the text: 'Diagnosis using Design Advisors: Stats Advisor, Index Advisor, Access Path Advisor, and Query Advisor'. The interface also shows a table with columns 'ST_ELAP' and 'SIMT_TEXT' and a row containing the value '0208168179381' and a SQL query snippet 'SELECT N_NAME, L_EXTENDEDP... AS REVENUE FROM CUSTOMER...'. The status bar at the bottom indicates 'Retrieve SQL stateme... table: (0%)'.



Ready for Problem Resolution – Expert Users

The screenshot displays the IBM DB2 Optimization Expert for z/OS application window. The interface includes a Project Navigator on the left, a main workspace with tabs for 'Configure Subsystem', 'View Queries', 'View Workloads', and 'View Monitor', and a status bar at the bottom. The 'View Queries' tab is active, showing the 'Subsystem Context' and 'View Queries List' sections. The 'View Queries List' section displays a table with columns 'STAT_EXEC' and 'STAT'. A context menu is open over the table, listing several diagnostic tools: 'Query Annotation', 'Access Plan Graph', 'Visual Plan Hint', 'Query Reports', and 'Gather Service Information'. A blue callout box points to this menu with the text: 'Diagnosis using Expert Tuning Tools: Annotation, report, index evaluator, explain, visual hint, service SQL'. The status bar at the bottom indicates 'Retrieve SQL statements... table: (0%)'.





Problem Resolution with Design Advisors



- Quick way to identify the cause of performance problem
- Recommended for all users
- Four Advisors
 - Statistics Advisor
 - Index Advisor
 - Query Advisor
 - Access Path Advisor



Tuning Problem Query with Design Advisors

The screenshot shows the IBM DB2 Optimization Expert for z/OS interface. The 'Source' section is active, with 'Text' selected as the query source. The 'Query text' section contains the following SQL query:

```
EXPLAIN opt...
SELECT N_NA...
WHERE C_CU...
AND C_NATIO...
AND R_NAME...
AND O_ORDERDATE < DATE('1994-01-01') + 1 YEAR;
```

A red box highlights the 'Advisors' menu, which includes the following options:

- Run All Advisors
- Run Statistics Advisor
- Run Query Advisor
- Run Access Path Advisor
- Run Index Advisor
- Show Advisor Options

A blue callout box with the text "Advisors make tuning recommendations" points to the 'Advisors' menu.



Tuning Problem Query with Stats Advisor

Recommendation Summary

Recommendations

Select a recommendation to view more details.

Number	Priority	Recommendation	Description
1	High	Run repairing RUNSTATS	Repair statistics problems for this query. Ga...
2	Low	Run complete RUNSTATS	Gather and recollect all of relevant statistic...

Selected Recommendations: No.1: Run repairing RUNSTATS

View more details for this recommendation, run the statement, or copy the statement for later use.

RUNSTATS Control Statements

```
RUNSTATS INDEX(SYSADM.PXR@RKNM)
SHRLEVEL CHANGE REPORT YES

RUNSTATS INDEX
(SYSADM.PXL@OKSDRFSKEPDC KEYCARD,
SYSADM.SXL@PKSKOKEPDSQN
KEYCARD)
SHRLEVEL CHANGE REPORT YES

RUNSTATS INDEX(SYSADM.PXS@SKNK)
SHRLEVEL CHANGE REPORT YES
```

Description

Repair statistics problems for this query. Gather missing statistics. Recollect conflicting statistics and potential obsolete statistics. Collect statistics for potential data skew and data correlation problems.

Actions

Details...
Run...
Copy
Save

RUNSTATS commands

Recommendation Explanation

Project Query **Statistics Advisor**

Run the Query Advisor



Tuning Problem Query with Stats Advisor

Detailed Recommendation Explanation

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Project Navigator

- Welcome
- Configure Subsystems
- View Monitors
- View Queries
- View Workloads
- New Project1
 - Project
 - Query
 - Statistics Advisor

Configure Subsystems | View Queries | View Workloads | View Monitors | * New Project1

This page explains the reasons for the selected RUNSTATS recommendation. To run the statement or copy it for later use, return to the Statistics Advisor page.

Table, index, column, and column group details

Statistics Advisor Detail Report
Analysis Begin Time: 2007-02-14 15:54:41.986
Analysis End Time: 2007-02-14 15:54:42.747

=====

TABLE SYSADM.REGION
Cardinality: 5.0
Collection Time: 0001-01-01 00:00:00.0
Statistics Status: OK

INDEXES:
SYSADM.PXR@RKNM (R_REGIONKEY,R_NAME)
First Key Cardinality: 5.0
Full Key Cardinality: 5.0
Data Repetition Factor: -1.0
Collection Time: 0001-01-01 00:00:00.0
Statistics Status: missing

Interesting Columns:
R_REGIONKEY
Cardinality: 5.0
Uniform Statistics Collection Time: 0001-01-01 00:00:00.0

Conflicts detail

Project | Query | Statistics Advisor | Statistics Advisor Details

Run the Query Advisor



Tuning Problem Query with Query Advisor

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The main window is titled "Query Advisor Recommendations List" and contains a table with the following data:

No	Severity	Confidence	Line Number	Description
1	Low	Medium	16	Consider rewriting the following predi

Below the table, the "SQL Text" section shows the following query:

```
AND SYSADM.ORDER.O_ORDERKEY = SYSADM.LINEITEM.L_ORDERKEY  
AND SYSADM.NATION.N_REGIONKEY = SYSADM.REGION.R_REGIONKEY  
AND SYSADM.CUSTOMER.C_NATIONKEY = SYSADM.SUPPLIER.S_NATIONKEY  
AND SYSADM.LINEITEM.L_SUPPKEY = SYSADM.SUPPLIER.S_SUPPKEY  
AND SYSADM.SUPPLIER.S_NATIONKEY = SYSADM.NATION.N_NATIONKEY  
AND SYSADM.CUSTOMER.C_NAME = 'IBM'  
AND YEAR( SYSADM.LINEITEM.L_SHIPDATE ) = 1994
```

The "Selected Recommendation:" section is divided into two columns:

- Description:** Consider rewriting the following predicate so that it is either an indexable or a stage-1 predicate: YEAR(SYSADM.LINEITEM.L_SHIPDATE) = 1994 so that it might filter out unnecessary rows earlier. Check the explanation for this warning for more details about possible impact and examples.
- Explanation:** The specified SQL statement might perform faster if you rewrite the stage 2 predicate as an indexable predicate or as a stage 1 predicate. Stage 1 predicates are better than stage 2 predicates because they disqualify rows earlier and reduce the amount of processing that DB2 needs to perform during later stages of evaluation. Also, because processing of stage 2 predicates can take many CPU cycles, these predicates are generally slower than stage 1 predicates. Indexable predicates, which are

Callouts in the image point to the "Recommendation Summary" (top right), "Recommendation Detail" (left), and "Recommendation Explanation" (right). A yellow circle highlights the "Query Advisor" tab in the bottom navigation bar.

Recommendation Summary

Recommendation Detail

Recommendation Explanation



Tuning Problem Query with Index Advisor

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The main window shows the 'Index Advisor Recommendations' section, which lists several index recommendations for a selected query. A red box highlights this section. A blue callout box labeled 'Index Recommendation' points to the table. Below the table, a 'Run Selected DDL Statements' dialog box is open, showing the SQL statements for creating the recommended indexes. A blue callout box labeled 'DDL for creating recommended index' points to this dialog. The 'Index Advisor' tab is selected in the bottom navigation bar, which is circled in yellow.

Index Recommendation

Feature Details	Object Name	Columns	Estimated Disk Space
<input checked="" type="checkbox"/> CUSTOMER			
<input checked="" type="checkbox"/> Index	CUSTOMER_VIRT...	C_NAME(ASC), C_NATIONKEY(...)	206.0234375 M
<input checked="" type="checkbox"/> REGION			
<input checked="" type="checkbox"/> Index	REGION_VIRT...	R_NAME(ASC), R_REGIONKEY(...)	0.0234375 M
<input checked="" type="checkbox"/> ORDER			
<input checked="" type="checkbox"/> Index	ORDER_VIRT...	O_CUSTKEY(ASC), O_ORDERD...	1024.0703125 M
<input checked="" type="checkbox"/> NATION			
<input checked="" type="checkbox"/> Index	NATION_VIRT...	N_REGIONKEY(ASC), N_NATIO...	0.0234375 M

DDL for creating recommended index

```
CREATE INDEX DB2OE.CUSTOMER_VIRT_IDX_1171438532131 ON
SYSADM.CUSTOMER ( C_NAME ASC, C_NATIONKEY ASC, C_CUSTKEY
ASC) FREEPAGE 0 PCTFREE 10;

CREATE INDEX DB2OE.REGION_VIRT_IDX_1171438532105 ON
SYSADM.REGION ( R_NAME ASC, R_REGIONKEY ASC) FREEPAGE 0
PCTFREE 10;

CREATE INDEX DB2OE.ORDER_VIRT_IDX_1171438563818 ON
SYSADM.ORDER ( O_CUSTKEY ASC, O_ORDERDATE ASC, O_ORDERKEY
ASC) FREEPAGE 0 PCTFREE 10;

CREATE INDEX DB2OE.NATION_VIRT_IDX_1171438588804 ON
SYSADM.NATION ( N_REGIONKEY ASC, N_NATIONKEY ASC, N_NAME
ASC) FREEPAGE 0 PCTFREE 10;
```



Tuning Problem Query with Access Path Advisor

Access Path Warning List

Access Path Warning List

The following list shows potential problems in the selected statement's access path. Select a warning to view a more detailed description. Make the necessary changes to avoid this warning in the future.

Severity	Query Block Number	Plan Number	Description
APA_HIGH_SEVERITY	1	2	The inner table SYSADM.OR...

Warning Description

Access Path Warning Details

Description	Explanation
The inner table SYSADM.ORDER in the nested loop join is accessed by a relational scan. When a large number of records are returned after the outer table is accessed, DB2 might be using an inefficient access path. Check the explanation for this warning for	When the DB2 optimizer chooses a nested loop join, DB2 first scans the outer table and then scans the inner table one time for each qualifying row in the outer table. The DB2 optimizer might choose to access the inner table by using a table space scan.

Corresponding record in PLAN_TABLE

PLAN_TABLE record

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

QBLOCKNO	PLANNO	MIXOPSEQ	METHOD	CREATOR	TNAME	CORRE...	ACCESST...	PROCESS...
1	2	0	1	SYSADM	ORDER		R	
1	1	0	0	SYSADM	LINEITEM		R	

Warning Explanation

Access Path Advisor





Problem Resolution with Tuning Tools



-
- Deep-dive into the root cause of the problem
 - Recommended for expert users
 - Expert Tuning Tools
 - Query Annotation
 - Query Report
 - Visual Explain
 - Visual Plan Hint
 - Service SQL



Tuning Problem Query with Tuning Tools

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The window title is "IBM DB2 Optimization Expert for z/OS". The interface includes a "Project Navigator" on the left, a "Tools" menu, and a main workspace. A blue callout bubble with the text "Expert tools can be invoked in the same Query Project" is overlaid on the "Source" section. The "Query text" section shows a list of tools: "Query Annotation", "Access Plan Graph", "Visual Plan Hint", "Query Reports", and "Gather Service Information". The "Query Annotation" tool is highlighted with a red box. Below the tools list, a SQL query is visible, starting with "SELECT N_NAME, L_EXTENDE...".

Expert tools can be invoked in the same Query Project

Tools available in the Query text section:

- Query Annotation
- Access Plan Graph
- Visual Plan Hint
- Query Reports
- Gather Service Information

```
SELECT N_NAME, L_EXTENDE...  
WHERE C_CUSTKEY = O_CU...  
AND C_NATIONKEY = S_NATI...  
AND R_NAME = 'ASIA' AND C_NAME = IBM AND YEAR(L_SHIPDATE) = 1994 AND O_ORDERDATE >= DATE('1994-01-01')  
AND O_ORDERDATE < DATE('1994-01-01') + 1 YEAR;
```



Understanding Query with Query Annotation

Original and transformed Query

Formatted, reorganized query text

Annotations (catalog stats, cost estimation)

Formatted Query

```

SELECT SYSADM.NATION.N_NAME
, SYSADM.LINEITEM.L_EXTENDEDPRICE AS REVENUE
FROM SYSADM.REGION
, SYSADM.CUSTOMER
, SYSADM.NATION
, SYSADM.SUPPLIER
, SYSADM.ORDER
, SYSADM.LINEITEM
WHERE ( SYSADM.CUSTOMER.C_CUSTKEY = SYSADM.ORDER.O_CUSTKEY
AND SYSADM.ORDER.O_ORDERKEY = SYSADM.LINEITEM.L_ORDERKEY
AND SYSADM.NATION.N_REGIONKEY = SYSADM.REGION.R_REGIONKEY
AND SYSADM.CUSTOMER.C_NATIONKEY = SYSADM.SUPPLIER.S_NATIONKEY
AND SYSADM.LINEITEM.L_SUPPKEY = SYSADM.SUPPLIER.S_SUPPKEY
AND SYSADM.SUPPLIER.S_NATIONKEY = SYSADM.NATION.N_NATIONKEY
AND SYSADM.CUSTOMER.C_NAME = 'IBM'
AND YEAR( SYSADM.LINEITEM.L_SHIPDATE ) = 1994
AND SYSADM.ORDER.O_ORDERDATE < ( DATE( '1994-01-01' ) + 1 YEAR
AND SYSADM.ORDER.O_ORDERDATE >= DATE( '1994-01-01' )
AND SYSADM.REGION.R_NAME = 'ASIA'
)
    
```

Annotation

CARDF=5	QUALIFIED_ROWS=0.99999994	NPAG
CARDF=4,500,000	QUALIFIED_ROWS=9.999995	
CARDF=25	QUALIFIED_ROWS=25.0	NPAG
CARDF=300,000	QUALIFIED_ROWS=300,000.0	
CARDF=45,000,000	QUALIFIED_ROWS=6,823,580.	
CARDF=179,998,372	QUALIFIED_ROWS=7,199,935	
COLCARDF=4,500,000/3,000,000	MAX_FREQ=/	FF=
COLCARDF=45,000,000/45,000,000	MAX_FREQ=/	FI
COLCARDF=5/5	MAX_FREQ=/	FF=0.1995
COLCARDF=25/25	MAX_FREQ=/	FF=0.035
COLCARDF=303,104/300,000	MAX_FREQ=/	FF=:
COLCARDF=25/25	MAX_FREQ=/	FF=0.035
COLCARDF=450,000	MAX_FREQ=	FF=2.22
COLCARDF=450,000	MAX_FREQ=	FF=2.222220650292E-6
COLCARDF=2,304	MAX_FREQ=	FF=0.455
COLCARDF=2,304	MAX_FREQ=	FF=0.695
COLCARDF=5	MAX_FREQ=	FF=0.19995



Understanding Query with Query Annotation

Query Annotation

Indicate which annotation to display and customize your view. Selecting a row will highlight all of the relevant rows for this table. Selecting a join predicate will highlight all of the join predicate rows in both of the joined tables. Click Reset Text to return to the original text view.

Annotation to display: All

Expand All Collapse All Customize Save

Formatted Query	Annotation
FROM SYSADM.REGION	CARDF=5 QUALIFIED_ROWS=0.99999994 NF
, SYSADM.CUSTOMER	CARDF=4,500,000 QUALIFIED_ROWS=9.999995
, SYSADM.NATION	CARDF=25 QUALIFIED_ROWS=25.0 NPAG
, SYSADM.SUPPLIER	CARDF=300,000 QUALIFIED_ROWS=300,000.0
, SYSADM.ORDER	CARDF=45,000,000 QUALIFIED_ROWS=6,823,580.
, SYSADM.LINEITEM	CARDF=179,998,372 QUALIFIED_ROWS=7,199,935
WHERE (SYSADM.CUSTOMER.C_CUSTKEY = SYSADM.ORDER.O_CUSTKEY	COLCARDF=4,500,000/3,000,000 MAX_FREQ=/ FF=
AND SYSADM.ORDER.O_ORDERKEY = SYSADM.LINEITEM.L_ORDERKEY	COLCARDF=45,000,000/45,000,000 MAX_FREQ=/ FI
AND SYSADM.NATION.N_REGIONKEY = SYSADM.REGION.R_REGIONKEY	COLCARDF=5/5 MAX_FREQ=/ FF=0.1995
AND SYSADM.CUSTOMER.C_NATIONKEY = SYSADM.SUPPLIER.S_NATIONKEY	COLCARDF=25/25 MAX_FREQ=/ FF=0.035
AND SYSADM.LINEITEM.L_SUPPKEY = SYSADM.SUPPLIER.S_SUPPKEY	COLCARDF=303,104/300,000 MAX_FREQ=/ FF=:
AND SYSADM.SUPPLIER.S_NATIONKEY = SYSADM.NATION.N_NATIONKEY	COLCARDF=25/25 MAX_FREQ=/ FF=0.035
AND SYSADM.CUSTOMER.C_NAME = 'IBM'	COLCARDF=450,000 MAX_FREQ= FF=2.22E-6
AND YEAR(SYSADM.LINEITEM.L_SHIPDATE) = 1994	COLCARDF=450,000 MAX_FREQ= FF=2.222220650292E-6
AND SYSADM.ORDER.O_ORDERDATE < (DATE('1994-01-01') + 1 YEARS	COLCARDF=2,304 MAX_FREQ= FF=0.455
AND SYSADM.ORDER.O_ORDERDATE >= DATE('1994-01-01')	COLCARDF=2,304 MAX_FREQ= FF=0.695
AND SYSADM.REGION.N_NAME = 'ASIA'	COLCARDF=5 MAX_FREQ= FF=0.19995

Two tables with small number of qualified rows. Either one could be the leading table.

The most selective local predicate. (FF=2.22E-6)



Query Report – Table Report

Legend for column names that have been truncated	
PARTS	-- Number of partitions in tablespace
QUALROWS	-- Optimizer's estimate for how many rows qualify if this table were the outer table
CLU	-- Is the index the clustering index?
UR	-- Unique rule
CR	-- Clusterratiof

TABLE_SPACE	NACTIVEF	PARTS	SEGSIZE	PG_SIZE						
DB4REGN.TSREGION	900.0	5	0	4						
		TABLE	CARDF	NPAGESF	TABNO	QUALROWS				
		SYSADM.REGION	5	0	6	0.99999994				
		INDEX	CLU	UR	NLEAF	NLEVEL	CR	KEYCOLNAME	COLCARDF	MCARDF
		SYSADM.PXR@RKNM	Y	U	5	2	1.0	R_REGIONKEY	5	5
								R_NAME	5	5
		COLUMN_GROUP	MCARDF							
DB4NATN.TSNATION	4500.0	25	0	4						
		TABLE	CARDF	NPAGESF	TABNO	QUALROWS				
		SYSADM.NATION	25	0	5	25.0				
		INDEX								
		SYSADM.PXN@NKNM								
		COLUMN_GROUP								
DB4CUST.TSCUST	192569.0	60	0	4						
		TABLE	CARDF	NPAGESF	TABNO	QUALROWS				
		SYSADM.CUSTOMER	4500000	0	1	9.999995				
		INDEX	CLU	UR	NLEAF	NLEVEL	CR	KEYCOLNAME	COLCARDF	MCARDF
		SYSADM.PXC@CKNKMS	Y	U	29040	3	1.0	C_CUSTKEY	4500000	4500000
								C_NATIONKEY	25	-1
								C_MKTSEGMENT	5	4500000
		SYSADM.UXC@NKCK	N	U	17858	3	0.999	C_NATIONKEY	25	25
								C_CUSTKEY	4500000	4500000
		COLUMN_GROUP	MCARDF							

Two indexes on the table 'CUSTOMER', but neither contains the column key 'C_NAME'.



Query Report – Predicate Report

Legend for column names that have been truncated											
FF	--Filter factor										
BT	--Whether this predicate is a boolean term predicate										
S1	--Whether the predicate is a stage 1 predicate										
JN	--Whether this predicate is a join predicate										
AJ	--Whether this predicate is an aggregation predicate										
PTC	--Whether the predicate belongs to a query selection when the set of predicates										
MARKE	--Whether this predicate is a marked predicate										
PREDNO	--The predicate number										

TABLE	TABNO	COLNAME	COLNO	COLCARD	HIGH2KEY	LOW2KEY	MAX FREQ	TYPE	OTH TABLE	OTH TABNO	OTH COLNAME	
SYSADM.CUSTOMER	1	C_CUSTKEY	1	4500000	8044AA1F40404040	8000000240404040		EQUAL	SYSADM.ORDER	2	O_CUSTKEY	
		C_NATIONKEY	4	25	8000001740404040	8000000140404040		EQUAL	SYSADM.SUPPLIER	4	S_NATIONKEY	
									EQUAL	SYSADM.NATION	5	N_NATIONKEY
		C_NAME	2	450000	Customer	Customer		EQUAL		0	VALUE	
SYSADM.ORDER	2	O_CUSTKEY	2	3000000	8044AA1E40404040	8000000240404040		EQUAL	SYSADM.CUSTOMER	1	C_CUSTKEY	
		O_ORDERKEY	1	45000000	8ABA94E740404040	8000000240404040		EQUAL	SYSADM.LINEITEM	3	L_ORDERKEY	
		O_ORDERDATE	5	2304	1998-08-02	1992-01-01		RANGE		0	VALUE	
									RANGE		0	VALUE
SYSADM.LINEITEM	3	L_ORDERKEY	1	45000000	8ABA94E740404040	8000000240404040		EQUAL	SYSADM.ORDER	2	O_ORDERKEY	
SYSADM.SUPPLIER	4	S_SUPPLIERKEY	4	45000000	8ABA94E740404040	8000000140404040		EQUAL	SYSADM.CUSTOMER	1	C_NATIONKEY	
SYSADM.NATION	5	N_NATIONKEY	5	45000000	8ABA94E740404040	8000000140404040		EQUAL	SYSADM.SUPPLIER	4	S_NATIONKEY	
SYSADM.REGION	6	R_REGIONKEY	1	5	8000000340404040	8000000140404040		EQUAL	SYSADM.NATION	5	N_REGIONKEY	
		R_NAME	2	5	MIDDLE E	AMERICA		EQUAL		0	VALUE	

Three EQUAL join predicates and one EQUAL local predicate on the table 'CUSTOMER'

One EQUAL join predicate and one EQUAL local predicate on the table 'REGION'



Query Report – Index Report

TABLE	CORR_NAME				
SYSADM.REGION					
	INDEX	INDEX_ONLY	ONE_FETCH	EQUAL_UNIQUE	GB_OB_DISTINCT
	SYSADM.PXR@RKNM	Y	N	N	N
	KEYCOL	ORDER	COLUMN_CARD	MULTI_COL_CARD	PRED
	R_REGIONKEY	ASCENDING	5.0	5.0	SYSADM.NATION_N_REGIONKEY=SYSADM.REGION_R_REGIONKEY (FF:0.19999998807907104)
	R_NAME	ASCENDING	5.0	5.0	SYSADM.REGION_R_NAME='ASIA' (FF:0.19999998807907104)

TABLE	CORR_NAME				
SYSADM.NATION					
	INDEX				
	SYSADM.PXN@NKNMRK				
	KEYCOL	ORDER	COLUMN_CARD	MULTI_COL_CARD	PRED
	N_NATIONKEY	ASCENDING	25.0	25.0	SYSADM.SUPPLIER_S_NATIONKEY=SYSADM.NATION_N_NATIONKEY
					N_NATIONKEY
	N_REGIONKEY	ASCENDING	5.0	25.0	SYSADM.NATION_N_REGIONKEY=SYSADM.REGION_R_REGIONKEY (FF:0.19999998807907104)

This index supports index-only scan on the table 'REGION'.

These two indexes supports index + fetch scan on the table 'CUSTOMER'.

QB: PLAN:3

TABLE	CORR_NAME				
SYSADM.CUSTOMER					
	INDEX	INDEX_ONLY	ONE_FETCH	EQUAL_UNIQUE	GB_OB_DISTINCT
	SYSADM.PXC@CKNKMS	N	N	N	N
	KEYCOL	ORDER	COLUMN_CARD	MULTI_COL_CARD	PRED
	C_CUSTKEY	ASCENDING	4500000.0	4500000.0	SYSADM.CUSTOMER_C_CUSTKEY=SYSADM.ORDER_O_CUSTKEY (FF:2.2222218376555247E-7)
	C_NATIONKEY	ASCENDING	25.0	-1.0	SYSADM.CUSTOMER_C_NATIONKEY=SYSADM.SUPPLIER_S_NATIONKEY (FF:0.03999999910593033)
					SYSADM.CUSTOMER_C_NATIONKEY=SYSADM.NATION_N_NATIONKEY (FF:0.03999999910593033)
	C_MKTSEGMENT	ASCENDING	5.0	4500000.0	
	INDEX	INDEX_ONLY	ONE_FETCH	EQUAL_UNIQUE	GB_OB_DISTINCT
	SYSADM.UXC@NKCK	N	N	N	N
	KEYCOL	ORDER	COLUMN_CARD	MULTI_COL_CARD	PRED



Understanding the Access Path with Visual Explain

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. A blue callout box labeled "Access Plan Graph" points to the main diagram area. A red box highlights the "Node Descriptor" window on the left and the "Access Plan Graph" area on the right.

Node Descriptor Window:

Query: QUERY
 @ query

Show attribute explanation Views: st_es

Name	Value
Type	SELECT
CPU Cost (ms)	6181
CPU Cost (su)	16055
Cost Category	A
Reason	
Timestamp	2007-02-15 19:32:36.18

Attribute explanation:

Buttons: Save as ... Print ...

Access Plan Graph:

The graph shows a hierarchical structure of query operations. At the top is a green circle labeled "QUERY". Below it is a blue circle "QOB1 17", followed by a blue octagon "NLJOIN 17.8252". The graph branches into several paths, including "NLJOIN 18.0096", "NLJOIN 15.00", "WFSCAN 5.8", "WORK FILE 5.8", "SORT 5.8", "NLJOIN 5.8", "XONLYSCAN 2.3647", "XO@CKOKODSP 4500000", "XONLYSCAN 4", "PX@OKSDRFSKEPD 179998372", "NLJOIN 5", "SCAN", "XONLYSCAN 5", "XSCAN 180000", and "CUSTOMER 4500000".

Hierarchical
Descriptor
Structure

Detailed
Descriptor
Attributes



Implementing Emergency Solution with Visual Plan Hint

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Configure Subsystems View Queries View Workloads View Monitors * New Project1 X

Query Blocks: Query block 1 Graph Report Validate Hints Deploy Hints Zoom In Zoom Out

Show Local Predicates Show Join Predicates Create Join Node Default Join Sequence Delete Selected Nodes

Clear All

Default Join Sequence

Criteria Edit Hint Criteria Delete Hint Criteria

PROJECT CREATOR TNAME CORRELATION_NAME ACCESTYPE ACCESSCREATOR ACCESSNAME PREFETCH PAGE_RANGE METHOD SORTN_JOIN

Project Query Query Annotation Report Access Plan Graph Visual Plan Hint Annotate Query



Implementing Emergency Solution with Visual Plan Hint

Hint Customization Rule

CREATOR	SYSADM	
TNAME	CUSTOMER	
CORRELATION_NAME	NULL	
ACCESSTYPE	INDEX	NULL
ACCESSCREATOR	DB2OE	NULL
ACCESSNAME	CUSTOMER_VIRT_IDX_11716082	NULL CUSTOMER_VIRT_IDX_11716082 PXC@CKNKMS UXC@NKCK
PREFETCH		
PAGE_RANGE		
METHOD	NULL	
SORTN_JOIN	N	NULL
SORTC_JOIN	N	NULL
PARALLELISM_MODE	NULL	NULL
ACCESS_DEGREE	NULL	NULL
JOIN_DEGREE	NULL	NULL
PRIMARY_ACCESSTYPE		NULL

table

OK Cancel

Project Query Query Annotation Report Access Plan Graph Visual Plan Hint

Annotate Query

Customize the access path using hint



Automatic Collection of Relevant Information for PMR Analysis

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Service SQL

Generate Options

Specify options for generating a report about the query. You can also specify an FTP server to send the report to IBM Software Support.

Options

Convert to version: No conversion

Generate CREATE statements Generate statistics Parallelism Edited

PMR or ETR Number (xxxxx,yyy,zzz): , , Version:

C:/Program Files/IBM/OE Browse...

Generate Report Send Report to IBM

Include analysis results

FTP Server Settings

Server name: testcase.software.ibm.com Port: 21

User: anonymous Password: *****

Directory: /toibm/im Proxy Settings...

Upload file list: Add... Remove

< Back Next > Finish Cancel

Send Service Information

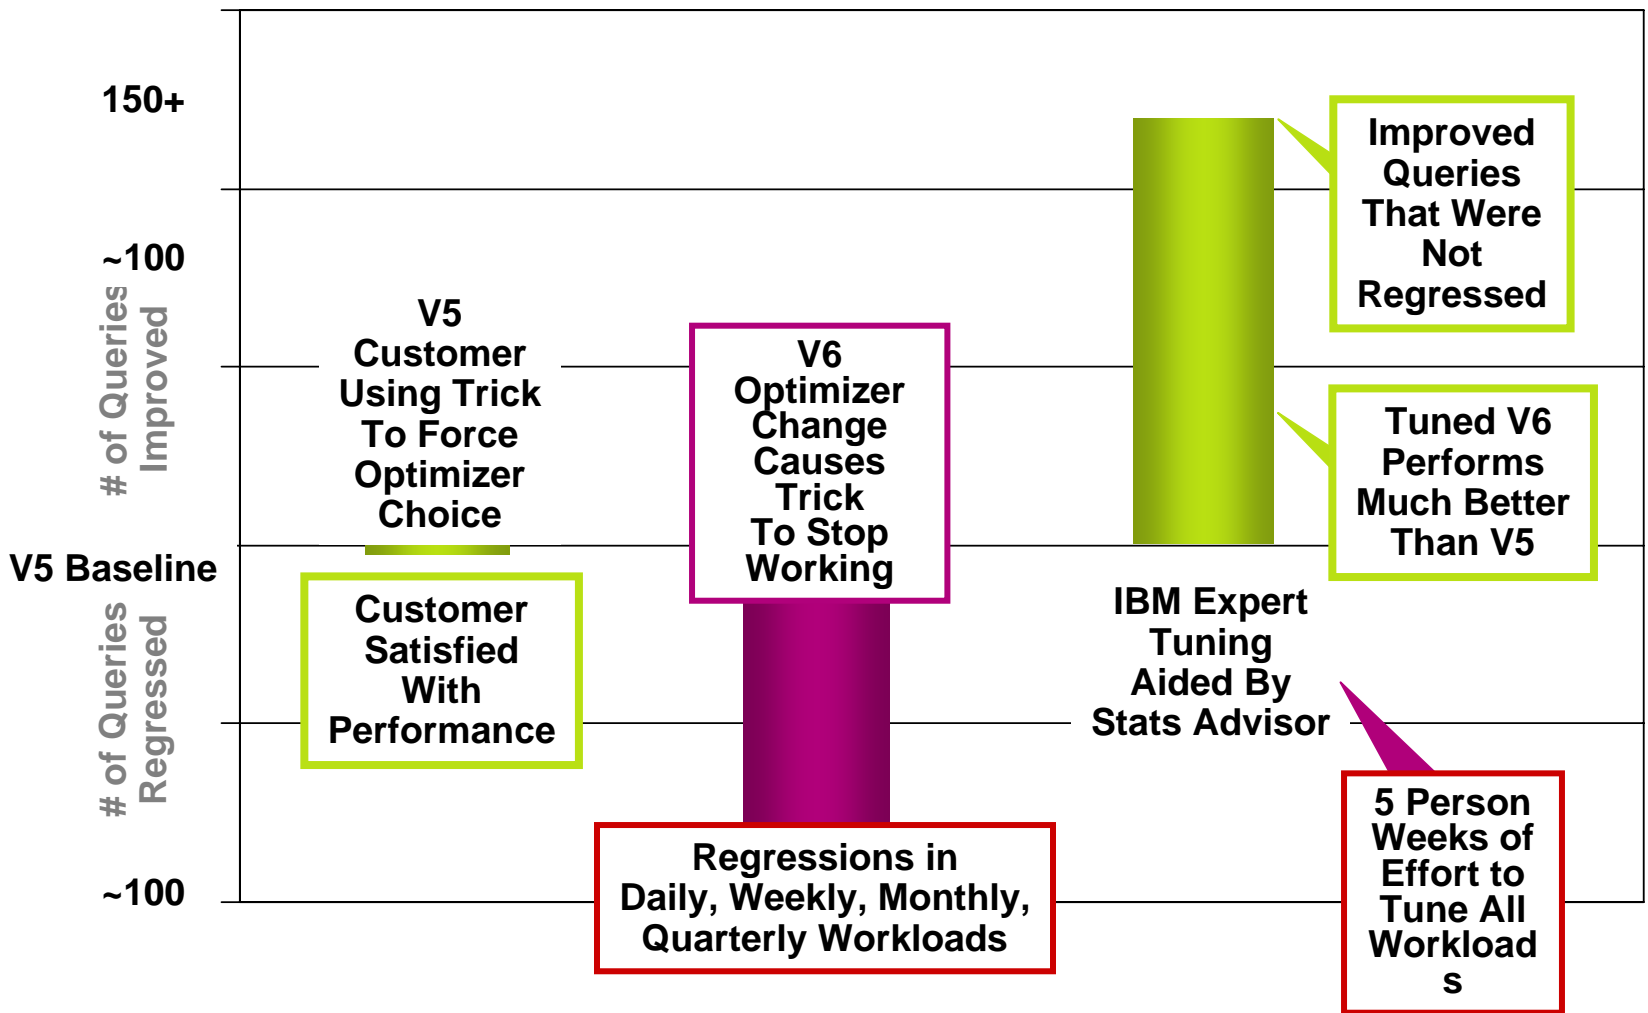




Capturing Application Workloads



Scenario: Version to version migration failure



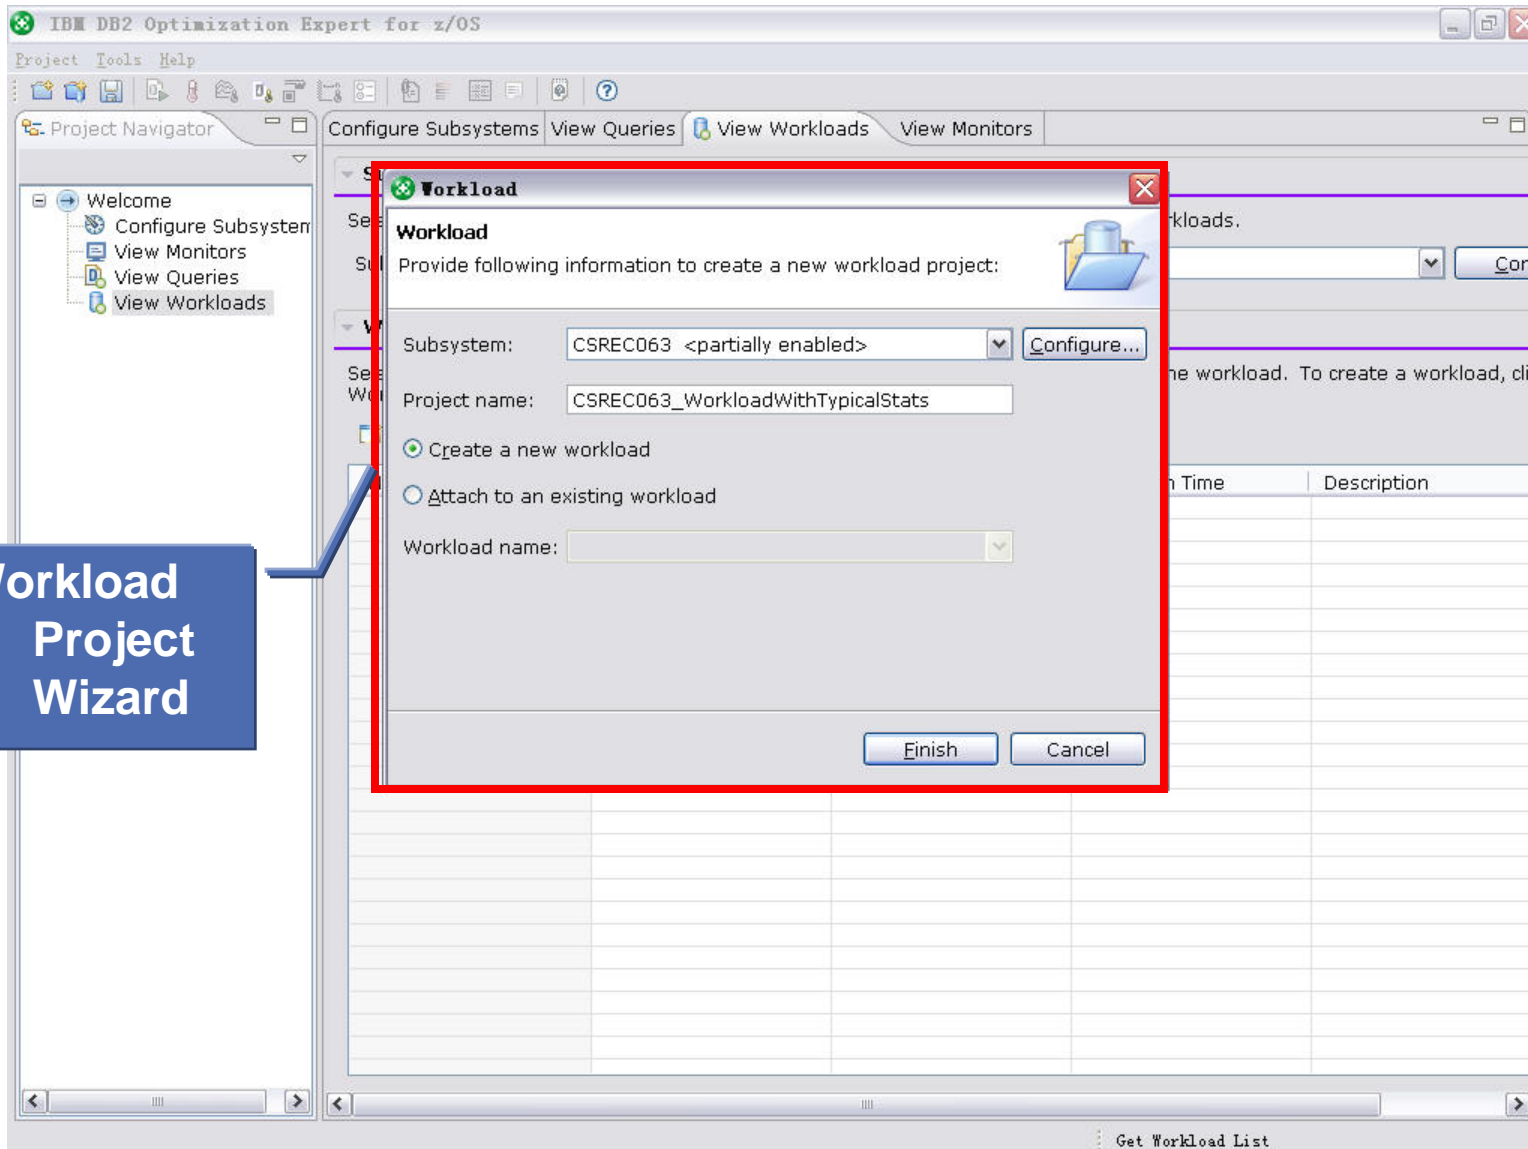
Scenario: Version to version migration failure

Post-mortem Analysis

- Customer satisfied with V5 performance, but it was actually sub-optimal
- Complexity of problem tested the limits of expert analysis
 - ✓ 30 table joins nested behind several layers of views
 - ✓ Column names changed in views
 - ✓ Days to analyze a single query
 - ✓ 100s of queries in multiple workloads
- Customer costs
 - ✓ 6-8 month delay in migration
 - ✓ Person weeks of effort to resolve problems
 - ✓ Inability to take advantage of V6 enhancements



Capturing Workload Queries



**Workload
Project
Wizard**



Capturing Workload Queries – General Info.

Workload Wizard

Workload
Provide the following information to create a workload.

Steps

1. Workload
2. Source
3. Filter
4. Capture

Workload name: WorkloadWithTypicalStats

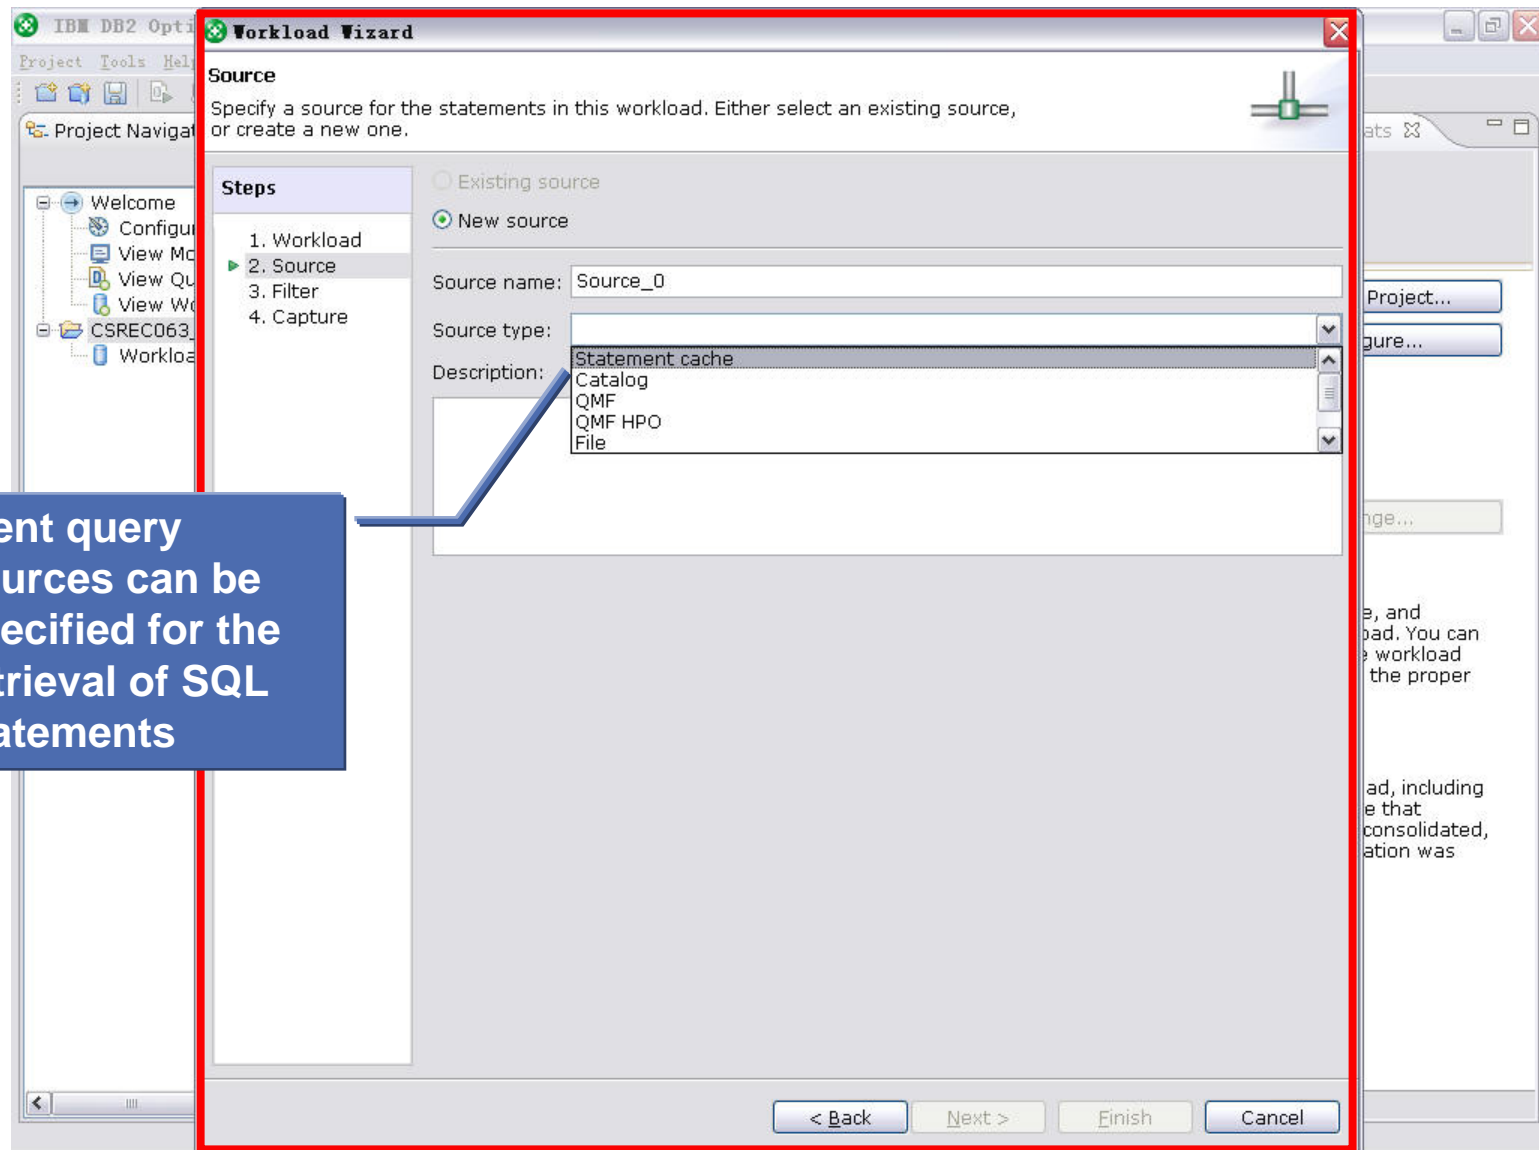
Owner: SYSADM

Description:

< Back Next > Finish Cancel



Capturing Workload Queries – Query Source



Different query sources can be specified for the retrieval of SQL statements



Capturing Workload Queries – Selection Criteria

Source Filter

Use the Operator and Value columns to define filtering criteria. Only those statements that satisfy the criteria are included in this workload.

Steps

1. Workload
2. Source
3. Filter
4. Capture

Column Name	Operator	Value	Comment
PRIMAUTH	=	ADMFO01	The primary authorization ID that did the in
CURSQLID	=		The CURRENT SQLID that did the initial PREF
SCHEMA	=		The value of the CURRENT SCHEMA special r
BIND_QUALIFIER	=		The BIND qualifier. For unqualified table nar
BIND_ISO	=		The value of the ISOLATION BIND option tha
BIND_CDATA	=		The value of the CURRENTDATA BIND option
BIND_DYNRL	=		The value of the DYNAMICRULES BIND optio
BIND_DEGREE	=		The value of the CURRENT DEGREE special r
BIND_SQLRL	=		The value of the CURRENT RULES special re
BIND_CHOLD	=		The value of the WITH HOLD attribute of the
STAT_EXEC	=		The number of times this statement has be
STAT_GPAG	=		The number of getpage operations that are
STAT_SYNR	=		The number of synchronous buffer reads th
STAT_WRIT	=		The number of buffer write operations that
STAT_EROW	=		The number of rows that are examined for t
STAT_PROW	=		The number of rows that are processed for
STAT_SORT	=		The number of sorts that are performed for
STAT_INDX	=		The number of index scans that are perform
STAT_RSCN	=		The number of tablespace scans that are pe
STAT_PGRP	=		The number of parallel groups that are crea
STAT_RIDLIMT	=		The number of times a RID list was not use
STAT_RIDSTOR	=		The number of times a RID list was not use
AVG_STAT_GPAG	=		The average number of getpage operations
AVG_STAT_SYNR	=		The average number of synchronous buffer
AVG_STAT_WRIT	=		The average number of buffer write operati
AVG_STAT_ER...	=		The average number of rows that are exam
AVG_STAT_PR...	=		The average number of rows that are proce
AVG_STAT_SORT	=		The average number of sorts.
AVG_STAT_INDX	=		The average number of index scans.
AVG_STAT_RSCN	=		The average number of tablespace scans.
AVG_STAT_PGRP	=		The average number of parallel groups that

Selection criteria

< Back Next > Finish Cancel



Capturing Workload Queries – When & How

Workload Wizard

Capture Type
Specify any related options. To customize the capture profile, select the Customize profile to choose from more options.

Steps	Capture profile	Description
1. Workload	Immediately	
2. Source	One Time	
3. Filter	Time Period	
4. Capture	Periodic Sampling	
	Customize	

Profile Details

< Back Next > Finish Cancel

Capture the workload now or later, once, multiple times or periodically



Capturing Workload Queries – Capturing

Capturing

IBM DB2 Optimization Expert for z/OS

Project Navigator

- Welcome
- Configure Subsystem
- View Monitors
- View Queries
- View Workloads
- CSRECO63_WorkloadWi
- Workload

Configure Subsystems | View Queries | View Workloads [BUSY] | View Monitors | CSRECO63_WorkloadWithTypic...

Workload Project

To tune a workload, specify the workload name and choose one of the following actions.

Project: CSRECO63_WorkloadWithTypicalStats Rename Project...

Subsystem: CSRECO63 <partially enabled> Configure...

Define Workload...

Define Workload... Elapsed time: 15.2 seconds

Define capture task for the workload

Run in Background Cancel Details >>

Run Advisors

Get recommendations for workloads that could improve workload performance. Schedule workload analysis for a later time.

History

View the history of this workload, including when it was created, each time that statements were captured or consolidated, and each time EXPLAIN information was gathered.

Schedule Tasks

Schedule when to capture statements, consolidate statements, and gather EXPLAIN information.

Workload

Define Workload ...



Capturing Workload Queries – Completion

Queries captured!

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Project Navigator

- Welcome
- Configure Subsystem
- View Monitors
- View Queries
- View Workloads
- CSRECO63_WorkloadWi
 - Workload

Configure Subsystems View Queries View Workloads View Monitors CSRECO63_WorkloadWithTypicalStats

Workload Project

To tune a workload, specify the workload name and choose one of the following actions.

Project: CSRECO63_WorkloadWithTypicalStats [Rename Project...](#)

Subsystem: CSRECO63 <partially enabled> [Configure...](#)

Workload Name: WorkloadWithTypicalStats

Workload Owner: SYSADM

Summary Status: CAPTURED [Change...](#)

Description:

Workload Statements
Capture workload statements, recommendations from the advisors, and use tools to tune an individual query.

Users
Grant or revoke owner, update, and read-only access to the workload. You can create authorization IDs in the workload control center only if you have the proper authority.

Run Advisors
Get recommendations for workloads that could improve workload performance. Schedule workload analysis for a later time.

History
View the history of this workload, including when it was created, each time that statements were captured or consolidated, and each time EXPLAIN information was gathered.

Schedule Tasks
Schedule when to capture statements, consolidate statements, and gather EXPLAIN information.

Workload

Define Workload ...



Capturing Workload Queries – Browsing

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The main window shows a list of workload statements under the heading "Workload Statements". A red box highlights a table of captured statements. The table has columns for Source, Accumulated Elapsed Time, Average Elapsed Time, Accumulated CPU Time, Average CPU Time, and Statement Text. The statements are SQL queries, including SELECT and INSERT statements.

Workload Statements

Immediately capture statements or multiple sources to this workload, launch workload advisors, use tools to tune selected queries from the workload, or schedule tasks for capture, consolidation, and analysis.

Capture Workload Tools Schedule Remove Query Tools Refresh

All of the rows are displayed. The number of rows is 10.

Source	Accumulated El...	Average Elapse...	Accumulated C...	Average CPU Ti...	Statement Text
CHE	0.001967156	0.001967156	0.0019474978	0.0019474978	SELECT C_NAME,C_ADDRESS,N...
CHE	0.005028826	0.005028826	0.004934781	0.004934781	SELECT O_ORDERKEY,O_CUSTKE...
CHE	0.0029194686	0.0029194686	0.0028823551	0.0028823551	SELECT S_SUPPKEY,S_NAME, SU...
CHE	0.0029387153	0.0029387153	0.0018932667	0.0018932667	SELECT C_NAME,C_ADDRESS,N...
CHE	0.0047716247	0.0047716247	0.004721173	0.004721173	SELECT O_ORDERKEY,O_CUSTKE...
CHE	0.0034057186	0.0034057186	0.0033768355	0.0033768355	SELECT S_SUPPKEY,S_NAME, SU...
CHE	0.0020138593	0.0020138593	0.0019925022	0.0019925022	SELECT C_NAME,C_ADDRESS,N...
CHE	0.004906592	0.004906592	0.0046221213	0.0046221213	SELECT O_ORDERKEY,O_CUSTKE...
CHE	0.0027890936	0.0027890936	0.0025451551	0.0025451551	SELECT S_SUPPKEY,S_NAME, SU...
CHE	0.004266531	0.004266531	0.0040141614	0.0040141614	INSERT INTO REGION (SELECT N...

Workload Statements

Captured Statements



Capture Workload Queries with Monitor

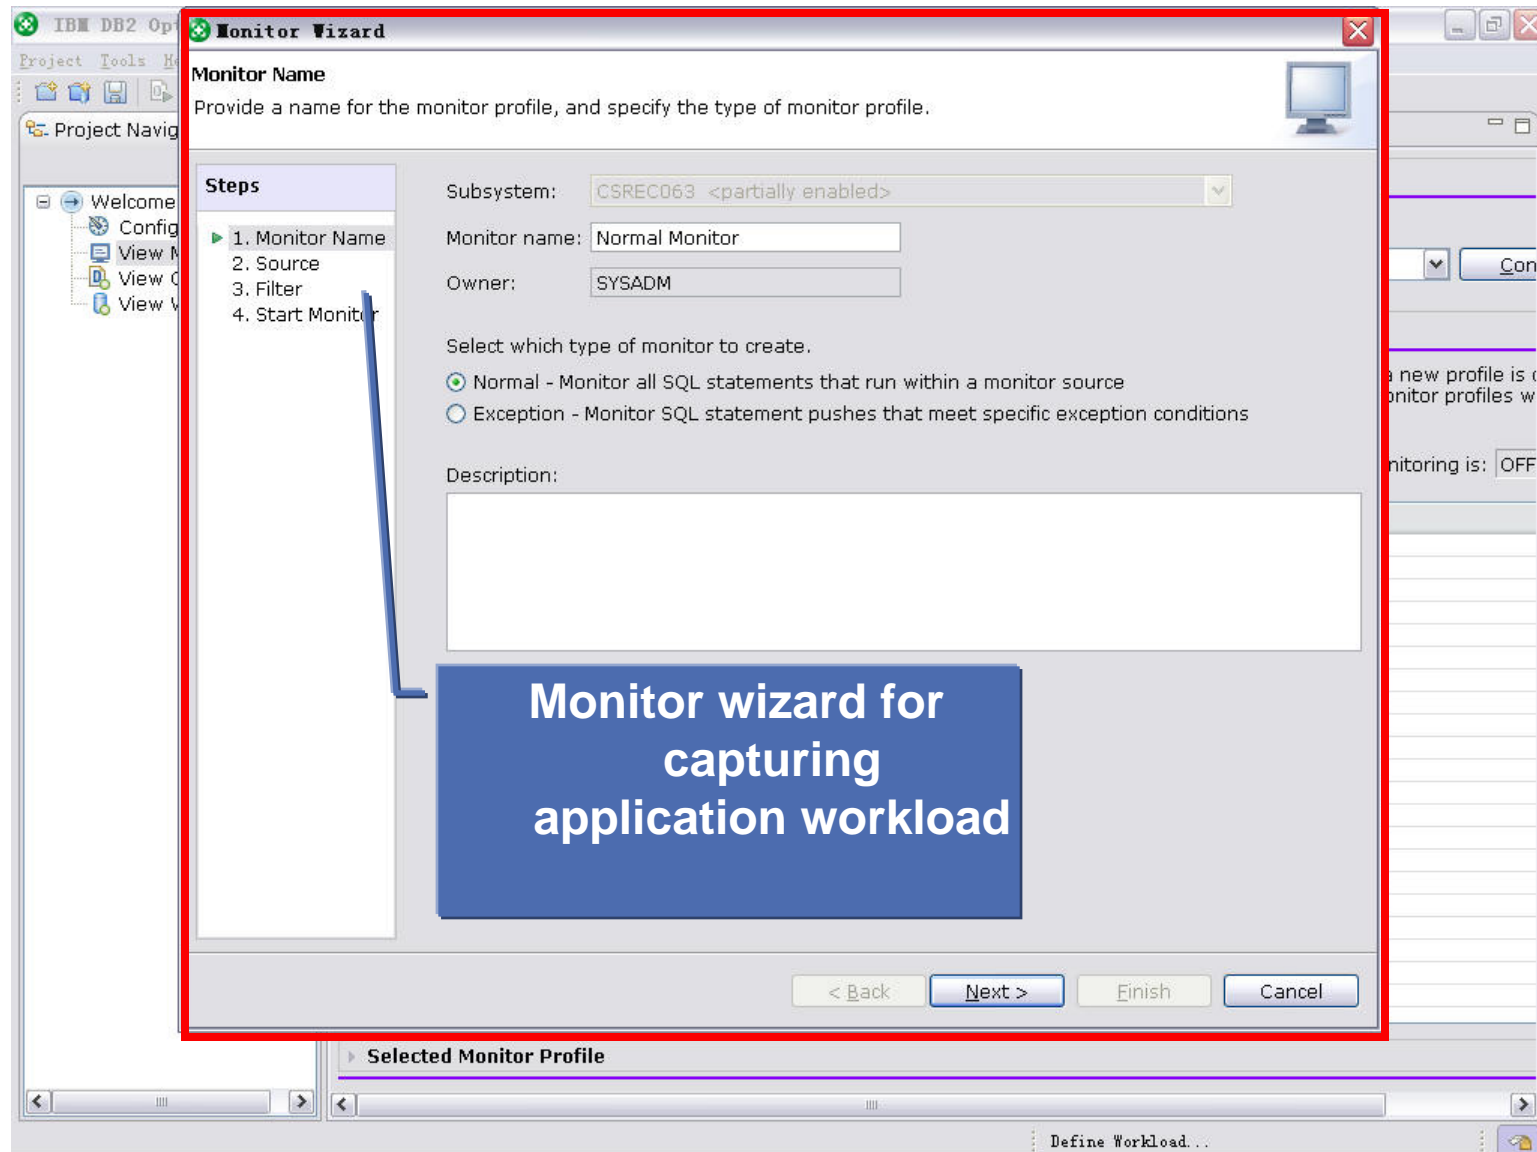
The screenshot shows the IBM DB2 Optimization Expert for z/OS application window. The title bar reads "IBM DB2 Optimization Expert for z/OS". Below the title bar is a menu bar with "Project", "Tools", and "Help". A "Welcome" tab is active, showing a banner with the product name and the IBM logo. The main content area is titled "WELCOME" and contains a paragraph: "Welcome! To get started with the DB2 Optimization Expert (OE), you must first configure a connection to a DB2 for z/OS subsystem. Then you can create a new project to tune a problem query or an entire query workload." Below this are six task cards arranged in a 3x2 grid:

- Configure DB2 Subsystems**: Connect to DB2 subsystems, enable OE, and grant EXPLAIN authorizations.
- View Workloads**: View the status of workloads, open existing workloads, and archive the workloads on a subsystem.
- View Query Activity**: View and sort dynamic and static queries to find potential problems.
- Tune a Workload**: Use OE advisors and advanced tools to capture and analyze and improve the performance of query workloads.
- Tune a Single Query**: Use OE advisors and advanced tools to analyze and improve the performance of a query.
- View Monitor Profiles**: View the status of all monitor profiles on a subsystem, create new monitor profiles, and open existing monitor profiles.

The "View Monitor Profiles" card is circled in red. At the bottom right of the window, there is a button labeled "Define Workload..."



Capture Workload Queries with Monitor



Capture Workload Queries with Monitor

Settings
Define a filter to specify when the monitor profile pushes out information about a statement execution and which information is recorded at each statement push.

Steps

1. Monitor Name
2. Source
3. Filter
4. Start Monitor

Monitor type: Normal execution
Description: Monitor a normal query exception.

General Settings

- Push out EXPLAIN information

Granularity:

- Push out number of executions and accumulated CPU time (recommended for minimal effects to performance)
- Push out complete runtime information

Limit for statement pushes that leave the cache: 5000

< Back **Next >** Finish Cancel

Selected Monitor Profile

Define Workload...

Capturing runtime and explain information for up to 5000 statements



Capture Workload Queries with Monitor

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The 'Workload Statements' panel is active, showing a list of captured queries. A blue callout box highlights the authentication details: 'Auth. ID: ADMF001' and 'IP: 9.181.133.73'. A red box highlights the table of captured statements.

Statements captured
Auth. ID: ADMF001
IP: 9.181.133.73

All of the rows are displayed. The number of rows is 8.

	Accumulated Ela...	Average Elaps...	Accumulated C...	Average CPU Ti...	Statement Text
R	1.1260938E-4	1.1260938E-4	1.11054695E-4	1.11054695E-4	SELECT 1 FROM SYSIBM.SYSDUMMY1
R	0.0048566284	0.0024283142	0.0021274262	0.0010637131	SELECT C_NAME,C_ADDRESS,N_NAM
R	0.007253484	0.003626742	0.0050335685	0.0025167842	SELECT O_ORDERKEY,O_CUSTKEY,O
R	0.0029460466	0.0014730233	0.0028902674	0.0014451337	SELECT S_SUPPKEY,S_NAME, SUM(L
R	0.0047927657	0.0023963829	0.004579288	0.002289644	SELECT O_ORDERKEY,O_CUSTKEY,O
R	0.003707945	0.0018539726	0.0034317768	0.0017158884	SELECT S_SUPPKEY,S_NAME, SUM(L
R	0.0052859383	0.0026429691	0.004823123	0.0024115616	SELECT O_ORDERKEY,O_CUSTKEY,O
R	0.0021096093	0.0010548047	0.0020319007	0.0010159503	SELECT S_SUPPKEY,S_NAME, SUM(L





Performing Health Check for Application Workloads with Design Advisors



Design Advisors

- Stats Advisor – stats recommendation for access path selection
- Index Advisor – recommendation for index design
- Query Advisor – recommendation for query design



Tune Query Workload

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The main window is titled "Workload Project" and shows configuration details for a workload named "DemoWorkload". The configuration includes the project name, subsystem (V9EC03), workload owner (SYSADM), and summary status (EXPLAINING). A red box highlights the "Run Advisors" section, which contains a list of advisor options: "Run All Advisors", "Run Workload Statistics Advisor", "Run Workload Index Advisor", "Run Workload Query Advisor", and "Show Advisor Options". A blue callout box labeled "Workload-based Advisors" points to this list. The interface also includes a Project Navigator on the left and a bottom status bar indicating "Retrieving PLAN_TABLE Records".

Workload-based Advisors

- Run All Advisors
- Run Workload Statistics Advisor
- Run Workload Index Advisor
- Run Workload Query Advisor
- Show Advisor Options



Tune Query Workload – Statistics Advisor

Recommendation Summary

Number	Priority	Recommendation	Description
1	High	Run complete RUNSTATS	Gather or Recollect all relevant statistics for...
2	Low	Run partial RUNSTATS	Repair the statistics problems within this w...

RUNSTATS command

RUNSTATS Control Statements

```

RUNSTATS TABLESPACE DB4PART.TSPART
TABLE(SYSADM.PART) SAMPLE 5
COLUMN
(P_MFGR,P_SIZE,P_RETAILPRICE,P_TYPE,P_PAR
TKEY,
P_NAME)
COLGROUP(P_NAME) HISTOGRAM
NUMQUANTILES 20
COLGROUP(P_TYPE) HISTOGRAM
NUMQUANTILES 20
COLGROUP(P_MFGR) FREQVAL COUNT 10
SORTDEVT SYSDA
INDEX(SYSADM.UXP@SZPTPKMF KEYCARD
FREQVAL NUMCOLS 1 COUNT 10,
    
```

Description

Gather or Recollect all relevant statistics for the entire workload. Periodical statistics collection brings all relevant statistics up to date and consistent, and avoids the workload performance degradation due to obsolete statistics.

Actions

- Details...
- Run...
- Copy
- Save

Recommendation Explanation



Tune Query Workload – Index Advisor

Workload Index Advisor Recommendations

The following information shows the index recommendations for this workload. You can view the performance improvement when all recommendations are applied. There is the option to run index analysis again with different values to see if there are better recommendations.

Workload performance improvement is an estimate based on applying all recommendations.

Estimated performance improvement: 36.39 %
Disk space required(DASD space): 44626.41 MB

Feature Details	Action	Object...	Columns	Estimated Di
<input checked="" type="checkbox"/> Index	Create	LINEITE...	L_DISCOUNT(ASC), L...	1796.082031
<input checked="" type="checkbox"/> Index	Create	LINEITE...	L_DISCOUNT(ASC), L...	4056.003906

Run Selected DDL Statements

```
CREATE INDEX "DB2OE"."SUPPLIER_VIRT_IDX_1171449271453" ON "SYSADM"."SUPPLIER" ("S_SUPPKEY" ASC, "S_NATIONKEY" ASC, "S_NAME" ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;  
CREATE INDEX "DB2OE"."SUPPLIER_VIRT_IDX_1171448270672" ON "SYSADM"."SUPPLIER" ("S_ACCTBAL" ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;  
CREATE INDEX "DB2OE"."SUPPLIER_VIRT_IDX_1171448270675" ON "SYSADM"."SUPPLIER" ("S_SUPPKEY" ASC, "S_ACCTBAL" ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;  
CREATE INDEX "DB2OE"."REGION_VIRT_IDX_1171448180065" ON "SYSADM"."REGION" ("R_REGIONKEY" ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;
```



Tune Query Workload – Query Advisor

Recommendation Summary

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	94
Statements with Warnings	34
Number of High Severity Warnings	0
Number of Medium Severity Warnings	4
Number of Low Severity Warnings	50
Statements with High Severity Warnings	0
Statements with Medium Severity Warnings	4
Statements with Low Severity Warnings	31

View statements that meet the following criteria:

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

[Restore Defaults](#) [Save as Defaults](#) [View Statements](#)

Workload | Statements | Advisors | **Query Advisor Summary**



Tune Query Workload – Query Advisor

Recommendations List for one query

Recommendation Description

Recommendation Explanation

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The main window shows the 'Query Advisor Recommendations List' for query 332. A table lists one recommendation with a severity of 'Low' and a confidence of 'Low'. The description of this recommendation is 'Consider replacing the asterisk (*) or'. Below this, the 'SQL Text' is shown as 'SELECT * FROM LINEITEM WHERE L_COMMENT='BNQ2y5xz1BBC06n4' for fetch only'. The 'Selected Recommendation' section is expanded to show a detailed 'Description' and an 'Explanation'. The 'Description' explains that using asterisks or long column lists in the SELECT clause is a bad practice because it causes DB2 to return unnecessary long rows, increasing CPU cost and overhead. The 'Explanation' further details that extraneous columns increase the cost of performing the sort and might discourage the access path that performs the sort, and that extraneous columns also increase the sort data length. The 'Query Advisor' tab is highlighted in the bottom navigation bar.

No	Severity	Confidence	Line Number	Description
1	Low	Low		Consider replacing the asterisk (*) or

Description

Consider replacing the asterisk (*) or the long column list of table SYSADM.LINEITEM in the SELECT list with the names of only the required columns in table SYSADM.LINEITEM. Check the explanation for this warning for more details about possible impact and examples.

Explanation

Using asterisks in the select list of an SQL statement is generally considered a bad practice and should be avoided. Extraneous columns cause DB2 to return unnecessarily long rows, thereby increasing CPU cost and overhead as the qualified rows are returned to the client. In addition, if a query performs a sort, the long select column list increases the cost of performing the sort and might discourage the access path that performs the sort. Extraneous columns also increase the sort data length, and might





Monitoring Workload Exceptions



Defining Performance Exceptions

- Two kinds of exceptions
 - CPU time exception
 - CPU spike exception
- Notify performance exception as soon as it occurs



Capture Workload Exceptions with Monitor

Monitor Name
Provide a name for the monitor profile, and specify the type of monitor profile.

Steps

1. Monitor Name
2. Source
3. Filter
4. Start Monitor

Subsystem: CSRECO63 <partially enabled>

Monitor name: CPU Time Exception Monitor

Owner: SYSADM

Select which type of monitor to create.

- Normal - Monitor all SQL statements that run within a monitor source
- Exception - Monitor SQL statement pushes that meet specific exception conditions

Description:

< Back Next > Finish Cancel

Number of push out for all statements when the statement leaves the cache: 5000

Number of push out for each statement when exception threshold is met:

70



Capture Workload Exceptions with Monitor

Monitor Source
Define sources for the monitor profile. You can specify more than one statement source in a single monitor profile. Specify Authorization ID and IP address for each dynamic statement source, and specify

Steps

1. Monitor Name
2. Source
3. Filter
4. Start Monitor

Dynamic statements

Authorization ID: ADMF001

IP address: 9.181.133.73

Embedded statements

Plan name:

Collection ID:

Package name:

Add Remove

Source List

Authorizati...	IP Address	Plan Name	Collection ID	Package N...
ADMF001	9.181.133.73			

< Back Next > Finish Cancel

Number of push out for all statements when the statement leaves the cache: 5000

Number of push out for each statement when execution threshold is met:

Define Workload...

**Scope of the exception monitor –
Auth ID: ADMF001
IP: 9.181.133.73**



Capture Workload Exceptions with Monitor

Settings
Define a filter to specify when the monitor profile pushes out information about a statement execution and which information is recorded at each statement push.

Steps

1. Monitor Name
2. Source
3. Filter
4. Start Monitor

Monitor type:

- CPU time exception
- Relative CPU time exce

Description:

Monitor Settings
Specify the type of normal monitor information to push out, or specify threshold criteria for the selected exception monitor types.

CPU time threshold: seconds

Relative CPU time threshold: %

General Settings

- Push out EXPLAIN information

Granularity:

- Push out number of executions and accumulated CPU time (recommended for minimal effects to performance)
- Push out complete runtime information

Limit for individual statement push when exception threshold is met:

Limit for all statement pushes when exception threshold is met:

< Back Next > Finish Cancel

Number of push out for all statements when the statement leaves the cache:

Number of push out for each statement when exception threshold is met:

Define Workload...

Raise the exception whenever CPU time for a single execution exceeds 0.002 second



Capture Workload Exceptions with Monitor

Start Monitor
Start or schedule the monitor profile, and specify monitoring details.

Steps

1. Monitor Name
2. Source
3. Filter
- ▶ 4. Start Monitor

Start monitor profile: Start Now
Description: The monitor profile starts immediately when you finish the monitor wizard.

Profile Details
Interval for consolidating statement pushes: 30 minutes

Post new exceptions, if there is any, every 30 minutes

< Back Next > Finish Cancel

Number of push out for all statements when the statement leaves the cache: 5000
Number of push out for each statement when execution threshold is met:



Capture Workload Exceptions with Monitor

The screenshot shows the IBM DB2 Optimization Expert for z/OS interface. The main window displays the 'Workload Statements' section for a monitor named 'CSREC063_CPU Time Exception Monitor'. A table lists captured exceptions, with a red box highlighting the first four rows. A blue callout box points to the table with the text: 'Exceptions captured Auth ID: ADMF001 IP: 9.181.133.73'. The table has columns for Execution ID, Source, Accumulated Elapsed Time, Average Elapsed Time, Accumulated CPU Time, Average CPU Time, and Statement Text.

Exe...	Source	Accumulated Ela...	Average Elap...	Accumulat...	Average CP...	Statement Text
1	MONITOR	4.2859367E-5	4.2859367E-5	4.2703123...	4.2703123E...	SELECT 1 FROM SYSIB
10	MONITOR	0.009633453	9.633453E-4	0.0089606...	8.960639E-4	SELECT * FROM SCQA
42	MONITOR	0.0075379997	1.7947618E-4	0.0072187...	1.7187552E...	INSERT INTO DSNACC.
116	MONITOR	0.024593571	2.1201355E-4	0.0193262...	1.6660566E...	INSERT INTO DSNACC.



Tuning Problem Queries

The screenshot shows the IBM DB2 Optimization Expert for z/OS interface. The main window is titled 'View Queries' and displays the 'Subsystem Context' and 'Queries List' sections. The 'Subsystem Context' section shows the selected subsystem as 'CSRECO63 <partially enabled>'. The 'Queries List' section shows the selected query source as 'Monitor' and the monitor profile as 'CPU Time Exception Monitor'. A red box highlights the 'Advisors' menu, which includes options like 'Run All Advisors', 'Run Statistics Advisor', 'Run Query Advisor', 'Run Access Path Advisor', 'Run Index Advisor', and 'Show Advisor Options'. A blue callout box points to this menu with the text: 'Running design advisors or expert expert tools for problem query'. Below the menu, a table displays query execution statistics.

	LAST_UPDATE_TS	LAST_EXPLAIN_TS	CACHED_TS
...NACC.JSRE...	2007-02-14 00:54:19...	2007-02-14 00:54:1...	2007-02-14 00:54:19.74
...NACC.JSRE...	2007-02-14 01:01:58...	2007-02-14 01:01:5...	2007-02-14 01:01:58.66
...SYSIBM.SY...	2007-02-14 01:26:20...	2007-02-14 01:26:2...	2007-02-14 01:26:20.65
...SCQA0000...	2007-02-14 01:26:22...	2007-02-14 01:26:2...	2007-02-14 01:26:22.41



Thank you for your interest in Optimization Service Center & DB2 Optimization Expert

- Gene Fuh, fuh@us.ibm.com
- Jay Bruce, jmbruce@us.ibm.com
- Patrick Bossman, bossman@us.ibm.com
- Kun Peng Ren, renkunp@cn.ibm.com
- Xiao Yi Wang, wangxyi@cn.ibm.com

