



Session: Z32

DB2 for z/OS Visual Explain V8

Patrick Bossman

IBM DB2 Information Management
Technical Conference

Sept. 20-24, 2004

Las Vegas, NV

Visual Explain for DB2 z/OS V8

- Visual Explain overview
 - Basic capabilities
 - Service SQL function
- Enhanced capabilities of Visual Explain
 - Qualified rows estimate
 - Wealth of predicate information
 - Limited partition scan information
 - Parallelism details
- Statistics Advisor

Interface overview

- Basic GUI functions
 - Connect lists all catalogued databases
 - Enable/maintain Visual Explain
- Browse subsystem parameters
- Browse external plan tables
 - PLAN_TABLE, DSN_STATEMENT_TABLE,
DSN_FUNCTION_TABLE

Connect

The screenshot shows the Visual Explain application window. At the top, there is a menu bar with 'Subsystem', 'Tools', 'Properties', 'Windows', and 'Help'. Below the menu bar is a toolbar with various icons. The main area is titled 'List Databases' and contains a table with the following data:

Database Alias	Status	Database Name	Comment	User ID	SQLID
JESSICA	✓	STLEC1			
LITALIEN	✓	STLEC1			
MOLERA	✓	STLEC1			
TERRY	✓	STLEC1			
UTEC492A	✓	STLEC1			
UTEC492B	✓	STLEC1B			
UTEC780A	⚡	STLEC1		ADMF001	ADMF001
UTEC780B	✓	STLEC1B			
TUTORIAL	✓	DSNVEDEMO	Visual Explain Tut...		

To the right of the table are several buttons: 'Connect...', 'Disconnect', 'Refresh', 'Update...', 'Catalog...', 'Uncatalog', and 'Help'. At the bottom left, there is a 'Database Properties' section with the following information:

Host name: v7ec022.svl.ibm.com Protocol: tcpip
Node name: NDEC9316 Port number: 446

A blue arrow points from the text 'Click to Connect / disconnect' to the 'UTEC780A' row in the table.

Visual overview

The screenshot shows the 'Visual Explain' application window. At the top is a menu bar with 'Subsystem', 'Tools', 'Properties', 'Windows', and 'Help'. Below the menu is a toolbar with icons for database operations. The main area displays a table titled 'List Database' with the following data:

Database Alias	Status	Database Name	Comment	User ID	SQLID
JESSICA	✓	STLEC1			
LITALIEN	✓	STLEC1			
MOLEPA	✓	STLEC1			
TERRY	✓	STLEC1			
UTEC492A	✓	STLEC1			
UTEC492B	✓	STLEC1B			
UTEC780A	⚡	STLEC1		ADMF001	ADMF001
UTEC780B	✓	STLEC1B			
TUTORIAL	✓	DSNVEDEMO	Visual Explain Tut...		

On the right side of the table, there are several buttons: 'Connect...', 'Disconnect', 'Refresh', 'Update...', 'Catalog...', and 'Uncatalog'. At the bottom of the window, there is a 'Database Properties' section with the following information:

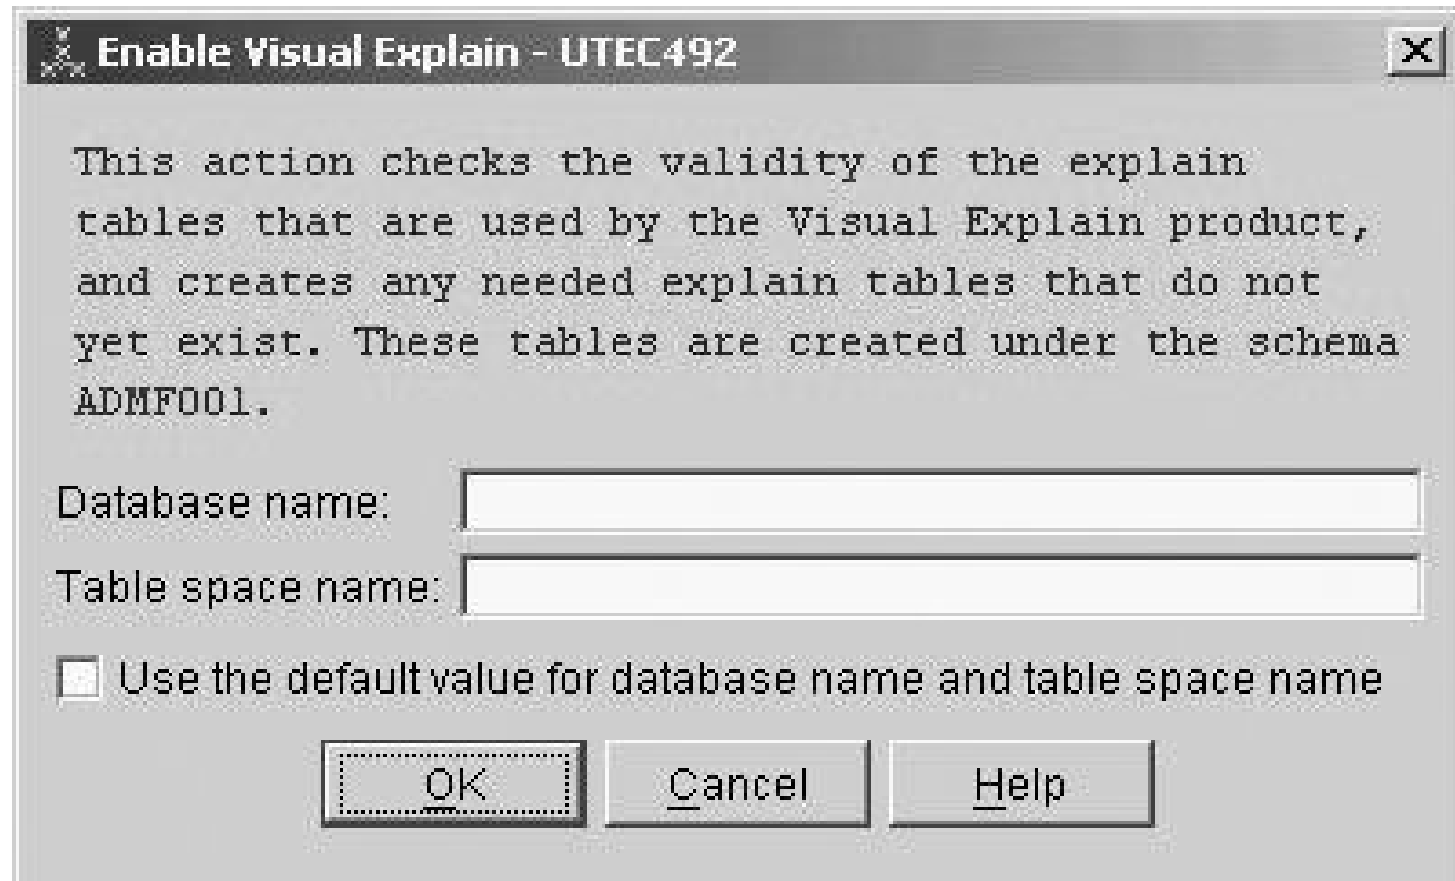
Host name: v7ec022.svl.ibm.com Protocol: tcpip
Node name: NDEC9316 Port number: 446

Four blue arrows point from text labels to specific icons in the toolbar:

- 'Enable IVE' points to the first icon (a green circle with a white checkmark).
- 'Maintain IVE' points to the second icon (a grid).
- 'View external explain tables' points to the third icon (a document with a magnifying glass).
- 'Browse zparms' points to the fourth icon (a document with a lightning bolt).

At the bottom right of the window, there is a 'Help' button.

Enable IVE



Enable Visual Explain - UTEC492 [X]

This action checks the validity of the explain tables that are used by the Visual Explain product, and creates any needed explain tables that do not yet exist. These tables are created under the schema ADMFOO1.

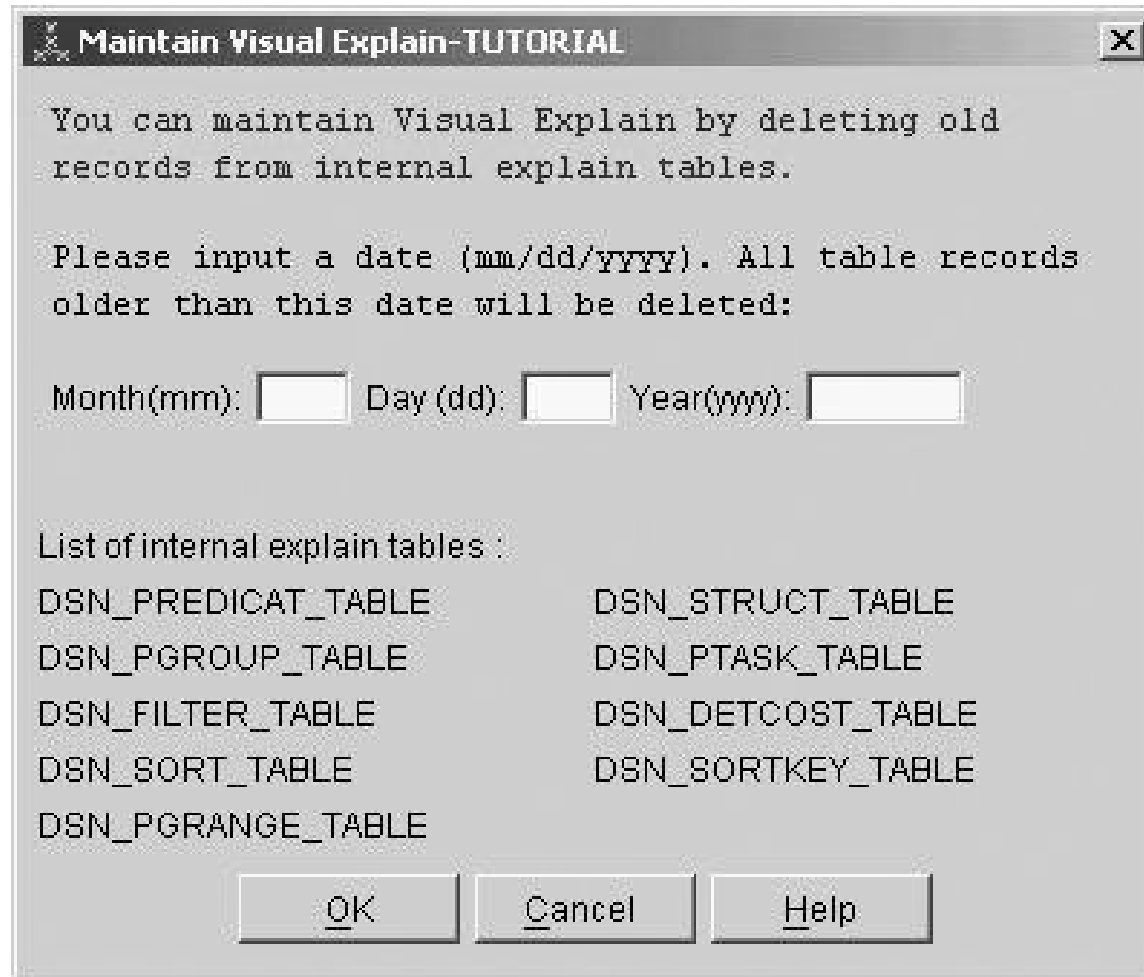
Database name:

Table space name:

Use the default value for database name and table space name

- **Creation of necessary objects within tool**
 - **Select database / table space**

Maintain IVE



- Date based purge for internal explain tables

Browse subsystem parameters

- Options to list parameters
 - Alphabetically by parameter name
 - Installation field name
 - Installation panel name
- Search by parameter name also available

Browse ZPARMs

The screenshot shows a window titled "Subsystem Parameters - TUTORIAL" with a search bar at the top. Below the search bar is a "List of parameter names" containing various parameters such as MAINTYPE, MAXARCH, MAXDBAT, MAXKEEPD, MAXRBLK, MAXRTU, MAXTYPE1, MAX_OPT_CPU, MAX_OPT_ELAP, MAX_OPT_STOR, MCCSID, MEMBNAME, and MINDVSCL. A red oval highlights the "MAXRBLK" entry in this list, with the text "Parameter List" written inside the oval. To the right of the list, the details for the selected parameter are shown: Value: 0000000063, Installation panel: DSNTIPC, Installation field: RID POOL SIZE, Macro: DSN6SPRM, and Parameter: MAXRBLK. A red arrow points from the text "ZPARM information" to the "Installation field" value. Below these details is a "Description:" section with a text area containing the following text: "This field specifies the amount of storage needed for the RID pool as calculated by the CLIST. Acceptable values are 0, or 16K-1000000K. The default size is based on calculations. You can estimate the storage required for the RID pool by multiplying the number of concurrent RID processing activities by the average number of RIDs times 2 times 4 bytes per RID. A value of 0 disables the use of the RID pool. If 0 is the value, DB2 does not use access paths or join methods that depend on RID pool storage." A red rounded rectangle with the text "Parameter Description" is placed below this text area. At the bottom right of the window is a "Help" button.

Parameter list close-up



Parameter setting close-up

Value:	SYSADM
Installation panel:	DSNTIPP
Installation field:	SYSTEM ADMIN 1
Macro:	DSN6SPRM
Parameter:	SYSADM
Description:	

Service SQL

- Provides problem recreation information to service team
 - Uses SQL statement or PLAN_TABLE as input
 - SQL statement (if SQL used as input)
 - DDL
 - Catalog statistics
 - Zparms (if DSNWZP stored procedure available)
 - Environment specific information
 - CPU speed
 - Bufferpool, ridpool, sortpool sizes
 - Number of processors

Service SQL (cont)

- Client version of host based DB2PLI8
 - Easier to install
 - Easier to use
 - Documentation from Service SQL more ready for use
- Eliminate common problems / frustrations
 - Input pmr number
 - SQL statement
 - SQL explained
 - Click to generate documentation
 - File names based on pmr
 - Click to FTP documentation
 - Appropriate FTP settings already set

Launch Service SQL

The screenshot shows the Visual Explain application window. The title bar reads "Visual Explain". Below the title bar is a menu bar with "Subsystem", "Tools", "Properties", "Windows", and "Help". A toolbar contains several icons, including a blue arrow pointing upwards to the "Tools" menu. The main window is titled "List Databases" and contains a table with the following data:

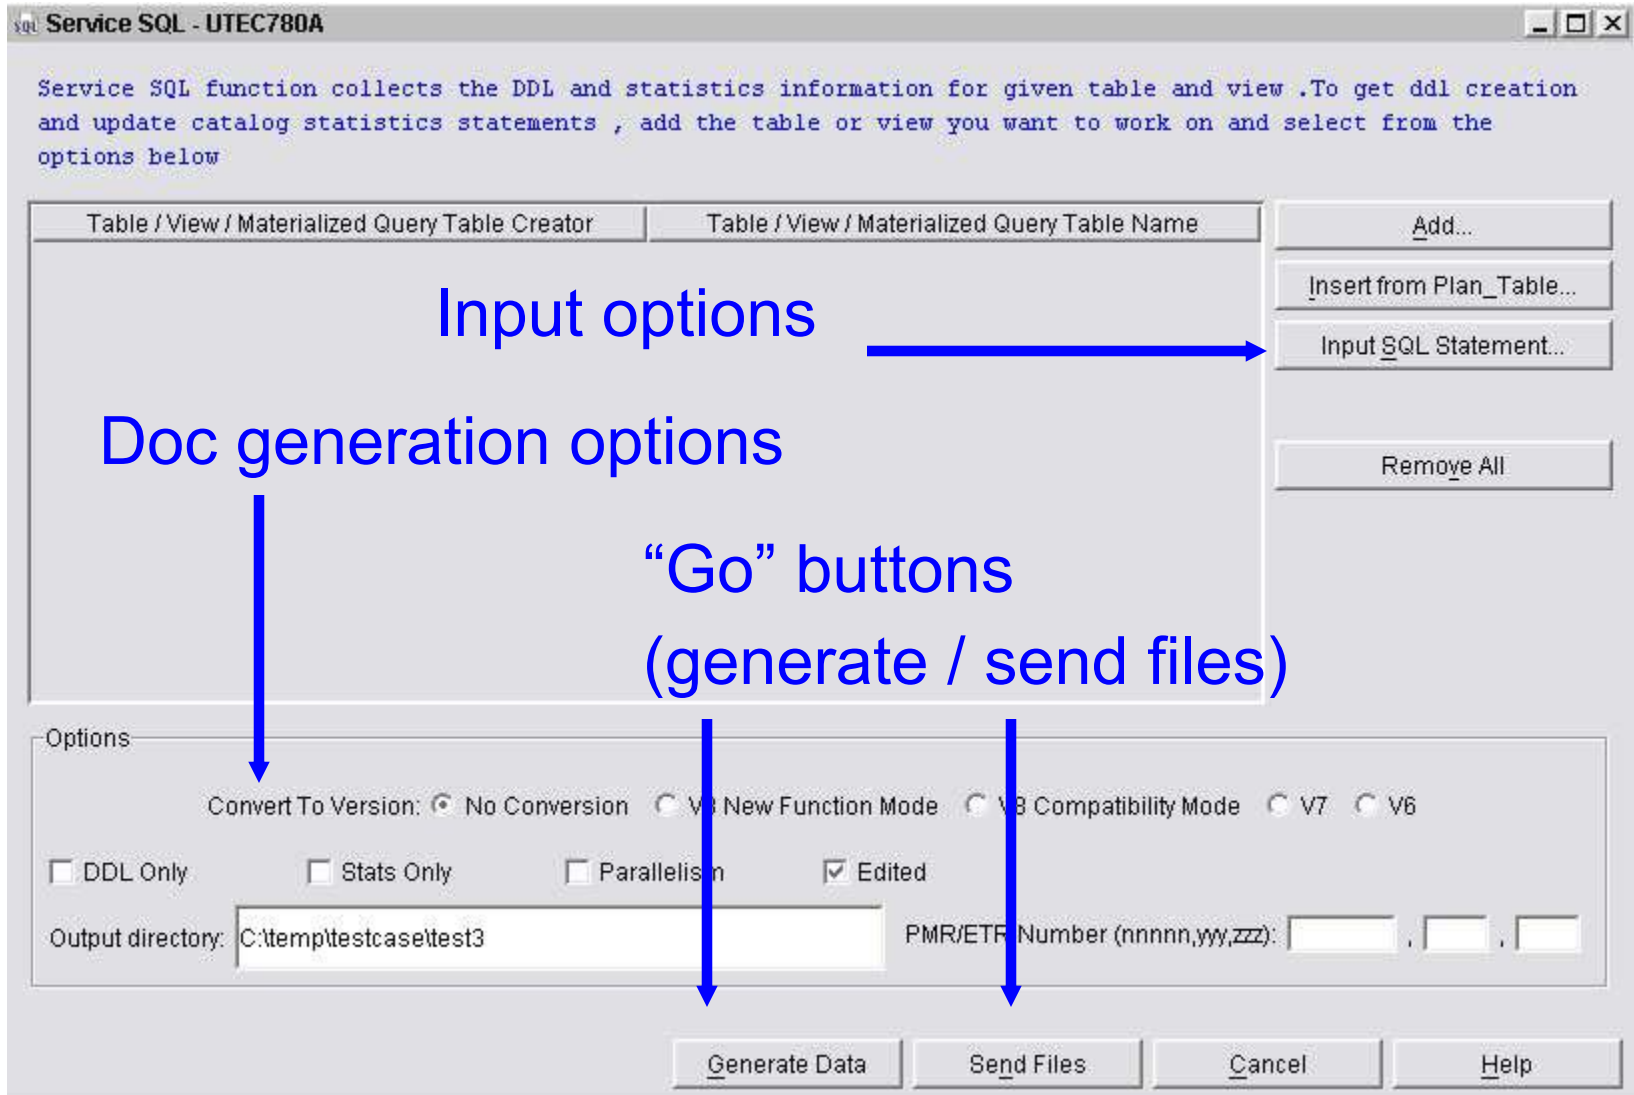
Database Alias	Status	Database Name	Comment	User ID	SQLID
JESSICA	✓	STLEC1			
LITALIEN	✓	STLEC1			
MOLERA	✓	STLEC1			
TERRY	✓	STLEC1			
UTEC492A	✓	STLEC1			
UTEC492B	✓	STLEC1B			
UTEC780A	⚡	STLEC1		ADMF001	ADMF001
UTEC780B	✓	STLEC1B			
TUTORIAL	✓	DSNVEDEMO	Visual Explain Tut...		

On the right side of the window, there are several buttons: "Connect...", "Disconnect", "Refresh", "Update...", "Catalog...", "Uncatalog", and "Help". At the bottom left, there is a "Database Properties" section with the following information:

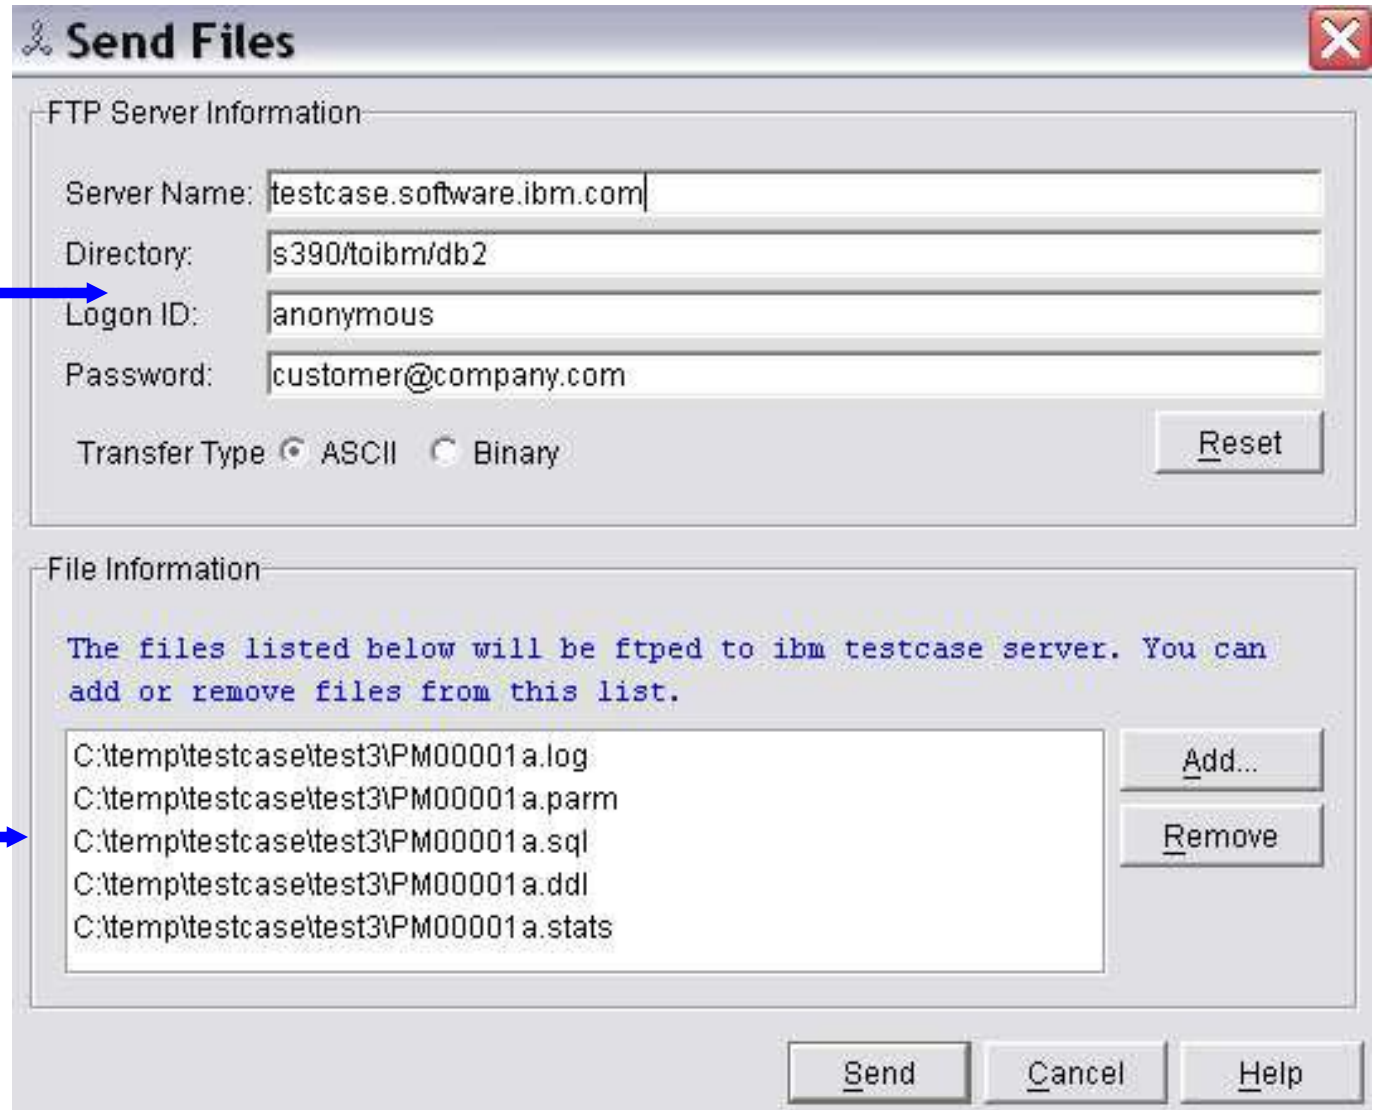
Host name: v7ec022.svl.ibm.com Protocol: tcpip
Node name: NDEC9316 Port number: 446

A blue arrow points from the text "Option off of Tools menu" to the Tools menu in the menu bar.

Service SQL main screen



Send files screen



The image shows a 'Send Files' dialog box with two main sections: 'FTP Server Information' and 'File Information'. The 'FTP Server Information' section contains fields for 'Server Name', 'Directory', 'Logon ID', and 'Password', along with radio buttons for 'Transfer Type' (ASCII and Binary) and a 'Reset' button. The 'File Information' section contains a text area with a list of files and two buttons: 'Add...' and 'Remove...'. At the bottom of the dialog are 'Send', 'Cancel', and 'Help' buttons.

Send Files

FTP Server Information

Server Name:

Directory:

Logon ID:

Password:

Transfer Type ASCII Binary

File Information

The files listed below will be ftped to ibm testcase server. You can add or remove files from this list.

C:\temp\testcase\test3\PM00001a.parm
C:\temp\testcase\test3\PM00001a.sql
C:\temp\testcase\test3\PM00001a.ddl
C:\temp\testcase\test3\PM00001a.stats

Server info



Files



Service SQL summary

- Improve serviceability for SQL performance problems
 - Make doc collection for SQL problems
 - Fast & easy
 - Complete
 - Provide more usable doc to SVL
 - Documentation almost ready to use
 - Recreate process much shorter for L2
 - Service SQL works with DB2 V6, V7, V8

Visual Explain for DB2 z/OS V8

- Visual Explain overview
 - Basic capabilities
 - Service SQL function
- Enhanced capabilities of Visual Explain
 - Qualified rows estimate
 - Wealth of predicate information
 - Limited partition scan information
 - Parallelism details
- Statistics Advisor

Visual Explain capability

- Graphically display access path
- Easy navigation of statistics
- Detailed explain information
 - Detailed explain information requires DB2 for z/OS V8 connection
 - Wealth of predicate information
 - Qualified row estimates (single table estimate, within explain)
 - Limited partition scan details
 - Parallelism details

Tune SQL option

The screenshot shows the 'Tune SQL - TUTORIAL' application window. The interface includes a menu bar (File, Edit, Action, View), a toolbar with icons for file operations, and a main workspace divided into several sections:

- Query Information:** QueryNo: 1, SQLID: ADMF007
- Command History:** A table listing various SQL queries and their execution details.
- Configuration:** Current degree: System default, Current refresh age: 0, Current maintained table types: None.
- Input SQL:** A text area containing a SQL query: `SELECT SUM(N_NATIONKEY), SUM(N_NATIONKEY), SUM(N_NATIONKEY), SUM(N_NATIONKEY), SUM(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), COUNT(*) FROM NATION_NP`
- Messages for Execution and Explain:** A large empty text area.
- Buttons:** Explain, Execute, Plan Hint, Help.

Annotations on the image:

- A blue arrow points from the 'Saved SQL' text to the 'Input SQL' field.
- A red arrow points from the 'Click to Explain' text to the 'Explain' button.

Name	SQL Statement	
Q1	SELECT SUM(...	index only
Q2	SELECT SUM(L...	index+fetch
Q3	SELECT L_OR...	partitioning index page range s
Q4	SELECT C_NA...	in-list index scan
Q5	SELECT COUN...	sequential prefetch
Q6	SELECT L_CO...	list prefetch
Q7	SELECT COUN...	page range scan
Q8	SELECT N_CO...	page range scan
Q9	SELECT L_QU...	group by

Look and feel of graph

The screenshot displays the DB2 Explain interface. At the top, the status bar reads: "DB2 Platform: Z/OS DB2 Version: v8 Explain Time: 2003-08-26 10:43:20.37".

The left pane, titled "query", contains a table with the following data:

Name	Value
Timestamp	2003-08-26 10:43:20.37
Type	SELECT
CPU Cost (ms)	964
CPU Cost (su)	10246
Cost Category	A
Reason	
Group Member	

Below the table, the "Attribute explanation:" section shows: "CPU Cost (ms): Estimated CPU cost (in milliseconds)".

The right pane displays a query plan graph. The nodes are:

- (1) QUERY (30029.727)
- (2) QB1 (30029.727)
- (3) FETCH (30029.727)
- (4) SORTRID (30030)
- (5) IXSCAN (30030)
- (6) SXL@PKSKOKEPDSN (179998372)
- (7) LINEITEM (179998372)

Annotations on the image:

- A blue arrow labeled "Related info" points to the "query" pane.
- A blue arrow labeled "Graph" points to the query plan graph.
- A blue arrow labeled "Details" points to the "CPU Cost (ms)" row in the table.
- Red arrows labeled "Nodes" point to the nodes (4) SORTRID, (5) IXSCAN, (6) SXL@PKSKOKEPDSN, and (7) LINEITEM.

Node information

- Show important information for node
 - Index --> FULLKEYCARDF
 - Table --> CARDF
 - Access methods (index scan, R-scan, sort, ...)
 - Qualified rows
 - Join methods
 - Qualified rows

Nodes

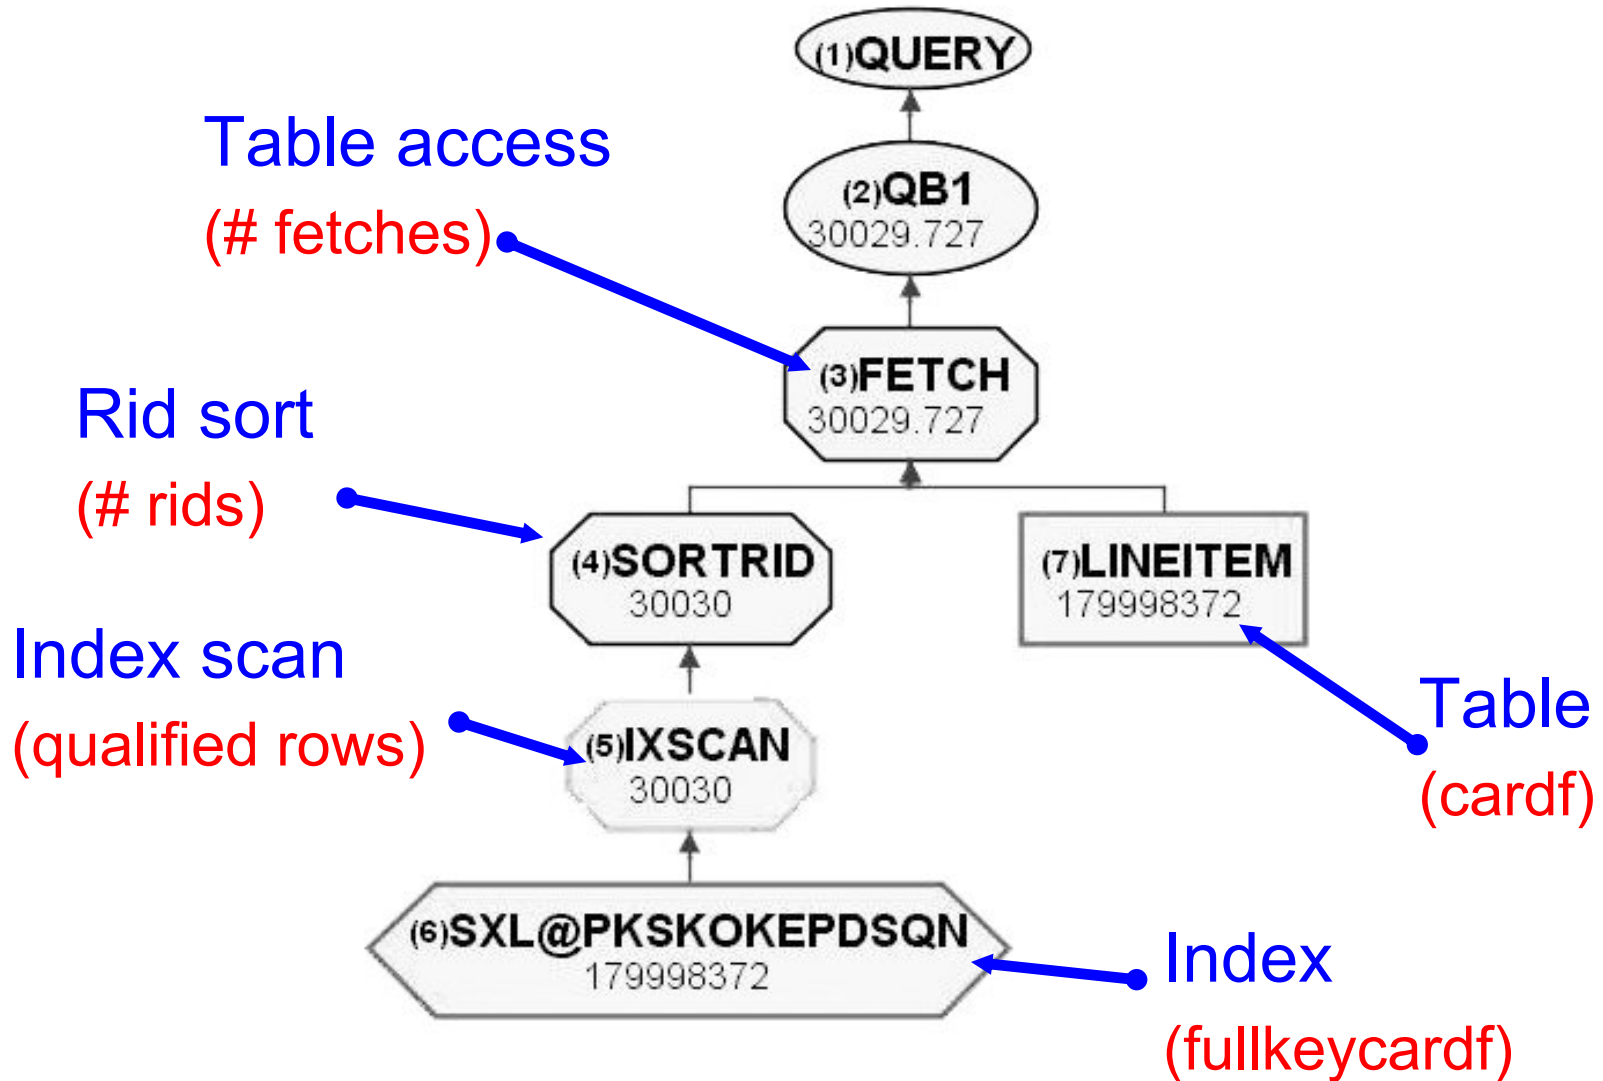


Table details

- Details shows closer look
 - Table statistics
 - Related information in tabbed format
 - Columns
 - Indexes
 - ...
 - Qualifying rows
 - Single table row estimate
 - As though table were outer table of join
 - Poor join sequence / method often due to inaccurate single table row estimates

Table graph

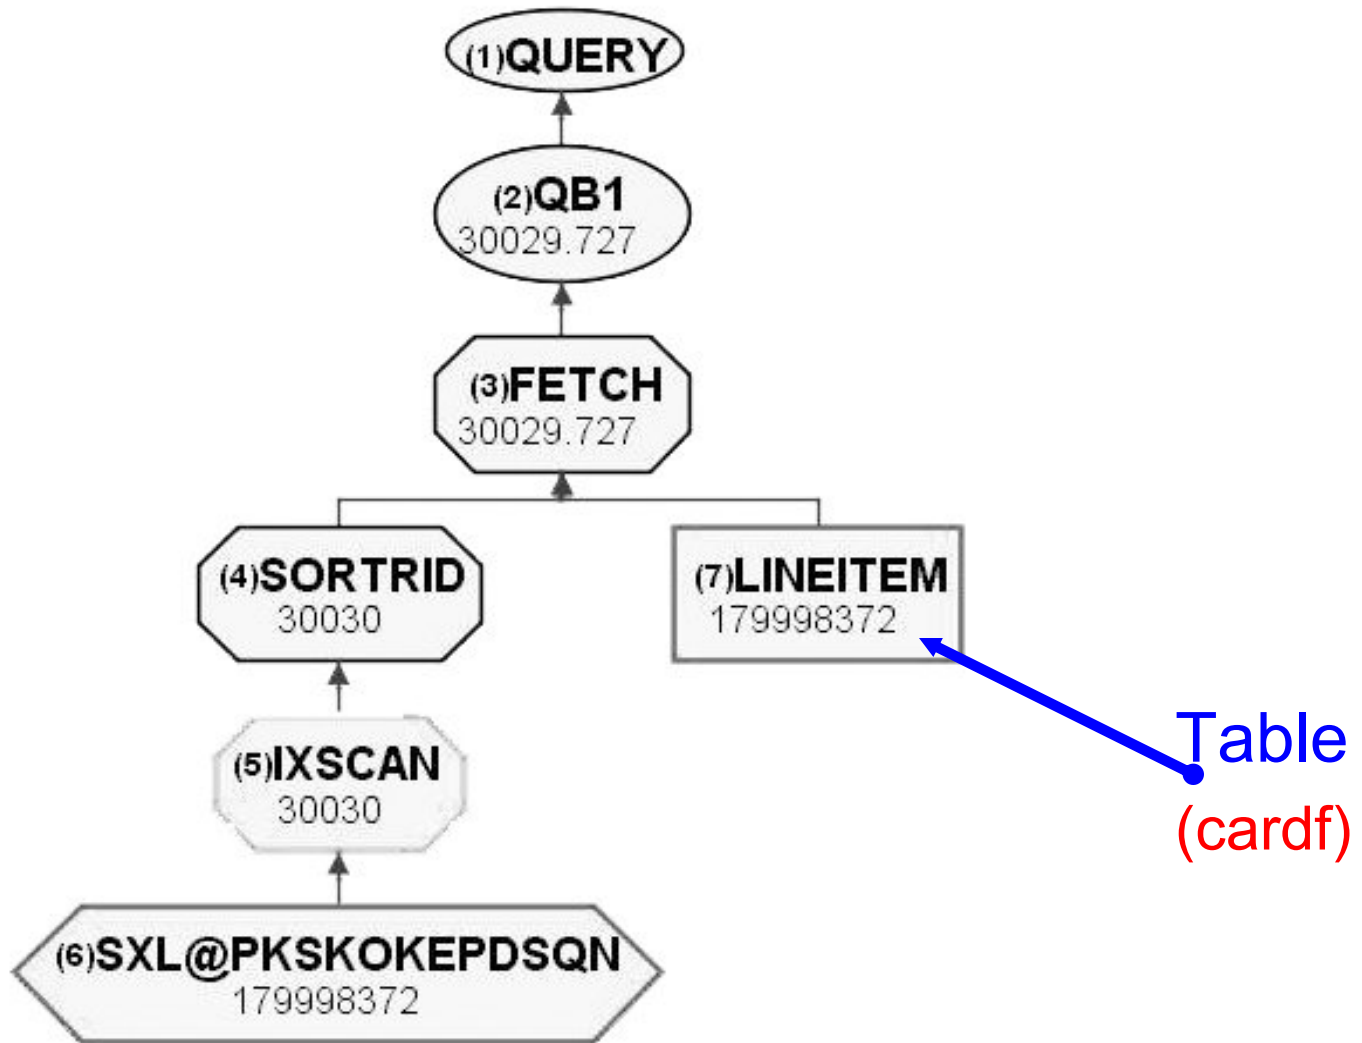


Table details

Descriptor - [7] table(LINEITEM)

table(LINEITEM)

- Columns
- Indexes
- Tablespace
- Table Partitions

Show attribute explanation Views: cost estimation

Name	Value
Name	LINEITEM
Creator	SYSADM
Correlation Name	
Qualifying Rows	30029.727
Rows	1.79998372E8
Pages	5996372
Compressed Row Percentage	0
Timestamp	1-01-01 00:00:00.0
Explain Time	2003-08-26 10:43:20.37

Table related info

- Columns
- Indexes
- Tablespace...

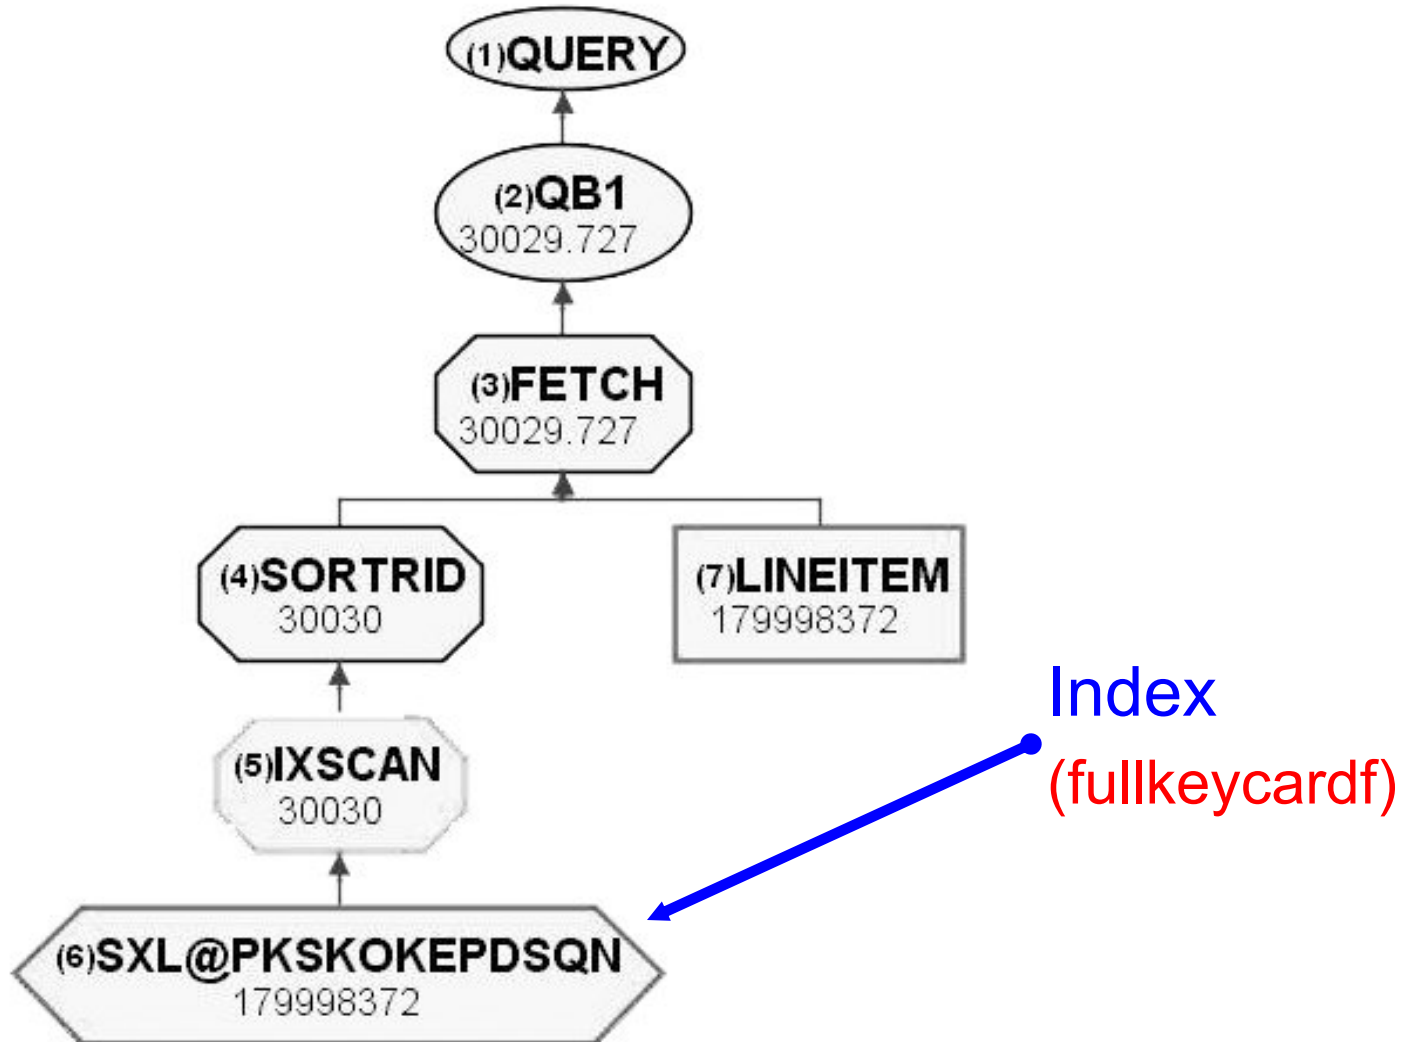
Qualifying rows

Table statistics

Index details

- Index information
 - Statistics
 - FIRSTKEYCARDF, FULLKEYCARDF
 - CLUSTERRATIOF
 - NLEAF
 - NLEVELS
 - Index key columns

Index graph



Index details

index(SXL@PKSKOKEPDSQN)

- Keys
 - indexkey(SXL@PKSKOKEPDSQN key1:L_PARTKEY)
 - indexkey(SXL@PKSKOKEPDSQN key2:L_SUPPKEY)
 - indexkey(SXL@PKSKOKEPDSQN key3:L_ORDERKEY)
 - indexkey(SXL@PKSKOKEPDSQN key4:L_EXTENDEDPR
 - indexkey(SXL@PKSKOKEPDSQN key5:L_DISCOUNT)
 - indexkey(SXL@PKSKOKEPDSQN key6:L_QUANTITY)
- Table

Show attribute explanation Views: cost estimation

Name	Value
Table Name	LINEITEM
Name	SXL@PKSKOKEPDSQN
Creator	SYSADM
Unique Rule	D
Clustering	N
Cluster Ratio	0.5
First Key Cardinality	6000000
Full Key Cardinality	1.79998372E8
Leaf Pages	2195103
Levels	4
Clustered	N
Type	2
Padded	
Timestamp	1-01-01 00:00:00.0
Explain Time	2003-08-26 10:43:20.37

Attribute explanation:
Table Name: Table name which index is created on. Double click this attribute to show table information.

Save As... Print... Suggesti... Help Close

Related info

- Key columns
- Table info

Properties

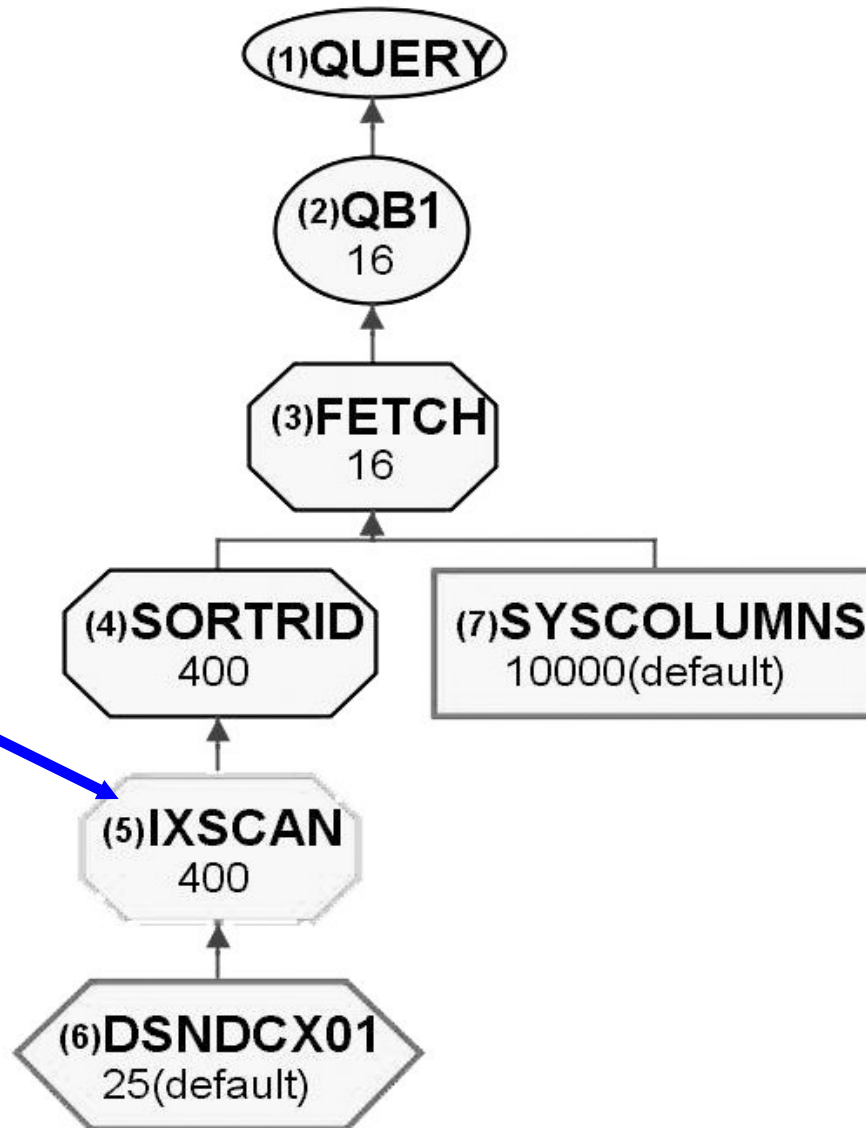
- Unique
- Clustering

Statistics


Index scan details

- Some really useful costing information...
 - Folder organizes predicates
 - Matching predicates
 - **Screening predicates**
 - Index cost estimates
 - Leaf pages scanned
 - Qualified RIDs
 - Filter factor estimates
 - Individual predicates
 - Total Matching filtering
 - Total index filtering

Index scan details



Index scan
(qualified rows)



Index scan details

Descriptor - [5] iscan

iscan

- Matching Predicates
 - SYSCOLUMNS.TBCREATOR=(EXPR)
- Screening Predicates
 - SYSCOLUMNS.NAME=(EXPR)

Show attribute explanation Views: cost estimation

Name	Value
Input RIDs	10000
Index Leaf Pages	34(default)
Matching Predicates	Filter Factor
SYSCOLUMNS.TBCREATOR=(EXPR)	0.04
Scanned Leaf Pages	2
Screening Predicates	Filter Factor
SYSCOLUMNS.NAME=(EXPR)	0.04
Output RIDs	400
Total Filter Factor	0.0016
Prefetch	
Matching Columns	1

Attribute explanation:
Input RIDs: Number of ROWIDs in the index

Save As... Print... Suggesti... Help Close

Predicate info

- Matching
- Screening

Scanned leaf pages

Output RIDs

Total filter factor

Index predicate details

- You can also see predicate information
 - Matching or screening
 - Type
 - Equal, range, etc.
 - Filter factor
 - Marker flag
 - Y = parameter marker, special register or host variable

Index predicate info

Descriptor - [5] iscan

iscan

- Matching Predicates
 - SYSCOLUMNS.TBCREATOR=(EXPR)**
- Screening Predicates
 - SYSCOLUMNS.NAME=(EXPR)

Show attribute explanation Views: cost estimation

Name	Value
Filter Factor	0.04
Type	EQUAL
Stage	MATCHING
Order	1
Marker	Y
LHS Text	TBCREATOR
RHS Text	VALUE

Attribute explanation:

Filter Factor:

Save As... Print... Suggesti... Help Close

Predicate info

- Matching
- Screening

Filter factor

Type

Marker?

Tuning comments

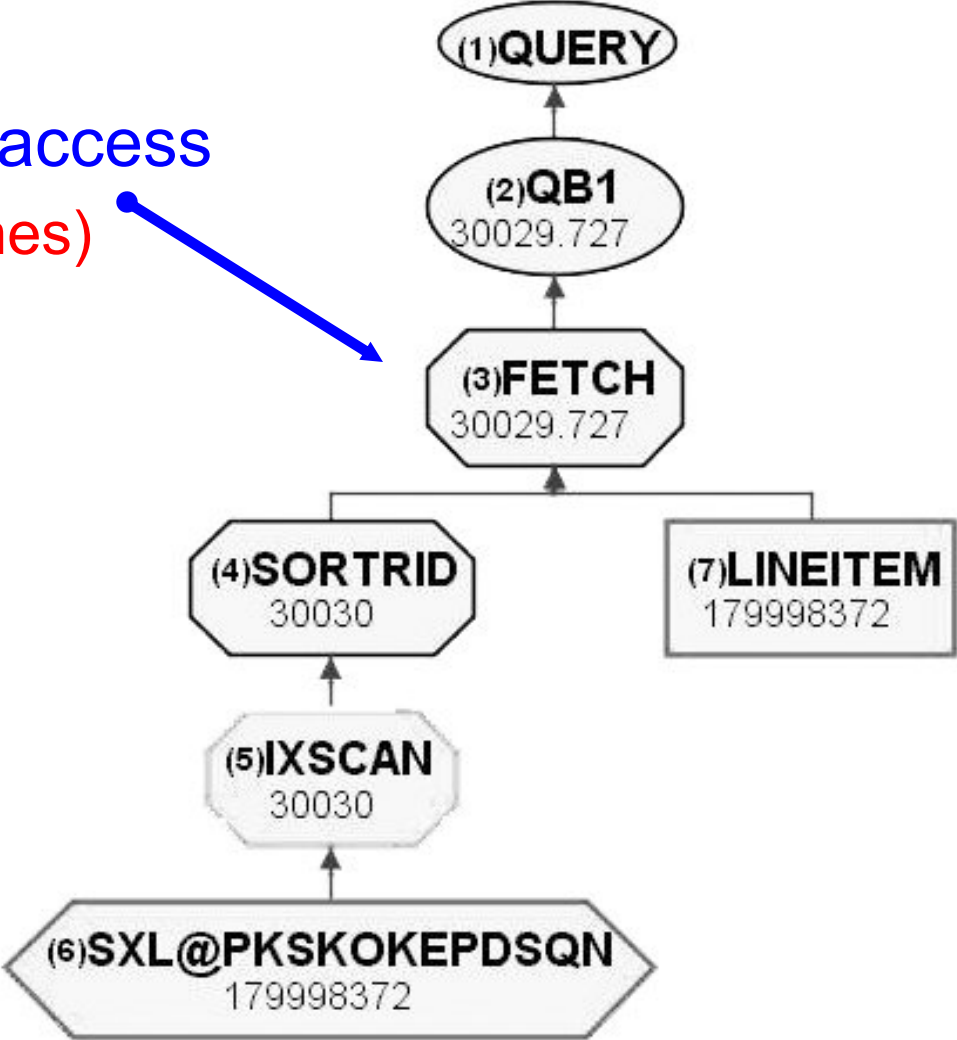
- Answer questions you've always had...
 - Am I getting screening?
 - On what predicates?
 - How costly is this index to use?
 - Leaf page scan estimate
 - Visibility into total index filtering estimate
 - Index qualify more / less rows than estimate?
 - More statistics, REOPT?

Fetch details

- Identify stage of predicates
- Row estimates
 - Stage 1
 - Stage 2
- Prefetch indicator
- Page Range indicator

Fetch details

Table access
(# fetches)



Fetch details

The screenshot shows a window titled "Descriptor - [6] fetch". The top part is a tree view showing a folder "fetch" containing "Stage 1 Predicates" and "Stage 2 Predicates". Under "Stage 1 Predicates", there is a selected item "O.O_ORDERPRIORITY='1-URGENT'". Below the tree view, there is a checkbox "Show attribute explanation" which is checked, and a "Views:" dropdown menu set to "cost estimation". The main area is a table with two columns: "Name" and "Value".

Name	Value
Input Cardinality	1500
Scanned Rows	1500
Stage 1 Predicates	Filter Factor
O.O_ORDERPRIORITY='1-URGENT'	0.2
Stage 1 Returned Rows	300
Stage 2 Predicates	Filter Factor
(O.O_ORDERKEY,O.O_ORDERSTATUS...	7.786864443914965E-6
Stage 2 Returned Rows	0.0023
Output Cardinality	0.0023
Stage 1 Columns	5
Page Range	
Prefetch	L

Below the table, there is a section for "Attribute explanation:" with the text "Input Cardinality: Number of input rows". At the bottom of the window, there are buttons for "Save As...", "Print...", "Suggesti...", "Help", and "Close".

Predicates

- Stage 1
- Stage 2

Access information

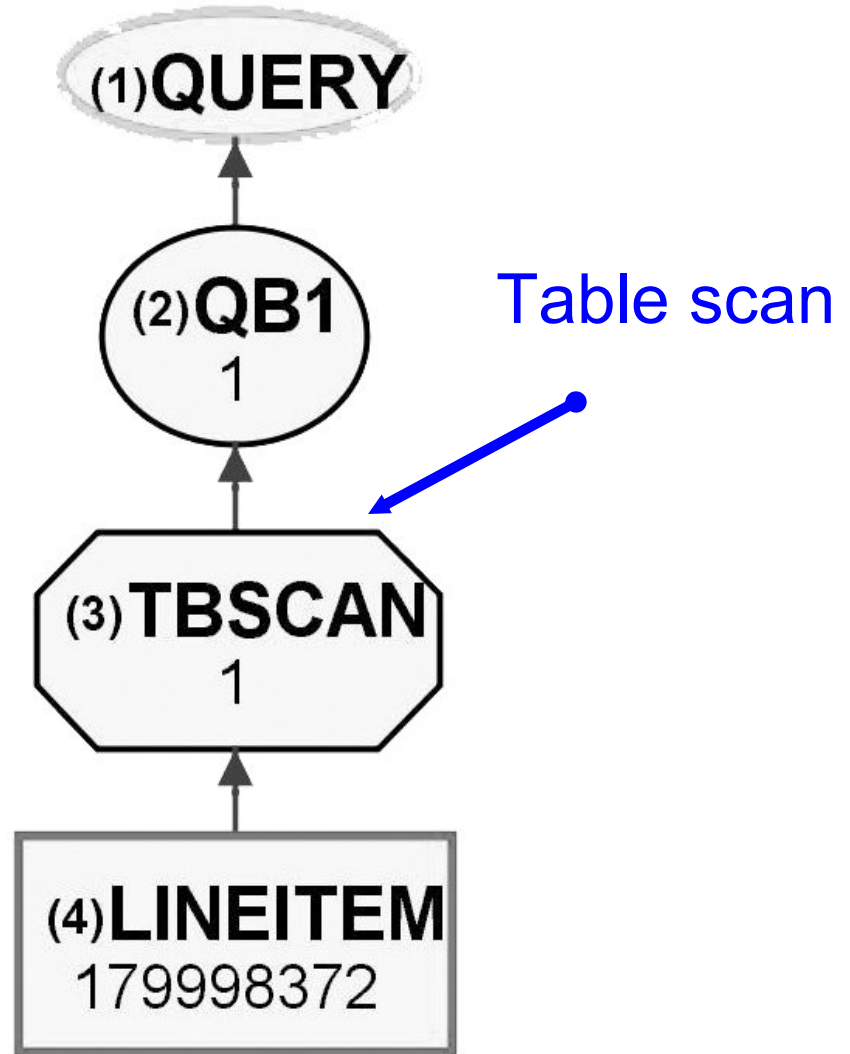
- Stage 1 rows
- Stage 2 rows
- Result rows
- Page range flag
- Prefetch flag

Page range example

- Simple table space scan graph
- Page range information
 - Number of page ranges
 - Number of qualified partitions
 - Page range details

Page range graph

```
SQL
SELECT COUNT(*)
FROM LINEITEM
WHERE L_ORDERKEY > 12000000
AND L_ORDERKEY < 15000000
AND L_TAX < 0.2
```



Page range details

The screenshot shows a window titled "Descriptor - [3] rscan". The tree view on the left shows a folder "rscan" containing "Stage 1 Predicates" and "Page Ranges", with "range1" selected under "Page Ranges". The main area displays a table of statistics under the view "cost estimation".

Name	Value
Input Cardinality	2998597
Scanned Rows	2998597
Stage 1 Predicates	Filter Factor
LINEITEM.L_ORDERKEY>12000000	0.9333
LINEITEM.L_ORDERKEY<15000000	0.0833
LINEITEM.L_TAX<0.2	0.8889
Stage 1 Returned Rows	1.2441756E7
Stage 2 Returned Rows	1.2441756E7
Output Cardinality	1
Stage 1 Columns	0
Prefetch	S
Page Range	Y
Qualified Partitions	1
Total Partitions	60
Qualified Partition Range	5

Attribute explanation:
Input Cardinality: Number of input rows

Buttons: Save As..., Print..., Suggesti..., Help, Close

Page range details

Page Range info

- Flag
- # partitions qualified
- Specific partitions

Page range details

The screenshot shows a window titled "Descriptor - [3] rscan". The tree view on the left shows a folder structure: "rscan" containing "Stage 1 Predicates" and "Page Ranges", with "range1" selected under "Page Ranges". Below the tree, there is a checkbox for "Show attribute explanation" (checked) and a "Views:" dropdown menu set to "cost estimation". A table displays the following data:

Name	Value
Qualified Partitions	1
First Partition	5
Last Partition	5

Below the table, there is a text field for "Attribute explanation:" and a label "Input Cardinality: Number of input rows". At the bottom, there are five buttons: "Save As...", "Print...", "Suggesti...", "Help", and "Close".

Per range details

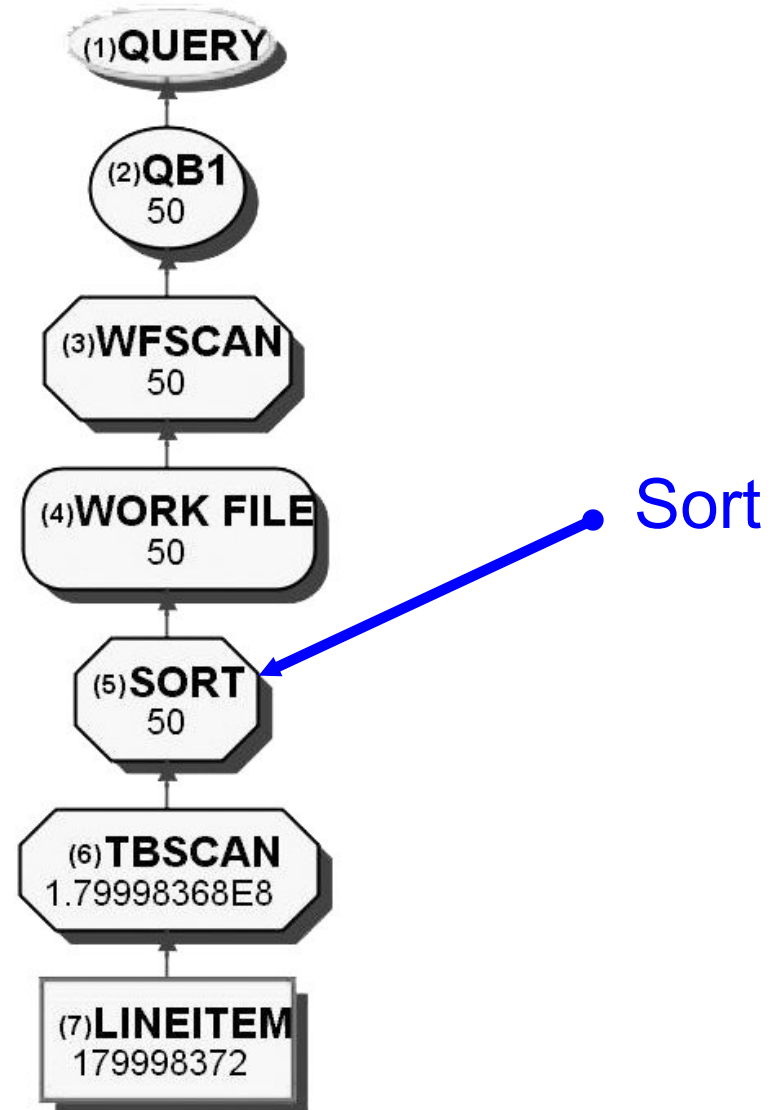
- # partitions in range
- First partition
- Last partition

Sort example

- Simple table space scan + sort
- Sort information
 - Show sort key
 - Sort key length
 - Sorted record length
- Why is this useful?
 - Input to workfile sizing
 - Insight into sort costing

Sort example

```
SQL
SELECT L_QUANTITY, MIN(L_TAX)
FROM LINEITEM
GROUP BY L_QUANTITY
```



Sort details

The screenshot shows a window titled "Descriptor - [5] sort". The top section is a tree view showing a folder "sort" containing a sub-folder "Sortkeys", which in turn contains a file "LINEITEM.L_QUANTITY". Below the tree view, there is a checkbox labeled "Show attribute explanation" which is checked, and a "Views:" dropdown menu set to "All". The main area of the window contains a table with two columns: "Name" and "Value". The table lists various sort-related metrics. At the bottom of the window, there is a section for "Attribute explanation:" with the text "Input Cardinality: Number of input rows". The bottom-most part of the window features a row of five buttons: "Save As...", "Print...", "Suggesti...", "Help", and "Close".

Name	Value
Input Cardinality	1.79998368E8
Output Cardinality	50
Pages	1975591
Runs	16079
Merges	1
Columns	3
Record Size	24
Key Size	8

← Sort keys, details

Sort information

- Estimated rows
- Estimated pages
- Sort key length
- Sort record length
- ...

Nested loop join details

- Simple nested loop join
- Outer table access
 - Nodes showing single table access
- Inner table access
 - Nodes showing how inner table accessed
- Join information
 - Join columns
 - Result cardinality

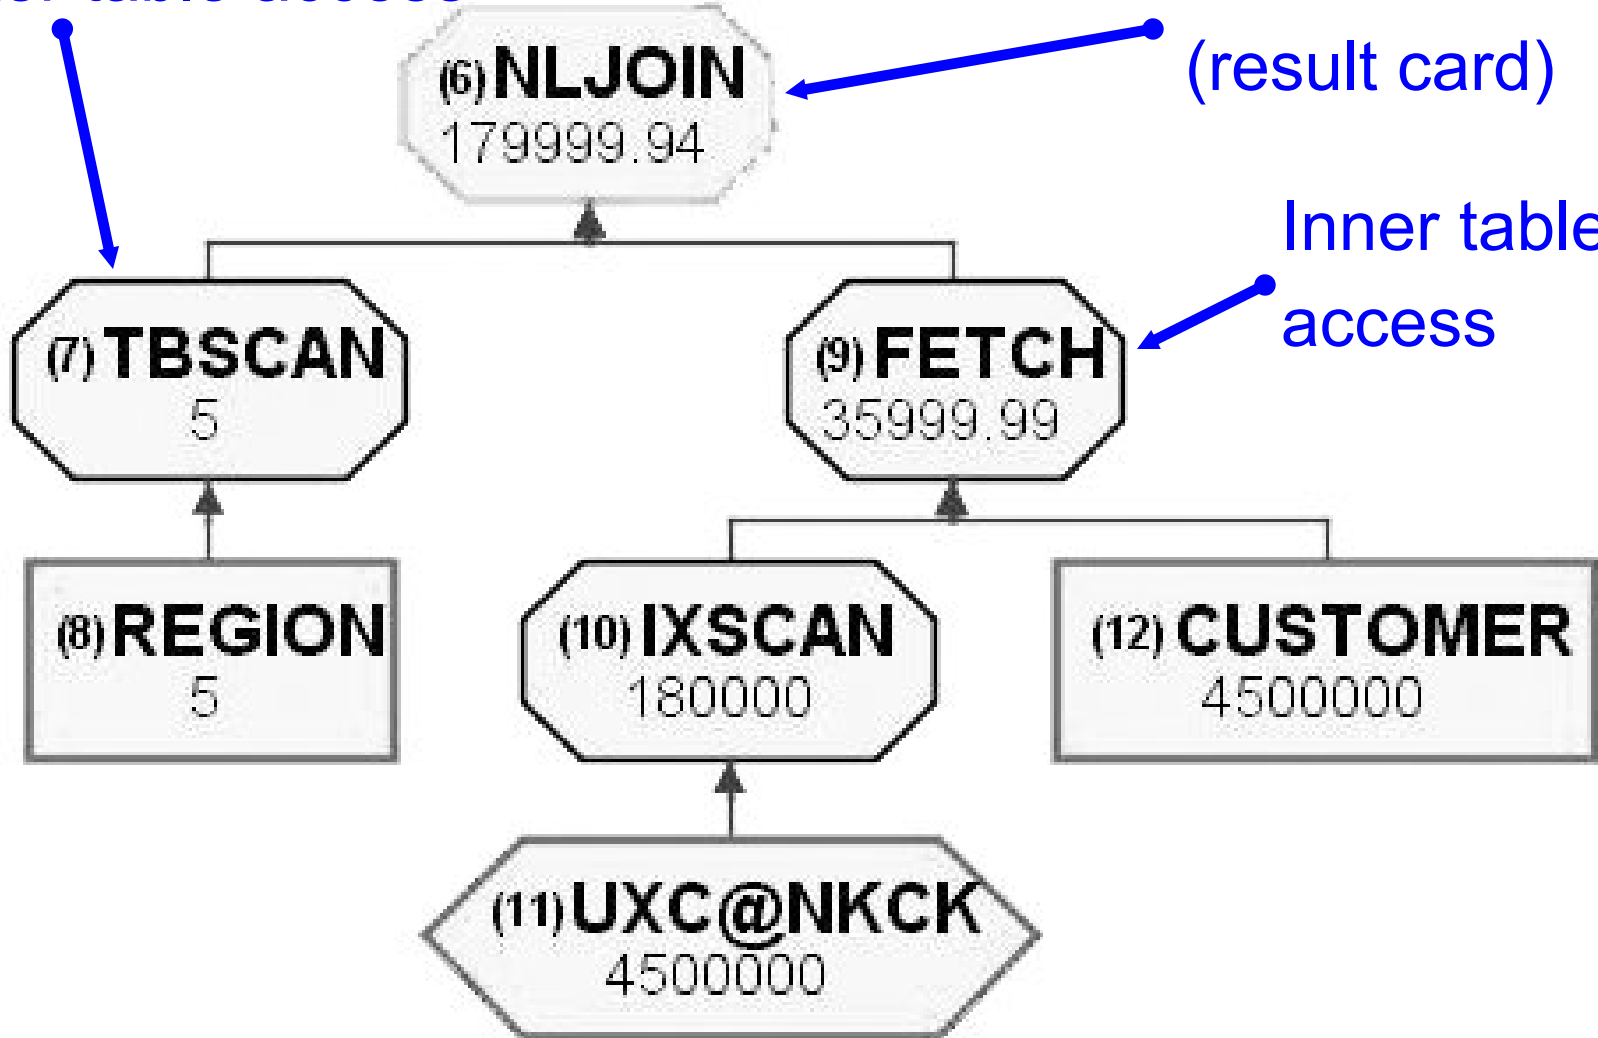
Nested loop join graph

Outer table access

Join

(result card)

Inner table access



Nested loop join details

Descriptor - [6] join

join

- Join Predicates
 - CUSTOMER.C_NATIONKEY=REGION.R_REGIONKEY

Show attribute explanation Views: cost estimation

Name	Value
Outer Input Cardinality	5.0
Inner Input Cardinality	35999.99
Join Predicates	Filter Factor
CUSTOMER.C_NATIONKEY=REGION.R_REGIONKEY	0.04
Output Cardinality	179999.94

Attribute explanation:

Outer Input Cardinality: Number of rows from outer table data stream

Save As... Print... Suggestions Help Close

Join column(s)

Join information

- Outer cardinality
- Result cardinality
- Join predicates

New information review

- Estimated number of records
 - Know single table qualified row estimates
 - What tables have worst estimates?
 - Can affect join sequence selected.
 - Join size estimate
 - Over estimation early in query can cause problems later with join sequence, join method

New information review (cont.)

- Predicate information
 - Stage of predicate application
 - Matching
 - Screening
 - Non-indexed stage 1
 - Non-indexed stage 2
 - Filter factor estimation
 - Single predicate filter factor
 - Applicable bounds also apply

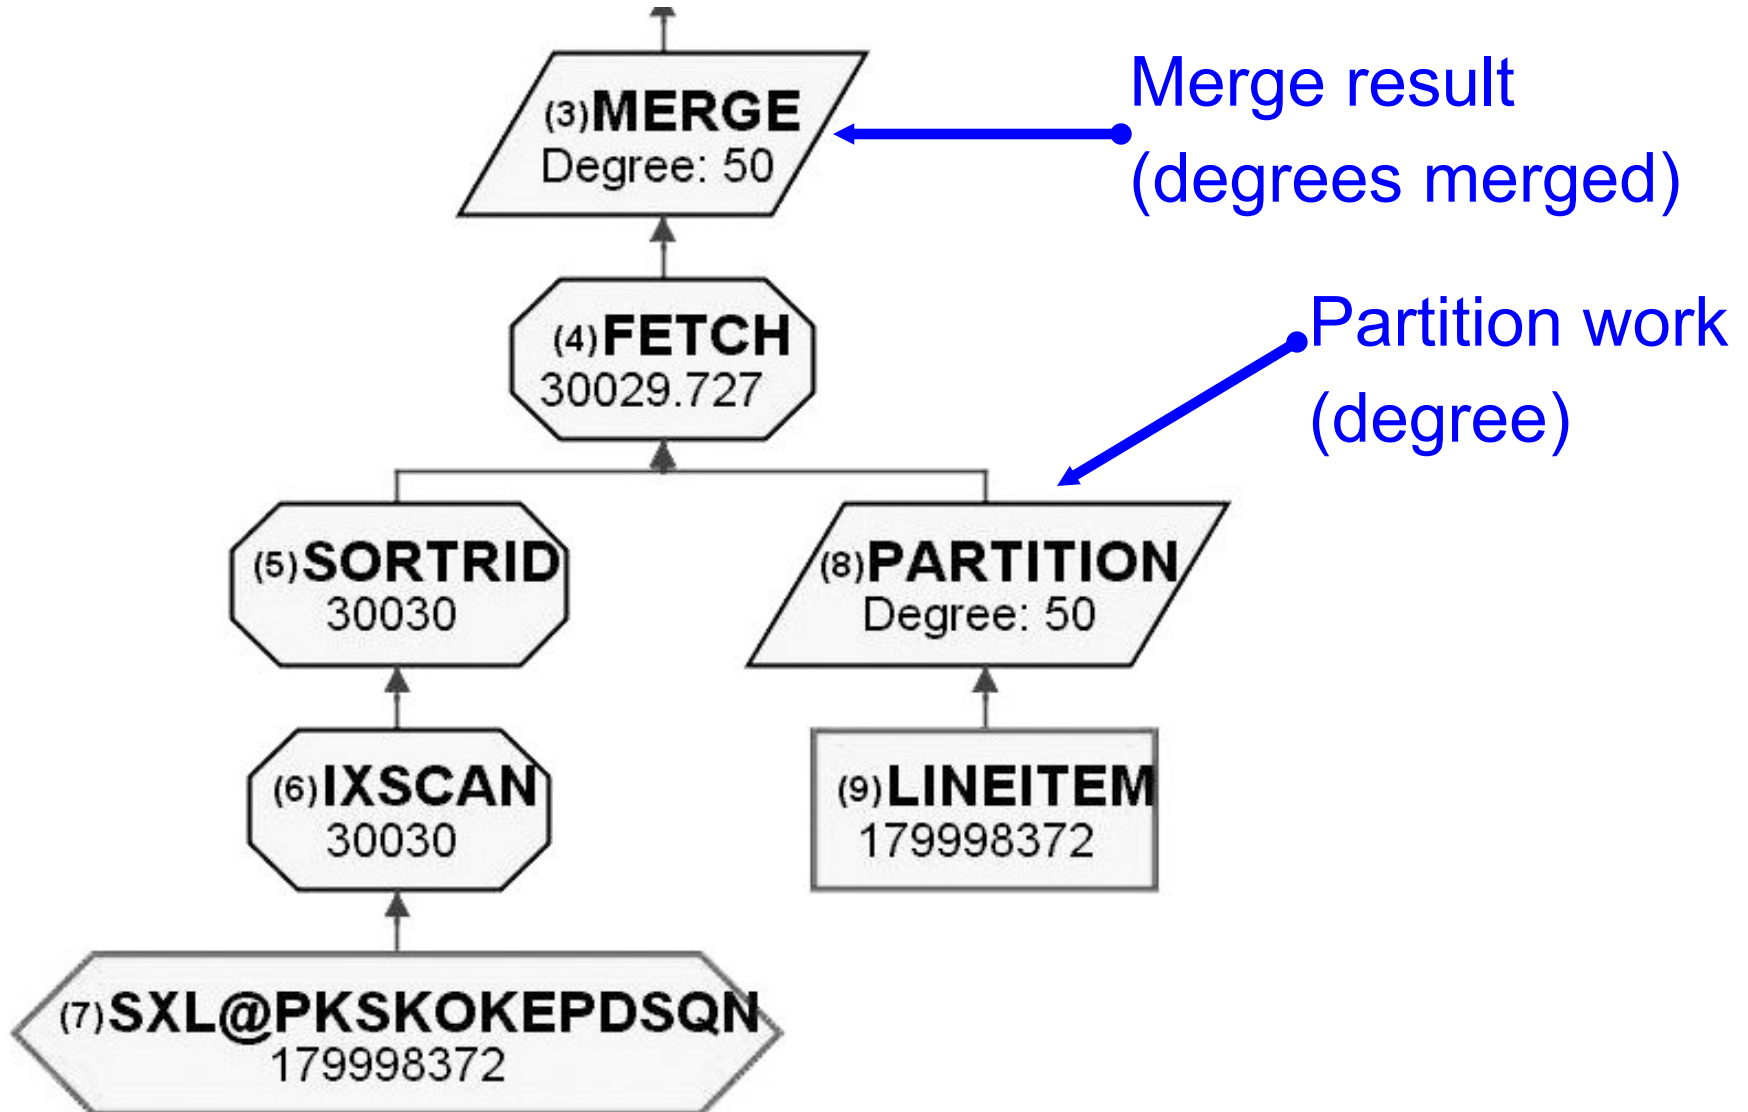
New information review (cont.)

- Limited partition scan information
 - What partitions are scanned?
 - How many page ranges?
 - How many partitions in specific range?
- Sort information
 - Sort key columns
 - Sort key length
 - Sort record length
 - Estimated sort records
 - Estimated pages scanned

Parallelism details

- Single table access parallelism graph
 - Partitioning of work node
 - Merge node
 - Degrees of parallelism
- Parallelism details
 - Mode of parallelism (I/O, CPU, SYSPLEX)
 - Type of parallelism
 - Page range
 - Key range
 - Parallel task details

Single table parallelism graph



Single table parallelism details

The screenshot shows a window titled "Descriptor - [8] parallepart". At the top, there is a tree view under "Parallel Tasks" containing four items: "paralleltask 1", "paralleltask 2", "paralleltask 3", and "paralleltask 4". Below this is a section with a checked "Show attribute explanation" checkbox and a "Views: All" dropdown menu. The main area is a table with the following data:

Name	Value
Parallel Group ID	1
Low Key	
High Key	
Key Columns	
First Page	2
Last Page	1048575
Parallel Tasks	50
Partition Type	P
Input and Output Style	RIRO
Partitioned Table	LINEITEM
Parallelism Mode	C

At the bottom, there is an "Attribute explanation:" section with the text: "Partition Type: Indicator for how partitioning should be done (K:Key range; P:Page range)". The window has buttons for "Save As...", "Print...", "Suggesti...", "Help", and "Close".

Parallel task details

Parallelism details

- Mode of parallelism
- Type of parallelism
- Parallel tasks
- Pages accessed

Parallel task details

The screenshot shows a window titled "Descriptor - [8] parallepart". It features a tree view on the left under "Parallel Tasks" with four items: "paralleltask 1", "paralleltask 2", "paralleltask 3" (selected), and "paralleltask 4". Below the tree is a checkbox for "Show attribute explanation" and a "Views:" dropdown set to "cost estimation". A table displays task details for the selected task. The table has two columns: "Name" and "Value". The rows are: "Parallel Group ID" (1), "Parallel Task Number" (3), "Key Column", "Low Key", "High Key", "Key Data Type", "Key Precision", "Key Scale", "Low Page" (39934), "High Page" (59965), "Low Partition" (3), and "High Partition" (4). At the bottom, there is an "Attribute explanation:" section with the text "Parallel Group ID: Parallel group ID" and a row of buttons: "Save As...", "Print...", "Suggesti...", "Help", and "Close".

Name	Value
Parallel Group ID	1
Parallel Task Number	3
Key Column	
Low Key	
High Key	
Key Data Type	
Key Precision	
Key Scale	
Low Page	39934
High Page	59965
Low Partition	3
High Partition	4

• List of tasks

• Parallel task details

• Page ranges for task

• or Key range for task

• Partitions for task

Parallelism review

- More specific information on parallelism
- Questions answered:
 - Mode of parallelism?
 - How many degrees?
 - Divided on page range or key range?
 - What range is each parallel task accessing?
 - What partitions is each task accessing?

Visual Explain for DB2 z/OS V8

- Visual Explain overview
 - Basic capabilities
 - Service SQL function
- Enhanced capabilities of Visual Explain
 - Qualified rows estimate
 - Wealth of predicate information
 - Limited partition scan information
 - Parallelism details
- Statistics Advisor

Statistics Advisor (SA)

- Statistics Advisor
 - Tool to automate the determination of statistics for a specific query
 - Input SQL statement
 - Output RUNSTATS commands

Why Statistics Advisor?

- Common problem
 - Most common cause of poor SQL performance
 - Insufficient statistics leads to inaccurate cost estimate
 - Efficient and inefficient access path estimates too close
 - "Did you run RUNSTATS?"
 - Question too vague

Why Statistics Advisor?

- Solution
 - **“Have you collected the right set of statistics, are they consistent and accurate?”**
 - Table space, table, index statistics
 - Column statistics
 - All columns used as predicates
 - Appropriate correlation / skew statistic
 - Where correlation / skew exists
 - Where statistics useful based on predicate type
 - **Difficult / time consuming to manually do this**

Statistics Advisor capabilities

- Analyze SQL statement
 - Predicate analysis
 - Column (groups) used as predicates
 - Type of predicates
 - Statistics analysis
 - Missing statistics (default)
 - Conflicting statistics
 - Missing appropriate correlation / skew statistics

Statistics Advisor capabilities

- Suggest, collect, educate
 - One-click to preliminary suggestions
 - Generates RUNSTATS statements
 - Two clicks from collection
 - Button on suggestion page to run RUNSTATS
 - Uses SYSPROC.DSNUTILS stored procedure to execute RUNSTATS utility
 - Educate
 - Explanation tab describes why statistics suggested
 - Conflict tab enumerates statistics conflicts

Input SQL, Click start

The screenshot shows a software window titled "New Analysis Session". At the top, there are fields for "Query No:" (containing "1") and "SQLID:" (containing "ADMFO01"). Below these is a large text area labeled "SQL Text" containing a complex SQL query with multiple table joins and a WHERE clause. At the bottom of the window, there are several buttons: "Start", "Cancel", and three icons (document, folder, floppy disk). A blue arrow points from the text "One click to analyze" to the "Start" button.

Query No: 1 SQLID: ADMFO01

SQL Text

```
T1.X_NEW_POSTN_ID,  
T1.PR_PER_ADDR_ID,  
T1.CREATED,  
T14.FST_NAME  
FROM  
  TSIEBEL.S_CONTACT T1  
  INNER JOIN TSIEBEL.S_POSTN T2 ON  
T1.PR_POSTN_ID = T2.ROW_ID  
  INNER JOIN TSIEBEL.S_POSTN_CON T3 ON  
T1.PR_POSTN_ID = T3.POSTN_ID AND T1.ROW_ID = T3.CON_ID  
  LEFT OUTER JOIN TSIEBEL.S_ORG_EXT T4 ON  
T1.PR_DEPT_OU_ID = T4.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_POSTN T5 ON  
T4.PR_POSTN_ID = T5.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_POSTN T6 ON  
T1.PR_POSTN_ID = T6.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_EMPLOYEE T7 ON  
T1.EMP_ID = T7.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_EMPLOYEE T8 ON  
T5.PR_EMP_ID = T8.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_EMPLOYEE T9 ON  
T6.PR_EMP_ID = T9.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_ADDR_PER T10 ON  
T1.PR_PER_ADDR_ID = T10.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_POSTN T11 ON  
T1.X_NEW_POSTN_ID = T11.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_ASSET T12 ON  
T1.X_NEW_CUSTOMER_ID = T12.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_POSTN T13 ON  
T1.PR_POSTN_ID = T13.ROW_ID  
  LEFT OUTER JOIN TSIEBEL.S_EMPLOYEE T14 ON  
T2.PR_EMP_ID = T14.ROW_ID  
WHERE  
  ((T1.BU_ID = ?) AND  
  (T1.X_CATEGORY_FLG = ?)) AND  
  (T1.X_POS_CHG_SIGN = ?)
```

Many tables, many predicates.
Difficult analysis.

One click to analyze

Start Cancel

VE GUI Example

The screenshot shows the Tune SQL - V8EC238 application window. The interface includes a menu bar (File, Edit, Action, View), a toolbar with icons for file operations, and a main workspace divided into several sections:

- QueryNo:** 3
- SQLID:** SYSADM
- Command history category:** default
- Command history table:** A table listing various SQL queries with columns for Name, SQL Statement, and Comment. The first row (Q1) is highlighted.
- Current degree:** System default
- Current refresh age:** 0
- Current maintained table types:** None
- Table qualifier for EXPLAIN stored procedure:** (empty)
- Input a SQL statement or select one from command history:** A text area containing a complex SQL query: `SELECT SUM(N_NATIONKEY), SUM(N_NATIONKEY), SUM(N_NATIONKEY), SUM(N_NATIONKEY), SUM(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), AVG(N_NATIONKEY), COUNT(*) FROM NATION_NP`. A blue arrow points from the text "SQL Statement" to this text area.
- Messages for Execution and Explain:** A large empty text area. A blue arrow points from the text "Click to invoke SA" to the "Analyze Statistics" button below it.
- Buttons:** "Analyze Statistics", "Execute", "Explain", and "Help".
- Checkboxes:** "Automatically save SQL in current command history category" (checked).

Name	SQL Statement	Comment
Q1	SELECT EMPN...	This query was saved at Jun 7, 2004 5:07:21 PM
Q1	SELECT EMPN...	This query was saved at Jun 8, 2004 12:18:55 PM
Q1	SELECT EMPN...	This query was saved at Jun 8, 2004 12:19:20 PM
Q1	SELECT EMPN...	This query was saved at Jun 8, 2004 12:21:30 PM
Q2	SELECT QUER...	This query was saved at Jun 8, 2004 5:27:14 PM
Q2	SELECT * FRO...	This query was saved at Jun 8, 2004 5:27:31 PM
Q2	SELECT EMPN...	This query was saved at Jun 10, 2004 10:54:40 AM
Q2	SELECT EMPN...	This query was saved at Jun 11, 2004 8:39:16 AM
Q2	select count(*) f...	This query was saved at Jun 11, 2004 9:01:27 AM
Q1	select count(*) f...	This query was saved at Jun 11, 2004 9:16:22 AM
Q1	select count(*) f...	This query was saved at Jun 11, 2004 9:17:25 AM
TestSP	SELECT EMPN...	This query was saved at Jun 14, 2004 8:30:44 AM
Q2	SELECT EMPN...	This query was saved at Jun 14, 2004 8:35:01 AM
Q2	SELECT EMPN...	This query was saved at Jun 14, 2004 8:35:51 AM
Q2	SELECT EMPN...	This query was saved at Jun 14, 2004 8:50:08 AM
Q2	SELECT EMPN...	This query was saved at Jun 14, 2004 8:50:14 AM
Q3	SELECT EMPN...	This query was saved at Jun 16, 2004 9:25:12 AM
Q3	SELECT EMPN...	This query was saved at Jun 16, 2004 9:25:17 AM
Q2	SELECT NAME,...	This query was saved at Jun 18, 2004 7:40:33 AM
Q2	SELECT NAME,...	This query was saved at Jun 18, 2004 7:43:38 AM

Suggestions for one Siebel query

The screenshot shows the DB2 Statistics Advisor Beta interface. The window title is "DB2 for OS/390 and z/OS Statistics Advisor Beta [DB2] v14ec004.svl.ibm.com [UserID] ADMF001 [SQLID] ADMF001". The interface includes a menu bar with "File" and "Tools", a toolbar with icons for SQL, tables, and search, and a main text area with tabs for "Tasks", "Explanation", and "Conflict Report".

The main text area contains the following SQL statements:

```
/* RUNSTATS statements =>
RUNSTATS TABLESPACE SIBDS020.SIBSS020
  TABLE(TSIEBEL.S_POSTN)
  INDEX(TSIEBEL.S_POSTN_P1,TSIEBEL.S_POSTN_V6,
        TSIEBEL.S_POSTN_V5,TSIEBEL.S_POSTN_F4,
        TSIEBEL.S_POSTN_M50,TSIEBEL.S_POSTN_V4,
        TSIEBEL.S_POSTN_F3,TSIEBEL.S_POSTN_V3,
        TSIEBEL.S_POSTN_U1,TSIEBEL.S_POSTN_F2,
        TSIEBEL.S_POSTN_V2,TSIEBEL.S_POSTN_V1,
        TSIEBEL.S_POSTN_M3,TSIEBEL.S_POSTN_M2,
        TSIEBEL.S_POSTN_M1)
SHRLEVEL CHANGE REPORT YES

RUNSTATS TABLESPACE SIBDS021.SIBSS021
  TABLE(TSIEBEL.S_POSTN_CON)
  INDEX(TSIEBEL.S_POSTN_CON_M2,TSIEBEL.S_POSTN_CON_M1,
        TSIEBEL.S_POSTN_CON_M50,TSIEBEL.S_POSTN_CON_P1,
        TSIEBEL.S_POSTN_CON_F50,TSIEBEL.S_POSTN_CON_U1 KEYCARD)
SHRLEVEL CHANGE REPORT YES

RUNSTATS INDEX (TSIEBEL.S_ADDR_PER_F50,TSIEBEL.S_ADDR_PER_M5,TSIEBEL.S_ADDR_PER_M4,
                TSIEBEL.S_ADDR_PER_M3,TSIEBEL.S_ADDR_PER_M2,TSIEBEL.S_ADDR_PER_M1,
                TSIEBEL.S_ADDR_PER_P1,TSIEBEL.S_ADDR_PER_U1,TSIEBEL.S_ADDR_PER_F2)
SHRLEVEL CHANGE REPORT YES

/* DSTATS statements =>
CARDINALITY IFLOW
VALUES 10,0
TSIEBEL.S_CONTACT.X_POS_CHG_SIGN
```

Two blue arrows point to the SQL statements. The first arrow points to the first two statements and is accompanied by the text "Targeted statistics based on SQL". The second arrow points to the "Execute RUNSTATS..." button at the bottom of the window and is accompanied by the text "Two clicks to collect statistics".

Explanation of suggestions

DB2 for OS/390 and z/OS Statistics Advisor Beta [DB2] v14ec004.svl.ibm.com [UserID] ADMF001 [SQLID] ADMF001

File Tools

Express Expert

Tasks Explanation Conflict Report

Preliminary Report

Statistics Analysis:

Table TSIEBEL.S_CONTACT (2988092) has conflicting statistics.

1 Column X_POS_CHG_SIGN (2) has low column cardinality relative to table cardinality.
Experience shows that columns with low cardinality relative to table cardinality are more likely to be skewed.

Table TSIEBEL.S_POSTN (6630) has conflicting statistics.

Index TSIEBEL.S_POSTN_M50 (11621) has conflicting statistics.
Index TSIEBEL.S_POSTN_U1 (6633) has conflicting statistics.

Table TSIEBEL.S_POSTN_CON (144221) has conflicting statistics.

2 Column COM_ID (144275) has conflicting uniform statistics.
Join Column Group (POSTN_ID, COM_ID) has missing cardinality.

3 Table TSIEBEL.S_ADDR_PER (0):

Index TSIEBEL.S_ADDR_PER_F50 (-1) has default statistics.
Index TSIEBEL.S_ADDR_PER_M5 (-1) has default statistics.
Index TSIEBEL.S_ADDR_PER_M4 (-1) has default statistics.
Index TSIEBEL.S_ADDR_PER_M3 (-1) has default statistics.
Index TSIEBEL.S_ADDR_PER_M2 (-1) has default statistics.

Explain why suggestion made (educate)

Statistics Advisor (cont.)

1 Suggesting frequency statistic

- COLCARDF low relative to CARDF
- Columns often skewed (eg. indicator values, flags)
- Table cardinality 2,988,092
- Column cardinality 2

2 Join column group missing cardinality

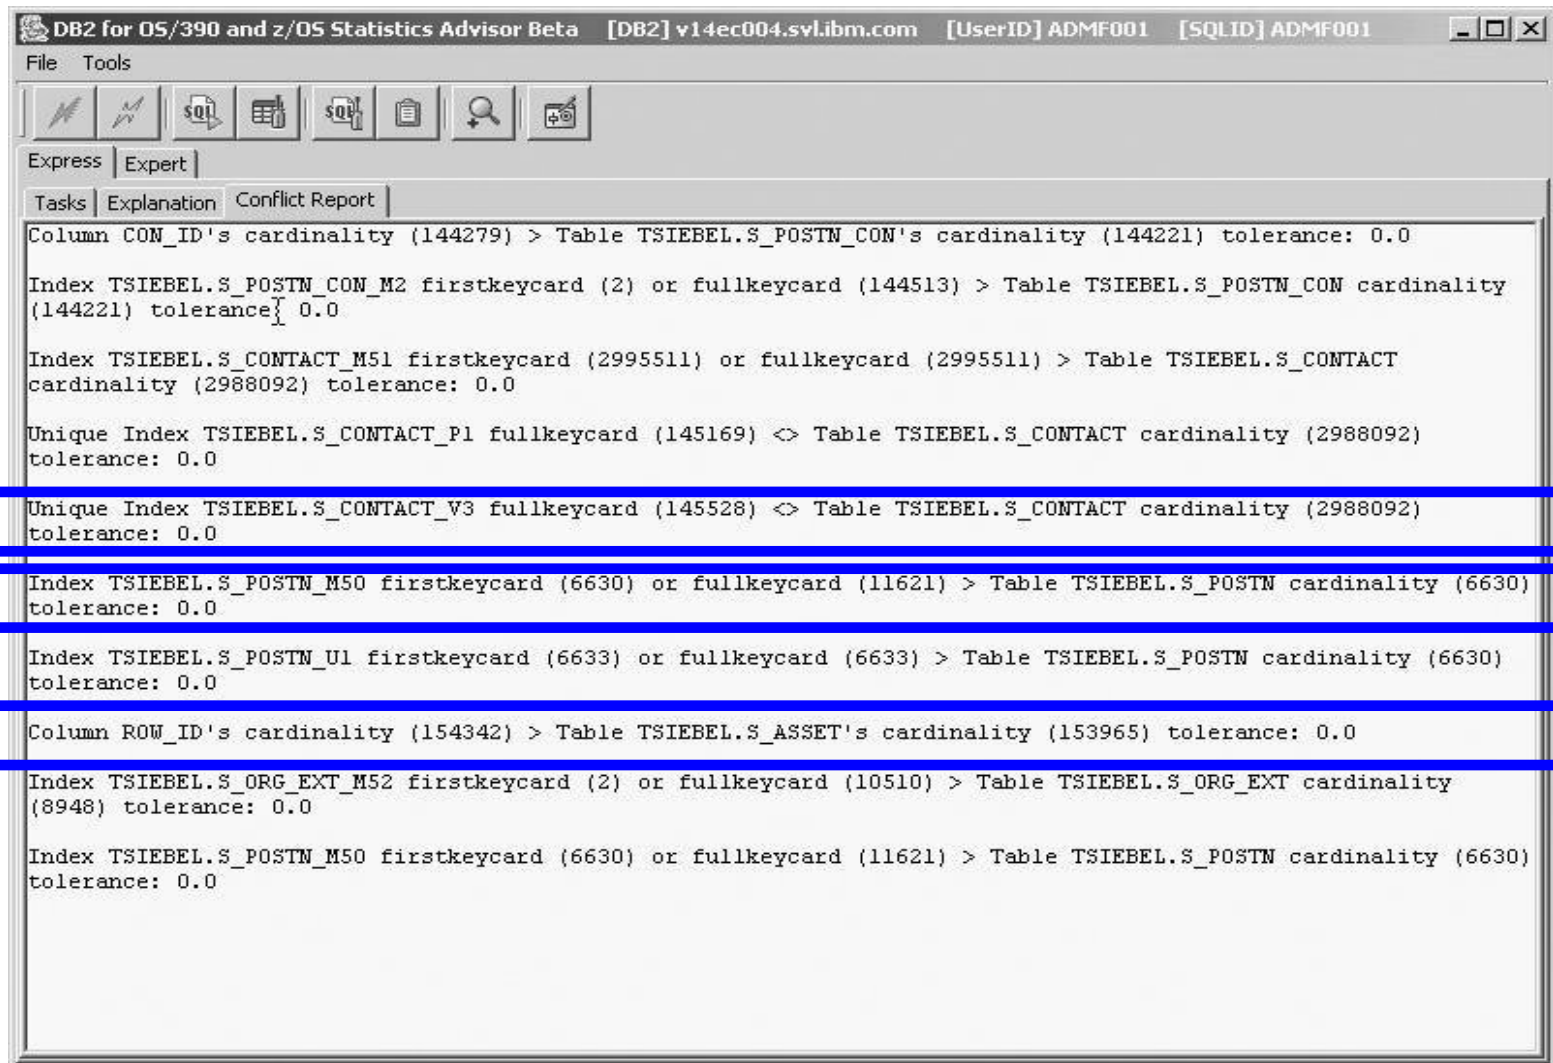
- Join on non-leading indexed columns
- No lower bound often leads to over-estimation of filtering
- This condition can be very difficult to detect manually
 - Complicated MCARD statistics
 - Analyze columns in every index on every table based on join predicates...

Statistics Advisor (cont.)

3 Default statistics

- Table has CARDF of 0, likely incorrect
- Columns have COLCARDF of -1

Details of conflicting statistics



The screenshot shows the DB2 Statistics Advisor Beta interface. The window title is "DB2 for OS/390 and z/OS Statistics Advisor Beta [DB2] v14ec004.svl.ibm.com [UserID] ADMF001 [SQLID] ADMF001". The interface includes a menu bar with "File" and "Tools", a toolbar with icons for SQL, tables, and search, and a main display area with tabs for "Express", "Expert", "Tasks", "Explanation", and "Conflict Report". The "Conflict Report" tab is active, displaying a list of statistics conflicts. Three entries are highlighted with blue boxes and numbered 1, 2, and 3 on the left side of the image.

```
Column CON_ID's cardinality (144279) > Table TSIEBEL.S_POSTN_CON's cardinality (144221) tolerance: 0.0

Index TSIEBEL.S_POSTN_CON_M2 firstkeycard (2) or fullkeycard (144513) > Table TSIEBEL.S_POSTN_CON cardinality (144221) tolerance: 0.0

Index TSIEBEL.S_CONTACT_M51 firstkeycard (2995511) or fullkeycard (2995511) > Table TSIEBEL.S_CONTACT cardinality (2988092) tolerance: 0.0

Unique Index TSIEBEL.S_CONTACT_P1 fullkeycard (145169) <> Table TSIEBEL.S_CONTACT cardinality (2988092) tolerance: 0.0

1 Unique Index TSIEBEL.S_CONTACT_V3 fullkeycard (145528) <> Table TSIEBEL.S_CONTACT cardinality (2988092) tolerance: 0.0

2 Index TSIEBEL.S_POSTN_M50 firstkeycard (6630) or fullkeycard (11621) > Table TSIEBEL.S_POSTN cardinality (6630) tolerance: 0.0

Index TSIEBEL.S_POSTN_U1 firstkeycard (6633) or fullkeycard (6633) > Table TSIEBEL.S_POSTN cardinality (6630) tolerance: 0.0

3 Column ROW_ID's cardinality (154342) > Table TSIEBEL.S_ASSET's cardinality (153965) tolerance: 0.0

Index TSIEBEL.S_ORG_EXT_M52 firstkeycard (2) or fullkeycard (10510) > Table TSIEBEL.S_ORG_EXT cardinality (8948) tolerance: 0.0

Index TSIEBEL.S_POSTN_M50 firstkeycard (6630) or fullkeycard (11621) > Table TSIEBEL.S_POSTN cardinality (6630) tolerance: 0.0
```

Conflicting statistics

1 Unique Index on table

- FULLKEYCARDF should equal CARDF
- FULLKEYCARDF = 145,528
- CARDF = 2,988,092

2 Index on table

- FULLKEYCARDF should be \leq CARDF for table
- FULLKEYCARDF = 11,621
- CARDF = 6630

Conflicting statistics (cont.)

- Some marginal cases
 - 3 Column on table
 - COLCARDF should be \leq CARDF for table
 - COLCARDF = 154,342
 - CARDF = 153,965
 - This discrepancy is close enough to not cause concern
 - Tolerance for error is configurable

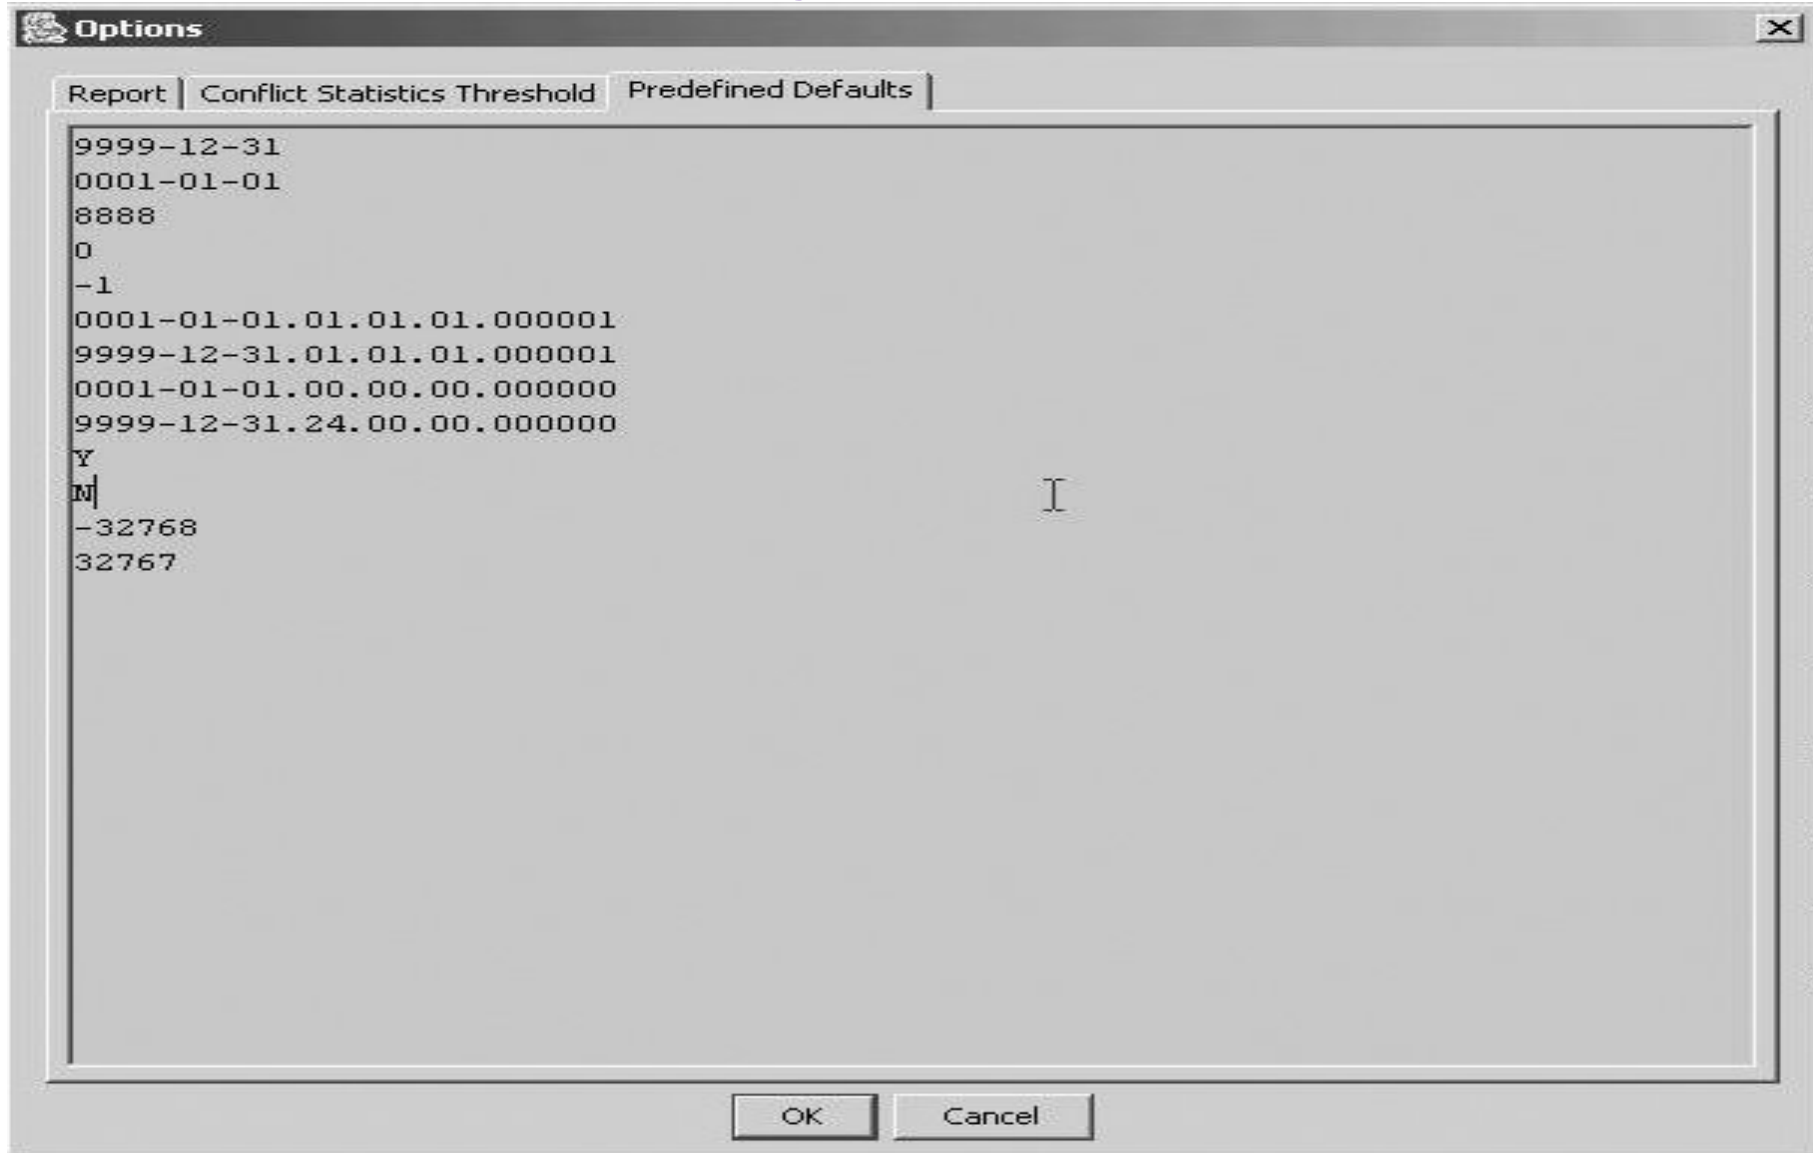
Statistics Advisor checks

- Checks highlighted / listed are only a sampling....
- Statistics analysis
 - Missing /conflicting statistics (default)
 - Table space, table, index statistics
 - Column statistics
 - Frequency statistics
 - Appropriate skew
 - Low COLCARDF relative to CARDF
 - Searching for null, default, typical default
 - Appropriate correlation
 - Join column groups
 - High COLCARDF columns

SA Configurable

- Configurable
 - User configurable "typical defaults"
 - Add typical defaults to make SA smarter
 - Columns with high COLCARDF often primarily skewed on single default value
 - Eg. Blanks, NULL, '9999-12-31'
 - Drive collection, drive collection of single value
 - Conflict thresholds also configurable
 - Small difference in some statistics not that bad
 - Ignore insignificant differences
 - Draw focus to significant discrepancies

Add / remove typical default values



Configurable thresholds

The screenshot shows a window titled "Options" with a close button (X) in the top right corner. The window has three tabs: "Report", "Conflict Statistics Threshold", and "Predefined Defaults". The "Conflict Statistics Threshold" tab is selected. The main area contains a list of 12 threshold settings, each with a text label and a numeric input field. All input fields are currently set to "0.0". At the bottom of the window are "OK" and "Cancel" buttons.

Threshold Label	Value
TABCARD < COLCARD	0.0
TABCARD < INDEX KEYCARD	0.0
TABCARD <> UNIQUE INDEX FULLKEYCARD	0.0
MCARD (C1,C2,C3) > MCARD (C1,C2,C3,C4)	0.0
MCARD (C1,C2,C3) > TABCARD	0.0
INDEX FIRSTKEYCARD > FULLKEYCARD	0.0
1-COL INDEX FIRSTKEYCARD <> FULLKEYCARD	0.0
SUM(FREQUENCY) > 1	0.0
MULTIPLE MCARD SOURCES FROM INDEXES	0.0
MULTIPLE MCARD SOURCES FROM INDEXES AND SYSCOLDIST	0.0
COUNT(FREQUENCY) > TABCARD	0.0
COLCARD(C1) * COLCARD(C2) < MCARD (C1,C2)	0.0
High colcard columns recollection threshold	0.0

SA Commentary

- Automated statistics determination
 - Often query may have inefficient OR unstable performance due to lack of statistics
 - SA automates the analysis of what statistics are required for an SQL statement
- Goal
 - Automate the SOLUTION to many common SQL performance problems
 - Solve SQL performance problems
 - Fast & easy

Visual Explain for DB2 z/OS V8

- Visual Explain overview
 - Basic capabilities
 - Service SQL function
- Enhanced capabilities of Visual Explain
 - Qualified rows estimate
 - Wealth of predicate information
 - Limited partition scan information
 - Parallelism details
- Statistics Advisor



Thank you for attending!!!

Session Z32

DB2 for z/OS Visual Explain V8

Patrick Bossman

E-mail: bossman@us.ibm.com

<http://www.ibm.com/software/data/db2/zos/osc/ve/index.html>

**IBM DB2 Information Management
Technical Conference**

Sept. 20-24, 2004

Las Vegas, NV