precisely

Syncsort Optimize DB2

(EZ-DB2)

Trend Analysis Tables Reference

Release 9.10

Version 9.1.2

July 2018

Syncsort DB2 Trend Analysis Tables Reference
(formerly EZ-DB2 Tables Reference Guide)

NOTICE

Contents

THIS PAGE IS INTENTIONALLY LEFT BLANK (if necessary)

Trend Analysis Database

The Trend Analysis Database is an optional database that can be defined and populated following the completion of an EZ-Tracer or EZ-Cache trace.

The database allows the use to query and produce analytics on SQL execution metrics across an extended period of time (depending on the frequency and setup of the EZ-DB2 tracing).

TRC_SUMMARY

One row for each Trace executed for each workload.

INDEXED BY

XTTTSPD1 ON (WORKLOAD,TRACE_NO)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
WORKLOAD_TYPE	CHAR(1) NOT NULL WITH DEFAULT 'U'	Workload Type: C – EZ-Cache T – EZ-Tracer U - Unknown
TADB_VERSION	SMALLINT NOT NULL WITH DEFAULT 1	Trend DB Version when trace loaded: 1 – Original TADB (V1) 2 – Enhanced TADB (V2)
TRACE_SUBSYS	CHAR(4) NOT NULL WITH DEFAULT	Db2 Subsystem traced
TRACE_SUBSYS_REL	CHAR(4) NOT NULL WITH DEFAULT	Db2 Release when traced
TRACE_SUBSYS_MOD	CHAR(3) NOT NULL WITH DEFAULT	Db2 Modification Level when traced (Db2 12)
TRACE_DATE	DATE NOT NULL	The start date of the trace.
TRACE_TIME	TIME	The start time of the trace.

	NOT NULL	
TRACE_DURATION	SMALLINT NOT NULL	The trace duration of the trace in minutes.

TRC_SUMMARY_LIST

One row for every Program, Plan, Collection, Authid, Table etc. encountered.

If the TRACE_NO is zero, then the statistics are the summarized totals for the applicable TYPE for ALL traces executed under a particular Workload

INDEXED BY

XTTSLPD1 ON (WORKLOAD,TRACE_NO,TYPE,NAME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
TYPE	CHAR(3) NOT NULL	Identifies the entity type to which this row relates:- DBN - Data relates to a Database TAB - Data relates to a Table PRG - Data relates to a Program PLN - Data relates to a Plan AUT - Data relates to an Authid COL - Data relates to a Collection
NAME	VARCHAR(128) NOT NULL	The Entity name. See Type above.
SQL	FLOAT NOT NULL	The number of distinct SQL statements executed against this entity.
SQL FREQ	FLOAT NOT NULL	The SQL frequency of all of the SQL for this entity.
CPU	FLOAT NOT NULL	The total CPU for all SQL for this entity.
CPU_PCT	SMALLINT NOT NULL	The Percentage CPU for this object compared to total for all objects for this entity type.
DBN_TABLES	FLOAT	If the entry is a Database then the number of tables referenced in this database
PRG_COLLID	VARCHAR(128)	If the entry is program then then this

		field contains the COLLID (if it is a package).

TRC_SQL_TEXT

One or more rows for each SQL statement executed for each workload.

INDEXED BY

XTSTXPD1 ON (WORKLOAD,SQL_HASH_KEY,STMT_SEQ_NO)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
SQL_HASH_KEY	CHAR (10) NOT NULL	The SQL Hash Key. This is an internally generated key used to uniquely identify matching SQL statements.
SELLIST_HASH_KEY	VARCHAR(10) NOT NULL	The SQL Select List Hash Key. This is an internally generated key used to uniquely identify SQL statements with matching Select Lists.
STMT_SEQ_NO	SMALLINT NOT NULL	The sequence number of this row where Statement text spans multiple rows. Starts at 1.
CONSOLIDATED_SQL	INTEGER NOT NULL WITH DEFAULT	Count of consolidated SQL, that are ‘like’ this SQL text
EXPLAIN_IND	CHAR(1) NOT NULL WITH DEFAULT 'Y'	Indicates whether this SQL can be explained. Y/N
SQL_TABLE_IND	CHAR(1) NOT NULL WITH DEFAULT 'N'	
ALERT_OPT_FOR	CHAR(1) NOT NULL WITH DEFAULT	Set to ‘Y’ if ‘OPTIMIZE FOR’ alert was set.
ALERT_HAS_INLIST	CHAR(1) NOT NULL WITH DEFAULT	Set to ‘Y’ if ‘SQL contains an INLIST’ alert was set.

ALERT_HAS_BETWEEN	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if 'SQL contains an BETWEEN' alert was set.
ALERT_NEW_V8_WORDS	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if SQL uses new V8 words.
ALERT_HAS_UPPER	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if 'SQL uses UPPER' alert was set.
ALERT_HAS_IN_SEL	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if 'SQL has an in subselect' alert was set.
ALERT_DISTINCT_JOIN	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if 'SQL uses a DISTINCT join' alert was set.
ALERT_ADDED_CASTS	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if 'SQL uses CAST' alert was set.
ALERT_LIT_CONCAT	CHAR(1) NOT NULL WITH DEFAULT	Set to 'Y' if 'SQL uses CONCAT' alert was set.
STMT_TEXT	VARCHAR(254) NOT NULL	The SQL statement text. If the SQL text is too long to fit in one row, there will be as many rows as required ordered by STMT_SEQ_NO.

TRC_SUMMARY_TABLES

One row for every table reference by SQL captured during a particular trace.

INDEXED BY

XTTSTPD1 ON (WORKLOAD,TRACE_NO,TABLE_CREATOR,TABLE_NAME,DBNAME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
TABLE_CREATOR	VARCHAR(128) NOT NULL	The table creator name.
TABLE_NAME	VARCHAR(128) NOT NULL	The Table name.
DBNAME	VARCHAR(24) NOT NULL	The database name.
SQL	FLOAT NOT NULL	The number of distinct SQL that reference this table.
SQL FREQ	FLOAT NOT NULL	The total frequency of all of the SQL that reference this table.
CPU	FLOAT NOT NULL	The total CPU for all SQL that reference this table.
CPU_PCT	SMALLINT NOT NULL	The percentage CPU of the SQL for this table against the total of the CPU for all tables.

TRC_SQL_TABLE

Contains one row for each table for each distinct SQL in the workload.

INDEXED BY

XTSTBPD1 ON (WORKLOAD,SQL_HASH_KEY,TABLE_CREATOR,TABLE_NAME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
SQL_HASH_KEY	CHAR (10) NOT NULL	The SQL Hash Key. This is an internally generated key used to uniquely identify matching SQL statements.
DBNAME	VARCHAR(24) NOT NULL	The applicable database name.
TABLE_CREATOR	VARCHAR(128 NOT NULL	Authorization ID of the owner of the table, view or Alias.
TABLE_NAME	VARCHAR(128) NOT NULL	Name of the table view or Alias.

TRC_SQL_TRACE

One row per SQL statement..

Table contains Summary Totals for the Workload, and each Trace for a given workload.

If TRACE_NO is zero then row represents the summary for the workload.

INDEXED BY

XTTRCPD1 ON (WORKLOAD,TRACE_NO,SQL_NO)

XTTRCAD4 ON (WORKLOAD,TRACE_NO,SQL_HASH_KEY)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
SQL_NO	INTEGER NOT NULL	Internally generated numeric identifier for the statement
SQL_HASH_KEY	CHAR(10) NOT NULL	The SQL Hash Key. This is an internally generated key used to uniquely identify matching SQL statements.
SELLIST_HASH_KEY	VARCHAR(254) NOT NULL	The SQL Select List Hash Key. This is an internally generated key used to uniquely identify SQL statements with matching Select Lists.
PROGRAM_NAME	VARCHAR(128) NOT NULL	The Program name that contains this SQL statement.
PLAN_NAME	VARCHAR(24) NOT NULL	The Plan name that the Package was bound in.
CURSOR_NAME	VARCHAR(128) NOT NULL	Name of the Cursor associated with the statement.
STATEMENT_NO	INTEGER NOT NULL	The Pre-Compiler generated statement number.
PREPARE_STMT_NO	INTEGER NOT NULL	The statement number of the related Prepare for this statement.
CONTOKEN	CHAR(8) NOT NULL	Consistency token for the package from a DBRM; the level as specified by the LEVEL option at precompile or the timestamp indicating when package

		program was precompiled.
STMT_TYPE	CHAR(6) NOT NULL	The Statement type OPEN SELECT INSERT DELETE
STATIC_DYNAMIC	CHAR(1) NOT NULL	Static/Dynamic SQL indicator:- S Statement is static SQL D Statement is Dynamic SQL.
SQL_LITERALS	CHAR(1) NOT NULL WITH DEFAULT **>N=**	
REOPT_VARS	CHAR(1) NOT NULL WITH DEFAULT **>N=**	Whether the program was bound with REOPT(VARS) Y – Yes N - No
WORKSTATION	CHAR(18) NOT NULL WITH DEFAULT	The Workstation name if available.
LAST_START_TIME	CHAR(16) NOT NULL WITH DEFAULT	The start time of the last execution of this statement in the format YYYY-MM-DD HH:MM
RID_LIST_FAILURES	INTEGER NOT NULL WITH DEFAULT 0	The number of RID LIST failures for this statement.
PREPARE_COUNT	INTEGER NOT NULL WITH DEFAULT 0	The total number of Prepares for this SQL statement in the trace.
THREAD_COUNT	INTEGER NOT NULL WITH DEFAULT 0	The total number of Threads for this SQL in the trace.
TIMERON_COST	FLOAT WITH DEFAULT 0	The total cost of the SQL in DB2 Timerons for the trace.
FREQUENCY	FLOAT WITH DEFAULT 0	The execution frequency of the SQL in the trace.
FETCH_FREQ	FLOAT WITH DEFAULT 0	The number of rows fetched by the SQL in the trace.
CPU_MS	FLOAT WITH DEFAULT 0	The CPU time of this statement in milliseconds.
ELAPSED_MS	FLOAT WITH DEFAULT 0	The Elapsed time of this statement in milliseconds.

PREPARE_CPU	FLOAT WITH DEFAULT 0	The CPU time of the Prepare in milliseconds – only dynamic SQL.
PREPARE_ELAPSED	FLOAT WITH DEFAULT 0	The Elapsed time of the Prepare in milliseconds – only dynamic SQL.
FETCH_CPU_MS	FLOAT WITH DEFAULT 0	The CPU time of the Fetch in milliseconds.
FETCH_ELAPSED_MS	FLOAT WITH DEFAULT 0	The Elapsed time of the Fetch in milliseconds.
OPEN_CPU	FLOAT WITH DEFAULT 0	The CPU time of the Open in milliseconds – only for cursor.
OPEN_CLOCK	FLOAT WITH DEFAULT 0	The Elapsed time of the Open in milliseconds – only for cursor.
CLOSE_CPU	FLOAT WITH DEFAULT 0	The CPU time of the Close in milliseconds – only for cursor.
CLOSE_CLOCK	FLOAT WITH DEFAULT 0	The Elapsed time of the Close in milliseconds – only for cursor.
ZIIP_CPU	FLOAT WITH DEFAULT 0	The CPU time executed on an IBM SE engine in milliseconds[1]
PARA_CPU	FLOAT WITH DEFAULT 0	The CPU executed on parallel processes in milliseconds.
CACHE_HITS	FLOAT WITH DEFAULT 0	Number of times this SQL was found in cache (EZ-Cache only)
CACHE_LOADS	FLOAT WITH DEFAULT 0	Number of times this SQL was loaded in cache (EZ-Cache only)
DBRM_VERSION	VARCHAR(122) NOT NULL WITH DEFAULT	First traced DBRM Version of the package that executed the SQL.
CONN_TYPE	INTEGER NOT NULL WITH DEFAULT 0	DB2 Connection Type. From QWHCATYP: 1 – TSO 2 – CAF 3 - DL/I BATCH 4 - CICS ATTACH 5 - IMS ATTACH BMP 6 - IMS ATTACH MPP 7 - DB2 PRIVATE PROTOCOL 8 - DRDA PROTOCOL 9 - IMS CONTROL REGION 10 - IMS TRANSACTION BMP 11 - DB2 UTILITIES 12 - RRSAF ATTACH

[1] For DRDA Cursor statements, SE times cannot be relied on, due to the way IBM roll up certain IFCIDs for DRDA Cursors. There is currently NO way to accurately report the SE time for an individual DRDA Cursor. Where SE Time is reported for a DRDA Cursor, this can be used to indicate that an SE engine was utilized, but will generally show a higher than actual figure used.

TRC_SQL_AUTHID

One row per SQL statement for each trace for each Workload broken down by Plan and Authid

INDEXED BY

XTXSAPD1 ON (WORKLOAD,TRACE_NO,SQL_NO,AUTHID,PLAN_NAME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
SQL_NO	INTEGER NOT NULL	Internally generated numeric identifier for the statement
AUTHID	VARCHAR(128)	The Authid that executed the statement.
PLAN NAME	VARCHAR(24) NOT NULL	The name of the Plan
CURSOR_NAME	VARCHAR(128) NOT NULL	Name of the Cursor associated with the statement.
COLLID	VARCHAR(128) NOT NULL	Collection Name.
CORRID	CHAR(12) NOT NULL	Correlation ID
LAST_START_TIME	CHAR(16) NOT NULL WITH DEFAULT	The start time of the last execution of this statement in the format YYYY-MM-DD HH:MM
RID_LIST_FAILURES	INTEGER NOT NULL WITH DEFAULT 0	The number of RID LIST failures for this statement.
PREPARE_COUNT	INTEGER NOT NULL WITH DEFAULT 0	The total number of Prepares for this SQL statement in the trace.
THREAD_COUNT	INTEGER NOT NULL WITH DEFAULT 0	The total number of Threads for this SQL in the trace.
FREQUENCY	FLOAT WITH DEFAULT 1	The execution frequency of the SQL in the trace.
FETCH_FREQ	FLOAT WITH DEFAULT 0	The number of rows fetched by the SQL in the trace.

CPU_MS	FLOAT WITH DEFAULT 0	The CPU time of this statement in milliseconds.
ELAPSED_MS	FLOAT WITH DEFAULT 0	The Elapsed time of this statement in milliseconds.
PREPARE_CPU	FLOAT WITH DEFAULT 0	The CPU time of the Prepare in milliseconds – only dynamic SQL.
PREPARE_ELAPSED	FLOAT WITH DEFAULT 0	The Elapsed time of the Prepare in milliseconds – only dynamic SQL.
FETCH_CPU_MS	FLOAT WITH DEFAULT 0	The CPU time of the Fetch in milliseconds.
FETCH_ELAPSED_MS	FLOAT WITH DEFAULT 0	The Elapsed time of the Fetch in milliseconds.
OPEN_CPU	FLOAT WITH DEFAULT 0	The CPU time of the Open in milliseconds – only for cursor.
OPEN_CLOCK	FLOAT WITH DEFAULT 0	The Elapsed time of the Open in milliseconds – only for cursor.
CLOSE_CPU	FLOAT WITH DEFAULT 0	The CPU time of the Close in milliseconds – only for cursor.
CLOSE_CLOCK	FLOAT WITH DEFAULT 0	The Elapsed time of the Close in milliseconds – only for cursor.
ZIIP_CPU	FLOAT WITH DEFAULT 0	The CPU time executed on an IBM SE engine in milliseconds[1]
PARA_CPU	FLOAT WITH DEFAULT 0	The CPU executed on parallel processes in milliseconds.
CACHE_HITS	FLOAT WITH DEFAULT 0	Number of times this SQL was found in cache (EZ-Cache only)
CACHE_LOADS	FLOAT WITH DEFAULT 0	Number of times this SQL was loaded in cache (EZ-Cache only)
DBRM_VERSION	VARCHAR(122) NOT NULL WITH DEFAULT	First traced DBRM Version of the package that executed the SQL.
CONN_TYPE	INTEGER NOT NULL WITH DEFAULT 0	DB2 Connection Type. From QWHCATYP: 1 – TSO 2 – CAF 3 - DL/I BATCH 4 - CICS ATTACH 5 - IMS ATTACH BMP 6 - IMS ATTACH MPP 7 - DB2 PRIVATE PROTOCOL 8 - DRDA PROTOCOL 9 - IMS CONTROL REGION

[1] For DRDA Cursor statements, SE times cannot be relied on, due to the way IBM roll up certain IFCIDs for DRDA Cursors. There is currently NO way to accurately report the SE time for an individual DRDA Cursor. Where SE Time is reported for a DRDA Cursor, this can be used to indicate that an SE engine was utilized, but will generally show a higher than actual figure used.

		10 - IMS TRANSACTION BMP 11 - DB2 UTILITIES 12 - RRSAF ATTACH

TRC_SQL_STATISTICS

One or more rows for each SQL statement for each workload. Rows contain Row level statistics.

INDEXED BY

XTSSTPD1 ON (WORKLOAD,TRACE_NO,SQL_NO,SQL_OPEN-FETCH,TYPE)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
SQL_NO	INTEGER NOT NULL	Internally generated numeric identifier for the statement
SQL_OPEN_FETCH	CHAR(1)NOT NULL	Statistics relate to SQL Select, Open or Fetch.
TYPE	CHAR(4) NOT NULL	The Row level statistic type
ROWS_PROCESSED	FLOAT NOT NULL	The total number of rows processed.
ROWS_LOOKED_AT	FLOAT NOT NULL	The total number of rows looked at.
ROWS_DM_QUALIFIED	FLOAT NOT NULL	The total number of rows processed by Data manager (Stage 1.
ROWS_RDS_QUALIFIED	FLOAT NOT NULL	The total number of rows processed by RDS (Stage 2).
ROWS_UPDATED	FLOAT NOT NULL	The total number of rows updated.
GETPAGES	FLOAT NOT NULL	The total number of Get Pages.

TRC_SQL_AUTH_STATS

One or more rows for each SQL statement for each workload. Rows contain Row level statistics broken down by Plan and Authid.

INDEXED BY

XTSASPD1 ON (WORKLOAD,TRACE_NO,SQL_NO,AUTHID,PLAN_NAME, SQL_OPEN_FETCH,TYPE)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
SQL_NO	INTEGER NOT NULL	Internally generated numeric identifier for the statement
AUTHID	VARCHAR(128) NOT NULL	The Authid that executed the statement.
PLAN_NAME	VARCHAR(24) NOT NULL	The name of the Plan
SQL_OPEN_FETCH	CHAR(1)NOT NULL	Statistics relate to SQL Select, Open or Fetch.
TYPE	CHAR(4) NOT NULL	The Row level statistic type
ROWS_PROCESSED	FLOAT NOT NULL	The total number of rows processed.
ROWS_LOOKED_AT	FLOAT NOT NULL	The total number of rows looked at.
ROWS_DM_QUALIFIED	FLOAT NOT NULL	The total number of rows processed by Data manager (Stage 1.
ROWS_RDS_QUALIFIED	FLOAT NOT NULL	The total number of rows processed by RDS (Stage 2).
ROWS_UPDATED	FLOAT NOT NULL	The total number of rows updated.
GETPAGES	FLOAT NOT NULL	The total number of Get Pages.

TRC_USER_NOTE

One row for each 'note' added for recording user specified comments for SQL or Workload.

SQL_NO = zero indicates note is for the workload.

INDEXED BY

XTTUNPD1 ON (WORKLOAD,TRACE_NO,SQL_NO,NOTE_TIME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
SQL_NO	INTEGER NOT NULL	SQL_NO this note relates to.
NOTE_TIME	TIMESTAMP NOT NULL	Date and time Note added.
USERID	CHAR(8) NOT NULL WITH DEFAULT	Userid adding note.
NOTE	CLOB(1M) NOT NULL	Note text.

TRC_USER_FLAG

One row for each ‘Flag’ added for recording user specified flags for SQL.

INDEXED BY

XTTUFPD1 ON (WORKLOAD,SQL_NO)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
SQL_NO	INTEGER NOT NULL	SQL_NO this note relates to.
FLAG	CHAR(1) NOT NULL	User Flag ID

TRC_IRONSTREAM

One row for each trace sent to Ironstream

INDEXED BY

XTTIRPD1 ON (WORKLOAD,TRACE_NO)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
TRACE_NO	INTEGER NOT NULL	The Trace number for this trace.
IRONSTRM_TIME	TIMESTAMP NOT NULL	Timestamp when events sent to Ironstream
IRONSTRM_SENT	CHAR(1) NOT NULL WITH DEFAULT 'N'	Indicator for data sent to Ironstream.

TRC_PLAN_TABLE

At least one row for each table referenced in each SQL statement in the current Workload being explained.

INDEXED BY

XTPLTPD1 ON (WORKLOAD,VERSION,SQL_NO,BIND_TIME,EXPLAIN_TIME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
VERSION	CHAR(1) NOT NULL WITH DEFAULT 'O'	The Cost Indicator associated with this Explain data. O - Original Costs P - Predicted Costs F - Final Costs Any other alphanumeric character identifies a User defined variation, eg. A refined index configuration or a Adelta@ version Id associated with Impact Analyzer.
BIND_TIME	CHAR(26) NOT NULL WITH DEFAULT ' '	Date and Time access path generated at BIND with EXPLAIN=YES
EXPLAIN_TIME	TIMESTAMP NOT NULL WITH DEFAULT	Date and Time access path generated from Explain
SQL_NO	INTEGER NOT NULL	Internally generated numeric identifier for the statement
QBLOCKNO	SMALLINT NOT NULL	The position of the query in the statement being explained (1 for the outermost query, 2 for the next query, and so forth).For better performance, DB2 might merge a query block into another query block. When that happens, the position number of the merged query block will not be in QBLOCKNO.
PLANNO	SMALLINT NOT NULL	The number of the step in which the query indicated in QBLOCKNO was processed. This column indicates the order in which the steps were executed.
METHOD	SMALLINT NOT NULL	A number (0, 1, 2, 3, or 4) that indicates the join method used for the

Column Name	Data Type	Description
		step: 0 First table accessed, continuation of previous table accessed, or not used. 1 Nested loop join. For each row of the present composite table, matching rows of a new table are found and joined. 2 Merge scan join. The present composite table and the new table are scanned in the order of the join columns, and matching rows are joined. 3 Sorts needed by ORDER BY, GROUP BY, SELECT DISTINCT, UNION, a quantified predicate, or an IN predicate. This step does not access a new table. 4 Hybrid join. The current composite table is scanned in the order of the join-column rows of the new table. The new table is accessed using list prefetch.
STEPCOST	FLOAT NOT NULL WITH DEFAULT 0	The cost of this step.
COMPCOST	FLOAT NOT NULL WITH DEFAULT 0	The total cost.
COMPCARD	FLOAT NOT NULL WITH DEFAULT 0	The total composite cardinality.
OPENCOST	FLOAT NOT NULL WITH DEFAULT 0	The do-at-open total cost for non-correlated sub query.
ONECOMPROWS	FLOAT NOT NULL WITH DEFAULT 0	The number of rows qualified after applying local predicates.
DMROWS	FLOAT NOT NULL WITH DEFAULT 0	The number of data manager rows returned (after all stage 1 predicates are applied).
RDSROW	FLOAT NOT NULL WITH DEFAULT 0	The number of RDS rows returned (after all stage 1 and stage 2 predicates are applied).

Column Name	Data Type	Description
TIMES	FLOAT NOT NULL WITH DEFAULT 0	The estimated number rows returned by data manager; also the estimated number of times this query block isexecuted.
ROWCOUNT	INTEGER NOT NULL WITH DEFAULT 0	The estimated number rows returned by RDS (Query cardinality).
ATOPEN	CHAR(1) NOT NULL WITH DEFAULT	Whether the query block is moved up for do-at-open processing. Y if done at open, N otherwise.
DOATOPEN_PARENT	SMALLINT NOT NULL WITH DEFAULT 0	The parent query block number of the current query block. Do-at-open-parent if the query block is done-at-open, this may be different from the PARENT_QBLOCKNO in the PLAN_TABLE.
TABLE_CREATOR	VARCHAR(128) NOT NULL	The creator of the new table accessed in this step, blank if METHOD is 3.
TABLE_NAME	VARCHAR(128) NOT NULL	The name of a table, created or declared temporary table, materialized view, or materialized table expression. The value is blank if METHOD is 3. The column can also contain the name of a table in the form DSNWFQB(qblockno). DSNWFQB(qblockno) is used to represent the intermediate result of a UNION ALL or an outer join that is materialized. If a view is merged, the name of the view does not appear. A value of Q in TABLE_TYPE for the name of a view or nested table expression indicates that the materialization was virtual and not actual. Materialization can be virtual when the view or nested table expression definition contains a UNION ALL that is not distributed.
TABNO	SMALLINT NOT NULL	The Sequence of the table in the FROM clause. Value is 0 if METHOD is 3. For IBM use.
ACCESSTYPE	CHAR(2) NOT NULL	The method of accessing the new table: I By an index (identified in ACCESSCREATOR and ACCESSNAME) I1 By a one-fetch index scan N By an index scan when the matching predicate contains the IN keyword R By a table space scan

Column Name	Data Type	Description
		M By a multiple index scan (followed by MX, MI, or MU) MX By an index scan on the index named in ACCESSNAME MI By an intersection of multiple indexes MU By a union of multiple indexes T By a sparse index (star join work files) blank Not applicable to the current row
MATCHCOLS	SMALLINT NOT NULL	For ACCESSTYPE I, I1, N, or MX, the number of index keys used in an index scan; otherwise, 0.
INDEX_CREATOR	VARCHAR(128) NOT NULL	For ACCESSTYPE I, I1, N, or MX, the creator of the index; otherwise, blank.
INDEX_NAME	VARCHAR(128) NOT NULL	For ACCESSTYPE I, I1, N, or MX, the name of the index; otherwise, blank.
INDEXONLY	CHAR(1) NOT NULL	Whether access to an index alone is enough to carry out the step, or whether data too must be accessed. Y=Yes; N=No.
SORTN_UNIQ	CHAR(1) NOT NULL	Whether the new table is sorted to remove duplicate rows. Y=Yes; N=No.
SORTN_JOIN	CHAR(1 NOT NULL	Whether the new table is sorted for join method 2 or 4. Y=Yes; N=No.
SORTN_ORDERBY	CHAR(1) NOT NULL	Whether the new table is sorted for ORDER BY. Y=Yes; N=No.
SORTN_GROUPBY	CHAR(1 NOT NULL	Whether the new table is sorted for GROUP BY. Y=Yes; N=No.
SORTC_UNIQ	CHAR(1) NOT NULL	Whether the composite table is sorted to remove duplicate rows. Y=Yes; N=No.
SORTC_JOIN	CHAR(1 NOT NULL	Whether the composite (outer) table is sorted for join method 1, 2 or 4. Y=Yes; N=No.
SORTC_ORDERBY	CHAR(1) NOT NULL	Whether the composite table is sorted for an ORDER BY clause or a quantified predicate. Y=Yes; N=No.
SORTC_GROUPBY	CHAR(1) NOT NULL	Whether the composite table is sorted for a GROUP BY clause. Y=Yes; N=No.
TSLOCKMODE	CHAR(3) NOT NULL	An indication of the mode of lock to be acquired on either the new table or its table space or table space partitions. If the isolation can be determined at

Column Name	Data Type	Description
		bind time, the values are: IS Intent share lock SIX Intent exclusive lock S Share lock U Update lock X Exclusive lock SIX Share intent exclusive lock N UR isolation; no lock If the isolation cannot be determined at bind time, then the lock mode determined by the isolation at run time is shown by the following values. NS For UR isolation, no lock; for CS, RS, or RR, an S lock. NIS For UR isolation, no lock; for CS, RS, or RR, an IS lock. NSS For UR isolation, no lock; for CS or RS, an IS lock; for RR, an S lock. SS For UR, CS, or RS isolation, an IS lock; for RR, an S lock. The data in this column is right justified. For example, IX appears as a blank followed by I followed by X. If the column contains a blank, then no lock is acquired.
PREFETCH	CHAR(1) NOT NULL	Whether data pages are to be read in advance by prefetch. S = pure sequential prefetch; L = prefetch through a page list; blank = unknown or no prefetch.
COLUMN_FN_EVAL	CHAR(1) NOT NULL	When an SQL column function is evaluated. R = while the data is being read from the table or index; S = while performing a sort to satisfy a GROUP BY clause; blank = after data retrieval and after any sorts.
MIXOPSEQ	SMALLINT NOT NULL	The sequence number of a step in a multiple index operation. 1, 2, ... n For the steps of the multiple index procedure (ACCESSTYPE is MX, MI, or MU.) 0 For any other rows (ACCESSTYPE is I, I1, M, N,

Column Name	Data Type	Description
		R, or blank.)
ACCESS_DEGREE	SMALLINT	The number of parallel tasks or operations activated by a query. This value is determined at bind time; the actual number of parallel operations used at execution time could be different. This column contains 0 if there is a host variable.
ACCESS_PGROUP_ID	SMALLINT	The identifier of the parallel group for accessing the new table. A parallel group is a set of consecutive operations, executed in parallel, that have the same number of parallel tasks. This value is determined at bind time; it could change at execution time.
JOIN_DEGREE	SMALLINT	The number of parallel operations or tasks used in joining the composite table with the new table. This value is determined at bind time and can be 0 if there is a host variable. The actual number of parallel operations or tasks used at execution time could be different.
JOIN_PGROUP_ID	SMALLINT	The identifier of the parallel group for joining the composite table with the new table. This value is determined at bind time; it could change at execution time.
SORTC_PGROUP_ID	SMALLINT	The parallel group identifier for the parallel sort of the composite table.
SORTN_PGROUP_ID	SMALLINT	The parallel group identifier for the parallel sort of the new table.
PARALLELISM_MODE	CHAR(1)	The kind of parallelism, if any, that is used at bind time: I Query I/O parallelism C Query CP parallelism X Sysplex query parallelism
MERGE_JOIN_COLS	SMALLINT	The number of columns that are joined during a merge scan join(Method=2).
PAGE_RANGE	CHAR(1) NOT NULL WITH DEFAULT	Whether the table qualifies for page range screening, so that plans scan only the partitions that are needed. Y = Yes; blank = No.
JOIN_TYPE	CHAR(1) NOT NULL WITH	The type of join:

Column Name	Data Type	Description
	DEFAULT	F FULL OUTER JOIN L LEFT OUTER JOIN S STAR JOIN blank INNER JOIN or no join RIGHT OUTER JOIN converts to a LEFT OUTER JOIN when you use it, so that JOIN_TYPE contains L.
GROUP_MEMBER	CHAR(8) NOT NULL WITH DEFAULT	The member name of the DB2 that executed EXPLAIN. The column is blank if the DB2 subsystem was not in a data sharing environment when EXPLAIN was executed.
WHEN_OPTIMIZE	CHAR(1) NOT NULL WITH DEFAULT	When the access path was determined: blank At bind time, using a default filter factor for any host variables, parameter markers, or special registers. B At bind time, using a default filter factor for any host variables, parameter markers, or special registers; however, the statement is reoptimized at run time using input variable values for input host variables, parameter markers, or special registers. The bind option REOPT(VARS) must be specified for reoptimization to occur. R At run time, using input variables for any host variables, parameter markers, or special registers. The bind option REOPT(VARS) must be this to occur.
QBLOCK_TYPE	CHAR(6) NOT NULL WITH DEFAULT	For each query block, an indication of the type of SQL operation performed. For the outermost query, this column identifies the statement type. Possible values: SELECT SELECT INSERT INSERT UPDATE UPDATE DELETE DELETE SELUPD SELECT with FOR UPDATE OF

Column Name	Data Type	Description
		DELCUR DELETE WHERE CURRENT OF CURSOR UPDCUR UPDATE WHERE CURRENT OF CURSOR CORSUB Correlated subquery NCOSUB Noncorrelated subquery TABLEX Table expression UNION UNION UNIONA UNION ALL
OPTHINT	VARCHAR(128) NOT NULL WITH DEFAULT	A string that you use to identify this row as an optimization hint for DB2. DB2 uses this row as input when choosing an access path.
HINT_USED	VARCHAR(128) NOT NULL WITH DEFAULT	If DB2 used one of your optimization hints, it puts the identifier for that hint (the value in OPTHINT) in this column.
PRIMARY_ACCESSTYPE	CHAR(1) NOT NULL WITH DEFAULT	Indicates whether direct row access will be attempted first: D DB2 will try to use direct row access. If DB2 cannot use direct row access at runtime, it uses the access path described in the ACCESSTYPE column of PLAN_TABLE. blank DB2 will not try to use direct row access.
PARENT_QBLOCKNO	SMALLINT NOT NULL WITH DEFAULT	A number that indicates the QBLOCKNO of the parent query block.
TABLE_TYPE	CHAR(1)	The type of new table: F Table function Q Temporary intermediate result table (not materialized) T Table W Work file The value of the column is null if the query uses GROUP BY, ORDER BY, or DISTINCT, which requires an implicit sort.
DBNAME	VARCHAR(24) NOT NULL WITH DEFAULT	

Column Name	Data Type	Description
SUBSYS	CHAR(4) NOT NULL WITH DEFAULT	
SCREENCOLS	SMALLINT NOT NULL WITH DEFAULT	

TRC_PREDICAT_TABLE

One row for each predicate referenced in the SQL statement being explained.

INDEXED BY

XTPRDPD1 ON (WORKLOAD,VERSION,SQL_NO,BIND_TIME,EXPLAIN_TIME)

Column Name	Data Type	Description
WORKLOAD	CHAR(8) NOT NULL	The Workload name.
VERSION	CHAR(1) NOT NULL WITH DEFAULT 'O'	The Cost Indicator associated with this Explain data. O - Original Costs P - Predicted Costs F - Final Costs Any other alphanumeric character identifies a User defined variation, eg. A refined index configuration or a **A**delta@ version Id associated with Impact Analyzer.
BIND_TIME	CHAR(26) NOT NULL WITH DEFAULT ' '	Date and Time access path generated at BIND with EXPLAIN=YES
EXPLAIN_TIME	TIMESTAMP NOT NULL WITH DEFAULT	Date and Time access path generated from Explain
SQL_NO	INTEGER NOT NULL	Date and Time access path generated at BIND with EXPLAIN=YES
QBLOCKNO	SMALLINT NOT NULL	Date and Time access path generated from Explain
PLANNO	SMALLINT NOT NULL	Date and Time access path generated at BIND with EXPLAIN=YES
PREDNO	INTEGER NOT NULL	Predicate Number.
ORDERNO	INTEGER NOT NULL	Predicate Order.
FILTER_FACTOR	FLOAT(8) NOT NULL WITH DEFAULT 0	Estimated Filter Factor.
STAGE	CHAR(9)	Stage at which predicate is evaluated: Matching

Column Name	Data Type	Description
	NOT NULL	Screening Pagerange Stage 1 Stage 2
PUSHDOWN	CHAR(1) NOT NULL	Indicates if the predicate is pushed down the Index Manager or Data Manager for evaluation: I – Index Manager D – Data Manager blank – No Pushdown
BOOLEAN_TERM	CHAR(1) NOT NULL	Predicate can be used to determine the truth value of the whole WHERE clause
SEARCHARG	CHAR(1) NOT NULL	Predicate can be processed by DM vs RDS.
JOIN	CHAR(1) NOT NULL	Whether the predicate can be used as simple join predicate.
AFTER_JOIN	CHAR(1) NOT NULL	Predicate evaluation phase: A – After Join D – During Join blank – Not applicable
ADDED_PRED	CHAR(1) NOT NULL	Predicate generated by transitive closure: B – Bubble Up C – Correlation J – Join K – Like for expression based IX L – Localization P – Push Down R – Page Range S – Simplification T – Transitive Closure blank – Db2 did not add predicate
REDUNDANT_PRED	CHAR(1) NOT NULL	Redundant Predicate
DIRECT_ACCESS	CHAR(1) NOT NULL	Predicate is direct access (ie Uses ROWID)
KEYFIELD	CHAR(1) NOT NULL	Predicate includes the key column
NEGATION	CHAR(1) NOT NULL	Predicate is negated by NOT
ORIGIN	CHAR(1) NOT NULL	Origin of Predicate: C – Column Mask R – Row permission

Column Name	Data Type	Description
		U – Specified by user
CLAUSE	CHAR(8) NOT NULL	Clause where predicate exists: HAVING – Having Clause ON – On Clause WHERE – Where Clause SELECT – Select Clause
UNCERTAINTY	FLOAT(4) NOT NULL WITH DEFAULT 0	Uncertainty factor of filter factor. Higher means more uncertain
TEXT	VARCHAR(2000) NOT NULL	Transformed predicate text