

[Send feedback](#) | [Rate this page](#)

QAQQINI query options

The following table summarizes the query options that can be specified on the QAQQINI command:



Table 28. Query Options Specified on QAQQINI Command

Parameter	Value	Description
APPLY_REMOTE	*DEFAULT	The default value is set to *YES.
	*NO	The CHGQRYA attributes for the job are not applied to the remote jobs. The remote jobs will use the attributes associated to them on their servers.
	*YES	The query attributes for the job are applied to the remote jobs used in processing database queries involving distributed tables. For attributes where *SYSVAL is specified, the system value on the remote server is used for the remote job. This option requires that, if CHGQRYA was used for this job, the remote jobs must have authority to use the CHGQRYA command.
ASYNC_JOB_USAGE	*DEFAULT	The default value is set to *LOCAL.
	*LOCAL	Asynchronous jobs may be used for database queries that involve only tables local to the server where the database queries are being run. In addition, for queries involving distributed tables, this option allows the communications required to be asynchronous. This allows each server involved in the query of the distributed tables to run its portion of the query at the same time (in parallel) as the other servers.
	*DIST	Asynchronous jobs may be used for database queries that involve distributed tables.
	*ANY	Asynchronous jobs may be used for any database query.
	*NONE	No asynchronous jobs are allowed to be used for database query processing. In addition, all processing for queries involving distributed tables occurs synchronously. Therefore, no inter-system parallel processing will occur.
COMMITMENT_CONTROL_LOCK_LIMIT	*DEFAULT	*DEFAULT is equivalent to 500,000,000.
	Integer Value	The maximum number of records that can be locked to a commit transaction initiated after setting the new value. The valid integer value is 1–500,000,000.
	*DEFAULT	The default is set to *NO.
	*NO	Allow the optimizer to re-order join tables.
	*SQL	Only force the join order for those queries that use the SQL JOIN syntax. This mimics the behavior for the optimizer before V4R4M0.

FORCE_JOIN_ORDER	*PRIMARY nnn	Only force the join position for the file listed by the numeric value nnn (nnn is optional and will default to 1) into the primary position (or dial) for the join. The optimizer will then determine the join order for all of the remaining files based upon cost.
	*YES	Do not allow the query optimizer to re-order join tables as part of its optimization process. The join will occur in the order in which the tables were specified in the query.
IGNORE_DERIVED_INDEX	*DEFAULT	The default value is the same as *NO.
	*YES	<p>Allow the SQE optimizer to ignore the derived index and process the query. The resulting query plan will be created without any regard to the existence of the derived index(s). The index types that are ignored include:</p> <ul style="list-style-type: none"> • Keyed logical files defined with select or omit criteria and with the DYNSTL keyword omitted • Keyed logical files built over multiple physical file members (V5R2 restriction, not a restriction for V5R3) • Keyed logical files where one or more keys reference an intermediate derivation in the DDS. Exceptions to this are: 1. when the intermediate definition is simply defining the field in the DDS so that shows up in the logical's format and 2. RENAME of a field (these two exceptions do not make the key derived) • Keyed logical files with K *NONE specified. • Keyed logical files with Alternate Collating Sequence (ACS) specified • SQL indexes created when the sort sequence active at the time of creation requires a weighting (translation) of the key to occur. This is true when any of several non-US language IDs are specified. It also occurs if language ID shared weight is specified, even for language US.
	*NO	Do not ignore the derived index. If a derived index exists, have CQE process the query.
IGNORE_LIKE_REDUNDANT_SHIFTS	*DEFAULT	The default value is set to *OPTIMIZE.
	*ALWAYS	When processing the SQL LIKE predicate or OPNQRYF command %WLDCRD built-in function, redundant shift characters are ignored for DBCS-Open operands. Note that this option restricts the query optimizer from using an index to perform key row positioning for SQL LIKE or OPNQRYF %WLDCRD predicates involving DBCS-Open, DBCS-Either, or DBCS-Only operands.
		When processing the SQL LIKE predicate or the OPNQRYF command %WLDCRD built-in function, redundant shift characters may or may not be ignored for

	*OPTIMIZE	DBCS-Open operands depending on whether an index is used to perform key row positioning for these predicates. Note that this option will enable the query optimizer to consider key row positioning for SQL LIKE or OPNQRYP %WLDCRD predicates involving DBCS-Open, DBCS-Either, or DBCS-Only operands.
LOB_LOCATOR_THRESHOLD	*DEFAULT	The default value is set to 0. This indicates that the database will take no action to free locators.
	Integer Value	If the value is 0, then the database will take no action to free locators. For values 1 through 250,000, on a FETCH request, the database will compare the active LOB locator count for the job against the threshold value. If the locator count is greater than or equal to the threshold, the database will free host server created locators that have been retrieved. This option applies to all host server jobs (QZDASOINIT) and has no impact to other jobs.
MATERIALIZED_QUERY_TABLE_REFRESH_AGE:	*DEFAULT	The default value is *NONE.
	*NONE	Materialized Query Tables may not be used in query optimization and implementation.
	*ALL	User-maintained refresh-deferred query tables may be used.
	*USER	User-maintained materialized query tables may be used.
MATERIALIZED_QUERY_TABLE_USAGE:	*DEFAULT	The default value is 0.
	0	No materialized query tables may be used.
	*ANY	Any tables indicated by the MATERIALIZED_QUERY_TABLE_USAGE QAQQINI parameter may be used. Equivalent to specifying 9999 99 99 99 99 99 (which is 9999 years, 99 months, 99 days, 99 hours, 99 minutes, 99 seconds). If the materialized query table has never been refreshed by the REFRESH TABLE SQL statement, but the table should be considered, then the MATERIALIZED_QUERY_TABLE_REFRESH_AGE QAQQINI option must be set to *ANY.
	Timestamp_duration	Only tables indicated by the MATERIALIZED_QUERY_TABLE_USAGE QAQQINI option which have a REFRESH TABLE performed within the specified timestamp duration will be used. This is a DECIMAL(20,6) number which indicates a timestamp duration since the last REFRESH TABLE was done.
MESSAGES_DEBUG	*DEFAULT	The default is set to *NO.
	*NO	No debug messages are to be displayed.
	*YES	Issue all debug messages that would be generated for STRDBG.
	*DEFAULT	The default is set to *NO.
		Unicode constants, host variables, parameter markers,

NORMALIZE_DATA	*NO	and expressions that combine strings will not be normalized.
	*YES	Unicode constants, host variables, parameter markers, and expressions that combine strings will be normalized
OPEN_CURSOR_CLOSE_COUNT	*DEFAULT	*DEFAULT is equivalent to 0. See Integer Value for details.
	Integer Value	OPEN_CURSOR_CLOSE_COUNT is used in conjunction with OPEN_CURSOR_THRESHOLD to manage the number of open cursors within a job. If the number of open cursors, which includes open cursors and pseudo-closed cursors, reaches the value specified by the OPEN_CURSOR_THRESHOLD, pseudo-closed cursors are hard (fully) closed with the least recently used cursors being closed first. This value determines the number of cursors to be closed. The valid values for this parameter are 1 - 65536. The value for this parameter should be less than or equal to the number in the OPEN_CURSOR_THRESHOLD parameter. This value is ignored if OPEN_CURSOR_THRESHOLD is *DEFAULT. If OPEN_CURSOR_THRESHOLD is specified and this value is *DEFAULT, the number of cursors closed is equal to OPEN_CURSOR_THRESHOLD multiplied by 10 percent and rounded up to the next integer value.
OPEN_CURSOR_THRESHOLD	*DEFAULT	*DEFAULT is equivalent to 0. See Integer Value for details.
	Integer Value	OPEN_CURSOR_THRESHOLD is used in conjunction with OPEN_CURSOR_CLOSE_COUNT to manage the number of open cursors within a job. If the number of open cursors, which includes open cursors and pseudo-closed cursors, reaches this threshold value, pseudo-closed cursors are hard (fully) closed with the least recently used cursors being closed first. The number of cursors to be closed is determined by OPEN_CURSOR_CLOSE_COUNT. The valid user-entered values for this parameter are 1 - 65536. Having a value of 0 (default value) indicates that there is no threshold and hard closes will not be forced on the basis of the number of open cursors within a job.
OPTIMIZATION_GOAL	*DEFAULT	Optimization goal is determined by the interface (ODBC, SQL precompiler options, OPTIMIZE FOR nnn ROWS clause).
	*FIRSTIO	All queries will be optimized with the goal of returning the first page of output as fast as possible. This goal works well when the control of the output is controlled by a user who is most likely to cancel the query after viewing the first page of output data. Queries coded with an OPTIMIZE FOR nnn ROWS clause will honor the goal specified by the clause.

	*ALLIO	All queries will be optimized with the goal of running the entire query to completion in the shortest amount of elapsed time. This is a good option for when the output of a query is being written to a file or report, or the interface is queuing the output data. Queries coded with an OPTIMIZE FOR nnn ROWS clause will honor the goal specified by the clause.
OPTIMIZE_STATISTIC_LIMITATION (See note)	*DEFAULT	The amount of time spent in gathering index statistics is determined by the query optimizer.
	*NO	No index statistics will be gathered by the query optimizer. Default statistics will be used for optimization. (Use this option sparingly.)
	*PERCENTAGE integer value	Specifies the maximum percentage of the index that will be searched while gathering statistics. Valid values for are 1 to 99.
	*MAX_NUMBER_OF_RECORDS_ALLOWED integer value	Specifies the largest table size, in number of rows, for which gathering statistics is allowed. For tables with more rows than the specified value, the optimizer will not gather statistics and will use default values.
PARALLEL_DEGREE	*DEFAULT	The default value is set to *SYSVAL.
	*SYSVAL	The processing option used is set to the current value of the system value, QQRYDEGREE.
	*IO	Any number of tasks can be used when the database query optimizer chooses to use I/O parallel processing for queries. SMP parallel processing is not allowed.
	*OPTIMIZE	The query optimizer can choose to use any number of tasks for either I/O or SMP parallel processing to process the query or database file keyed access path build, rebuild, or maintenance. SMP parallel processing is used only if the system feature, DB2® Symmetric Multiprocessing for OS/400®, is installed. Use of parallel processing and the number of tasks used is determined with respect to the number of processors available in the server, this job has a share of the amount of active memory available in the pool in which the job is run, and whether the expected elapsed time for the query or database file keyed access path build or rebuild is limited by CPU processing or I/O resources. The query optimizer chooses an implementation that minimizes elapsed time based on the job has a share of the memory in the pool.
	*MAX	The query optimizer chooses to use either I/O or SMP parallel processing to process the query. SMP parallel processing will only be used if the system feature, DB2 Symmetric Multiprocessing for OS/400, is installed. The choices made by the query optimizer are similar to those made for parameter value *OPTIMIZE except the optimizer assumes that all active memory in the pool can

		be used to process the query or database file keyed access path build, rebuild, or maintenance.
	*NONE	No parallel processing is allowed for database query processing or database table index build, rebuild, or maintenance.
	*NUMBER_OF_TASKS nn	Indicates the maximum number of tasks that can be used for a single query. The number of tasks will be capped off at either this value or the number of disk arms associated with the table.
PARAMETER_MARKER_CONVERSION	*DEFAULT	The default value is set to *YES.
	*NO	Constants cannot be implemented as parameter markers.
	*YES	Constants can be implemented as parameter markers.
QUERY_TIME_LIMIT	*DEFAULT	The default value is set to *SYSVAL.
	*SYSVAL	The query time limit for this job will be obtained from the system value, QQRYSYTLMT.
	*NOMAX	There is no maximum number of estimated elapsed seconds.
	integer value	Specifies the maximum value that is checked against the estimated number of elapsed seconds required to run a query. If the estimated elapsed seconds is greater than this value, the query is not started. Valid values range from 0 through 2147352578.
REOPTIMIZE_ACCESS_PLAN	*DEFAULT	The default value is set to *NO.
	*NO	Do not force the existing query to be reoptimized. However, if the optimizer determines that optimization is necessary, the query will be reoptimized.
	*YES	Force the existing query to be reoptimized.
	*FORCE	Force the existing query to be reoptimized.
	*ONLY_REQUIRED	Do not allow the plan to be reoptimized for any subjective reasons. For these cases, continue to use the existing plan since it is still a valid workable plan. This may mean that you may not get all of the performance benefits that a reoptimization plan may derive. Subjective reasons include, file size changes, new indexes, and so on. Non-subjective reasons include, deletion of an index used by existing access plan, query file being deleted and recreated, and so on.
SQLSTANDARDS_MIXED_CONSTANT	*DEFAULT	The default value is set to *YES.
	*YES	SQL IGC constants will be treated as IGC-OPEN constants.
	*NO	If the data in the IGC constant only contains shift-out DBCS-data shift-in, then the constant will be treated as IGC-ONLY, otherwise it will be treated as IGC-OPEN.
	*DEFAULT	The default value is set to *NO.
		Examine the SQLCODE in the SQLCA after execution of a statement. If the SQLCODE is + 30, then alter the

SQL_SUPPRESS_WARNINGS	*YES	SQLCA so that no warning is returned to the caller. Set the SQLCODE to 0, the SQLSTATE to '00000' and SQLWARN to ''.
	*NO	Specifies that SQL warnings will be returned to the caller.
SQL_TRANSLATE_ASCII_TO_JOB	*DEFAULT	The default value is set to *NO.
	*YES	Translate ASCII SQL statement text to the CCSID of the iSeries™ job.
	*NO	Translate ASCII SQL statement text to the EBCDIC CCSID associated with the ASCII CCSID.
STAR_JOIN (see note)	*DEFAULT	The default value is set to *NO
	*NO	The EVI Star Join optimization support is not enabled.
	*COST	Allow query optimization to consider (cost) the usage of EVI Star Join support. The determination of whether the Distinct List selection is used will be determined by the optimizer based on how much benefit can be derived from using that selection.
SYSTEM_SQL_STATEMENT_CACHE	*DEFAULT	The default value is set to *YES.
	*YES	Examine the SQL system-wide statement cache when an SQL prepare request is processed. If a matching statement already exists in the cache, use the results of that prepare. This allows the application to potentially have better performing prepares.
	*NO	Specifies that the SQL system-wide statement cache should not be examined when processing an SQL prepare request.
UDF_TIME_OUT (see note)	*DEFAULT	The amount of time to wait is determined by the database. The default is 30 seconds.
	*MAX	The maximum amount of time that the database will wait for the UDF to finish.
	integer value	Specify the number of seconds that the database should wait for a UDF to finish. If the value given exceeds the database maximum wait time, the maximum wait time will be used by the database. Minimum value is 1 and maximum value is system defined.
VARIABLE_LENGTH_OPTIMIZATION	*DEFAULT	The default value is set to *YES.
	*YES	Allow aggressive optimization of variable length columns. Allows index only access for the column(s). It also allows constant value substitution when an equal predicate is present against the column(s). As a consequence, the length of the data returned for the variable length column may not include any trailing blanks that existed in the original data.
		Do not allow aggressive optimization of variable length

*NO columns.

**Note:**

Only modifies the environment for the Classic Query Engine.