

IBM i
7.3

*Database
Performance and Query Optimization*



Note

Before using this information and the product it supports, read the information in [“Notices” on page 1059.](#)

This document may contain references to Licensed Internal Code. Licensed Internal Code is Machine Code and is licensed to you under the terms of the IBM License Agreement for Machine Code.

© **Copyright International Business Machines Corporation 1998, 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

- Performance and query optimization.....1**
- What's new for IBM i 7.3..... 1
- Print PDF..... 11
- Query engine overview..... 11
 - SQE and CQE engines..... 12
 - Query dispatcher..... 13
 - Statistics manager..... 13
 - Global Statistics Cache..... 14
 - Plan cache..... 14
- Data access methods..... 16
 - Permanent objects & access methods..... 17
 - Table..... 17
 - Table scan..... 17
 - Table probe..... 18
 - Radix index..... 19
 - Radix index scan..... 20
 - Radix index probe..... 21
 - Encoded vector index..... 23
 - Encoded vector index RRN probe..... 24
 - EVI probe..... 25
 - EVI index-only access..... 26
 - EVI symbol table scan..... 27
 - EVI symbol table probe..... 30
 - Temporary objects & access methods..... 32
 - Temporary hash table..... 32
 - Hash table scan..... 33
 - Hash table probe..... 34
 - Temporary sorted list..... 35
 - Sorted list scan..... 36
 - Sorted list probe..... 37
 - Temporary distinct sorted list..... 38
 - Sorted list scan..... 38
 - Temporary list..... 39
 - List scan..... 39
 - Temporary values list..... 40
 - Values list scan..... 41
 - Temporary row number list..... 41
 - Row number list scan..... 41
 - Row number list probe..... 43
 - Temporary bitmap..... 44
 - Bitmap scan..... 45
 - Bitmap probe..... 46
 - Temporary index..... 48
 - Temporary index scan..... 48
 - Temporary index probe..... 50
 - Temporary buffer..... 51
 - Buffer scan..... 51
 - Queue..... 52
 - Enqueue..... 53
 - Dequeue..... 54
 - Array unnest temporary table..... 55

Array unnest temporary table scan.....	56
Temporary Indexed List.....	56
Temporary Indexed List scan and Index Merge.....	57
Window.....	59
Window scan.....	60
Objects processed in parallel.....	60
Spreading data automatically.....	61
Processing queries: Overview.....	61
Query optimizer.....	62
Query optimization tips.....	62
Access plan validation.....	63
Single table optimization.....	63
Solid State Drives.....	64
Memory preference controls.....	65
Join optimization.....	66
Nested loop join.....	66
Join optimization.....	68
Join order optimization.....	69
Full outer join.....	70
Join cost & index selection.....	71
Transitive closure predicates.....	72
LPG.....	72
CQE Join performance tips.....	73
Multiple join types.....	73
Join performance problems.....	74
Join performance tips.....	75
Distinct optimization.....	76
Grouping optimization.....	77
Hash grouping.....	77
Index grouping.....	78
Eliminate grouping columns.....	80
Add grouping columns.....	81
Index skip key processing.....	81
Read trigger considerations.....	82
Grouping set optimization.....	83
Ordering optimization.....	83
View implementation.....	85
View composite.....	85
View materialization.....	86
MQT optimization.....	87
MQT supported function.....	87
Using MQTs.....	89
MQT examples.....	89
MQT matching.....	91
Determining MQT usage.....	95
MQT recommendations.....	95
Recursive query optimization.....	96
Example.....	96
Multiple initialization & iterative fullselects.....	98
Predicate pushing.....	100
SEARCH considerations.....	100
CYCLE considerations.....	101
SMP & recursive queries.....	102
System-period temporal tables.....	105
Adaptive Query Processing.....	106
How AQP works.....	107
AQP example.....	107
AQP join order.....	108

Database Monitor additions.....	109
Row and column access control.....	111
Indexing strategy and RCAC.....	111
Materialized query tables and RCAC.....	114
Tools.....	118
Health Center.....	118
Navigator view.....	118
SQL procedures.....	118
Health_Database_Overview.....	118
Health_Activity.....	123
Health_Design_Limits.....	128
Health_Size_Limits.....	131
Health_Environmental_Limits.....	135
Reset_Environmental_Limits.....	139
Database Monitor.....	140
Start.....	141
End.....	144
Performance rows.....	146
Examples.....	146
Application with table scans.....	146
Queries with table scans.....	147
Table scan detail.....	148
Additional examples.....	149
Navigator monitors.....	151
Start.....	151
Analyze data.....	153
Compare data.....	154
View statements.....	155
Import.....	155
Index advisor.....	155
Index advice and OR predicates.....	156
Index advice for Encoded Vector Index RRN Probe Plans.....	158
Display information.....	160
System table.....	161
Column descriptions.....	163
Database monitor view.....	164
Condense advice.....	165
Visual Explain.....	165
Start.....	165
Information available.....	166
Adaptive Query Processing in VE.....	168
SQL Plan Cache.....	169
Show Statements.....	170
Column descriptions.....	173
Properties.....	174
Creating snapshots.....	181
Event monitor.....	183
SQL stored procedures.....	183
Verify performance.....	183
View debug messages.....	184
Print SQL Information.....	185
Tool comparison.....	185
Change query attributes.....	186
QAQQINI.....	186
CHGQRYA.....	188
Creating the QAQQINI query options file.....	188
QAQQINI override support.....	189
File format.....	189

Update.....	189
Authority requirements.....	189
System-supplied triggers.....	190
Query options.....	190
SQL_XML_DATA_CCSID option.....	204
Query Supervisor.....	205
Using Query Supervisor to monitor query resource usage.....	206
Query Supervisor configuration and operation.....	206
Managing Query Supervisor activity.....	207
Using DETECTION_FREQUENCY to protect system resources.....	208
Writing a Query Supervisor Exit Program.....	209
Query Supervisor example exit programs.....	210
Exit program to send message to QSYSOPR.....	210
Exit program to end query.....	214
Exit program to dump plan cache information for query.....	217
Exit program to log information using a data queue.....	220
Predictive Query Governor.....	236
How to use.....	237
Cancel a query.....	238
Control the reply.....	238
Test performance.....	239
Time limit examples.....	239
Test temporary storage use.....	240
Storage limit examples.....	240
Parallel processing.....	241
System-wide.....	241
Job level.....	242
Statistics manager.....	243
Automatic collection.....	244
Automatic refresh.....	245
View requests.....	245
Indexes and column statistics.....	245
Background collection.....	247
Replicate column statistics.....	247
View column statistics.....	248
Manual collection and refresh	248
APIs.....	249
Display MQT columns.....	249
Check pending constraints.....	251
Creating an index strategy.....	252
Binary radix indexes.....	252
Derived key index.....	253
Sparse indexes.....	253
Optimization.....	254
Matching algorithm.....	254
Sparse index examples.....	256
Specify PAGESIZE.....	259
Index maintenance.....	259
Encoded vector indexes.....	260
How the EVI works.....	261
When to create.....	262
Maintenance.....	263
Recommendations.....	264
Compare radix & EVIs.....	267
Indexes & the optimizer.....	267
Index not used.....	267
Display indexes for a table.....	268
Determine unnecessary indexes.....	272

Reset usage counts.....	272
View index builds.....	273
Manage index rebuilds.....	273
Indexing strategy.....	278
Reactive approach.....	279
Proactive approach.....	279
Coding for effective indexes.....	280
Avoid numeric conversions.....	280
Avoid arithmetic expressions.....	280
Avoid character string padding.....	281
LIKE considerations.....	281
Derived indexes.....	282
Sparse indexes.....	282
Indexes with sort sequence.....	283
Selection, joins, or grouping.....	283
Ordering.....	284
Index examples.....	284
Equal selection, no sort sequence.....	284
Equal selection, unique-weight sort sequence.....	284
Equal selection, shared-weight sort sequence.....	284
Greater than selection, unique-weight sort sequence.....	285
Join selection, unique-weight sort sequence.....	285
Join selection, shared-weight sort sequence.....	285
Order, no sort sequence.....	286
Order, unique-weight sort sequence.....	286
Order, shared-weight sort sequence.....	286
Order, ALWCOPYDTA(*OPTIMIZE), unique-weight sort sequence.....	286
Group, no sort sequence.....	287
Group, unique-weight sort sequence.....	287
Group, shared-weight sort sequence.....	287
Order & group on same columns, unique-weight sort sequence.....	288
Order & group on same columns, ALWCOPYDTA(*OPTIMIZE), unique-weight sort sequence....	288
Order & group on same columns, shared-weight sort sequence.....	288
Order & group on same columns, ALWCOPYDTA(*OPTIMIZE), shared-weight sort sequence....	289
Order & group on different columns, unique-weight sort sequence.....	289
Order & group on different columns, ALWCOPYDTA(*OPTIMIZE), unique-weight sort sequence.....	289
Order & group on different columns, ALWCOPYDTA(*OPTIMIZE), shared-weight sort sequence.....	290
Sparse index examples.....	290
Application design tips.....	293
Live data.....	293
Reduce open operations.....	294
Retain cursor positions.....	297
Non-ILE program calls.....	297
ILE program calls.....	297
General rules.....	298
Programming techniques.....	299
Use the OPTIMIZE clause.....	299
Use FETCH FOR n ROWS.....	300
Improve SQL blocking performance.....	301
Use INSERT n ROWS.....	301
Control database manager blocking.....	301
Optimize columns selected.....	303
PREPARE considerations.....	303
REFRESH(*FORWARD) considerations.....	303
Improve concurrency.....	304
Use SELECTIVITY to supply missing information.....	305

Performance considerations.....	306
Long object names.....	306
Precompile options.....	307
ALWCOPYDTA.....	308
VARCHAR and VARGRAPHIC.....	309
Field procedures.....	310
Examples.....	312
Db2 for i Services.....	313
Application Services.....	313
DELIMIT_NAME scalar function.....	313
OVERRIDE_QAQQINI procedure.....	314
OVERRIDE_TABLE procedure.....	315
PARSE_STATEMENT table function.....	315
WLM_SET_CLIENT_INFO procedure.....	320
Performance Services.....	321
ACT_ON_INDEX_ADVICE procedure.....	321
ACTIVE_QUERY_INFO table function.....	322
ADD_QUERY_THRESHOLD procedure.....	324
DATABASE_MONITOR_INFO view.....	327
HARVEST_INDEX_ADVICE procedure.....	331
MTI_INFO table function.....	332
QUERY_SUPERVISOR view.....	334
REMOVE_INDEXES procedure.....	335
REMOVE_QUERY_THRESHOLD procedure.....	336
RESET_TABLE_INDEX_STATISTICS procedure.....	336
Plan Cache Services.....	338
CHANGE_PLAN_CACHE_SIZE procedure.....	338
CLEAR_PLAN_CACHE procedure.....	338
DUMP_PLAN_CACHE procedure.....	340
DUMP_PLAN_CACHE_PROPERTIES procedure.....	341
DUMP_PLAN_CACHE_TOPN procedure.....	341
DUMP_SNAP_SHOT_PROPERTIES procedure.....	342
END_ALL_PLAN_CACHE_EVENT_MONITORS procedure.....	343
END_PLAN_CACHE_EVENT_MONITOR procedure.....	343
IMPORT_PC_EVENT_MONITOR procedure.....	344
IMPORT_PC_SNAPSHOT procedure.....	344
REMOVE_PC_EVENT_MONITOR procedure.....	344
REMOVE_PC_SNAPSHOT procedure.....	345
REMOVE_PERFORMANCE_MONITOR procedure.....	345
START_PLAN_CACHE_EVENT_MONITOR procedure.....	345
Utility Services.....	346
ANALYZE_CATALOG table function.....	346
CANCEL_SQL procedure.....	348
CHECK_SYSCST procedure.....	349
CHECK_SYSROUTINE procedure.....	350
DUMP_SQL_CURSORS procedure.....	351
END_IDLE_SQE_THREADS procedure.....	352
EXTRACT_STATEMENTS procedure.....	353
FIND_AND_CANCEL_QSQSRVR_SQL procedure.....	354
FIND_QSQSRVR_JOBS procedure.....	355
GENERATE_SQL procedure.....	356
GENERATE_SQL_OBJECTS procedure.....	365
RELATED_OBJECTS table function.....	376
RESTART_IDENTITY procedure.....	378
SWAP_DYNUSRPRF procedure.....	379
SYSFILES view.....	380
VALIDATE_DATA, VALIDATE_DATA_FILE, and VALIDATE_DATA_LIBRARY table functions.....	390
IBM i Services.....	391

Application Services.....	391
ACTIVATION_GROUP_INFO table function.....	391
BINDING_DIRECTORY_INFO view.....	393
BOUND_MODULE_INFO view.....	394
BOUND_SRVPGM_INFO view.....	403
CLEAR_DATA_QUEUE procedure.....	404
DATA_AREA_INFO table function.....	406
DATA_AREA_INFO view.....	407
DATA_QUEUE_ENTRIES table function.....	408
DATA_QUEUE_INFO view.....	411
DB_TRANSACTION_INFO view.....	414
ENVIRONMENT_VARIABLE_INFO view.....	431
EXIT_POINT_INFO view.....	432
EXIT_PROGRAM_INFO view.....	434
LPRINTF procedure.....	436
PROGRAM_EXPORT_IMPORT_INFO view.....	437
PROGRAM_INFO view.....	438
QCMDEXC procedure.....	452
QCMDEXC scalar function.....	452
RECEIVE_DATA_QUEUE table function.....	453
SEND_DATA_QUEUE, SEND_DATA_QUEUE_BINARY, and SEND_DATA_QUEUE_UTF8 procedures.....	455
SERVICES_INFO table.....	456
SET_PASE_SHELL_INFO procedure.....	459
SPLIT table function.....	460
STACK_INFO table function.....	461
USER_INDEX_ENTRIES table function.....	465
USER_INDEX_INFO view.....	466
USER_SPACE table function.....	467
USER_SPACE_INFO view.....	468
WATCH_DETAIL table function.....	469
WATCH_INFO view.....	474
Backup, Recovery, and Media Services (BRMS) Services.....	476
Communication Services.....	476
ACTIVE_DB_CONNECTIONS table function.....	477
DNS_LOOKUP table function.....	479
ENV_SYS_INFO view.....	480
HTTP_SERVER_INFO view.....	481
NETSTAT_INFO view.....	482
NETSTAT_INTERFACE_INFO view.....	489
NETSTAT_JOB_INFO view.....	498
NETSTAT_ROUTE_INFO view.....	499
SERVER_SBS_CONFIGURATION view.....	506
SERVER_SBS_ROUTING view.....	507
SET_SERVER_SBS_ROUTING procedure.....	510
TCPIP_INFO view.....	515
IFS Services.....	516
IFS_JOB_INFO table function.....	516
IFS_OBJECT_LOCK_INFO table function.....	519
IFS_OBJECT_PRIVILEGES table function.....	522
IFS_OBJECT_REFERENCES_INFO table function.....	525
IFS_OBJECT_STATISTICS table function.....	528
IFS_READ, IFS_READ_BINARY, and IFS_READ_UTF8 table functions.....	540
IFS_WRITE, IFS_WRITE_BINARY, and IFS_WRITE_UTF8 procedures.....	542
SERVER_SHARE_INFO view.....	544
Java Services.....	547
JVM_INFO view.....	547
SET_JVM procedure.....	548

Journal Services.....	549
ASSOCIATE_JOURNAL_RECEIVER table function.....	549
Audit journal entry services.....	552
AUDIT_JOURNAL table function common information.....	552
AUDIT_JOURNAL_AF table function.....	555
AUDIT_JOURNAL_CA table function.....	558
AUDIT_JOURNAL_CD table function.....	563
AUDIT_JOURNAL_CO table function.....	564
AUDIT_JOURNAL_CP table function.....	565
AUDIT_JOURNAL_DO table function.....	576
AUDIT_JOURNAL_EV table function.....	577
AUDIT_JOURNAL_GR table function.....	578
AUDIT_JOURNAL_JS table function.....	582
AUDIT_JOURNAL_OM table function.....	587
AUDIT_JOURNAL_OW table function.....	589
AUDIT_JOURNAL_PW table function.....	591
AUDIT_JOURNAL_ST table function.....	592
AUDIT_JOURNAL_SV table function.....	597
DISPLAY_JOURNAL table function.....	598
JOURNAL_INFO view.....	612
JOURNAL_RECEIVER_INFO view.....	620
JOURNALED_OBJECTS view.....	628
REMOTE_JOURNAL_INFO view.....	630
Librarian Services.....	635
JOURNAL_INHERIT_RULES view.....	635
LIBRARY_INFO table function.....	639
LIBRARY_LIST_INFO view.....	642
OBJECT_STATISTICS table function.....	643
Message Handling Services.....	650
HISTORY_LOG_INFO table function.....	650
JOBLOG_INFO table function.....	655
MESSAGE_FILE_DATA view.....	657
MESSAGE_QUEUE_INFO table function.....	660
MESSAGE_QUEUE_INFO view.....	663
REPLY_LIST_INFO view.....	664
SEND_MESSAGE procedure	665
Performance Services.....	666
COLLECTION_SERVICES_INFO view.....	666
PowerHA Services.....	669
Product Services.....	669
LICENSE_EXPIRATION_CHECK procedure.....	670
LICENSE_INFO view.....	670
SOFTWARE_PRODUCT_INFO view.....	673
PTF Services.....	678
ELECTRONIC_SERVICE_AGENT_INFO view.....	678
FIRMWARE_CURRENCY view.....	680
GROUP_PTF_CURRENCY view.....	682
GROUP_PTF_DETAILS view.....	683
GROUP_PTF_INFO view.....	686
PTF_INFO view.....	688
Security Services.....	691
AUTHORITY_COLLECTION view.....	691
AUTHORIZATION_LIST_INFO view.....	691
AUTHORIZATION_LIST_USER_INFO view.....	693
CERTIFICATE_INFO table function.....	695
CHANGE_USER_PROFILE table function.....	698
DRDA_AUTHENTICATION_ENTRY_INFO view.....	701
FUNCTION_INFO view.....	702

FUNCTION_USAGE view.....	703
GROUP_PROFILE_ENTRIES view.....	704
OBJECT_OWNERSHIP view.....	704
OBJECT_PRIVILEGES table function.....	706
OBJECT_PRIVILEGES view.....	709
SECURITY_INFO view.....	712
SET_COLUMN_ATTRIBUTE procedure.....	715
SQL_CHECK_AUTHORITY scalar function.....	716
SQL_CHECK_FUNCTION_USAGE scalar function.....	716
SQL_CHECK_SPECIAL_AUTHORITY scalar function.....	717
USER_INFO view.....	717
USER_INFO_BASIC view.....	726
Spool Services.....	735
DELETE_OLD_SPOOLED_FILES procedure.....	735
GENERATE_PDF scalar function.....	737
OUTPUT_QUEUE_ENTRIES table function.....	738
OUTPUT_QUEUE_ENTRIES view.....	742
OUTPUT_QUEUE_ENTRIES_BASIC view.....	747
OUTPUT_QUEUE_INFO view.....	749
SPOOLED_FILE_DATA table function.....	754
SPOOLED_FILE_INFO table function.....	755
Storage Services.....	759
ASP_INFO view.....	760
ASP_JOB_INFO view.....	767
ASP_VARY_INFO view.....	768
MEDIA_LIBRARY_INFO view.....	770
SYSDISKSTAT table function.....	772
SYSDISKSTAT view.....	775
SYSTMPSTG view.....	781
USER_STORAGE view.....	782
System Health Services.....	782
System limit alerts.....	785
SYSLIMTBL table.....	786
SYSLIMITS view.....	788
SYSLIMITS_BASIC view.....	790
QIBM_SYSTEM_LIMITS global variables.....	792
Work Management Services.....	793
ACTIVE_JOB_INFO table function.....	793
AUTOSTART_JOB_INFO view.....	810
COMMUNICATIONS_ENTRY_INFO view.....	810
GET_JOB_INFO table function.....	812
JOB_DESCRIPTION_INFO view.....	814
JOB_INFO table function.....	821
JOB_LOCK_INFO table function.....	833
JOB_QUEUE_INFO view.....	836
MEMORY_POOL table function.....	840
MEMORY_POOL_INFO view.....	842
OBJECT_LOCK_INFO view.....	844
OPEN_FILES table function.....	846
PRESTART_JOB_INFO view.....	848
PRESTART_JOB_STATISTICS table function.....	851
RECORD_LOCK_INFO view.....	854
ROUTING_ENTRY_INFO view.....	855
SCHEDULED_JOB_INFO view.....	857
SUBSYSTEM_INFO view.....	860
SUBSYSTEM_POOL_INFO view.....	862
SYSTEM_ACTIVITY_INFO table function.....	863
SYSTEM_STATUS table function.....	864

SYSTEM_STATUS_INFO view.....	872
SYSTEM_STATUS_INFO_BASIC view.....	878
SYSTEM_VALUE_INFO view.....	885
WORKLOAD_GROUP_INFO view.....	885
WORKSTATION_INFO view.....	886
SYSTOOLS.....	888
Using SYSTOOLS.....	888
HTTP function overview.....	890
BASE64DECODE scalar function.....	895
BASE64ENCODE scalar function.....	896
HTTPBLOB and HTTPCLOB scalar functions.....	896
HTTPBLOBVERBOSE and HTTPCLOBVERBOSE table functions.....	897
HTTPDELETEBLOB and HTTPDELETECLOB scalar functions.....	898
HTTPDELETEBLOBVERBOSE and HTTPDELETECLOBVERBOSE table functions.....	899
HTTPGETBLOB and HTTPGETCLOB scalar functions.....	899
HTTPGETBLOBVERBOSE and HTTPGETCLOBVERBOSE table functions.....	900
HTTPHEAD scalar function.....	901
HTTPPOSTBLOB and HTTPPOSTCLOB scalar functions.....	901
HTTPPOSTBLOBVERBOSE and HTTPPOSTCLOBVERBOSE table functions.....	902
HTTPPUTBLOB and HTTPPUTCLOB scalar functions.....	902
HTTPPUTBLOBVERBOSE and HTTPPUTCLOBVERBOSE table functions.....	903
URLDECODE scalar function.....	904
URLENCODE scalar function.....	904
Database monitor formats.....	904
SQL table.....	904
SQL view.....	910
1000 - SQL Information.....	910
3000 - Table Scan.....	928
3001 - Index Used.....	932
3002 - Index Created.....	938
3003 - Query Sort.....	944
3004 - Temp Table.....	948
3005 - Table Locked.....	954
3006 - Access Plan Rebuilt.....	956
3007 - Optimizer Timed Out.....	960
3008 - Subquery Processing.....	965
3010 - Host Variable, ODP Implementation.....	966
3011 - Array Host Variables.....	967
3012 - Global Variables.....	969
3014 - Generic QQ Information.....	970
3015 - Statistics Information.....	980
3018 - STRDBMON/ENDDDBMON.....	982
3019 - Rows retrieved.....	985
3020 - Index advised (SQE).....	987
3021 - Bitmap Created.....	990
3022 - Bitmap Merge.....	992
3023 - Temp Hash Table Created.....	995
3025 - Distinct Processing.....	998
3026 - Set operation.....	1000
3027 - Subquery Merge.....	1002
3028 - Grouping.....	1006
3030 - Materialized query tables.....	1009
3031 - Recursive common table expressions.....	1012
Messages reference.....	1014
Performance information.....	1014
Open data paths.....	1039
PRTSQLINF.....	1045

Notices	1059
Programming interface information.....	1060
Trademarks.....	1060
Terms and conditions.....	1061

Database performance and query optimization

The goal of database performance tuning is to minimize the response time of your queries by making the best use of your system resources. The best use of these resources involves minimizing network traffic, disk I/O, and CPU time. This goal can only be achieved by understanding the logical and physical structure of your data, the applications used on your system, and how the conflicting uses of your database might affect performance.

The best way to avoid performance problems is to ensure that performance issues are part of your ongoing development activities. Many of the most significant performance improvements are realized through careful design at the beginning of the database development cycle. To most effectively optimize performance, you must identify the areas that yield the largest performance increases over the widest variety of situations. Focus your analysis on these areas.

Many of the examples within this publication illustrate a query written through either an SQL or an OPNQRYF query interface. The interface chosen for a particular example does not indicate an operation exclusive to that query interface, unless explicitly noted. It is only an illustration of one possible query interface. Most examples can be easily rewritten into whatever query interface that you prefer.

Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

What's new for IBM i 7.3

The following information was added or updated in this release of the information:

- Support for System-period temporal tables: [“System-period temporal tables”](#) on page 105
- Support for EVI RRN probe: [“Encoded vector index RRN probe”](#) on page 24
- Index advice for EVI RRN probe: [“Index advice for Encoded Vector Index RRN Probe Plans”](#) on page 158
- Support for Window temporary result objects and window scans: [“Window”](#) on page 59
- Support for Temporary indexed lists: [“Temporary Indexed List”](#) on page 56
- Support for temporary index list scan and index merge: [“Temporary Indexed List scan and Index Merge”](#) on page 57
- SQL plan cache properties topic: [“SQL plan cache properties”](#) on page 174 has been updated.
- Indexes and column statistics topic: [“Indexes and column statistics”](#) on page 245 has been updated.
- Start database monitor (STRDBMON) command topic: [“Start Database Monitor \(STRDBMON\) command”](#) on page 141 has been updated.
- Environmental attributes that you can modify through the QAQQINI file : [“Controlling queries dynamically with the query options file QAQQINI”](#) on page 186 has been updated.
- [“QAQQINI query options”](#) on page 190 has been updated.
- Database monitor records that have additions in 7.3:
 - [“Database monitor view 1000 - SQL Information”](#) on page 910
 - [“Database monitor view 3001 - Index Used”](#) on page 932
 - [“Database monitor view 3014 - Generic QQ Information”](#) on page 970
 - [“Database monitor view 3018 - STRDBMON/ENDDDBMON”](#) on page 982
- **New services**
 - AUTHORITY_COLLECTION view: [“AUTHORITY_COLLECTION view”](#) on page 691
 - ENVIRONMENT_VARIABLE_INFO view: [“ENVIRONMENT_VARIABLE_INFO view”](#) on page 431
 - OUTPUT_QUEUE_INFO view: [“OUTPUT_QUEUE_INFO view”](#) on page 749

- SERVICES_INFO table and Db2 PTF Group level dependency information: [“SERVICES_INFO table” on page 456](#)
- **Updated services**
 - DISPLAY_JOURNAL table function accepts ending values as input parameters to limit the entries returned: [“DISPLAY_JOURNAL table function” on page 598](#)
 - NETSTAT_INFO view has been updated to return more information: [“NETSTAT_INFO view” on page 482](#)
 - NETSTAT_INTERFACE_INFO view has been updated to return more information: [“NETSTAT_INTERFACE_INFO view” on page 489](#)
 - NETSTAT_JOB_INFO view has been updated to return more information: [“NETSTAT_JOB_INFO view” on page 498](#)
 - NETSTAT_ROUTE_INFO view has been updated to return more information: [“NETSTAT_ROUTE_INFO view” on page 499](#)
 - OBJECT_STATISTICS table function has been updated to return more information: [“OBJECT_STATISTICS table function” on page 643](#)
 - PTF_INFO view has been updated to return more information: [“PTF_INFO view” on page 688](#)
 - SERVER_SBS_ROUTING view shows information about more servers: [“SERVER_SBS_ROUTING view” on page 507](#)
 - SET_SERVER_SBS_ROUTING procedure allows you to configuring more servers: [“SET_SERVER_SBS_ROUTING procedure” on page 510](#)
 - USER_INFO view has new columns for authority collection details: [“USER_INFO view” on page 717](#)
 - SYSLIMITS view returns more information about each object: [“SYSLIMITS view” on page 788](#)
 - An additional limit is tracked: Maximum extended dynamic package size: [“System Health Services” on page 782](#)

What’s new since the first 7.3 publication

The following revisions or additions have been made to the Performance and query optimization documentation since the first 7.3 publication:

• May 2022 update

- Query Supervisor example exit programs written in CL: [“Query Supervisor example exit programs” on page 210](#)
- **New services**
 - ACTIVATION_GROUP_INFO table function: [“ACTIVATION_GROUP_INFO table function” on page 391](#)
 - ASSOCIATE_JOURNAL_RECEIVER table function: [“ASSOCIATE_JOURNAL_RECEIVER table function” on page 549](#)
 - AUDIT_JOURNAL_JS, AUDIT_JOURNAL_OM, and AUDIT_JOURNAL_ST table functions return entry specific data: [“AUDIT_JOURNAL_JS table function” on page 582](#), [“AUDIT_JOURNAL_OM table function” on page 587](#), and [“AUDIT_JOURNAL_ST table function” on page 592](#)
 - BINDING_DIRECTORY_INFO view: [“BINDING_DIRECTORY_INFO view” on page 393](#)
 - DNS_LOOKUP table function: [“DNS_LOOKUP table function” on page 479](#)
 - ELECTRONIC_SERVICE_AGENT_INFO view: [“ELECTRONIC_SERVICE_AGENT_INFO view” on page 678](#)
 - JOURNAL_RECEIVER_INFO view: [“JOURNAL_RECEIVER_INFO view” on page 620](#)
 - MTI_INFO table function: [“MTI_INFO table function” on page 332](#)
 - REMOTE_JOURNAL_INFO view: [“REMOTE_JOURNAL_INFO view” on page 630](#)
 - SPOOLED_FILE_INFO table function: [“SPOOLED_FILE_INFO table function” on page 755](#)

- SQL_CHECK_FUNCTION_USAGE scalar function: [“SQL_CHECK_FUNCTION_USAGE scalar function” on page 716](#)
- SQL_CHECK_SPECIAL_AUTHORITY scalar function: [“SQL_CHECK_SPECIAL_AUTHORITY scalar function” on page 717](#)
- SYSTEM_ACTIVITY_INFO table function: [“SYSTEM_ACTIVITY_INFO table function” on page 863](#)
- **Updated services**
 - ACTIVE_JOB_INFO table function can return a WORK subset of data: [“ACTIVE_JOB_INFO table function” on page 793](#)
 - GENERATE_PDF scalar function supports *LAST option: [“GENERATE_PDF scalar function” on page 737](#)
 - GENERATE_SQL and GENERATE_SQL_OBJECTS procedures have a new option for generating comments for tables: [“GENERATE_SQL procedure” on page 356](#) and [“GENERATE_SQL_OBJECTS procedure” on page 365](#)
 - JOBLOG_INFO table function returns the qualified job name: [“JOBLOG_INFO table function” on page 655](#)
 - SYSTEM_STATUS table function, SYSTEM_STATUS_INFO view, and SYSTEM_STATUS_INFO_BASIC view no longer return information for the AVERAGE_CPU_RATE, AVERAGE_CPU_UTILIZATION, MINIMUM_CPU_UTILIZATION, and MAXIMUM_CPU_UTILIZATION columns due to performance concerns. This information is now returned through the SYSTEM_ACTIVITY_INFO table function. These columns can be explicitly requested by using an environment variable switch.
ADDENVVAR ENVVAR(QIBM_DB2_SYSTEM_STATUS_INFO) VALUE(FULL) LEVEL(*JOB)
[“SYSTEM_STATUS table function” on page 864](#), [“SYSTEM_STATUS_INFO view” on page 872](#), and [“SYSTEM_STATUS_INFO_BASIC view” on page 878](#)
- **September 2021 update**
 - New system limits global variables provide the ability to remove rows from the system limits table by the age of rows, in days: [“QIBM_SYSTEM_LIMITS global variables” on page 792](#)
 - **New services**
 - ACTIVE_QUERY_INFO table function returns information about active SQL Query Engine (SQE) queries: [“ACTIVE_QUERY_INFO table function” on page 322](#)
 - AUDIT_JOURNAL_CD, AUDIT_JOURNAL_CO, AUDIT_JOURNAL_CP, AUDIT_JOURNAL_DO, AUDIT_JOURNAL_EV, AUDIT_JOURNAL_GR, AUDIT_JOURNAL_SV table functions return the entry specific data for audit journal entries: [“Audit journal entry services” on page 552](#)
 - COLLECTION_SERVICES_INFO returns the configuration properties for Collection Services: [“COLLECTION_SERVICES_INFO view” on page 666](#)
 - SYSFILES view returns database file information: [“SYSFILES view” on page 380](#)
 - WORKLOAD_GROUP_INFO view returns information about workload groups: [“WORKLOAD_GROUP_INFO view” on page 885](#)
 - **Updated services**
 - ACTIVE_JOB_INFO includes a column for the workload group: [“ACTIVE_JOB_INFO table function” on page 793](#)
 - DISPLAY_JOURNAL table function adds a column with a new name to return the current user value: [“DISPLAY_JOURNAL table function” on page 598](#)
 - PARSE_STATEMENT returns the procedure name for a CALL statement: [“PARSE_STATEMENT table function” on page 315](#)
 - SECURITY_INFO view returns additional security-related system values: [“SECURITY_INFO view” on page 712](#)
- **April 2021 update**

- The Query Supervisor allows real time monitoring of resources used by queries: [“Query Supervisor” on page 205](#)
- **New services**
 - AUDIT_JOURNAL_AF, AUDIT_JOURNAL_CA, AUDIT_JOURNAL_OW, AUDIT_JOURNAL_PW table functions return the entry specific data for audit journal entries: [“Audit journal entry services” on page 552](#)
 - CHANGE_USER_PROFILE table function allows some user profile attributes to be changed: [“CHANGE_USER_PROFILE table function” on page 698](#)
 - END_IDLE_SQE_THREADS procedure ends SQE threads not being used by a job: [“END_IDLE_SQE_THREADS procedure ” on page 352](#)
 - GENERATE_PDF scalar function generates a PDF from a spooled file: [“GENERATE_PDF scalar function” on page 737](#)
 - MESSAGE_QUEUE_INFO table function allows filtering of the messages in a message queue: [“MESSAGE_QUEUE_INFO table function” on page 660](#)
 - QCMDXC scalar function executes a CL command from within a query: [“QCMDXC scalar function” on page 452](#)
 - SECURITY_INFO view provides system security settings: [“SECURITY_INFO view” on page 712](#)
 - SEND_MESSAGE procedure sends an informational message to the QSYSOPR message queue: [“SEND_MESSAGE procedure ” on page 665](#)
 - USER_INDEX_INFO view returns the attributes of a *USRIDX: [“USER_INDEX_INFO view” on page 466](#)
 - USER_INDEX_ENTRIES table function returns the contents of a *USRIDX: [“USER_INDEX_ENTRIES table function” on page 465](#)
 - USER_INFO_BASIC view is a faster version of USER_INFO, with fewer columns: [“USER_INFO_BASIC view” on page 726](#)
 - USER_SPACE_INFO view returns the attributes of a *USRSPC: [“USER_SPACE_INFO view” on page 468](#)
 - USER_SPACE table function returns the content of a *USRSPC object: [“USER_SPACE table function” on page 467](#)
- **Updated services**
 - ACTIVE_JOB_INFO and JOB_INFO table functions return individual columns for the parts of a qualified job name: [“ACTIVE_JOB_INFO table function” on page 793](#) and [“JOB_INFO table function” on page 821](#)
 - CLEAR_PLAN_CACHE procedure allows a plan for a query to be cleared: [“CLEAR_PLAN_CACHE procedure ” on page 338](#)
 - DISPLAY_JOURNAL table function returns more SYSLOG information: [“DISPLAY_JOURNAL table function” on page 598](#)
 - DUMP_PLAN_CACHE procedure allows a specific plan to be dumped: [“DUMP_PLAN_CACHE procedure” on page 340](#)
 - DUMP_PLAN_CACHE_TOPN procedure supports additional categories of plans to be dumped: [“DUMP_PLAN_CACHE_TOPN procedure” on page 341](#)
 - FIRMWARE_CURRENCY view includes information shown by the DSPFMWSTS command: [“FIRMWARE_CURRENCY view” on page 680](#)
 - GENERATE_SQL and GENERATE_SQL_OBJECTS procedures write output to a stream file: [“GENERATE_SQL procedure” on page 356](#) and [“GENERATE_SQL_OBJECTS procedure” on page 365](#)
 - IFS_OBJECT_STATISTICS table function return the symbolic link for an object: [“IFS_OBJECT_STATISTICS table function” on page 528](#)
 - OBJECT_STATISTICS table function supports a generic object name: [“OBJECT_STATISTICS table function” on page 643](#)

- SYSDISKSTAT view and SYSDISKSTAT table function have additional information about SSDs, including remaining lifetime: [“SYSDISKSTAT view” on page 775](#) and [“SYSDISKSTAT table function” on page 772](#)
- USER_INFO view indicates whether the user profile is disabled for IBM i NetServer use: [“USER_INFO view” on page 717](#)
- System limit alerting is instrumented for the maximum number of spooled files limit. In addition, global variables can be used to set the alerting level percent for each limit: [“System limit alerts” on page 785](#)
- **October 2020 update**
 - The SELECTIVITY clause allows predicate selectivity values to be assigned within a query: [“Use SELECTIVITY to supply missing information” on page 305](#)
 - **New services**
 - COMMUNICATIONS_ENTRY_INFO view: [“COMMUNICATIONS_ENTRY_INFO view” on page 810](#)
 - DATA_QUEUE_ENTRIES table function: [“DATA_QUEUE_ENTRIES table function” on page 408](#)
 - EXIT_POINT_INFO view: [“EXIT_POINT_INFO view” on page 432](#)
 - EXIT_PROGRAM_INFO view: [“EXIT_PROGRAM_INFO view” on page 434](#)
 - IFS_READ table function: [“IFS_READ, IFS_READ_BINARY, and IFS_READ_UTF8 table functions” on page 540](#)
 - IFS_WRITE procedure: [“IFS_WRITE, IFS_WRITE_BINARY, and IFS_WRITE_UTF8 procedures” on page 542](#)
 - JOURNAL_INHERIT_RULES view: [“JOURNAL_INHERIT_RULES view” on page 635](#)
 - JOURNALED_OBJECTS view: [“JOURNALED_OBJECTS view” on page 628](#)
 - OPEN_FILES table function: [“OPEN_FILES table function” on page 846](#)
 - RELATED_OBJECTS table function: [“RELATED_OBJECTS table function” on page 376](#)
 - SERVER_SHARE_INFO view: [“SERVER_SHARE_INFO view” on page 544](#)
 - SOFTWARE_PRODUCT_INFO view: [“SOFTWARE_PRODUCT_INFO view” on page 673](#)
 - SYSLIMITS_BASIC view: [“SYSLIMITS_BASIC view” on page 790](#)
 - WATCH_DETAIL table function: [“WATCH_DETAIL table function” on page 469](#)
 - WATCH_INFO view: [“WATCH_INFO view” on page 474](#)
 - **Updated services**
 - ACTIVE_JOB_INFO table function returns a column about open files: [“ACTIVE_JOB_INFO table function” on page 793](#)
 - ANALYZE_CATALOG table function has an option to return cross reference server status: [“ANALYZE_CATALOG table function” on page 346](#)
 - DB_TRANSACTION_INFO view returns additional columns for pending changes: [“DB_TRANSACTION_INFO view” on page 414](#)
 - DISPLAY_JOURNAL table function allows multiple object names and a fully qualified job name as input parameters, and returns more SYSLOG information: [“DISPLAY_JOURNAL table function” on page 598](#)
 - GROUP_PTF_DETAILS and GROUP_PTF_CURRENCY views return dates as a date data type column: [“GROUP_PTF_DETAILS view” on page 683](#) and [“GROUP_PTF_CURRENCY view” on page 682](#)
 - LIBRARY_INFO table function has a new detailed information parameter and returns additional journal columns: [“LIBRARY_INFO table function” on page 639](#)
 - SPLIT table function has an escape character parameter: [“SPLIT table function” on page 460](#)
 - SYSDISKSTAT view has additional information added, including statistics. The new SYSDISKSTAT table function allows the statistics to be reset: [“SYSDISKSTAT view” on page 775](#) and [“SYSDISKSTAT table function” on page 772](#)

- SYSTEM_STATUS table function and SYSTEM_STATUS_INFO view return job table information and system-wide journal information. A new view, SYSTEM_STATUS_INFO_BASIC does not return the job table columns: [“SYSTEM_STATUS table function” on page 864](#), [“SYSTEM_STATUS_INFO view” on page 872](#), and [“SYSTEM_STATUS_INFO_BASIC view” on page 878](#)
- **SYSTOOLS services are documented**
 - DELETE_OLD_SPOOLED_FILES procedure: [“DELETE_OLD_SPOOLED_FILES procedure” on page 735](#)
 - LPRINTF procedure: [“LPRINTF procedure” on page 436](#)
 - VALIDATE_DATA table functions: [“VALIDATE_DATA, VALIDATE_DATA_FILE, and VALIDATE_DATA_LIBRARY table functions” on page 390](#)
- **April 2020 update**
 - **New services**
 - ANALYZE_CATALOG table function: [“ANALYZE_CATALOG table function” on page 346](#)
 - AUTOSTART_JOB_INFO view: [“AUTOSTART_JOB_INFO view” on page 810](#)
 - CERTIFICATE_INFO table function: [“CERTIFICATE_INFO table function” on page 695](#)
 - DB_TRANSACTION_INFO view: [“DB_TRANSACTION_INFO view” on page 414](#)
 - HTTP_SERVER_INFO view: [“HTTP_SERVER_INFO view” on page 481](#)
 - IFS_OBJECT_STATISTICS table function: [“IFS_OBJECT_STATISTICS table function” on page 528](#)
 - JOB_LOCK_INFO table function: [“JOB_LOCK_INFO table function” on page 833](#)
 - LIBRARY_INFO table function: [“LIBRARY_INFO table function” on page 639](#)
 - PRESTART_JOB_INFO view: [“PRESTART_JOB_INFO view” on page 848](#)
 - PRESTART_JOB_STATISTICS table function: [“PRESTART_JOB_STATISTICS table function” on page 851](#)
 - ROUTING_ENTRY_INFO view: [“ROUTING_ENTRY_INFO view” on page 855](#)
 - SUBSYSTEM_INFO view: [“SUBSYSTEM_INFO view” on page 860](#)
 - SUBSYSTEM_POOL_INFO view: [“SUBSYSTEM_POOL_INFO view” on page 862](#)
 - WORKSTATION_INFO view: [“WORKSTATION_INFO view” on page 886](#)
 - **Updated services**
 - ACTIVE_JOB_INFO table function returns enhanced job type: [“ACTIVE_JOB_INFO table function” on page 793](#)
 - DISPLAY_JOURNAL table function returns more SYSLOG information: [“DISPLAY_JOURNAL table function” on page 598](#)
 - JOBLOG_INFO table function returns message key: [“JOBLOG_INFO table function” on page 655](#)
 - OBJECT_PRIVILEGES table function has been documented: [“OBJECT_PRIVILEGES table function” on page 706](#)
 - OBJECT_PRIVILEGES view has 2 new columns: [“OBJECT_PRIVILEGES view” on page 709](#)
 - OVERRIDE_QAQQINI procedure allows optional parameters: [“OVERRIDE_QAQQINI procedure” on page 314](#)
 - SYSTEM_STATUS table function and SYSTEM_STATUS_INFO view return additional columns: [“SYSTEM_STATUS table function” on page 864](#) and [“SYSTEM_STATUS_INFO view” on page 872](#)
 - USER_INFO view has new columns: [“USER_INFO view” on page 717](#)
- **October 2019 update**
 - New option to suppress alter table inquiry message. [“QAQQINI query options” on page 190](#)
 - **New services**
 - ACTIVE_DB_CONNECTIONS table function: [“ACTIVE_DB_CONNECTIONS table function” on page 477](#)

- DATA_QUEUE_INFO view: [“DATA_QUEUE_INFO view” on page 411](#)
- SEND_DATA_QUEUE procedure: [“SEND_DATA_QUEUE, SEND_DATA_QUEUE_BINARY, and SEND_DATA_QUEUE_UTF8 procedures” on page 455](#)
- RECEIVE_DATA_QUEUE table function: [“RECEIVE_DATA_QUEUE table function” on page 453](#)
- CLEAR_DATA_QUEUE procedure: [“CLEAR_DATA_QUEUE procedure” on page 404](#)
- BOUND_MODULE_INFO view: [“BOUND_MODULE_INFO view” on page 394](#)
- BOUND_SRVPGM_INFO view: [“BOUND_SRVPGM_INFO view” on page 403](#)
- PROGRAM_EXPORT_IMPORT_INFO view: [“PROGRAM_EXPORT_IMPORT_INFO view” on page 437](#)
- PROGRAM_INFO view: [“PROGRAM_INFO view” on page 438](#)
- IFS_JOB_INFO table function: [“IFS_JOB_INFO table function” on page 516](#)
- IFS_OBJECT_LOCK_INFO table function: [“IFS_OBJECT_LOCK_INFO table function” on page 519](#)
- IFS_OBJECT_REFERENCES_INFO table function: [“IFS_OBJECT_REFERENCES_INFO table function” on page 525](#)
- IFS_OBJECT_STATISTICS table function: [“IFS_OBJECT_STATISTICS table function” on page 528](#)
- OBJECT_OWNERSHIP view: [“OBJECT_OWNERSHIP view” on page 704](#)
- SERVER_SBS_CONFIGURATION view: [“SERVER_SBS_CONFIGURATION view” on page 506](#)
- SWAP_DYNUSRPRF procedure: [“SWAP_DYNUSRPRF procedure” on page 379](#)
- **Updated services**
 - OBJECT_PRIVILEGES view: [“OBJECT_PRIVILEGES view” on page 709](#)
 - OBJECT_STATISTICS table function returns additional columns: [“OBJECT_STATISTICS table function” on page 643](#)
 - Support for server routing by IP address added to SET_SERVER_SBS_ROUTING: [“SET_SERVER_SBS_ROUTING procedure” on page 510](#)
 - SYSTEM_STATUS table function and SYSTEM_STATUS_INFO view return additional columns: [“SYSTEM_STATUS table function” on page 864](#) and [“SYSTEM_STATUS_INFO view” on page 872](#)
 - JOBLOG_INFO table function has optional error handling parameter: [“JOBLOG_INFO table function” on page 655](#)
- **April 2019 update**
 - **New services**
 - ASP_JOB_INFO view: [“ASP_JOB_INFO view” on page 767](#)
 - DATA_AREA_INFO view and table function: [“DATA_AREA_INFO view” on page 407](#), [“DATA_AREA_INFO table function” on page 406](#)
 - FIRMWARE_CURRENCY view: [“FIRMWARE_CURRENCY view” on page 680](#)
 - MESSAGE_FILE_DATA view: [“MESSAGE_FILE_DATA view” on page 657](#)
 - SPLIT table function: [“SPLIT table function” on page 460](#)
 - SPOOLED_FILE_DATA table function: [“SPOOLED_FILE_DATA table function” on page 754](#)
 - The HTTP functions in SYSTOOLS have been documented: [“HTTP function overview” on page 890](#)
 - **Updated services**
 - The maximum table size has been added as a tracked system limit and as a limit that sends alerts: [“System Health Services” on page 782](#), [“System limit alerts” on page 785](#)
 - ASP_INFO returns the relational database name for SYSBASE: [“ASP_INFO view” on page 760](#)
 - GET_JOB_INFO returns additional columns: [“GET_JOB_INFO table function” on page 812](#)
 - OBJECT_PRIVILEGES view returns authorization list: [“OBJECT_PRIVILEGES view” on page 709](#)
 - PARSE_STATEMENT table function supports DROP statements and returns information about referential constraints: [“PARSE_STATEMENT table function” on page 315](#)

- WLM_SET_CLIENT_INFO procedure uses defaults for unspecified parameters: [“WLM_SET_CLIENT_INFO procedure” on page 320](#)
- **August 2018 update**
 - **New services**
 - GENERATE_SQL_OBJECTS procedure: [“GENERATE_SQL_OBJECTS procedure” on page 365](#)
 - JOB_DESCRIPTION_INFO view: [“JOB_DESCRIPTION_INFO view” on page 814](#)
 - OUTPUT_QUEUE_ENTRIES_BASIC view: [“OUTPUT_QUEUE_ENTRIES_BASIC view” on page 747](#)
 - **Updated services**
 - ACTIVE_JOB_INFO table function optionally returns more detailed information: [“ACTIVE_JOB_INFO table function” on page 793](#)
 - NETSTAT_INFO view and NETSTAT_JOB_INFO view return port names from service table entries: [“NETSTAT_INFO view” on page 482](#) and [“NETSTAT_JOB_INFO view” on page 498](#)
 - PARSE_STATEMENT table function supports some DDL references: [“PARSE_STATEMENT table function” on page 315](#)
- **October 2017 update**
 - **New services**
 - ASP_INFO view: [“ASP_INFO view” on page 760](#)
 - ASP_VARY_INFO view: [“ASP_VARY_INFO view” on page 768](#)
 - JOB_QUEUE_INFO view: [“JOB_QUEUE_INFO view” on page 836](#)
 - STACK_INFO table function: [“STACK_INFO table function” on page 461](#)
 - **Updated services**
 - DISPLAY_JOURNAL and HISTORY_LOG_INFO include syslog information: [“DISPLAY_JOURNAL table function” on page 598](#) and [“HISTORY_LOG_INFO table function” on page 650](#)
 - OVERRIDE_QAQQINI procedure has been fully documented: [“OVERRIDE_QAQQINI procedure” on page 314](#)
 - System limit notifications: [“System limit alerts” on page 785](#)
- **March 2017 update**
 - **New services**
 - AUTHORIZATION_LIST_INFO view: [“AUTHORIZATION_LIST_INFO view” on page 691](#)
 - AUTHORIZATION_LIST_USER_INFO view: [“AUTHORIZATION_LIST_USER_INFO view” on page 693](#)
 - OBJECT_PRIVILEGES view: [“OBJECT_PRIVILEGES view” on page 709](#)
 - MESSAGE_QUEUE_INFO view: [“MESSAGE_QUEUE_INFO view” on page 663](#)
 - LICENSE_EXPIRATION_CHECK procedure: [“LICENSE_EXPIRATION_CHECK procedure” on page 670](#)
 - SET_PASE_SHELL_INFO procedure: [“SET_PASE_SHELL_INFO procedure” on page 459](#)
 - **Updated services**
 - USER_INFO has new columns for supplemental group profile information and the PASE shell: [“USER_INFO view” on page 717](#)
 - LICENSE_INFO view has a new column indicating the install status: [“LICENSE_INFO view” on page 670](#)
 - RESET_TABLE_INDEX_STATISTICS procedure has a new option to remove rows from the index advice tracking table: [“RESET_TABLE_INDEX_STATISTICS procedure” on page 336](#)
- **November 2016 update**
 - STATEMENT DETERMINISTIC option has been added for functions. [“QAQQINI query options” on page 190](#)

– New services

- HISTORY_LOG_INFO table function: [“HISTORY_LOG_INFO table function” on page 650](#)
- JOB_INFO table function: [“JOB_INFO table function” on page 821](#)
- PARSE_STATEMENT table function: [“PARSE_STATEMENT table function” on page 315](#)

– Updated services

- DISPLAY_JOURNAL table function honors row and column access control: [“DISPLAY_JOURNAL table function” on page 598](#)
- GET_JOB_INFO table function has new columns for prestart job information: [“GET_JOB_INFO table function” on page 812](#)
- GROUP_PTF_CURRENCY view returns a new value to indicate PTFs will be current with the next IPL: [“GROUP_PTF_CURRENCY view” on page 682](#)
- GROUP_PTF_CURRENCY and GROUP_PTF_DETAILS views have been updated to access a new XML feed: [“GROUP_PTF_CURRENCY view” on page 682](#) and [“GROUP_PTF_DETAILS view” on page 683](#)
- OBJECT_STATISTICS table function added an option to efficiently return a list of libraries: [“OBJECT_STATISTICS table function” on page 643](#)

What’s new since the first 7.2 publication

The following revisions or additions have been made to the Performance and query optimization documentation since the first 7.2 publication:

• October 2015 update

– New services

- GROUP_PTF_DETAILS view: [“GROUP_PTF_DETAILS view” on page 683](#)
- LICENSE_INFO view: [“LICENSE_INFO view” on page 670](#)
- MEDIA_LIBRARY_INFO view: [“MEDIA_LIBRARY_INFO view” on page 770](#)
- MEMORY_POOL table function: [“MEMORY_POOL table function” on page 840](#)
- MEMORY_POOL_INFO view: [“MEMORY_POOL_INFO view” on page 842](#)
- NETSTAT_INFO view: [“NETSTAT_INFO view” on page 482](#)
- NETSTAT_INTERFACE_INFO view: [“NETSTAT_INTERFACE_INFO view” on page 489](#)
- NETSTAT_JOB_INFO view: [“NETSTAT_JOB_INFO view” on page 498](#)
- NETSTAT_ROUTE_INFO view: [“NETSTAT_ROUTE_INFO view” on page 499](#)
- OBJECT_LOCK_INFO view: [“OBJECT_LOCK_INFO view” on page 844](#)
- OUTPUT_QUEUE_ENTRIES table function: [“OUTPUT_QUEUE_ENTRIES table function” on page 738](#)
- OUTPUT_QUEUE_ENTRIES view: [“OUTPUT_QUEUE_ENTRIES view” on page 742](#)
- RECORD_LOCK_INFO view: [“RECORD_LOCK_INFO view” on page 854](#)
- SYSTEM_STATUS table function: [“SYSTEM_STATUS table function” on page 864](#)
- SYSTEM_STATUS_INFO view: [“SYSTEM_STATUS_INFO view” on page 872](#)

– Updated services

- ACTIVE_JOB_INFO table function has been updated to return elapsed time: [“ACTIVE_JOB_INFO table function” on page 793](#)
- DATABASE_MONITOR_INFO view has been updated to describe new filter values: [“DATABASE_MONITOR_INFO view” on page 327](#)
- ENV_SYS_INFO view has been updated to return the total configured memory: [“ENV_SYS_INFO view” on page 480](#)
- GET_JOB_INFO table function has been updated to return the client IP address: [“GET_JOB_INFO table function” on page 812](#)

- SET_SERVER_SBS_ROUTING procedure allows you to configuring the remote command server: [“SET_SERVER_SBS_ROUTING procedure” on page 510](#)
- **May 2015 update**
 - Additional information was added to QOI1 - Insert unique count in the database monitor 1000 record. For details, see: [“Database monitor view 1000 - SQL Information” on page 910](#)
 - Additional options were added to the QAQQINI query option Memory_Pool_Preference. For details, see: [“QAQQINI query options” on page 190](#)
 - CLEAR_PLAN_CACHE procedure. For details, see: [“CLEAR_PLAN_CACHE procedure ” on page 338](#)
 - **New services**
 - ACTIVE_JOB_INFO table function: [“ACTIVE_JOB_INFO table function” on page 793](#)
 - DATABASE_MONITOR_INFO view: [“DATABASE_MONITOR_INFO view” on page 327](#)
 - DRDA_AUTHENTICATION_ENTRY_INFO view: [“DRDA_AUTHENTICATION_ENTRY_INFO view” on page 701](#)
 - JVM_INFO view: [“JVM_INFO view” on page 547](#)
 - SCHEDULED_JOB_INFO view: [“SCHEDULED_JOB_INFO view” on page 857](#)
 - SERVER_SBS_ROUTING view: [“SERVER_SBS_ROUTING view” on page 507](#)
 - SET_JVM procedure: [“SET_JVM procedure” on page 548](#)
 - SET_SERVER_SBS_ROUTING procedure: [“SET_SERVER_SBS_ROUTING procedure” on page 510](#)
 - **Updated services**
 - GET_JOB_INFO table function has been updated to return additional SQL information for a job: [“GET_JOB_INFO table function” on page 812](#)
 - OBJECT_STATISTICS table function has a new optional parameter to specify the name of the object to return. It will also return the long SQL name for an object and has new columns to return the text, the long schema name, and the SQL type of an object: [“OBJECT_STATISTICS table function” on page 643](#)
 - System Health Services has been updated to track index limits: [“System Health Services” on page 782](#)
- **October 2014 update**
 - **Updates to the QAQQINI query options topic**



For details, see [“QAQQINI query options” on page 190](#).
 - **Memory preference controls enhanced for SQL**

For details, see [“Memory preference controls” on page 65](#)
 - The database monitor topic has been updated: [“Monitoring your queries using the Database Monitor ” on page 140](#)
 - The SQL Plan Cache topic has been updated: [“Optimizing performance using the Plan Cache” on page 169](#)
 - **New services**
 - LIBRARY_LIST_INFO view: [“LIBRARY_LIST_INFO view” on page 642](#)
 - REPLY_LIST_INFO view: [“REPLY_LIST_INFO view” on page 664](#)
 - JOURNAL_INFO view: [“JOURNAL_INFO view” on page 612](#)
 - GROUP_PTF_CURRENCY view: [“GROUP_PTF_CURRENCY view” on page 682](#)
 - JOBLOG_INFO table function: [“JOBLOG_INFO table function” on page 655](#)
 - **Tracking of additional file system limits**

For details, see [“System Health Services” on page 782](#)

How to see what's new or changed

To help you see where technical changes have been made, this information uses:

- The  image to mark where new or changed information begins.
- The  image to mark where new or changed information ends.

To find other information about what's new or changed this release, see the [Memo to users](#).



PDF file for Database performance and query optimization

View and print a PDF of this information.

To view or download the PDF version of this document, select [Database performance and query optimization](#).

Other information

You can also view or print any of the following PDF files:


- [Preparing for and Tuning the SQL Query Engine on DB2® for i5/OS](#) 
- [SQL Performance Diagnosis on IBM DB2 Universal Database for iSeries](#) 

Saving PDF files

To save a PDF on your workstation for viewing or printing:

1. Right-click the PDF in your browser (right-click the preceding link).
2. Click the option that saves the PDF locally.
3. Navigate to the directory in which you want to save the PDF.
4. Click **Save**.

Downloading Adobe Reader

You need Adobe Reader installed on your system to view or print these PDF files. You can download a free copy from [Adobe](http://get.adobe.com/reader/) (<http://get.adobe.com/reader/>) .

Query engine overview

IBM Db2 for i provides two query engines to process queries: Classic Query Engine (CQE) and SQL Query Engine (SQE).

SQL-based interfaces, such as ODBC, JDBC, CLI, Query Manager, Net.Data®, RUNSQLSTM, and embedded or interactive SQL, run through SQE. Also by default some non-SQL based interface such as OPNQRYF and Query/400 will run through SQE. The CQE processes queries originating from non-SQL interfaces: QQQQry API. For ease of use, the routing decision for processing the query by either CQE or SQE is pervasive and under the control of the system. The requesting user or application program cannot control or influence this behavior except for non-SQL interfaces through use of a QAQQINI. However, a better understanding of the engines and process that determines which path a query takes can give you a better understanding of query performance.

Within SQE, several more components were created and other existing components were updated. Additionally, new data access methods are possible with SQE that are not supported under CQE.

Related information

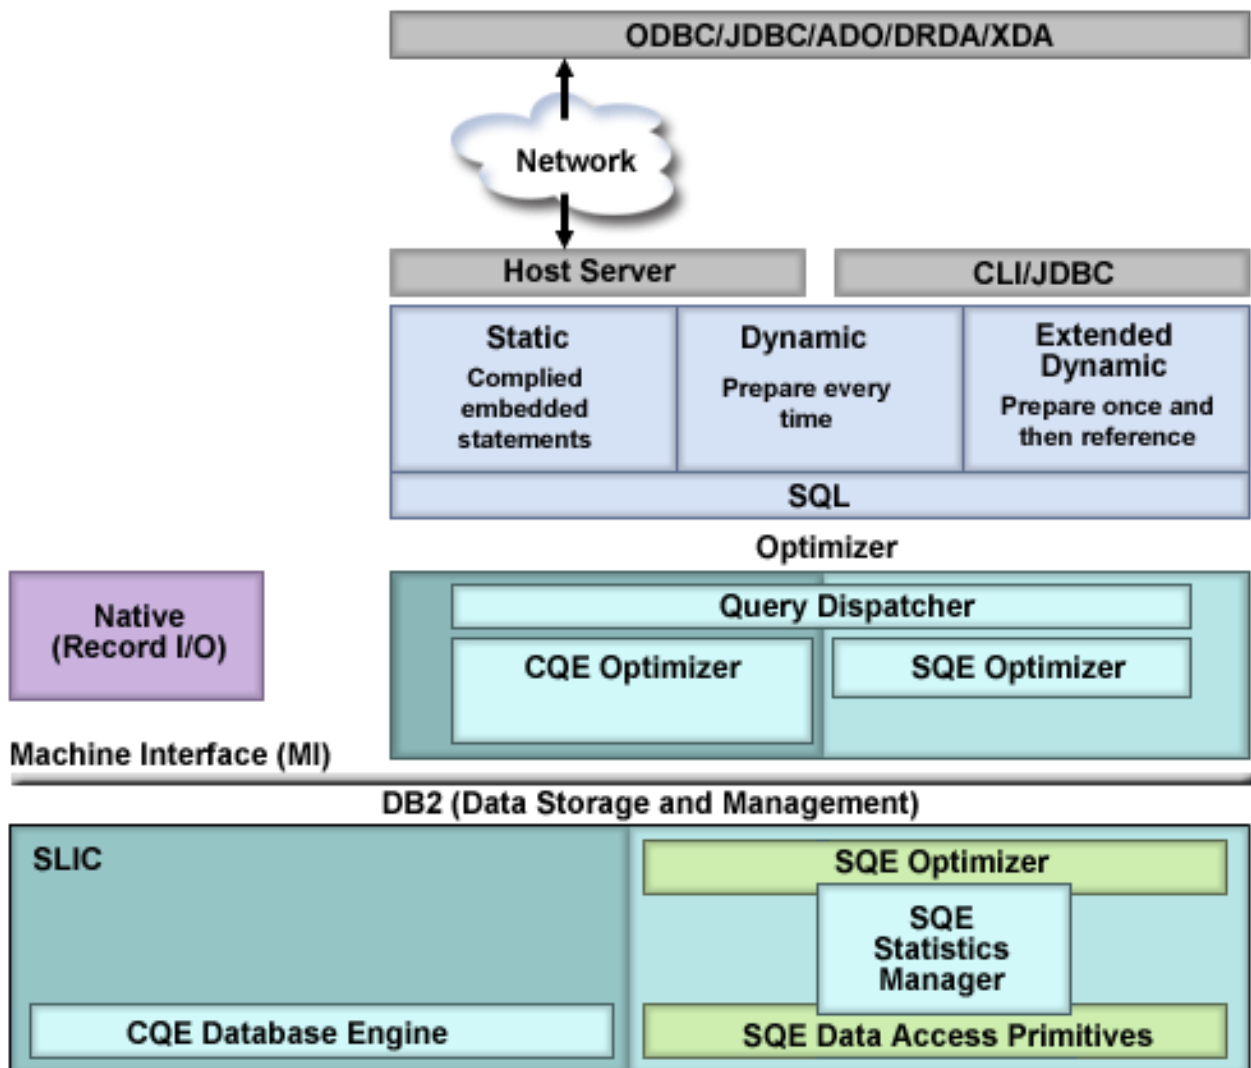
[Embedded SQL programming](#)

[SQL programming](#)
[Query \(QQQRY\) API](#)
[Open Query File \(OPNQRYF\) command](#)
[Run SQL Statements \(RUNSQLSTM\) command](#)

SQE and CQE engines

It is important to understand the implementation differences of query management and processing in CQE versus SQE.

The following figure shows an overview of the IBM Db2 for i architecture. It shows the delineation between CQE and SQE, how query processing is directed by the query dispatcher, and where each SQE component fits. The functional separation of each SQE component is clearly evident. This division of responsibility enables IBM to more easily deliver functional enhancements to the individual components of SQE, as and when required. Notice that most of the SQE Optimizer components are implemented below the MI. This implementation translates into enhanced performance efficiency.



As seen in the previous graphic, the query runs from any query interface to the optimizer and the query dispatcher. The query dispatcher determines whether the query is implemented with CQE or SQE.

Query dispatcher

The function of the dispatcher is to route the query request to either CQE or SQE, depending on the attributes of the query. All queries are processed by the dispatcher. It cannot be bypassed.

Currently, the dispatcher routes queries to SQE unless it finds that the query references or contains any of the following:

- INSERT WITH VALUES statement or the target of an INSERT with subselect statement
- Tables with Read triggers
- Read-only queries with more than 1000 dataspace, or updatable queries with more than 256 dataspace.
- Db2 Multisystem tables
- QQQQry API

For other non-SQL queries, for example Query/400 or OPNQRYF, the routing of the query can be controlled by the QAQQINI SQE_NATIVE_ACCESS option. See "table 46".

Related reference

MQT supported function

Although an MQT can contain almost any query, the optimizer only supports a limited set of query functions when matching MQTs to user specified queries. The user-specified query and the MQT query must both be supported by the SQE optimizer.

Statistics manager

In CQE, the retrieval of statistics is a function of the Optimizer. When the Optimizer needs to know information about a table, it looks at the table description to retrieve the row count and table size. If an index is available, the Optimizer might extract information about the data in the table. In SQE, the collection and management of statistics is handled by a separate component called the statistics manager. The statistics manager leverages all the same statistical sources as CQE, but adds more sources and capabilities.

The statistics manager does not actually run or optimize the query. Instead, it controls the access to the metadata and other information that is required to optimize the query. It uses this information to answer questions posed by the query optimizer. The statistics manager always provides answers to the optimizer. In cases where it cannot provide an answer based on actual existing statistics information, it is designed to provide a predefined answer.

The Statistics manager typically gathers and tracks the following information:

Cardinality of values	The number of unique or distinct occurrences of a specific value in a single column or multiple columns of a table.
Selectivity	Also known as a histogram, this information is an indication of how many rows are selected by any given selection predicate or combination of predicates. Using sampling techniques, it describes the selectivity and distribution of values in a given column of the table.
Frequent values	The top <i>nn</i> most frequent values of a column together with a count of how frequently each value occurs. This information is obtained by using statistical sampling techniques. Built-in algorithms eliminate the possibility of data skewing. For example, NULL values and default values that can influence the statistical values are not taken into account.
Metadata information	Includes the total number of rows in the table, indexes that exist over the table, and which indexes are useful for implementing the particular query.
Estimate of IO operation	An estimate of the amount of IO operations that are required to process the table or the identified index.

The Statistics manager uses a hybrid approach to manage database statistics. Most of this information can be obtained from existing indexes. In cases where the required statistics cannot be gathered from existing indexes, statistical information is constructed on single columns of a table and stored internally. By default, this information is collected automatically by the system, but you can manually control the collection of statistics. Unlike indexes, however, statistics are not maintained immediately as data in the tables change.

Related reference

[Collecting statistics with the statistics manager](#)

The collection of statistics is handled by a separate component called the statistics manager. Statistical information can be used by the query optimizer to determine the best access plan for a query. Since the query optimizer bases its choice of access plan on the statistical information found in the table, it is important that this information is current.

Global Statistics Cache

In SQE, the Db2 Statistics Manager stores actual row counts into a Global Statistics Cache. In this manner, the Statistics Manager refines its estimates over time as it learns where estimates have deviated from actual row counts.

Both completed queries and currently executing queries might be inspected by the “[Adaptive Query Processing](#)” on page 106 (AQP) task, which compares estimated row counts to actual row counts. If there are any significant discrepancies, the AQP task notifies the Db2 Statistics Manager (SM). The SM stores this actual row count (also called observed row count) into a Global Statistics Cache (GSC).

If the query which generated the observed statistic in the GSC is reoptimized, the actual row count estimate is used in determining a new query plan. Further, if a different query asks for the same or a similar row count, the SM could return the stored actual row count from the GSC. Faster query plans can be generated by the query optimizer.

Typically, observed statistics are for complex predicates such as with a join. A simple example is a query joining three files A, B, and C. There is a discrepancy between the estimate and actual row count of the join of A and B. The SM stores an observed statistic into the GSC. Later, if a different join query of A, B, and Z is submitted, the SM recalls the observed statistic of the A and B join. The SM considers that observed statistic in its estimate of the A, B, and Z join.

The Global Statistics Cache is an internal Db2 object, and the contents of it are not directly observable.

Plan cache

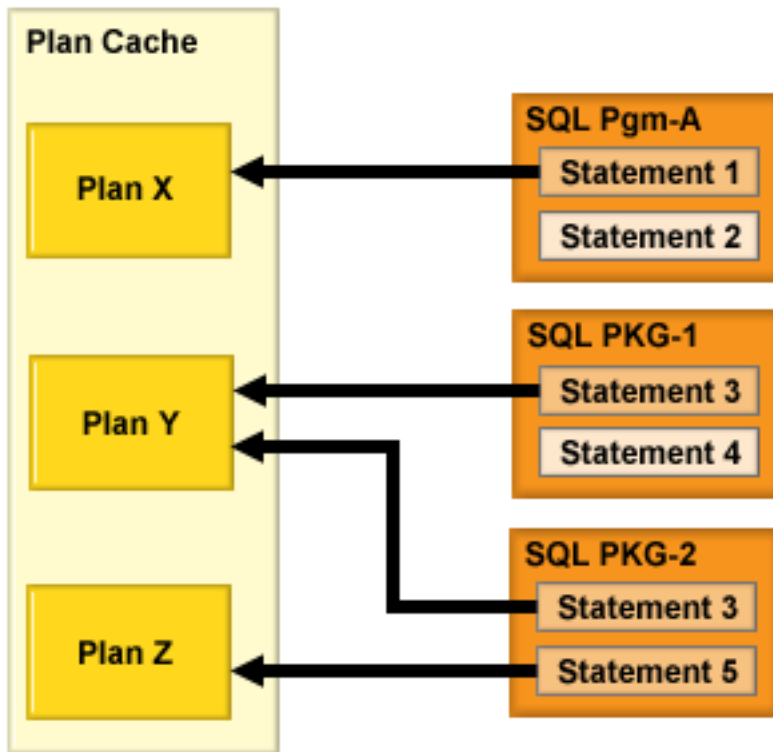
The plan cache is a repository that contains the access plans for queries that were optimized by SQE.

Access plans generated by CQE are not stored in the plan cache; instead, they are stored in SQL packages, the system-wide statement cache, and job cache. The purposes of the plan cache are to:

- Facilitate the reuse of a query access plan when the same query is re-executed
- Store runtime information for subsequent use in future query optimizations
- Provide performance information for analysis and tuning

Once an access plan is created, it is available for use by all users and all queries, regardless of where the query originates. Furthermore, when an access plan is tuned, for example, when creating an index, all queries can benefit from this updated access plan. This updated access plan eliminates the need to re-optimize the query, resulting in greater efficiency.

The following graphic shows the concept of re-usability of the query access plans stored in the plan cache:



As shown in the previous graphic, statements from packages and programs are stored in unique plans in the plan cache. If Statement 3 exists in both SQL package 1 and SQL package 2, the plan is stored once in the plan cache. The plan cache is interrogated each time a query is executed. If an access plan exists that satisfies the requirements of the query, it is used to implement the query. Otherwise a new access plan is created and stored in the plan cache for future use.

The plan cache is automatically updated with new query access plans as they are created. When new statistics or indexes become available, an existing plan is updated the next time the query is run. The plan cache is also automatically updated by the database with runtime information as the queries are run.

Each plan cache entry contains the original query, the optimized query access plan, and cumulative runtime information gathered during the runs of the query. In addition, several instances of query runtime objects are stored with a plan cache entry. These runtime objects are the real executable objects and temporary storage containers (hash tables, sorts, temporary indexes, and so on) used to run the query.

By default the SQE Plan Cache will auto adjust from an initial threshold size of 512 MB to an internally managed maximum. Automatic management of the SQL Plan Cache Threshold Size by the system will not take effect if the plan cache threshold size is explicitly set on the system. See the SQL plan cache properties topic for more information: [“SQL plan cache properties” on page 174](#)

- When processing is initiated to remove plans in the cache due to size constraint, the efficiency rating of the cache is checked. If the rating is too low, the database will automatically increase the plan cache size.
- The plan cache auto-sizing maximum size will not exceed a small percentage of free storage on the system.
- The plan cache auto-sizing will decrease the size if the temporary storage on the machine exceeds a certain percentage.
- The auto-sized adjusted threshold value does not survive an IPL. The default plan cache size is used after an IPL and auto sizing begins again.
- To reset an explicitly set plan cache size in order to allow auto-sizing to take effect, set the plan cache size to zero.

Example:

```
CALL qsys2.change_plan_cache_size(0)
```

When the plan cache exceeds its designated size, a background task is automatically scheduled to remove plans from the plan cache. Access plans are deleted based upon age, how frequently it is used, and how much cumulative resources (CPU/IO) were consumed.

The total number of access plans stored in the plan cache depends largely upon the complexity of the SQL statements that are being executed. The plan cache is cleared when a system Initial Program Load (IPL) is performed.

Multiple access plans for a single SQL statement can be maintained in the plan cache. Although the SQL statement is the primary key into the plan cache, different environmental settings can cause additional access plans to be stored. Examples of these environmental settings include:

- Different SMP Degree settings for the same query
- Different library lists specified for the query tables
- Different settings for the share of available memory for the job in the current pool
- Different ALWCOPYDTA settings
- Different selectivity based on changing host variable values used in selection (WHERE clause)

Currently, the plan cache can maintain a maximum of three different access plans for the same SQL statement. The uniqueness of an SQL statement is determined by an internal representation of the query which includes, but is not limited to, full schema qualification of all objects referenced in query and the data types of columns. As new access plans are created for the same SQL statement, older access plans are discarded to make room for the new access plans. There are, however, certain conditions that can cause an existing access plan to be invalidated. Examples of these conditions include:

- Specifying REOPTIMIZE_ACCESS_PLAN(*YES) or (*FORCE) in the QAQQINI table or in Run SQL Scripts
- Deleting or recreating the table that the access plan refers to
- Deleting an index that is used by the access plan

Related reference

[Effects of the ALWCOPYDTA parameter on database performance](#)

Some complex queries can perform better by using a sort or hashing method to evaluate the query instead of using or creating an index.

[Changing the attributes of your queries](#)

You can modify different types of query attributes for a job with the **Change Query Attributes (CHGQRYA)** CL command. You can also use the System i Navigator Change Query Attributes interface.

[Optimizing performance using the Plan Cache](#)

The SQL Plan Cache contains a wealth of information about the SQE queries being run through the database. Its contents are viewable through the System i Navigator GUI interface. Certain portions of the plan cache can also be modified.

Data access methods

Data access methods are used to process queries and access data.

In general, the query engine has two kinds of raw material with which to satisfy a query request:

- The database objects that contain the data to be queried
- The executable instructions or operations to retrieve and transform the data into usable information

There are only two types of permanent database objects that can be used as source material for a query — tables and indexes. Indexes include binary radix and encoded vector indexes.

In addition, the query engine might need to create temporary objects to hold interim results or references during the execution of an access plan. The Db2 Symmetric Multiprocessing feature provides

the optimizer with additional methods for retrieving data that include parallel processing. Finally, the optimizer uses certain methods to manipulate these objects.

Permanent objects and access methods

There are three basic types of access methods used to manipulate the permanent and temporary database objects -- Create, Scan, and Probe.

The following table lists each object and the access methods that can be performed against that object. The symbols shown in the table are the icons used by Visual Explain.

<i>Table 1. Permanent object data access methods</i>		
Permanent objects	Scan operations	Probe operations
Table	Table scan	Table probe
Radix index	Radix index scan	Radix index probe
Encoded vector index	Encoded vector index symbol table scan	Encoded vector index probe

Table

An SQL table or physical file is the base object for a query. It represents the source of the data used to produce the result set for the query. It is created by the user and specified in the FROM clause (or OPNQRYF FILE parameter).

The optimizer determines the most efficient way to extract the data from the table in order to satisfy the query. These ways could include scanning or probing the table or using an index to extract the data.

Visual explain icon:




Table scan

A table scan is the easiest and simplest operation that can be performed against a table. It sequentially processes all the rows in the table to determine if they satisfy the selection criteria specified in the query. It does this processing in a way to maximize the I/O throughput for the table.

A table scan operation requests large I/Os to bring as many rows as possible into main memory for processing. It also asynchronously pre-fetches the data to make sure that the table scan operation is never waiting for rows to be paged into memory. Table scan however, has a disadvantage in it has to process all the rows in order to satisfy the query. The scan operation itself is efficient if it does not need to perform the I/O synchronously.

<i>Table 2. Table scan attributes</i>	
Data access method	Table scan
Description	Reads all the rows from the table and applies the selection criteria to each of the rows within the table. The rows in the table are processed in no guaranteed order, but typically they are processed sequentially.
Advantages	<ul style="list-style-type: none"> • Minimizes page I/O operations through asynchronous pre-fetching of the rows since the pages are scanned sequentially • Requests a larger I/O to fetch the data efficiently

<i>Table 2. Table scan attributes (continued)</i>	
Data access method	Table scan
Considerations	<ul style="list-style-type: none"> • All rows in the table are examined regardless of the selectivity of the query • Rows marked as deleted are still paged into memory even though none are selected. You can reorganize the table to remove deleted rows.
Likely to be used	<ul style="list-style-type: none"> • When expecting many rows returned from the table • When the number of large I/Os needed to scan is fewer than the number of small I/Os required to probe the table
Example SQL statement	<pre>SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3000 - Table Scan
SMP parallel enabled	Yes
Also referred to as	Table Scan, Preload
Visual Explain icon	

Related concepts

Nested loop join implementation

Db2 for i provides a **nested loop** join method. For this method, the processing of the tables in the join are ordered. This order is called the **join order**. The first table in the final join order is called the **primary table**. The other tables are called **secondary tables**. Each join table position is called a **dial**.

Table probe


A table probe operation is used to retrieve a specific row from a table based upon its row number. The row number is provided to the table probe access method by some other operation that generates a row number for the table.

This can include index operations as well as temporary row number lists or bitmaps. The processing for a table probe is typically random. It requests a small I/O to retrieve only the row in question and does not attempt to bring in any extraneous rows. This method leads to efficient processing for smaller result sets because only rows needed to satisfy the query are processed, rather than scanning all rows.

However, since the sequence of the row numbers is not known in advance, little pre-fetching can be performed to bring the data into main memory. This randomness can result in most of the I/Os associated with table probe to be performed synchronously.

<i>Table 3. Table probe attributes</i>	
Data access method	Table probe
Description	Reads a single row from the table based upon a specific row number. A random I/O is performed against the table to extract the row.

Table 3. Table probe attributes (continued)

Data access method	Table probe
Advantages	<ul style="list-style-type: none"> • Requests smaller I/Os to prevent paging rows into memory that are not needed • Can be used with any access method that generates a row number for the table probe to process
Considerations	Because of the synchronous random I/O the probe can perform poorly when many rows are selected
Likely to be used	<ul style="list-style-type: none"> • When row numbers (from indexes or temporary row number lists) are used, but data from the underlying table is required for further processing of the query • When processing any remaining selection or projection of the values
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (LastName) SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' AND LastName IN ('Smith', 'Jones', 'Peterson') OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3001 Index Used, where QVC14 (Index_Only_Access) set to 'N' indicates that a table probe was used in conjunction with the index access operation.
SMP parallel enabled	Yes
Also referred to as	Table Probe, Preload
Visual Explain icon	

Radix index

An SQL index (or keyed sequence access path) is a permanent object that is created over a table. The index is used by the optimizer to provide a sequenced view of the data for a scan or probe operation.

The rows in the tables are sequenced in the index based upon the key columns specified on the creation of the index. When the optimizer matches a query to index key columns, it can use the index to help satisfy query selection, ordering, grouping, or join requirements.

Typically, using an index also includes a table probe to provide access to columns needed to satisfy the query that cannot be found as index keys. If all the columns necessary to satisfy the query can be found as index keys, then the table probe is not required. The query uses index-only access. Avoiding the table probe can be an important savings for a query. The I/O associated with a table probe is typically the more expensive synchronous random I/O.

Visual Explain icon:

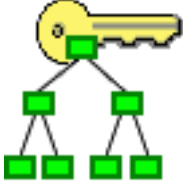


Radix index scan

A radix index scan operation is used to retrieve the rows from a table in a keyed sequence. Like a table scan, all the rows in the index are sequentially processed, but the resulting row numbers are sequenced based upon the key columns.

The sequenced rows can be used by the optimizer to satisfy a portion of the query request (such as ordering or grouping). They can also be used to provide faster throughput by performing selection against the index keys rather than all the rows in the table. Since the index I/Os only contain keys, typically more rows can be paged into memory in one I/O than rows in a table with many columns.

Table 4. Radix index scan attributes	
Data access method	Radix index scan
Description	Sequentially scan and process all the keys associated with the index. Any selection is applied to every key value of the index before a table row
Advantages	<ul style="list-style-type: none"> • Only those index entries that match any selection continue to be processed • Potential to extract all the data from the index key values, thus eliminating the need for a Table Probe • Returns the rows back in a sequence based upon the keys of the index
Considerations	Generally requires a Table Probe to be performed to extract any remaining columns required to satisfy the query. Can perform poorly when many rows are selected because of the random I/O associated with the Table Probe.
Likely to be used	<ul style="list-style-type: none"> • When asking for or expecting only a few rows to be returned from the index • When sequencing the rows is required for the query (for example, ordering or grouping) • When the selection columns cannot be matched against the leading key columns of the index
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (LastName, WorkDept) SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' ORDER BY LastName OPTIMIZE FOR 30 ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3001 Index Used, where QQKP (Index_Probe_Used) set to 'N' will indicate an index scan operation. Preload indicated by QVPARPL = 'Y'. Distinct Probe indicated by QVC11 = 'Y'.
SMP parallel enabled	Yes

<i>Table 4. Radix index scan attributes (continued)</i>	
Data access method	Radix index scan
Also referred to as	Index Scan Index Scan, Preload Index Scan, Distinct Index Scan Distinct, Preload Index Scan, Key Selection
Visual Explain icon	

Related reference

Effects of the ALWCPYDTA parameter on database performance
 Some complex queries can perform better by using a sort or hashing method to evaluate the query instead of using or creating an index.

Radix index probe

A radix index probe operation is used to retrieve the rows from a table in a keyed sequence. The main difference between the radix index probe and the scan is that the rows returned are first identified by a probe operation to subset them.

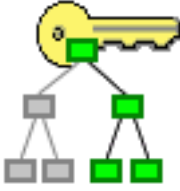
The optimizer attempts to match the columns used for some or all the selection against the leading keys of the index. It then rewrites the selection into a series of ranges that can be used to probe directly into the index key values. Only those keys from the series of ranges are paged into main memory.

The resulting row numbers generated by the probe can then be further processed by any remaining selection against the index keys or a table probe operation. This method provides for quick access to only the rows of the index that satisfy the selection.

The main function of a radix index probe is to provide quick selection against the index keys. In addition, the row sequencing can be used to satisfy other portions of the query, such as ordering or grouping. Since the index I/Os are only for rows that match the probe selection, no extraneous processing is performed on rows that do not match. This savings in I/Os against rows that are not a part of the result set is one of the primary advantages for this operation.

<i>Table 5. Radix index probe attributes</i>	
Data access method	Radix index probe
Description	The index is quickly probed based upon the selection criteria that were rewritten into a series of ranges. Only those keys that satisfy the selection are used to generate a table row number.
Advantages	<ul style="list-style-type: none"> • Only those index entries that match any selection continue to be processed • Provides quick access to the selected rows • Potential to extract all the data from the index key values, thus eliminating the need for a Table Probe • Returns the rows back in a sequence based upon the keys of the index

Table 5. Radix index probe attributes (continued)

Data access method	Radix index probe
Considerations	Generally requires a Table Probe to be performed to extract any remaining columns required to satisfy the query. Can perform poorly when many rows are selected because of the random I/O associated with the Table Probe.
Likely to be used	<ul style="list-style-type: none"> • When asking for or expecting only a few rows to be returned from the index • When sequencing the rows is required the query (for example, ordering or grouping) • When the selection columns match the leading key columns of the index
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (LastName, WorkDept) SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' AND LastName IN ('Smith', 'Jones', 'Peterson') OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	<p>QQRID 3001 Index Used where QQKP (Index_Probe_Used) set to 'Y' will indicate an index probe operation.</p> <p>Preload indicated by QVPARPL = 'Y'</p> <p>Distinct Probe indicated by QVC11 = 'Y'</p>
SMP parallel enabled	Yes
Also referred to as	<p>Index Probe</p> <p>Index Probe, Preload</p> <p>Index Probe, Distinct</p> <p>Index Probe Distinct, Preload</p> <p>Index Probe, Key Positioning</p> <p>Index Scan, Key Row Positioning</p>
Visual Explain icon	

The following example illustrates a query where the optimizer might choose the radix index probe access method:

```
CREATE INDEX X1 ON Employee (LastName, WorkDept)

SELECT * FROM Employee
WHERE WorkDept BETWEEN 'A01' AND 'E01'
AND LastName IN ('Smith', 'Jones', 'Peterson')
OPTIMIZE FOR ALL ROWS
```


In this example, index X1 is used to position to the first index entry that matches the selection built over both columns LastName and WorkDept. The selection is rewritten into a series of ranges that match all the leading key columns used from the index X1. The probe is then based upon the composite concatenated values for all the leading keys. The pseudo-SQL for this rewritten SQL might look as follows:

```
SELECT * FROM X1
WHERE X1.LeadingKeys BETWEEN 'JonesA01' AND 'JonesE01'
      OR X1.LeadingKeys BETWEEN 'PetersonA01' AND 'PetersonE01'
      OR X1.LeadingKeys BETWEEN 'SmithA01' AND 'SmithE01'
```

All the key entries that satisfy the probe operation are used to generate a row number for the table associated with the index (for example, Employee). The row number is used by a Table Probe operation to perform random I/O on the table to produce the results for the query. This processing continues until all the rows that satisfy the index probe operation have been processed. In this example, all the index entries processed and rows retrieved met the index probe criteria.

Additional selection might be added that cannot use an index probe, such as selection against columns which are not leading key columns of the index. Then the optimizer performs an index scan operation within the range of probed values. This process still allows for selection to be performed before the Table Probe operation.

Related concepts

[Nested loop join implementation](#)

Db2 for i provides a **nested loop** join method. For this method, the processing of the tables in the join are ordered. This order is called the **join order**. The first table in the final join order is called the **primary table**. The other tables are called **secondary tables**. Each join table position is called a **dial**.

Related reference

[Effects of the ALWCPYDTA parameter on database performance](#)

Some complex queries can perform better by using a sort or hashing method to evaluate the query instead of using or creating an index.

Encoded vector index

An encoded vector index is a permanent object that provides access to a table. This access is done by assigning codes to distinct key values and then representing those values in a vector.

The size of the vector matches the number of rows in the underlying table. Each vector entry represents the table row number in the same position. The codes generated to represent the distinct key values can be 1 byte, 2 bytes, or 4 bytes in length. The key length depends upon the number of distinct values that need to be represented in the vector. Because of their compact size and relative simplicity, the EVI can be used to process large amounts of data efficiently.

An encoded vector index is used to represent the values stored in a table. However, the index itself cannot be used to directly gain access to the table. Instead, the encoded vector index can only be used to generate either a temporary row number list or a temporary row number bitmap. These temporary objects can then be used with a table probe to specify the rows in the table that the query needs to process.

The main difference in the table probe using an encoded vector index vs. a radix index is that the I/O paging can be asynchronous. The I/O can now be scheduled more efficiently to take advantage of groups of selected rows. Large portions of the table can be skipped over where no rows are selected.

Visual explain icon:



Related concepts

Encoded vector indexes

An encoded vector index (EVI) is used to provide fast data access in decision support and query reporting environments.

EVI maintenance

There are unique challenges to maintaining EVIs. The following table shows a progression of how EVIs are maintained, the conditions under which EVIs are most effective, and where EVIs are least effective, based on the EVI maintenance characteristics.

Encoded vector index RRN probe


A table probe operation is used to retrieve a specific row from a table based upon its row number. The row number is provided to the table probe access method by some other operation that generates a row number for the table.

The encoded vector index (EVI) RRN probe is an index only access method that is used to provide selected columns by retrieving the value from the EVI instead of using a table probe to access the table. Retrieving the value from the EVI should provide better I/O characteristics than the random I/Os associated with a table probe operation.

This access method is used in conjunction with a radix index probe, radix index scan, or EVI probe operation. The radix index probe, radix index scan, or EVI probe operation is used to select the rows and then the RRN of the selected row is used to probe into EVIs to retrieve any selected values that were not provided by the index used for selection. The EVI RRN probe can access multiple EVIs to provide selected values.

<i>Table 6. EVI RRN probe attributes</i>	
Data access method	EVI RRN Probe
Description	The encoded vector index (EVI) is quickly probed based upon the RRNs provided by the underlying index access.
Advantages	<ul style="list-style-type: none"> • Potential to extract all the data from the EVI index key values, thus eliminating the need for a Table Probe • Provides better paging characteristics than a Table Probe.
Considerations	<ul style="list-style-type: none"> • Only single key EVIs are considered for this implementation • All selected columns must have a single column EVI created. • The EVIs must fit in the query's fair share of optimizer memory
Likely to be used	<ul style="list-style-type: none"> • When the table row size is wide, the number of select columns is small compared to the number of columns in the table and the query requires a table probe to retrieve columns
Example SQL statement	<pre>CREATE ENCODED VECTOR INDEX EVI1 ON Employee (WorkDept) CREATE ENCODED VECTOR INDEX EVI2 ON Employee (Salary)WITH 10000 DISTINCT VALUES CREATE ENCODED VECTOR INDEX EVI3 ON Employee (LASTNAME)WITH 100000 DISTINCT VALUES CREATE INDEX IX1 ON Employee (Job) SELECT LASTNAME, WORKDEPT, SALARY FROM EMPLOYEE WHERE JOB = 'ANALYST'</pre>
Database Monitor and Plan Cache record indicating use	A QQRID 3001 Index Used record for each EVI with QQRCOD = 'I8'
SMP parallel enabled	Yes

Table 6. EVI RRN probe attributes (continued)

Data access method	EVI RRN Probe
Also referred to as	Table Probe, Preload
Visual Explain icon	

Prior to encoded vector index only access (EOA), the recommendation had been to only create EVIs for column with low cardinality (small number of distinct values). This recommendation has now changed. EVI RRN Probe can be used for columns with high cardinality (large number of distinct values). However, when creating the EVI, the WITH integer DISTINCT VALUES clause should be used to set the initial size of the codes appropriately and to minimize maintenance time if the database manager needs to use a larger code. See the CREATE INDEX statement in the SQL Reference for more details.


Encoded vector index probe

The encoded vector index (EVI) is quickly probed based upon the selection criteria that were rewritten into a series of ranges. It produces either a temporary row number list or bitmap.

Table 7. Encoded vector index probe attributes

Data access method	Encoded vector index probe
Description	The encoded vector index (EVI) is quickly probed based upon the selection criteria that were rewritten into a series of ranges. It produces either a temporary row number list or bitmap.
Advantages	<ul style="list-style-type: none"> • Only those index entries that match any selection continue to be processed • Provides quick access to the selected rows • Returns the row numbers in ascending sequence so that the Table Probe can be more aggressive in pre-fetching the rows for its operation
Considerations	EVI is usually built over a single key. The more distinct the column is and the higher the overflow percentage, the less advantageous the encoded vector index becomes. EVIs always require a Table Probe to be performed on the result of the EVI probe operation.
Likely to be used	<ul style="list-style-type: none"> • When the selection columns match the leading key columns of the index • When an encoded vector index exists and savings in reduced I/O against the table justifies the extra cost. This cost includes probing the EVI and fully populating the temporary row number list.

Table 7. Encoded vector index probe attributes (continued)

Data access method	Encoded vector index probe
Example SQL statement	<pre>CREATE ENCODED VECTOR INDEX EVI1 ON Employee (WorkDept) CREATE ENCODED VECTOR INDEX EVI2 ON Employee (Salary) CREATE ENCODED VECTOR INDEX EVI3 ON Employee (Job) SELECT * FROM Employee WHERE WorkDept = 'E01' AND Job = 'CLERK' AND Salary = 5000 OPTIMIZE FOR 99999 ROWS</pre>
Database Monitor and Plan Cache record indicating use	<pre>QQRID 3001 Index Used with QQRCD='I5', QQRID 3021 Bitmap Created and optionally QQRID 3022 Bitmap Merge.</pre>
SMP parallel enabled	Yes
Also referred to as	Encoded Vector Index Probe, Preload
Visual Explain icon	

Using the example above, the optimizer chooses to create a temporary row number bitmap for each of the encoded vector indexes used by this query. Each bitmap only identifies those rows that match the selection on the key columns for that index.

These temporary row number bitmaps are then merged together to determine the intersection of the rows selected from each index. This intersection is used to form a final temporary row number bitmap used to help schedule the I/O paging against the table for the selected rows.

The optimizer might choose to perform an index probe with a binary radix tree index if an index existed over all three columns. The implementation choice is probably decided by the number of rows to be returned and the anticipated cost of the I/O associated with each plan.

If few rows are returned, the optimizer probably chooses the binary radix tree index and performs the random I/O against the table. However, selecting more rows causes the optimizer to use the EVIs, because of the savings from the more efficiently scheduled I/O against the table.

Encoded vector index index-symbol table only access

The encoded vector index can also be used for index-symbol table only access.

The EVI can be used for more than generating a bitmap or row number list to provide an asynchronous I/O map to the desired table rows. The EVI can also be used by two index-only access methods that can be applied specific to the symbol table itself. These two index-only access methods are the EVI symbol table scan and the EVI symbol table probe.

These two methods can be used with GROUP BY or DISTINCT queries that can be satisfied by the symbol table. This symbol table-only access can be further employed in aggregate queries by adding INCLUDE values to the encoded vector index.

The following information is a summary of the symbol table-only scan and probe access methods.

Use the following links to learn in-depth information.

Related concepts

Encoded vector indexes

An encoded vector index (EVI) is used to provide fast data access in decision support and query reporting environments.

How the EVI works

EVI works in different ways for costing and implementation.

Related reference

Index grouping implementation

There are two primary ways to implement grouping using an index: Ordered grouping and pre-summarized processing.

Encoded vector index symbol table scan

An encoded vector index symbol table scan operation is used to retrieve the entries from the symbol table portion of the index.

All entries (symbols) in the symbol table are sequentially scanned if a scan is chosen. The symbol table can be used by the optimizer to satisfy GROUP BY or DISTINCT portions of a query request.

Selection is applied to every entry in the symbol table. The selection must be applied to the symbol table keys unless the EVI was created as a sparse index, with a WHERE clause. In that case, a portion of the selection is applied as the symbol table is built and maintained. The query request must include matching predicates to use the sparse EVI.

All entries are retrieved directly from the symbol table portion of the index without any access to the vector portion of the index. There is also no access to the records in the associated table over which the EVI is built.

Encoded vector index INCLUDE aggregates

To enhance the ability of the EVI symbol table to provide aggregate answers, the symbol table can be created to contain additional INCLUDE values. These are ready-made numeric aggregate results, such as SUM, COUNT, AVG, or VARIANCE values requested over non-key data. These aggregates are specified using the INCLUDE keyword on the CREATE ENCODED VECTOR INDEX request.

These included aggregates are maintained in real time as rows are inserted, updated, or deleted from the corresponding table. The symbol table maintains these additional aggregate values in addendum to the EVI keys for each symbol table entry. Because these are numeric results and finite in size, the symbol table is still a desirable compact size.


These included aggregates are over non-key columns in the table where the grouping is over the corresponding EVI symbol table defined keys. The aggregate can be over a single column or a derivation.

Data access method	Encoded vector index symbol table scan
Description	Sequentially scan and process all the symbol table entries associated with the index. When there is selection (WHERE clause), it is applied to every entry in the symbol table. An exception is made in the case of a sparse EVI, where the selection is applied as the index is created and maintained. Selected entries are retrieved directly without any access to the vector or the associated table.

Table 8. Encoded vector index symbol table scan attributes (continued)

Data access method	Encoded vector index symbol table scan
Advantages	<ul style="list-style-type: none"> • Pre-summarized results are readily available • Only processes the unique values in the symbol table, avoiding processing table records. • Extract all the data from the index unique key values or INCLUDE values, thus eliminating the need for a Table Probe or vector scan. • With INCLUDE, provides ready-made numeric aggregates, eliminating the need to access corresponding table rows to perform the aggregation
Considerations	<p>Dramatic performance improvement for grouping queries where the resulting number of groups is relatively small compared to the number of records in the underlying table. Can perform poorly when there are many groups involved such that the symbol table is large. Poor performance is even more likely if a large portion of the symbol table has been put into the overflow area.</p> <p>Dramatic performance improvement for grouping queries when the aggregate is specified as an INCLUDE value of the symbol table.</p>
Likely to be used	<ul style="list-style-type: none"> • When asking for GROUP BY, DISTINCT, COUNT, or COUNT DISTINCT from a single table and the referenced columns are in the key definition. • When the number of unique values in the columns of the key definition is small relative to the number of records in the underlying table. • When there is no selection (WHERE clause) within the query or the selection does not reduce the result set much. • When the symbol table key satisfies the GROUP BY, and requested aggregates, like SUM or COUNT, are specified as INCLUDE values. • when the query is run with commitment control *NONE or *CHG.

Table 8. Encoded vector index symbol table scan attributes (continued)

Data access method	Encoded vector index symbol table scan
Example SQL statement	<pre data-bbox="626 254 1268 285">CREATE ENCODED VECTOR INDEX EVI1 ON Sales (Region)</pre> <p data-bbox="610 306 740 338">Example 1</p> <pre data-bbox="626 369 919 464">SELECT Region, count(*) FROM Sales GROUP BY Region OPTIMIZE FOR ALL ROWS</pre> <p data-bbox="610 485 740 516">Example 2</p> <pre data-bbox="626 548 911 621">SELECT DISTINCT Region FROM Sales OPTIMIZE FOR ALL ROWS</pre> <p data-bbox="610 642 740 674">Example 3</p> <pre data-bbox="626 705 1000 758">SELECT COUNT(DISTINCT Region) FROM Sales</pre> <p data-bbox="610 779 1446 842">Example 4 uses the INCLUDE option. The sums of revenue and cost of goods per sales region is maintained in real time.</p> <pre data-bbox="626 873 1252 989">CREATE ENCODED VECTOR INDEX EVI2 ON Sales(Region) INCLUDE(SUM(SALES)) SELECT Region, SUM(SALEs) FROM Sales GROUP BY Region</pre>
Database Monitor and Plan Cache record indicating use	<p data-bbox="610 1031 1081 1104">QQRID 3001 Index Used where QQC15 = 'E' AND QQRcod = 'I2'</p>
Also referred to as	<p data-bbox="610 1136 1211 1167">Encoded Vector Index Symbol Table Scan, Preload</p>
Visual Explain icon	

Related concepts

[Encoded vector indexes](#)

An encoded vector index (EVI) is used to provide fast data access in decision support and query reporting environments.

[How the EVI works](#)

EVI's work in different ways for costing and implementation.

Related reference

[Index grouping implementation](#)

There are two primary ways to implement grouping using an index: Ordered grouping and pre-summarized processing.

Related information

[SQL INCLUDE statement](#)

Encoded vector index symbol table probe

An encoded vector index symbol table probe operation is used to retrieve entries from the symbol table portion of the index. Scanning the entire symbol table is not necessary.

The symbol table can be used by the optimizer to satisfy GROUP BY or DISTINCT portions of a query request.

The optimizer attempts to match the columns used for some or all the selection against the leading keys of the EVI index. It then rewrites the selection into a series of ranges that can be used to probe directly into the symbol table. Only those symbol table pages from the series of ranges are paged into main memory.

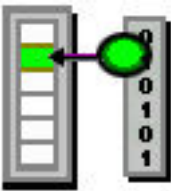
The resulting symbol table entries generated by the probe operation can then be further processed by any remaining selection against EVI keys. This strategy provides for quick access to only the entries of the symbol table that satisfy the selection.

Like an encoded vector symbol table scan, a symbol table probe can return ready-made aggregate results if INCLUDE is specified when the EVI is created.

All entries are retrieved directly from the symbol table portion of the index without any access to the vector portion of the index. In addition, it is unnecessary to access the records in the associated table over which the EVI is built.

Data access method	Encoded vector index symbol table probe
Description	
Advantages	Probe the symbol table entries associated with the index. When there is selection (WHERE clause), it is applied to every entry in the symbol table that meets the probe criteria. If there are sparse EVIs, the selection is applied as the EVI is created and maintained. Selected entries are retrieved directly without any access to the vector or the associated table.
Considerations	<ul style="list-style-type: none">• Pre-summarized results are readily available• Only processes the unique values in the symbol table, avoiding processing table records.• Extracts all the data from the index unique key values or include values, or both, thus eliminating the need for a table probe or vector scan• With INCLUDE, provides ready-made numeric aggregates, eliminating the need to access corresponding table rows to perform the aggregation

Table 9. Encoded vector index symbol table probe attributes (continued)

Data access method	Encoded vector index symbol table probe
Likely to be used	<ul style="list-style-type: none"> • When asking for GROUP BY, DISTINCT, COUNT, or COUNT DISTINCT from a single table and the referenced columns are in the key definition. • When the number of unique values in the columns of the key definition is small relative to the number of records in the underlying table. • When there is selection (WHERE clause) that reduces the selection from the Symbol Table and the WHERE clause involves leading, probable keys. • When the symbol table key satisfies the GROUP BY and the WHERE clause reduces selection to the leading keys, and aggregates are specified as INCLUDE values. • When the query is run with commitment control *NONE or *CHG.
Example SQL statement	<pre>CREATE ENCODED VECTOR INDEX EVI1 ON Sales (Region)</pre> <p>Example 1</p> <pre>SELECT Region, COUNT(*) FROM Sales WHERE Region in ('Quebec', 'Manitoba') GROUP BY Region OPTIMIZE FOR ALL ROWS</pre> <p>Example 2</p> <pre>CREATE ENCODED VECTOR INDEX EVI2 ON Sales(Region) INCLUDE(SUM(SALES))</pre> <pre>SELECT Region, SUM(SALES) FROM Sales WHERE Region = 'PACIFIC' GROUP BY Region</pre>
Database Monitor and Plan Cache record indicating use	<pre>QQRID 3001 Index Used where QQC15 = 'E' AND QQRCD = 'I2' AND QQKP = 'Y'</pre>
Also referred to as	Encoded Vector Index Table Probe, Preload
Visual Explain icon	

Related concepts

[Encoded vector indexes](#)

An encoded vector index (EVI) is used to provide fast data access in decision support and query reporting environments.

[How the EVI works](#)

EVI's work in different ways for costing and implementation.

Related reference

[Index grouping implementation](#)

There are two primary ways to implement grouping using an index: Ordered grouping and pre-summarized processing.

Related information

[SQL INCLUDE statement](#)

Temporary objects and access methods

Temporary objects are created by the optimizer in order to process a query. In general, these temporary objects are internal objects and cannot be accessed by a user.

Table 10. Temporary object data access methods

Temporary create objects	Scan operations	Probe operations
Temporary hash table	Hash table scan	Hash table probe
Temporary sorted list	Sorted list scan	Sorted list probe
Temporary distinct sorted list	Sorted list scan	N/A
Temporary list	List scan	N/A
Temporary values list	Values list scan	N/A
Temporary row number list	Row number list scan	Row number list probe
Temporary bitmap	Bitmap scan	Bitmap probe
Temporary index	Temporary index scan	Temporary index probe
Temporary buffer	Buffer scan	N/A
Queue	N/A	N/A
Array unnest temporary table	Temporary table scan	N/A
Temporary Indexed List	Temporary Indexed List Scan and Index Merge	N/A
Window	Window scan	N/A

Temporary hash table

The temporary hash table is a temporary object that allows the optimizer to collate the rows based upon a column or set of columns. The hash table can be either scanned or probed by the optimizer to satisfy different operations of the query.

A temporary hash table is an efficient data structure because the rows are organized for quick and easy retrieval after population has occurred. The hash table remains resident within main memory to avoid any I/Os associated with either the scan or probe against the temporary object. The optimizer determines the optimal hash table size based on the number of unique column combinations used as keys for the creation.

Additionally the hash table can be populated with all the necessary columns to satisfy any further processing. This population avoids any random I/Os associated with a table probe operation.

However, the optimizer can selectively include columns in the hash table when the calculated size exceeds the memory pool storage available for the query. In these cases, a table probe operation is required to recollect the missing columns from the hash table before the selected rows can be processed.

The optimizer also can populate the hash table with distinct values. If the query contains grouping or distinct processing, then all the rows with the same key value are not required in the hash table. The rows

are still collated, but the distinct processing is performed during the population of the hash table itself. This method allows a simple scan on the result in order to complete the grouping or distinct operation.

A temporary hash table is an internal data structure and can only be created by the database manager

Visual explain icon:




Hash table scan

During a hash table scan operation, the entire temporary hash table is scanned and all the entries contained within the hash table are processed.

The optimizer considers a hash table scan when the data values need to be collated together, but sequencing of the data is not required. A hash table scan allows the optimizer to generate a plan that takes advantage of any non-join selection while creating the temporary hash table.

An additional benefit is that the temporary hash table data structure will typically cause the table data to remain resident within main memory after creation. Resident table data reduces paging on the subsequent hash table scan operation.

<i>Table 11. Hash table scan attributes</i>	
Data access method	Hash table scan
Description	Read all the entries in a temporary hash table. The hash table can perform distinct processing to eliminate duplicates. Or the temporary hash table can collate all the rows with the same value together.
Advantages	<ul style="list-style-type: none"> • Reduces the random I/O to the table associated with longer running queries that might otherwise use an index to collate the data • Selection can be performed before generating the hash table to subset the number of rows in the temporary object
Considerations	Used for distinct or group by processing. Can perform poorly when the entire hash table does not stay resident in memory as it is being processed.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA) • When the data is required to be collated based upon a column or columns for distinct or grouping
Example SQL statement	<pre>SELECT COUNT(*), FirstNme FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' GROUP BY FirstNme</pre>
Database Monitor and Plan Cache record indicating use	<pre>QQRID 3023 Temp Hash Table Created where QVC1F (HashTable_ReasonCode) = 'G'</pre>
SMP parallel enabled	Yes

<i>Table 11. Hash table scan attributes (continued)</i>	
Data access method	Hash table scan
Also referred to as	Hash Scan, Preload Hash Table Scan Distinct Hash Table Scan Distinct, Preload
Visual Explain icon	


Hash table probe

A hash table probe operation is used to retrieve rows from a temporary hash table based upon a probe lookup operation.

The optimizer initially identifies the keys of the temporary hash table from the join criteria specified in the query. When the hash table is probed, the values used to probe into the hash table are extracted from the join-from criteria specified in the selection.

These values are sent through the same hashing algorithm used to populate the temporary hash table. They determine if any rows have a matching equal value. All the matching join rows are then returned to be further processed by the query.

<i>Table 12. Hash table probe attributes</i>	
Data access method	Hash table probe
Description	The temporary hash table is quickly probed based upon the join criteria.
Advantages	<ul style="list-style-type: none"> • Provides quick access to the selected rows that match probe criteria • Reduces the random I/O to the table associated with longer running queries that use an index to collate the data • Selection can be performed before generating the hash table to subset the number of rows in the temporary object
Considerations	Used to process equal join criteria. Can perform poorly when the entire hash table does not stay resident in memory as it is being processed.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA) • When the data is required to be collated based upon a column or columns for join processing • The join criteria was specified using an equals (=) operator
Example SQL statement	<pre>SELECT * FROM Employee XXX, Department YYY WHERE XXX.WorkDept = YYY.DeptNo OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	<pre>QQRID 3023 Temp Hash Table Created where QVC1F (HashTable_ReasonCode) = 'J'</pre>
SMP parallel enabled	Yes

<i>Table 12. Hash table probe attributes (continued)</i>	
Data access method	Hash table probe
Also referred to as	Hash Table Probe, Preload Hash Table Probe Distinct Hash Table Probe Distinct, Preload
Visual Explain icon	

The hash table probe access method is considered when determining the implementation for a secondary table of a join. The hash table is created with the key columns that match the equal selection or join criteria for the underlying table.

The hash table probe allows the optimizer to choose the most efficient implementation in selecting rows from the underlying table, without regard for join criteria. This single pass through the underlying table can now use a table scan or existing index to select the rows needed for the hash table population.

Since hash tables are constructed so that most of the hash table remains resident within main memory, the I/O associated with a hash probe is minimal. Additionally, if the hash table was populated with all necessary columns from the underlying table, no additional table probe is required to finish processing this table. This method causes further I/O savings.

Related concepts

Nested loop join implementation

Db2 for i provides a **nested loop** join method. For this method, the processing of the tables in the join are ordered. This order is called the **join order**. The first table in the final join order is called the **primary table**. The other tables are called **secondary tables**. Each join table position is called a **dial**.

Temporary sorted list

The temporary sorted list is a temporary object that allows the optimizer to sequence rows based upon a column or set of columns. The sorted list can be either scanned or probed by the optimizer to satisfy different operations of the query.

A temporary sorted list is a data structure where the rows are organized for quick and easy retrieval after population has occurred. During population, the rows are copied into the temporary object and then a second pass is made through the temporary object to perform the sort.

In order to optimize the creation of this temporary object, minimal data movement is performed while the sort is processed. It is not as efficient to probe a temporary sorted list as it is to probe a temporary hash table.

Additionally, the sorted list can be populated with all the necessary columns to satisfy any further processing. This population avoids any random I/Os associated with a table probe operation.

However, the optimizer can selectively include columns in the sorted list when the calculated size exceeds the memory pool storage available for this query. In those cases, a table probe operation is required to recollect the missing columns from the sorted list before the selected rows can be processed.

A temporary sorted list is an internal data structure and can only be created by the database manager.

Visual explain icon:



Sorted list scan


During a sorted list scan operation, the entire temporary sorted list is scanned and all the entries contained within the sorted list are processed.

A sorted list scan is considered when the data values need to be sequenced. A sorted list scan allows the optimizer to generate a plan that can take advantage of any non-join selection while creating the temporary sorted list.

An additional benefit is that the data structure will usually cause the table data within the sorted list to remain resident within main memory after creation. This resident data reduces paging on the subsequent sorted list scan operation.

<i>Table 13. Sorted list scan attributes</i>	
Data access method	Sorted list scan
Description	Read all the entries in a temporary sorted list. The sorted list can perform distinct processing to eliminate duplicate values or take advantage of the temporary sorted list to sequence all the rows.
Advantages	<ul style="list-style-type: none"> • Reduces the random I/O to the table associated with longer running queries that would otherwise use an index to sequence the data. • Selection can be performed prior to generating the sorted list to subset the number of rows in the temporary object
Considerations	Used to process ordering or distinct processing. Can perform poorly when the entire sorted list does not stay resident in memory as it is being populated and processed.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA) • When the data is required to be ordered based upon a column or columns for ordering or distinct processing
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (LastName, WorkDept) SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' ORDER BY FirstNme OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3003 Query Sort. There is no specific field that indicates whether or not the sorted list was used for a scan or a probe. Refer to Visual Explain diagram for query implementation details.
SMP parallel enabled	No
Also referred to as	Sorted List Scan, Preload Sorted List Scan Distinct Sorted List Scan Distinct, Preload

Table 13. Sorted list scan attributes (continued)

Data access method	Sorted list scan
Visual Explain icon	

Sorted list probe


A sorted list probe operation is used to retrieve rows from a temporary sorted list based upon a probe lookup operation.

The optimizer initially identifies the temporary sorted list keys from the join criteria specified in the query. The values used to probe into the temporary sorted list are extracted from the join-from criteria specified in the selection. Those values are used to position within the sorted list in order to determine if any rows have a matching value. All the matching join rows are then returned to be further processed by the query.

Table 14. Sorted list probe attributes

Data access method	Sorted list probe
Description	The temporary sorted list is quickly probed based upon the join criteria.
Advantages	<ul style="list-style-type: none"> • Provides quick access to the selected rows that match probe criteria • Reduces the random I/O to the table associated with longer running queries that otherwise use an index to collate the data • Selection can be performed before generating the sorted list to subset the number of rows in the temporary object
Considerations	Used to process non-equal join criteria. Can perform poorly when the entire sorted list does not stay resident in memory as it is being populated and processed.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA) • When the data is required to be collated based upon a column or columns for join processing • The join criteria was specified using a non-equals operator
Example SQL statement	<pre>SELECT * FROM Employee XXX, Department YYY WHERE XXX.WorkDept > YYY.DeptNo OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3003 Query Sort. There is no specific field that indicates whether or not the sorted list was used for a scan or a probe. Refer to Visual Explain diagram for query implementation details.
SMP parallel enabled	Yes
Also referred to as	Sorted List Probe, Preload Sorted List Probe Distinct Sorted List Probe Distinct, Preload

Table 14. Sorted list probe attributes (continued)

Data access method	Sorted list probe
Visual Explain icon	

The sorted list probe access method is considered when determining the implementation for a secondary table of a join. The sorted list is created with the key columns that match the non-equal join criteria for the underlying table. The optimizer chooses the most efficient implementation to select the rows from the underlying table without regard to any join criteria. This single pass through the underlying table can use a Table Scan or an existing index to select the rows needed to populate the sorted list.

Since sorted lists are constructed so that most of the temporary object remains resident within main memory, the sorted list I/O is minimal. If the sorted list was populated with all necessary table columns, no additional Table Probe is required to finish processing the table, causing further I/O savings.

Related concepts

Nested loop join implementation

Db2 for i provides a **nested loop** join method. For this method, the processing of the tables in the join are ordered. This order is called the **join order**. The first table in the final join order is called the **primary table**. The other tables are called **secondary tables**. Each join table position is called a **dial**.

Temporary distinct sorted list

A temporary distinct sorted list combines the features of the temporary hash table and the temporary sorted list.

Like the hash table, the temporary distinct sorted list allows the optimizer to collate the rows based on a column or set of columns. Like the sorted list, the temporary distinct sorted list also allows the optimizer to sequence the rows.

A temporary distinct sorted list contains a hash table data structure set up for efficient access to aggregate rows during population. In addition, a binary tree data structure is maintained over the hash table data structure so that the data can be accessed in sequence. The sorted aspect of the data structure allows for the efficient computation of super-aggregate rows in SQL statements that contain GROUP BY ROLLUP.

A temporary sorted aggregate hash table is an internal data structure and can only be created by the database manager.

Visual explain icon:




Sorted list scan

During the sorted list scan, the entire temporary distinct sorted list is scanned and all the entries contained within the temporary are processed.

The optimizer uses the sorted list scan when the data values need to be aggregated and sequenced. The optimizer generates this plan that can take advantage of any non-join selection while creating the temporary distinct sorted list. The data structure of the temporary distinct sorted list will typically cause

the table data to remain resident within main memory after creation. This memory-resident data reduces paging on the subsequent sorted list scan.

<i>Table 15. Sorted list scan attributes</i>	
Data access method	Sorted list scan
Description	Reads all the entries in a temporary distinct sorted list
Advantages	<ul style="list-style-type: none"> • Allows efficient computation of ROLLUP super-aggregate rows. • Reduces the random I/O to the table associated with longer running queries that might otherwise use an index to collate the data. • Selection can be performed before generating the distinct sorted list to subset the number of rows in the temporary object.
Considerations	Used for GROUP BY ROLLUP processing, Can perform poorly when the entire temporary object does not stay resident in memory as it is being processed.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed in the query environmental parameter (ALWCPYDTA) • When a GROUP BY ROLLUP is in the SQL statement
Messages indicating use	N/A
SMP parallel enabled	Yes
Also referred to as	N/A
Visual Explain icon	

Temporary list

The temporary list is a temporary object that allows the optimizer to store intermediate results of a query. The list is an unsorted data structure that is used to simplify the operation of the query. Since the list does not have any keys, the rows within the list can only be retrieved by a sequential scan operation.

The temporary list can be used for various reasons, some of which include an overly complex view or derived table, Symmetric Multiprocessing (SMP) or to prevent a portion of the query from being processed multiple times.

A temporary list is an internal data structure and can only be created by the database manager.


Visual explain icon:



List scan

The list scan operation is used when a portion of the query is processed multiple times, but no key columns can be identified. In these cases, that portion of the query is processed once and its results

are stored within the temporary list. The list can then be scanned for only those rows that satisfy any selection or processing contained within the temporary object.

<i>Table 16. List scan attributes</i>	
Data access method	List scan
Description	Sequentially scan and process all the rows in the temporary list.
Advantages	<ul style="list-style-type: none"> • The temporary list and list scan can be used by the optimizer to minimize repetition of an operation or to simplify the optimizer logic flow. • Selection can be performed before generating the list to subset the number of rows in the temporary object.
Considerations	Used to prevent portions of the query from being processed multiple times when no key columns are required to satisfy the request.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA). • When Db2 symmetric multiprocessing is used for the query.
Example SQL statement	<pre>SELECT * FROM Employee XXX, Department YYY WHERE XXX.LastName IN ('Smith', 'Jones', 'Peterson') AND YYY.DeptNo BETWEEN 'A01' AND 'E01' OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3004 Temp Table
SMP parallel enabled	Yes
Also referred to as	List Scan, Preload
Visual Explain icon	

Using the example above, the optimizer chose to create a temporary list to store the selected rows from the DEPARTMENT table. Since there is no join criteria, a Cartesian product join is performed between the two tables. To prevent the join from scanning all the rows of the DEPARTMENT table for each join possibility, the selection against the DEPARTMENT table is performed once. The results are stored in the temporary list. The temporary list is then scanned for the Cartesian product join.

Temporary values list

The temporary values list allows the optimizer to store rows of data specified in a VALUES clause of a SELECT or CREATE VIEW statement.

The list is an unsorted data structure that is used to simplify the operation of the query. Since the list does not have any keys, the rows within the list can only be retrieved by a sequential scan operation.


A temporary values list is an internal data structure and can only be created by the database manager.

Visual explain icon:



Values list scan

During a values list scan operation, the entire temporary values list is scanned and all the rows of data are processed.

Table 17. Values list scan attributes	
Data access method	Values list scan
Description	Sequentially scan and process all the rows of data in the temporary values list.
Advantages	The temporary values list and values list scan can be used by the optimizer to simplify the optimizer logic flow.
Likely to be used	When a VALUES clause is specified in the from-clause of an SQL fullselect
Example SQL statement	<pre>SELECT EMPNO, 'empproject' FROM EMPPROJECT WHERE PROJNO IN('MA2100', 'MA2110', 'MA2112') UNION VALUES ('NEWAAA', 'new'), ('NEWBBB', 'new')</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3000 where QVQTBL = '*VALUES'
SMP parallel enabled	Yes
Visual Explain icon	


Temporary row number list

The temporary row number list, also referred to as an RRN List, is a temporary object that allows the optimizer to sequence rows based upon their row address (their row number). The row number list can be either scanned or probed by the optimizer to satisfy different operations of the query.

A temporary row number list is a data structure where the rows are organized for quick and efficient retrieval. The row number list only contains the row number for the associated row. Since no table data is present, a table probe operation is typically associated with it in order to retrieve the underlying table data. Because the row numbers are sorted, the random I/O associated with the table probe operation is performed more efficiently. The database manager performs pre-fetch or look-ahead logic to determine if multiple rows are located on adjacent pages. If so, the table probe requests a larger I/O to bring the rows into main memory more efficiently.

A temporary row number list is an internal data structure and can only be created by the database manager.

Visual explain icon:



Row number list scan

The entire temporary row number list is scanned and all the row addresses contained within the row number list are processed. The optimizer considers this plan when there is an applicable encoded vector


index or if the index probe or scan random I/O can be reduced. The random I/O can be reduced by first preprocessing and sorting the row numbers associated with the Table Probe.

The use of a row number list scan allows the optimizer to generate a plan that can take advantage of multiple indexes to match up to different portions of the query.

An additional benefit is that the data structure of the temporary row number list guarantees that the row numbers are sorted. It closely mirrors the row number layout of the table data, ensuring that the table paging never visits the same page of data twice. This results in increased I/O savings for the query.

A row number list scan is identical to a bitmap scan operation. The only difference is that the list scan is over a list of row addresses while the bitmap scan is over a bitmap representing the addresses.

<i>Table 18. Row number list scan</i>	
Data access method	Row number list scan
Description	Sequentially scan and process all the row numbers in the temporary row number list. The sorted row numbers can be merged with other temporary row number lists or can be used as input into a Table Probe operation.
Advantages	<ul style="list-style-type: none"> • The temporary row number list only contains address, no data, so the temporary can be efficiently scanned within memory. • The row numbers contained within the temporary object are sorted to provide efficient I/O processing to access the underlying table. • Selection is performed as the row number list is generated to subset the number of rows in the temporary object.
Considerations	Since the row number list contains only the addresses of the selected rows in the table, a separate Table Probe fetches the table rows.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA). • When the cost of sorting of the row number is justified by the more efficient I/O that can be performed during the Table Probe operation. • When multiple indexes over the same table need to be combined in order to minimize the number of selected rows.
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (WorkDept) CREATE ENCODED VECTOR INDEX EVI2 ON Employee (Salary) CREATE ENCODED VECTOR INDEX EVI3 ON Employee (Job) SELECT * FROM Employee WHERE WorkDept = 'E01' AND Job = 'CLERK' AND Salary = 5000 OPTIMIZE FOR 99999 ROWS</pre>
Database Monitor and Plan Cache record indicating use	<p>QQRID 3001 and QQRID 3021 records for each index used.</p> <p>The QQC11 field in the 3021 record will be 'L'.</p> <p>A QQRID 3000 record with QQC11 (Skip_Sequential_Table_Scan) = 'Y'.</p> <p>Optionally, QQRID 3022 records if bitmap merging occurred.</p>
SMP parallel enabled	Yes
Also referred to as	Row Number List Scan, Preload; RRN Scan; RRN Scan, Preload

<i>Table 18. Row number list scan (continued)</i>	
Data access method	Row number list scan
Visual Explain icon	

Using the example above, the optimizer created a temporary row number list for each of the indexes used by this query. These indexes included a radix index and two encoded vector indexes. Each index row number list was scanned and merged into a final composite row number list representing the intersection of all the index row number lists. The final row number list is then used by the Table Probe to determine which rows are selected and processed for the query results.

Row number list probe


A row number list probe is used to test row numbers generated by a separate operation against the selected rows of a temporary row number list. The row numbers can be generated by any operation that constructs a row number for a table. That row number is then used to probe into a temporary row number list to determine if it matches the selection used to generate the list.

The use of a row number list probe operation allows the optimizer to generate a plan that can take advantage of any sequencing provided by an index, but still use the row number list to perform additional selection before any Table probe operations.

A row number list probe is identical to a bitmap probe operation. The only difference is that the list probe is over a list of row addresses while the bitmap probe is over a bitmap representing the addresses.

<i>Table 19. Row number list probe</i>	
Data access method	Row number list probe
Description	The temporary row number list is quickly probed based upon the row number generated by a separate operation.
Advantages	<ul style="list-style-type: none"> • The temporary row number list only contains a row address, no data, so the temporary can be efficiently probed within memory. • The row numbers represented within the row number list are sorted to provide efficient lookup processing to test the underlying table. • Selection is performed as the row number list is generated to subset the number of selected rows in the temporary object.
Considerations	Used when the query contains ordering and additional selection that can be satisfied by additional indexes. Since the row number list contains only the addresses of the selected rows in the table, a separate Table Probe fetches the table rows.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCOPYDTA). • When the cost of creating and probing the row number list is justified by reducing the number of Table Probe operations that must be performed. • When multiple indexes over the same table need to be combined in order to minimize the number of selected rows.

Table 19. Row number list probe (continued)

Data access method	Row number list probe
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (WorkDept) CREATE ENCODED VECTOR INDEX EVI2 ON Employee (Salary) CREATE ENCODED VECTOR INDEX EVI3 ON Employee (Job) SELECT * FROM Employee WHERE WorkDept = 'E01' AND Job = 'CLERK' AND Salary = 5000 ORDER BY WorkDept</pre>
Database Monitor and Plan Cache record indicating use	<p>QQRID 3001 and QQRID 3021 records for each index used.</p> <p>The QQC11 field in the 3021 record will be 'L'.</p> <p>Optionally, QQRID 3022 records if bitmap merging occurred.</p>
SMP parallel enabled	Yes
Also referred to as	Row Number List Probe, Preload; RRN Probe; RRN Probe, Preload
Visual Explain icon	

Using the example above, the optimizer created a temporary row number list for each of the encoded vector indexes. Additionally, an index probe operation was performed against the radix index X1 to satisfy the ordering requirement. Since the ORDER BY requires that the resulting rows be sequenced by the WorkDept column, the row number list cannot be scanned for the selected rows.

However, the temporary row number list can be probed using a row address extracted from the index X1 used to satisfy the ordering. By probing the list with the row address extracted from the index probe, the sequencing of the keys in the index X1 is preserved. The row can still be tested against the selected rows within the row number list.

Temporary bitmap


The temporary bitmap is a temporary object that allows the optimizer to sequence rows based upon their row address (their row number). The bitmap can be either scanned or probed by the optimizer to satisfy different operations of the query.

A temporary bitmap is a data structure that uses a bitmap to represent all the row numbers for a table. Since each row is represented by a separate bit, all the rows within a table can be represented in a fairly condensed form. When a row is selected, the bit within the bitmap that corresponds to the selected row is set on. After the temporary bitmap is populated, all the selected rows can be retrieved in a sorted manner for quick and efficient retrieval. The temporary bitmap only represents the row number for the associated selected rows.

No table data is present within the temporary bitmap. A table probe operation is typically associated with the bitmap in order to retrieve the underlying table data. Because the bitmap is by definition sorted, the random I/O associated with the table probe operation can be performed more efficiently. The database manager performs pre-fetch or look-ahead logic to determine if multiple rows are located on adjacent pages. If so, the table probe requests a larger I/O to bring the rows into main memory more efficiently.

A temporary bitmap is an internal data structure and can only be created by the database manager.

Visual explain icon:



Bitmap scan


During a bitmap scan operation, the entire temporary bitmap is scanned and all the row addresses contained within the bitmap are processed. The optimizer considers this plan when there is an applicable encoded vector index or if the index probe or scan random I/O can be reduced. The random I/O can be reduced by first preprocessing and sorting the row numbers associated with the Table Probe.

The use of a bitmap scan allows the optimizer to generate a plan that can take advantage of multiple indexes to match up to different portions of the query.

An additional benefit is that the data structure of the temporary bitmap guarantees that the row numbers are sorted. It closely mirrors the row number layout of the table data, ensuring that the table paging never visits the same page of data twice. This results in increased I/O savings for the query.

A bitmap scan is identical to a row number list scan operation. The only difference is that the list scan is over a list of row addresses while the bitmap scan is over a bitmap representing the addresses.

Table 20. Bitmap scan attributes	
Data access method	Bitmap scan attributes
Description	Sequentially scan and process all the row numbers in the temporary bitmap. The sorted row numbers can be merged with other temporary bitmaps or can be used as input into a Table Probe operation.
Advantages	<ul style="list-style-type: none"> • The temporary bitmap only contains a reference to a row address, no data, so the temporary can be efficiently scanned within memory. • The row numbers represented within the temporary object are sorted to provide efficient I/O processing to access the underlying table. • Selection is performed as the bitmap is generated to subset the number of selected rows in the temporary object.
Considerations	Since the bitmap contains only the addresses of the selected rows in the table, a separate Table Probe fetches the table rows.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCOPYDTA). • When the cost of sorting of the row numbers is justified by the more efficient I/O that can be performed during the Table Probe operation. • When multiple indexes over the same table need to be combined in order to minimize the number of selected rows.
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (WorkDept) CREATE ENCODED VECTOR INDEX EVI2 ON Employee (Salary) CREATE ENCODED VECTOR INDEX EVI3 ON Employee (Job) SELECT * FROM Employee WHERE WorkDept = 'E01' AND Job = 'CLERK' AND Salary = 5000 OPTIMIZE FOR 99999 ROWS</pre>

Data access method	Bitmap scan attributes
Database Monitor and Plan Cache record indicating use	<p>QQRID 3001 and QQRID 3021 records for each index used.</p> <p>The QQC11 field in the 3021 record will be 'B'.</p> <p>A QQRID 3000 record with QQC11 (Skip_Sequential_Table_Scan) = 'Y'.</p> <p>Optionally, QQRID 3022 records if bitmap merging occurred.</p>
SMP parallel enabled	Yes
Also referred to as	<p>Bitmap Scan, Preload</p> <p>Row Number Bitmap Scan</p> <p>Row Number Bitmap Scan, Preload</p> <p>Skip Sequential Scan</p>
Visual Explain icon	

Using the example above, the optimizer created a temporary bitmap for each of the indexes used by this query. These indexes included a radix index and two encoded vector indexes. Each index temporary bitmap was scanned and merged into a final composite bitmap representing the intersection of all the index temporary bitmaps. The final bitmap is then used by the Table Probe operation to determine which rows are selected and processed for the query results.

Bitmap probe

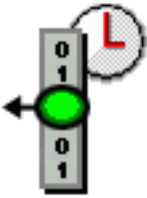
A bitmap probe operation is used to test row numbers generated by a separate operation against the selected rows of a temporary bitmap. The row numbers can be generated by any operation that constructs a row number for a table. That row number is then used to probe into a temporary bitmap to determine if it matches the selection used to generate the bitmap.

The use of a bitmap probe operation allows the optimizer to generate a plan that can take advantage of any sequencing provided by an index, but still use the bitmap to perform additional selection before any Table Probe operations.

A bitmap probe is identical to a row number list probe operation. The only difference is that the list probe is over a list of row addresses while the bitmap probe is over a bitmap representing the addresses.

Data access method	Bitmap probe attributes
Description	The temporary bitmap is quickly probed based upon the row number generated by a separate operation.
Advantages	<ul style="list-style-type: none"> • The temporary bitmap only contains a reference to a row address, no data, so the temporary can be efficiently probed within memory. • The row numbers represented within the bitmap are sorted to provide efficient lookup processing to test the underlying table. • Selection is performed as the bitmap is generated to subset the number of selected rows in the temporary object.

Table 21. Bitmap probe attributes (continued)

Data access method	Bitmap probe attributes
Considerations	Since the bitmap contains only the addresses of the selected rows in the table, a separate Table Probe fetches the table rows.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA). • When the cost of creating and probing the bitmap is justified by reducing the number of Table Probe operations that must be performed. • When multiple indexes over the same table need to be combined in order to minimize the number of selected rows.
Example SQL statement	<pre>CREATE INDEX X1 ON Employee (WorkDept) CREATE ENCODED VECTOR INDEX EVI2 ON Employee (Salary) CREATE ENCODED VECTOR INDEX EVI3 ON Employee (Job) SELECT * FROM Employee WHERE WorkDept = 'E01' AND Job = 'CLERK' AND Salary = 5000 ORDER BY WorkDept</pre>
Database Monitor and Plan Cache record indicating use	<p>QQRID 3001 and QQRID 3021 records for each index used.</p> <p>The QQC11 field in the 3021 record will be 'BL'.</p> <p>Optionally, QQRID 3022 records if bitmap merging occurred.</p>
SMP parallel enabled	Yes
Also referred to as	<p>Bitmap Probe, Preload</p> <p>Row Number Bitmap Probe</p> <p>Row Number Bitmap Probe, Preload</p>
Visual Explain icon	

Using the example above, the optimizer created a temporary bitmap for each of the encoded vector indexes. Additionally, an index probe operation was performed against the radix index X1 to satisfy the ordering requirement. Since the ORDER BY requires that the resulting rows be sequenced by the WorkDept column, the bitmap cannot be scanned for the selected rows.

However, the temporary bitmap can be probed using a row address extracted from the index X1 used to satisfy the ordering. By probing the bitmap with the row address extracted from the index probe, the sequencing of the keys in the index X1 is preserved. The row can still be tested against the selected rows within the bitmap.

Temporary index

A temporary index is a temporary object that allows the optimizer to create and use a radix index for a specific query. The temporary index has all the same attributes and benefits as a radix index created through the CREATE INDEX SQL statement or **Create Logical File (CRTLF)** CL command.

Additionally, the temporary index is optimized for use by the optimizer to satisfy a specific query request. This optimization includes setting the logical page size and applying any selection to the index to speed up its use after creation.

The temporary index can be used to satisfy various query requests:

- Ordering
- Grouping/Distinct
- Joins
- Record selection

Generally a temporary index is a more expensive temporary object to create than other temporary objects. It can be populated by a table scan, or by one or more index scans or probes. The optimizer considers all the methods available when determining which method to use to produce the rows for the index creation. This process is like the costing and selection of the other temporary objects used by the optimizer.

One significant advantage of the temporary index over other temporary objects is that it is the only temporary object maintained if the underlying table changes. The temporary index is identical to a radix index in that any inserts or updates against the table are reflected immediately through normal index maintenance.

SQE usage of temporary indexes is different from CQE usage in that SQE allows reuse. References to temporary indexes created and used by the SQE optimizer are kept in the system Plan Cache. A temporary index is saved for reuse by other instances of the same query or other instances of the same query running in a different job. It is also saved for potential reuse by a different query that can benefit from the use of the same temporary index.


By default, an SQE temporary index persists until the Plan Cache entry for the last referencing query plan is removed. With the SQE Plan Cache auto sizing capability, there is the potential for SQE temporary indexes to persist longer. You can control this behavior by setting the CACHE_RESULTS QAQQINI value. The default for this INI value allows the optimizer to keep temporary indexes around for reuse.

Changing the INI value to '*JOB' prevents the temporary index from being saved in the Plan Cache; the index does not survive a hard close. The *JOB option causes the SQE optimizer use of temporary indexes to behave more like the CQE optimizer. The temporary index has a shorter life, but is still shared as long as there are active queries using it. This behavior can be desirable in cases where there is concern about increased maintenance costs for temporary indexes that persist for reuse.

A SQE temporary index can also be used as a source of statistics.

A temporary index is an internal data structure and can only be created by the database manager.

Visual explain icon:

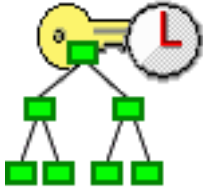


Temporary index scan

A temporary index scan operation is identical to the index scan operation that is performed upon the permanent radix index. It is still used to retrieve the rows from a table in a keyed sequence; however, the

temporary index object must first be created. All the rows in the index are sequentially processed, but the resulting row numbers are sequenced based upon the key columns.

The sequenced rows can be used by the optimizer to satisfy a portion of the query request (such as ordering or grouping).

<i>Table 22. Temporary index scan attributes</i>	
Data access method	Temporary index scan
Description	Sequentially scan and process all the keys associated with the temporary index.
Advantages	<ul style="list-style-type: none"> • Potential to extract all the data from the index key values, thus eliminating the need for a Table Probe • Returns the rows back in a sequence based upon the keys of the index
Considerations	Generally requires a Table Probe to be performed to extract any remaining columns required to satisfy the query. Can perform poorly when many rows are selected because of the random I/O associated with the Table Probe.
Likely to be used	<ul style="list-style-type: none"> • When sequencing the rows is required for the query (for example, ordering or grouping) • When the selection columns cannot be matched against the leading key columns of the index • When the overhead cost associated with the creation of the temporary index can be justified against other alternative methods to implement this query
Example SQL statement	<pre>SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' ORDER BY LastName OPTIMIZE FOR ALL ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3002 record and QQRID 3001 where QQKP(Index_Probe_Used) = 'N'.
SMP parallel enabled	Yes
Also referred to as	Index Scan Index Scan, Preload Index Scan, Distinct Index Scan Distinct, Preload Index Scan, Key Selection
Visual Explain icon	

Using the example above, the optimizer chose to create a temporary index to sequence the rows based upon the LastName column. A temporary index scan might then be performed to satisfy the ORDER BY clause in this query.

The optimizer determines where the selection against the WorkDept column best belongs. It can be performed as the temporary index itself is being created or it can be performed as a part of the temporary index scan. Adding the selection to the temporary index creation has the possibility of making the open data path (ODP) for this query non-reusable. This ODP reuse is considered when determining how selection is performed.

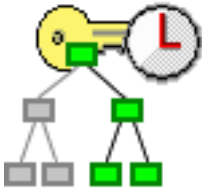
Temporary index probe

A temporary index probe operation is identical to the index probe operation that is performed on the permanent radix index. Its main function is to provide quick access against the index keys of the temporary index. However, it can still be used to retrieve the rows from a table in a keyed sequence.

The temporary index is used by the optimizer to satisfy the join portion of the query request.

<i>Table 23. Temporary index probe attributes</i>	
Data access method	Temporary index probe
Description	The index is quickly probed based upon the selection criteria that were rewritten into a series of ranges. Only those keys that satisfy the selection is used to generate a table row number.
Advantages	<ul style="list-style-type: none"> • Only those index entries that match any selection continue to be processed. Provides quick access to the selected rows • Potential to extract all the data from the index key values, thus eliminating the need for a Table Probe • Returns the rows back in a sequence based upon the keys of the index
Considerations	Generally requires a Table Probe to be performed to extract any remaining columns required to satisfy the query. Can perform poorly when many rows are selected because of the random I/O associated with the Table Probe.
Likely to be used	<ul style="list-style-type: none"> • When the ability to probe the rows required for the query (for example, joins) exists • When the selection columns cannot be matched against the leading key columns of the index • When the overhead cost associated with the creation of the temporary index can be justified against other alternative methods to implement this query
Example SQL statement	<pre> SELECT * FROM Employee XXX, Department YYY WHERE XXX.WorkDept = YYY.DeptNo OPTIMIZE FOR ALL ROWS </pre>
Database Monitor and Plan Cache record indicating use	QQRID 3002 record and QQRID 3001 where QQKP(Index_Probe_Used) = 'Y'.
SMP parallel enabled	Yes

Table 23. Temporary index probe attributes (continued)

Data access method	Temporary index probe
Also referred to as	Index Probe Index Probe, Preload Index Probe, Distinct Index Probe Distinct, Preload Index Probe, Key Selection
Visual Explain icon	

Using the example above, the optimizer chose to create a temporary index over the DeptNo column to help satisfy the join requirement against the DEPARTMENT table. A temporary index probe was then performed against the temporary index to process the join criteria between the two tables. In this particular case, there was no additional selection that might be applied against the DEPARTMENT table while the temporary index was being created.

Temporary buffer

The temporary buffer is a temporary object that is used to help facilitate operations such as parallelism. It is an unsorted data structure that is used to store intermediate rows of a query. The difference between a temporary buffer and a temporary list is that the buffer does not need to be fully populated before its results are processed.


The temporary buffer acts as a serialization point between parallel and non-parallel portions of a query. The operations used to populate the buffer cannot be performed in parallel, whereas the operations that fetch rows from the buffer can be performed in parallel.

The temporary buffer is required for SQE because the index scan and index probe operations are not SMP parallel-enabled for this engine. Unlike CQE, which performs these index operations in parallel, SQE does not subdivide the index operation work to take full advantage of parallel processing.

The buffer is used to allow a query to be processed under parallelism by serializing access to the index operations. Any remaining work within the query is processed in parallel.

A temporary buffer is an internal data structure and can only be created by the database manager.

Visual explain icon:




Buffer scan

The buffer scan is used when a query is processed using Db2 Symmetric Multiprocessing, yet a portion of the query is unable to be parallel processed. The buffer scan acts as a gateway to control access to rows between the parallel enabled portions of the query and the non-parallel portions.

Multiple threads can be used to fetch the selected rows from the buffer, allowing the query to perform any remaining processing in parallel. However, the buffer is populated in a non-parallel manner.

A buffer scan operation is identical to the list scan operation that is performed upon the temporary list object. The main difference is that a buffer does not need to be fully populated before the start of the scan operation. A temporary list requires that the list is fully populated before fetching any rows.

<i>Table 24. Buffer scan attributes</i>	
Data access method	Buffer scan
Description	Sequentially scan and process all the rows in the temporary buffer. Enables SMP parallelism to be performed over a non-parallel portion of the query.
Advantages	<ul style="list-style-type: none"> • The temporary buffer can be used to enable parallelism over a portion of a query that is non-parallel • The temporary buffer does not need to be fully populated in order to start fetching rows
Considerations	Used to prevent portions of the query from being processed multiple times when no key columns are required to satisfy the request.
Likely to be used	<ul style="list-style-type: none"> • When the query is attempting to take advantage of Db2 Symmetric Multiprocessing • When a portion of the query cannot be performed in parallel (for example, index scan or index probe)
Example SQL statement	<pre> CHGQRYA DEGREE(*OPTIMIZE) CREATE INDEX X1 ON Employee (LastName, WorkDept) SELECT * FROM Employee WHERE WorkDept BETWEEN 'A01' AND 'E01' AND LastName IN ('Smith', 'Jones', 'Peterson') OPTIMIZE FOR ALL ROWS </pre>
Database Monitor and Plan Cache record indicating use	QVPARU will be greater than 0 on the associated database monitor record.
SMP parallel enabled	Yes
Also referred to as	Not applicable
Visual Explain icon	

Using the example above, the optimizer chose to use the existing index X1 to perform an index probe operation against the table. In order to speed up the remaining Table Probe operation for this query, Db2 Symmetric Multiprocessing is used to perform the random probe into the table. Since the index probe is not SMP parallel-enabled for SQE, it is placed within a temporary buffer to control access to the selected index entries.

Queue

The Queue is a temporary object that the optimizer uses to feed recursion by putting data values needed for the recursion on it. This data typically includes those values used on the recursive join predicate, and other recursive data accumulated or manipulated during the recursive process.

The Queue has two operations allowed:

- Enqueue: puts data on the queue
- Dequeue: takes data off the queue

A queue is an efficient data structure because it contains only the data needed to feed the recursion or directly modified by the recursion process. Its size is managed by the optimizer.

Unlike other temporary objects created by the optimizer, the queue is not populated all at once by the underlying query node tree. It is a real-time temporary holding area for values feeding the recursion. In this regard, a queue is not considered temporary, as it does not prevent the query from running if ALWCPYDTA(*NO) was specified. The data can flow from the query at the same time the recursive values are inserted into the queue and used to retrieve additional join rows.

A queue is an internal data structure and can only be created by the database manager.

Visual explain icon:




Enqueue

During an enqueue operation, an entry is put on the queue. The entry contains key values used by the recursive join predicates or data manipulated as a part of the recursion process. The optimizer always supplies an enqueue operation to collect the required recursive data on the query node directly above the Union All.

<i>Table 25. Enqueue Attributes</i>	
Data Access Method	Enqueue
Description	Places an entry on the queue needed to cause further recursion
Advantages	<ul style="list-style-type: none"> • Required as a source for the recursion. Only enqueues required values for the recursion process. Each entry has short life span, until it is dequeued. • Each entry on the queue can seed multiple iterative fullselects that are recursive from the same RCTE or view.
Likely to be used	A required access method for recursive queries
Example SQL statement	<pre>WITH RPL (PART, SUBPART, QUANTITY) AS (SELECT ROOT.PART, ROOT.SUBPART, ROOT.QUANTITY FROM PARTLIST ROOT WHERE ROOT.PART = '01' UNION ALL SELECT CHILD.PART, CHILD.SUBPART, CHILD.QUANTITY FROM RPL PARENT, PARTLIST CHILD WHERE PARENT.SUBPART = CHILD.PART) SELECT DISTINCT PART, SUBPART, QUANTITY FROM RPL</pre>
Database Monitor and Plan Cache record indicating use	There are no explicit records that indicate the use of an enqueue
SMP parallel enabled	Yes
Also referred to as	Not applicable

Table 25. Enqueue Attributes (continued)

Data Access Method	Enqueue
Visual Explain icon	

Use the CYCLE option in the definition of the recursive query if the data reflecting the parent-child relationship could be cyclic, causing an infinite recursion loop. CYCLE prevents already visited recursive key values from being put on the queue again for a given set of related (ancestry chain) rows.

Use the SEARCH option in the definition of the recursive query to return the results of the recursion in the specified parent-child hierarchical ordering. The search choices are Depth or Breadth first. Depth first means that all the descendents of each immediate child are returned before the next child is returned. Breadth first means that each child is returned before their children are returned.

SEARCH requires not only the specification of the relationship keys, the columns which make up the parent-child relationship, and the search type of Depth or Breadth. It also requires an ORDER BY clause in the main query on the provided sequence column in order to fully implement the specified ordering.

Dequeue


During a dequeue operation, an entry is taken off the queue. Those values specified by recursive reference are fed back in to the recursive join process.

The optimizer always supplies a corresponding enqueue, dequeue pair of operations for each recursive common table expression or recursive view in the specifying query. Recursion ends when there are no more entries to pull off the queue.

Table 26. Dequeue Attributes

Data Access Method	Dequeue
Description	Removes an entry off the queue. Minimally, provides one side of the recursive join predicate that feeds the recursive join and other data values that are manipulated through the recursive process. The dequeue operation is always on the left side of the inner join with constraint, where the right side is the target child rows.
Advantages	<ul style="list-style-type: none"> • Provides quick access to recursive values • Allows for post selection of local predicate on recursive data values
Likely to be used	<ul style="list-style-type: none"> • A required access method for recursive queries • A single dequeued value can feed the recursion of multiple iterative fullselects that reference the same RCTE or view
Example SQL statement	<pre> WITH RPL (PART, SUBPART, QUANTITY) AS (SELECT ROOT.PART, ROOT.SUBPART, ROOT.QUANTITY FROM PARTLIST ROOT WHERE ROOT.PART = '01' UNION ALL SELECT CHILD.PART, CHILD.SUBPART, CHILD.QUANTITY FROM RPL PARENT, PARTLIST CHILD WHERE PARENT.SUBPART = CHILD.PART) SELECT DISTINCT PART, SUBPART, QUANTITY FROM RPL </pre>

Table 26. Dequeue Attributes (continued)

Data Access Method	Dequeue
Database Monitor and Plan Cache record indicating use	There are no explicit records that indicate the use of the dequeue operation.
SMP parallel enabled	Yes
Also referred to as	Not applicable
Visual Explain icon	

Array unnest temporary table


The array unnest temporary table is a temporary object that holds the output of an UNNEST of an array or a list of arrays. It can be viewed vertically, with each column of array values having the same format. The temporary table contains one or more arrays specified by the user in an UNNEST clause of a SELECT statement.

UNNEST creates a temporary table with the arrays specified as columns in the table. If more than one array is specified, the first array provides the first column in the result table. The second array provides the second column, and so on.

The arrays might be of different lengths. Shorter arrays are primed with nulls to match the length of the longest array in the list.

If WITH ORDINALITY is specified, an extra counter column of type BIGINT is appended to the temporary table. The ordinality column contains the index position of the elements in the arrays.

The array unnest temporary table is an internal data structure and can only be created by the database manager.



Visual explain icon:

Related reference

[QAQQINI query options](#)

There are different options available for parameters in the QAQQINI file.

Related information

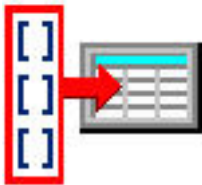
[Array support in SQL procedures](#)

[Debugging an SQL routine](#)

[table-reference](#)

Array unnest temporary table scan

During an array unnest temporary table scan operation, the temporary table is processed one row at a time.

Table 27. Array unnest temporary table scan operation	
Data access method	Array unnest temporary table scan
Description	Sequentially scan and process all the rows of data in the unnest temporary table.
Advantages	The array unnest temporary table and temporary table scan can be used to simplify the logic flow of the optimizer for processing arrays.
Likely to be used	When an UNNEST clause is specified in the from-clause of an SQL fullselect.
Example SQL statement	<pre>CREATE PROCEDURE processCustomers() BEGIN DECLARE ids INTARRAY; DECLARE names STRINGARRAY; set ids = ARRAY[5,6,7]; set names = ARRAY['Ann', 'Bob', 'Sue']; INSERT INTO customerTable(id, name, order) (SELECT Customers.id, Customers.name, Customers.order FROM UNNEST(ids, names) WITH ORDINALITY AS Customers(id, name, order)); END CALL processCustomers()</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3000 where QVQTBL = '*UNNEST'
SMP parallel enabled	Yes
Also referred to as	
Visual Explain icon	

Temporary Indexed List

The temporary indexed list is a temporary object that allows the optimizer to sequence rows based upon a column or set of columns. The temporary indexed list can be scanned by the optimizer to satisfy ordering or grouping requirements of the query.

A temporary indexed list is a data structure that contains a radix index object to provide ordering. This object is generally used when a query contains ordering and non-equal predicates and only a partial answer set is fetched. The object is initially populated with a distinct set of values by distinct scanning or probing an existing index. As processing continues and more rows are needed, additional rows will be inserted into the temporary indexed list for additional processing.

If the underlying index used to build the temporary indexed list contains all the necessary columns to satisfy any further processing, then these columns are built into the temporary indexed list. This population avoids any random I/Os associated with a table probe operation. If the underlying index does not contain all the necessary columns, then a table probe operation is required to recollect the missing columns from the temporary indexed list before the selected rows can be processed.

A temporary indexed list is an internal data structure and can only be created by the database manager. The temporary indexed list is always used in conjunction with the temporary indexed list scan and index merge data access methods.

Visual explain icon:



Temporary Indexed List scan and Index Merge

A temporary indexed list scan operation is similar to the index scan operation that is performed upon the permanent radix index. It is still used to retrieve the rows from a table in a keyed sequence; however, the temporary indexed list object must first be created. This operation is often performed in conjunction with an index merge operation. However, for some distinct and grouping queries, the temporary indexed list scan may be used without an associated index merge operation.

This access method is used when there is non-equal selection and is used in conjunction with an existing permanent radix index. A distinct index scan or a distinct index probe operation is performed against the index to get a unique set of rows. The rows are then inserted into the temporary indexed list to provide the appropriate ordering.


In general, for a radix index to be eligible for temporary list scan and Index Merge Ordering, the index should be created as follows:

- Any columns with equal predicates that are not in the ORDER BY list should be first keys in the index.
- Followed by any columns for non-equal predicates that are not in the ORDER BY list. The columns should be arranged so the most selective predicate column is the first column following the equal predicate columns.
- Followed by the ORDER BY columns in the same order as the ORDER BY clause of the query and with the same ASC or DESC attribute.
- Optionally, at the end you may include other selection columns or other selected columns for Index Only Access.

A temporary sorted list is an internal data structure and can only be created by the database manager.

<i>Table 28. Temporary indexed list scan</i>	
Data access method	Temporary indexed list scan
Description	Sequentially scan and process all the keys associated with the temporary indexed list.
Advantages	<ul style="list-style-type: none"> • Provides an alternate implementation to order data when the query contains non-equal where selection. • Returns the rows back in a sequence based upon the keys of the index. • Provides better paging characteristics than a sorted list scan in low memory environments. • Provides better performance characteristics than a Sorted List Scan when only a partial answer set is fetched.
Considerations	Used to process ordering, grouping or distinct processing for a single table query.

Table 28. Temporary indexed list scan (continued)

Data access method	Temporary indexed list scan
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA). • When the data is required to be ordered or grouped based upon a column or columns and the query contains non-equal where selection. • The query is being optimized for FIRST I/O. • The number of distinct values is low.
Example SQL statement	<pre>CREATE INDEX INDEX1 ON EMPLOYEE(SALARY, WORKDEPT) SELECT DISTINCT WORKDEPT FROM EMPLOYEE WHERE SALARY > 30000 OPTIMIZE FOR 30 ROWS</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3001 where QQRCOD = 'I7'
SMP parallel enabled	No
Also referred to as	
Visual Explain icon	

Using the example above, the optimizer chose to create a temporary indexed list of distinct WORKDEPT values from index INDEX1. The selection SALARY > 3000 was applied to the index probe distinct.

Index Merge:


An index merge operation is used to provide ordering in conjunction with a temporary indexed list scan and an index probe distinct or distinct index.

The sequenced rows can be used by the optimizer to satisfy a portion of the query request (such as ordering or grouping).

Table 29. Index Merge

Data access method	Index Merge
Description	Sequentially scan and process all the keys associated with the temporary indexed list.
Advantages	<ul style="list-style-type: none"> • Potential to extract all the data from the index key values, thus eliminating the need for a Table Probe. • Returns the rows back in a sequence based upon the keys of the index.

Table 29. Index Merge (continued)

Data access method	Index Merge
Considerations	Used to process ordering, grouping or distinct processing for a single table query.
Likely to be used	<ul style="list-style-type: none"> • When the use of temporary results is allowed by the query environmental parameter (ALWCPYDTA). • When the data is required to be ordered or grouped based upon a column or columns and the query contains non-equal where selection. • The query is being optimized for FIRST I/O. • The number of distinct values is low.
Example SQL statement	<pre>CREATE INDEX IX1 ON Sales(Sales_date, Region, Sales_person) SELECT * FROM Sales WHERE Sales_date BETWEEN '1996-03-29' AND '1996-04-29' AND Region IN ('Quebec','Manitoba') ORDER BY Sales_person;</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3001 where QQRCD = 'I7'
SMP parallel enabled	No
Also referred to as	
Visual Explain icon	

Window

The window is a temporary object that holds intermediate results needed to determine the result for On-Line Analytical Processing (OLAP) functions. The number of rows in the window may change as the data flows through it.


The window is an internal data structure and can only be created by the database manager.

Visual explain icon:



Window scan

During processing of an OLAP function the window may be processed multiple time in order to determine the result of the OLAP function for each row

Table 30. Window scan attributes	
Data access method	Window scan
Description	Sequentially scan and process all the rows needed to determine the result of the OLAP function.
Advantages	<ul style="list-style-type: none">Allows the database manager to process subsets of the data to calculate the OLAP result without processing the entire result set for each row.
Considerations	Required for OLAP specifications that require intermediate results to determine the value of the OLAP expression. The OLAP specifications that require a window include any of the aggregate OLAPs, LAG, LEAD, NTILE, and CUME_DIST.
Likely to be used	Required for OLAP specifications that require intermediate results to determine the value of the OLAP expression.
Example SQL statement	<pre>SELECT workdept, empno, salary, decimal(avg(salary) over(partition by workdept order by salary), 10,2) FROM employee</pre>
Database Monitor and Plan Cache record indicating use	QQRID 3004 where QQRCOD = 'H8'.
SMP parallel enabled	No
Also referred to as	
Visual Explain icon	

Objects processed in parallel

The Db2 Symmetric multiprocessing feature provides the optimizer with additional methods for retrieving data that include parallel processing. Symmetrical multiprocessing is a form of parallelism achieved on a single system where multiple CPU and I/O processors sharing memory and disk work simultaneously toward a single result.

This parallel processing means that the database manager can have more than one (or all) of the system processors working on a single query simultaneously. The performance of a CPU-bound query can be improved with this feature on multiple-processor systems by distributing the processor load across more than one processor.

The preceding tables indicate what data access methods are enabled to take advantage of the Db2 Symmetric Multiprocessing feature. An important thing to note, however, is that the parallel implementation differs for both the SQL Query Engine and the Classic Query Engine.

Processing requirements

Parallelism requires that SMP parallel processing must be enabled by one of the following methods:

- System value QQRVDEGREE

- Query option file
- DEGREE parameter on the **Change Query Attributes (CHGQRYA)** command
- SQL SET CURRENT DEGREE statement

Once parallelism has been enabled, a set of database system tasks or threads is created at system startup for use by the database manager. The database manager uses the tasks to process and retrieve data from different disk devices. Since these tasks can be run on multiple processors simultaneously, the elapsed time of a query can be reduced. Even though the tasks do much of the parallel I/O and CPU processing, the I/O and CPU resource accounting is transferred to the application job. The summarized I/O and CPU resources for this type of application continue to be accurately displayed by the **Work with Active Jobs (WRKACTJOB)** command.

The job must be run in a shared storage pool with the *CALC paging option, as this method causes more efficient use of active memory.

Related concepts

[Nested loop join implementation](#)

Db2 for i provides a **nested loop** join method. For this method, the processing of the tables in the join are ordered. This order is called the **join order**. The first table in the final join order is called the **primary table**. The other tables are called **secondary tables**. Each join table position is called a **dial**.

Related reference

[Changing the attributes of your queries](#)

You can modify different types of query attributes for a job with the **Change Query Attributes (CHGQRYA)** CL command. You can also use the System i Navigator Change Query Attributes interface.

Related information

[SET CURRENT DEGREE statement](#)

[Performance system values: Parallel processing for queries and indexes](#)

[Adjusting performance automatically](#)

[Work with Active Jobs \(WRKACTJOB\) command](#)

[Change Query Attributes \(CHGQRYA\) command](#)

[DB2 Symmetric Multiprocessing](#)

Spreading data automatically

Db2 for i automatically spreads the data across the disk devices available in the auxiliary storage pool (ASP) where the data is allocated. This process ensures that the data is spread without user intervention.

The spreading allows the database manager to easily process the blocks of rows on different disk devices in parallel. Even though Db2 for i spreads data across disk devices within an ASP, sometimes the allocation of the data extents (contiguous sets of data) might not be spread evenly. This unevenness occurs when there is uneven allocation of space on the devices, or when a new device is added to the ASP. The allocation of the table data space could be spread again by saving, deleting, and then restoring the table.

Maintaining an even distribution of data across all the disk devices can lead to better throughput on query processing. The number of disk devices used and how the data is spread across them is considered by the optimizer while costing the different plan permutations.

Processing queries: Overview

This overview of the query optimizer provides guidelines for designing queries that perform and use system resources more efficiently.

This overview covers queries that are optimized by the query optimizer and includes interfaces such as SQL, OPNQRYF, APIs (QQQRY), ODBC, and Query/400 queries. Whether you apply the guidelines, the query results are still correct.

Note: The information in this overview is complex. You might find it helpful to experiment with an IBM i product as you read this information to gain a better understanding of the concepts.

When you understand how Db2 for i processes queries, it is easier to understand the performance impacts of the guidelines discussed in this overview. There are two major components of Db2 for i query processing:

- How the system accesses data.

These methods are the algorithms that are used to retrieve data from the disk. The methods include index usage and row selection techniques. In addition, parallel access methods are available with the Db2 Symmetric Multiprocessing operating system feature.

- Query optimizer

The query optimizer identifies the valid techniques which can be used to implement the query and selects the most efficient technique.

How the query optimizer makes your queries more efficient

Data manipulation statements such as SELECT specify only what data the user wants, not how to retrieve that data. This path to the data is chosen by the optimizer and stored in the access plan. Understand the techniques employed by the query optimizer for performing this task.

The optimizer is an important part of Db2 for i because the optimizer:

- Makes the key decisions which affect database performance.
- Identifies the techniques which can be used to implement the query.
- Selects the most efficient technique.

General query optimization tips

Here are some tips to help your queries run as fast as possible.

- Create indexes whose leftmost key columns match your selection predicates to help supply the optimizer with selectivity values (key range estimates).
- For join queries, create indexes that match your join columns to help the optimizer determine the average number of matching rows.
- Minimize extraneous mapping by specifying only columns of interest on the query. For example, specify only the columns you need to query on the SQL SELECT statement instead of specifying SELECT *. Also, specify FOR FETCH ONLY if the columns do not need to be updated.
- If your queries often use table scan, use the **Reorganize Physical File Member (RGZPFM)** command to remove deleted rows from tables, or the **Change Physical File (CHGPF) REUSEDLT (*YES)** command to reuse deleted rows.

Consider using the following options:

- Specify ALWCOPYDTA(*OPTIMIZE) to allow the query optimizer to create temporary copies of data so better performance can be obtained. The IBM i Access ODBC driver and Query Management driver always use this mode. If ALWCOPYDTA(*YES) is specified, the query optimizer attempts to implement the query without copies of the data, but might create copies if required. If ALWCOPYDTA(*NO) is specified, copies of the data are not allowed. If the query optimizer cannot find a plan that does not use a temporary, then the query cannot be run.
- For SQL, use CLOSQLCSR(*ENDJOB) or CLOSQLCSR(*ENDACTGRP) to allow open data paths to remain open for future invocations.
- Specify DLYPRP(*YES) to delay SQL statement validation until an OPEN, EXECUTE, or DESCRIBE statement is run. This option improves performance by eliminating redundant validation.
- Use ALWBK(*ALLREAD) to allow row blocking for read-only cursors.

Related information

[Reorganize Physical File Member \(RGZPFM\) command](#)

Access plan validation

An access plan is a control structure that describes the actions necessary to satisfy each query request. It contains information about the data and how to extract it. For any query, whenever optimization occurs, the query optimizer develops an optimized plan of how to access the requested data.

To improve performance, an access plan is saved once it is built (see following exceptions), to be available for potentially future runs of the query. However, the optimizer has dynamic replan capability. This means that even if a previously built (and saved) plan is found, the optimizer could rebuild it if a more optimal plan is possible. This process allows for maximum flexibility while still taking advantage of saved plans.

- For dynamic SQL, an access plan is created at prepare or open time. However, optimization uses the host variable values to determine an optimal plan. Therefore, a plan built at prepare time could be rebuilt the first time the query is opened (when the host variable values are present).
- For an IBM i program that contains static embedded SQL, an access plan is initially created at compile time. Again, since optimization uses the host variable values to determine an optimal plan, the compile-time plan could be rebuilt the first time the query is opened.
- For Open Query File (OPNQRYF), an access plan is created but is not saved. A new access plan is created each time the OPNQRYF command is processed.
- For Query/400, an access plan is saved as part of the query definition object.

In all the preceding cases where a plan is saved, including static SQL, dynamic replan can still apply as the queries are run over time.

The access plan is validated when the query is opened. Validation includes the following:

- Verifying that the same tables are referenced in the query as in the access plan. For example, the tables were not deleted and recreated or that the tables resolved by using *LIBL have not changed.
- Verifying that the indexes used to implement the query, still exist.
- Verifying that the table size or predicate selectivity has not changed significantly.
- Verifying that QAQQINI options have not changed.

Single table optimization

At run time, the optimizer chooses an optimal access method for a query by calculating an *implementation cost* based on the current state of the database. The optimizer uses two costs in its decision: an I/O cost and a CPU cost. The goal of the optimizer is to minimize both I/O and CPU cost.

Improved query optimization I/O cost estimates

The time it takes to perform a disk I/O operation can vary according to the connecting infrastructure, the external or internal nature of the media and media type, spinning disk or Solid State Disk. Consequently, the total I/O cost associated with a particular query access method may vary from system to system.

In order to more accurately estimate these costs, the optimizer considers the performance of each disk unit individually. It does this by measuring the time it takes for read operations to complete across a sample of pages across the disk. This analysis is done at each IPL for disks in the system and user ASPs and at vary-on time for independent ASPs. With this information and with the additional knowledge about how database objects are spread across various disk units, the optimizer can make a reasonable estimate about the time it takes to perform I/O against a given database object. This means that no matter where your data resides, and even as it moves around, the optimizer can choose the most efficient plan to execute your queries.

Optimizing Access to each table

The optimizer uses a general set of guidelines to choose the best method for accessing data in each table. The optimizer:

- Determines the default filter factor for each predicate in the selection clause.
- Determines the true filter factor of the predicates by key range estimate when the selection predicates match the index left-most keys, or by available column statistics.
- Determines the cost of table scan processing if an index is not required.
- Determines the cost of creating an index over a table if an index is required. This index is created by performing either a table scan or creating an index-from-index.
- Determines the cost of using a sort routine or hashing method if appropriate.
- Determines the cost of using existing indexes using Index Probe or Index Scan
 - Orders the indexes. For SQE, the indexes are ordered in general such that the indexes that access the smallest number of entries are examined first. For CQE, the indexes are ordered from mostly recently created to oldest.
 - For each index available, the optimizer does the following:
 - Determines if the index meets the selection criteria.
 - Determines the cost of using the index by estimating the number of I/Os and CPU needed to Index Probe or Index Scan, and possible Table Probes.
 - Compares the cost of using this index with the previous cost (current best).
 - Picks the cheaper one.
 - Continues to search for best index until the optimizer decides to look at no more indexes.

SQE orders the indexes so that the best indexes are examined first. Once an index is found that is more expensive than the previously chosen best index, the search is ended.

For CQE, the *time limit* controls how much time the optimizer spends choosing an implementation. The time limit is based on how much time was spent so far and the current best implementation cost found. The idea is to prevent the optimizer from spending more time optimizing the query than it takes to actually execute the query. Dynamic SQL queries are subject to the optimizer time restrictions. Static SQL query optimization time is not limited. For OPNQRYF, if you specify OPTALLAP(*YES), the optimization time is not limited.

For small tables, the query optimizer spends little time in query optimization. For large tables, the query optimizer considers more indexes. For CQE, the optimizer generally considers five or six indexes for each table of a join before running out of optimization time. Because of this processing, it is normal for the optimizer to spend longer lengths of time analyzing queries against the tables.

- Determines the cost of using a temporary bitmap
 - Order the indexes that can be used for bit mapping. In general the indexes that select the smallest number of entries are examined first.
 - Determine the cost of using this index for bit mapping and the cost of merging this bitmap with any previously generated bitmaps.
 - If the cost of this bitmap plan is cheaper than the previous bitmap plan, continue searching for bitmap plans.
- After examining the possible methods of access the data for the table, the optimizer chooses the best plan from all the plans examined.

Solid State Drives

Solid State Drives (SSDs) offer a number of advantages over traditional hard disk drives (HDDs)

Solid State Drives

Solid State Drives (SSDs) offer a number of advantages over traditional hard disk drives (HDDs). With no seek time or rotational delays, SSDs can deliver substantially better I/O performance than HDDs. Capable of driving tens of thousands of I/O operations per second as opposed to hundreds for HDDs, SSDs break through performance bottlenecks of I/O-bound applications. Applications that require dozens and dozens

of “extra” HDDs for performance can meet their I/O performance requirements with far fewer SSDs, resulting in energy, space, and cost savings.

As IBM i has its own storage manager and Db2 for i built in, the integration of SSDs on IBM i is a fairly simple task. The functions provided for management of SSDs and adjusting their impact on Applications and Database are very simple and easy to use.

There are three basic methodologies to place data on SSD.

- ASP Balancer – Enhanced for SSDs
- Library and SSD Integration
- Db2 and SSD Integration

To allow you to specify what data should be allocated on SSD, Db2 has provided the capability to specify a “media preference” as an attribute of a database table, partition, or index. It should be noted that this attribute specifies that storage allocations on SSD are preferred, but if no SSD disks are available or if the SSD disks do not have enough space left to allocate the entire object, at least some part of the object will be allocated on traditional disks. See the **UNIT** parameter on **CRTPF** and **CRTL** or the **UNIT** clause on the **CREATE TABLE**, **CREATE INDEX** and **ALTER TABLE** SQL statements.

You should consider SSDs if your I/O demands have outpaced the performance capabilities of traditional HDDs, latencies associated with spinning platters and moving arms limit the speed of HDD data access. SSDs’ near instantaneous data access removes this I/O bottleneck, creating a paradigm shift in I/O performance. Applications throttled by poor I/O performance can benefit greatly from SSDs.

Memory preference controls

Memory preference controls can be used as a technique to maximize performance and utilization of resources.

Memory preference controls

Memory preference controls can be used against performance critical database tables, indexes, physical files, and logical files as a technique to maximize performance and utilization of resources. Several approaches are available for controlling the memory preference:

1. **Set Object Access (SETOBJACC)** command

One benefit of **SETOBJACC** is that you can carve out a separate memory pool that is not used by from any running applications or **MEMORY_POOL_PREFERENCE** and those objects will then not get paged out because neither applications nor SQE will be using that pool. If the target objects are primarily accessed using Native database I/O, **SETOBJACC** is the preferred approach. **SETOBJACC** uses a single thread to bring the object into memory.

2. **Change Physical File (CHGPF)** and **Change Logical File (CHGLF)** commands - Keep in memory (**KEEPINMEM**) parameter

When an object is changed to have Keep in memory set to *YES, the database will bring the object into memory and attempt to keep it in memory when it is accessed using SQL via SQE. Native database I/O (for example RPG CHAIN, READ, etc.) does not do this. **KEEPINMEM** has the ability to use parallel I/O to bring the object into memory.

- **CHGPF KEEPINMEM(*YES|*NO)**
- **CHGLF KEEPINMEM(*YES|*NO)**

3. The SQL memory-preference can be used as an alternative to the **KEEPINMEM** command parameter.

The behavior of SQL configured in memory objects matches the behavior described in the **KEEPINMEM** section.

KEEP IN MEMORY <NO/YES> is available on the following SQL statements:

- **ALTER TABLE**
- **CREATE INDEX**

- **CREATE TABLE**
- **DECLARE GLOBAL TEMPORARY TABLE**

Note: The QSYS2/SYSPARTITIONSTAT and SYSPARTITIONINDEXSTAT catalogs can be queried to determine the memory-preference for specific objects. When a memory-preference is specified for an object, the MEMORY_POOL_PREFERENCE QAQQINI option can be used to influence where we attempt to page objects. There is no guarantee that objects will remain in memory.

Join optimization

A join operation is a complex function that requires special attention in order to achieve good performance. This section describes how Db2 for i implements join queries and how optimization choices are made by the query optimizer. It also describes design tips and techniques which help avoid or solve performance problems.

Nested loop join implementation

Db2 for i provides a **nested loop** join method. For this method, the processing of the tables in the join are ordered. This order is called the **join order**. The first table in the final join order is called the **primary table**. The other tables are called **secondary tables**. Each join table position is called a **dial**.

The nested loop is implemented either using an index on secondary tables, a hash table, or a table scan (arrival sequence) on the secondary tables. In general, the join is implemented using either an index or a hash table.

Index nested loop join implementation

During the join, Db2 for i:

1. Accesses the first primary table row selected by the predicates local to the primary table.
2. Builds a key value from the join columns in the primary table.
3. Chooses the access to the first secondary table:
 - If using an index, Radix Index Probe is used to locate the first row satisfying the join condition for the secondary table. The probe uses an index with keys matching the join condition or local row selection columns of the secondary table.
 - Applies bitmap selection, if applicable.

All rows that satisfy the join condition from each secondary dial are located using an index. Rows are retrieved from secondary tables in random sequence. This random disk I/O time often accounts for a large percentage of the processing time of the query. Since a given secondary dial is searched once for each row selected from the primary and the preceding secondary dials that satisfy the join condition for each of the preceding secondary dials, many searches could be against the later dials. Any inefficiencies in the processing of the later dials can significantly inflate the query processing time. This reason is why attention to performance considerations for join queries can reduce the run time of a join query from hours to minutes.

If an efficient index cannot be found, a temporary index could be created. Some join queries build temporary indexes over secondary dials even when an index exists for all the join keys. Because efficiency is important for secondary dials of longer running queries, the optimizer could build a temporary index containing only entries with local row selection for that dial. This preprocessing of row selection allows the database manager to process row selection in one pass instead of each time rows are matched for a dial.

- If using a Hash Table Probe, a hash temporary result table is created containing all rows from local selection against the table on the first probe. The structure of the hash table is such that rows with the same join value are loaded into the same hash table partition (clustered). The location of the rows for any given join value can be found by applying a hashing function to the join value.

A nested loop join using a Hash Table Probe has several advantages over a nested loop join using an Index Probe:

- The structure of a hash temporary result table is simpler than the structure of an index. Less CPU processing is required to build and probe a hash table.
 - The rows in the hash result table contain all the data required by the query. There is no need to access the dataspace of the table with random I/O when probing the hash table.
 - Like join values are clustered, so all matching rows for a given join value can typically be accessed with a single I/O request.
 - The hash temporary result table can be built using SMP parallelism.
 - Unlike indexes, entries in hash tables are not updated to reflect changes of column values in the underlying table. The existence of a hash table does not affect the processing cost of other updating jobs in the system.
 - If using a Sorted List Probe, a sorted list result is created containing all the rows from local selection against the table on the first probe. The structure of the sorted list table is such that rows with the same join value are sorted together in the list. The location of the rows for any given join value can be found by probing using the join value.
 - If using a Table Scan, locate the first row that satisfies the join condition or local row selection columns of the secondary table. The join could be implemented with a table scan when the secondary table is a user-defined table function.
4. Determines if the row is selected by applying any remaining selection local to the first secondary dial. If the secondary dial row is not selected then the next row that satisfies the join condition is located. Steps 1 through 4 are repeated until a row that satisfies both the join condition and any remaining selection is selected from all secondary tables
5. Returns the result join row.
6. Processes the last secondary table again to find the next row that satisfies the join condition in that dial.
- During this processing, when no more rows satisfying the join condition can be selected, the processing backs up to the logical previous dial. It attempts to read the next row that satisfies its join condition.
7. Ends processing when all selected rows from the primary table are processed.

Note the following characteristics of a nested loop join:

- If ordering or grouping is specified, and all the columns are over a single table eligible to be the primary, then the optimizer costs the join with that table as the primary table, performing the grouping and ordering with an index.
- If ordering and grouping is specified on two or more tables or if temporary objects are allowed, Db2 for i breaks the processing of the query into two parts:
 1. Perform the join selection, omitting the ordering or grouping processing, and write the result rows to a temporary work table. This method allows the optimizer to consider any table of the join query as a candidate for the primary table.
 2. Perform the ordering or grouping on the data in the temporary work table.

Queries that cannot use hash join

Hash join cannot be used for queries that:

- Hash join cannot be used for queries involving physical files or tables that have read triggers.
- Require that the cursor position is restored as the result of the SQL ROLLBACK HOLD statement or the ROLLBACK CL command. For SQL applications using commitment control level other than *NONE, this method requires that *ALLREAD be specified as the value for the ALWBLK precompiler parameter.
- Hash join cannot be used for a table in a join query where the join condition something other than an equals operator.
- CQE does not support hash join if the query contains any of the following:

- Subqueries unless all subqueries in the query can be transformed to inner joins.
- UNION or UNION ALL
- Perform left outer or exception join.
- Use a DDS created join logical file.

Related concepts

Objects processed in parallel

The Db2 Symmetric multiprocessing feature provides the optimizer with additional methods for retrieving data that include parallel processing. Symmetrical multiprocessing is a form of parallelism achieved on a single system where multiple CPU and I/O processors sharing memory and disk work simultaneously toward a single result.

Related reference

Table scan

A table scan is the easiest and simplest operation that can be performed against a table. It sequentially processes all the rows in the table to determine if they satisfy the selection criteria specified in the query. It does this processing in a way to maximize the I/O throughput for the table.

Sorted list probe

A sorted list probe operation is used to retrieve rows from a temporary sorted list based upon a probe lookup operation.

Hash table probe

A hash table probe operation is used to retrieve rows from a temporary hash table based upon a probe lookup operation.

Radix index probe

A radix index probe operation is used to retrieve the rows from a table in a keyed sequence. The main difference between the radix index probe and the scan is that the rows returned are first identified by a probe operation to subset them.

Join optimization algorithm

The query optimizer must determine the join columns, join operators, local row selection, dial implementation, and dial ordering for a join query.

The join columns and join operators depend on the following situations:

- Join column specifications of the query
- Join order
- Interaction of join columns with other row selection

Join specifications not implemented for the dial are deferred until a later dial or, if an inner join, processed as row selection.

For a given dial, the only join specifications which are usable as join columns are those being joined to a *previous* dial. For example, the second dial can only use join specifications which reference columns in the primary dial. Likewise, the third dial can only use join specifications which reference columns in the primary and the second dials, and so on. Join specifications which reference later dials are deferred until the referenced dial is processed.

Note: For OPNQRYF, only one type of join operator is allowed for either a left outer or an exception join. That is, the join operator for all join conditions must be the same.

When looking for an existing index to access a secondary dial, the query optimizer looks at the left-most key columns of the index. For a given dial and index, the join specifications which use the left-most key columns can be used. For example:

```

DECLARE BROWSE2 CURSOR FOR
SELECT * FROM EMPLOYEE, EMP_ACT
WHERE EMPLOYEE.EMPNO = EMP_ACT.EMPNO
AND EMPLOYEE.HIREDATE = EMP_ACT.EMSTDATE
OPTIMIZE FOR 99999 ROWS

```

For the index over EMP_ACT with key columns EMPNO, PROJNO, and EMSTDATE, the join operation is performed only on column EMPNO. After the join is performed, index scan-key selection is done using column EMSTDATE.

The query optimizer also uses local row selection when choosing the best use of the index for the secondary dial. If the previous example had been expressed with a local predicate as:

```
DECLARE BROWSE2 CURSOR FOR
SELECT * FROM EMPLOYEE, EMP_ACT
WHERE EMPLOYEE.EMPNO = EMP_ACT.EMPNO
AND EMPLOYEE.HIREDATE = EMP_ACT.EMSTDATE
AND EMP_ACT.PROJNO = '123456'
OPTIMIZE FOR 99999 ROWS
```

The index with key columns EMPNO, PROJNO, and EMSTDATE are fully utilized by combining join and selection into one operation against all three key columns.

When creating a temporary index, the left-most key columns are the usable join columns in that dial position. All local row selection for that dial is processed when selecting entries for inclusion into the temporary index. A temporary index is like the index created for a select/omit keyed logical file. The temporary index for the previous example has key columns of EMPNO and EMSTDATE.

Since the optimizer tries a combination of join and local row selection, you can achieve almost all the advantages of a temporary index by using an existing index. In the preceding example, using either implementation, an existing index could be used or a temporary index could be created. A temporary index is built with the local row selection on PROJNO applied during the index creation. The temporary index has key columns of EMPNO and EMSTDATE to match the join selection.

If, instead, an existing index was used with key columns of EMPNO, PROJNO, EMSTDATE (or PROJNO, EMP_ACT, EMSTDATE), the local row selection can be applied **at the same time** as the join selection. This method contrasts to applying the local selection before the join selection, as happens when the temporary index is created. Or applying the local selection after the join selection, as happens when only the first key column of the index matches the join column.

The existing index implementation is more likely to provide faster performance because join and selection processing are combined without the overhead of building a temporary index. However, the existing index could have slightly slower I/O processing than the temporary index because the local selection is run many times rather than once. In general, create indexes with key columns for the combination of join and equal selection columns as the left-most keys.

Join order optimization

The SQE optimizer allows join reordering for a join logical file. However, the join order is fixed if CQE runs a query that references a join logical file. The join order is also fixed if the OPNQRYF JORDER(*FILE) parameter is specified. In addition, the join order is fixed if the query options file (QAQQINI) FORCE_JOIN_ORDER parameter is *YES

Otherwise, the following join ordering algorithm is used to determine the order of the tables:

1. Determine an access method for each individual table as candidates for the primary dial.
2. Estimate the number of rows returned for each table based on local row selection.

If the join query with ordering or grouping is processed in one step, the table with the ordering or grouping columns is the primary table.

3. Determine an access method, cost, and expected number of rows returned for each join combination of candidate tables as primary and first secondary tables.

The join order combinations estimated for a four table inner join would be:

```
1-2  2-1  1-3  3-1  1-4  4-1  2-3  3-2  2-4  4-2  3-4  4-3
```

4. Choose the combination with the lowest join cost and number of selected rows or both.
5. Determine the cost, access method, and expected number of rows for each remaining table joined to the previous secondary table.

6. Select an access method for each table that has the lowest cost for that table.
7. Choose the secondary table with the lowest join cost and number of selected rows or both.
8. Repeat steps 4 through 7 until the lowest cost join order is determined.

Note: After dial 32, the optimizer uses a different method to determine file join order, which might not be the lowest cost.

When a query contains a left or right outer join or a right exception join, the join order is not fixed. However, all from-columns of the ON clause must occur from dials previous to the left or right outer or exception join. For example:

```
FROM A INNER JOIN B ON A.C1=B.C1
LEFT OUTER JOIN C ON B. C2=C.C2
```

The allowable join order combinations for this query would be:

1-2-3, 2-1-3, or 2-3-1

Right outer or right exception joins are implemented as left outer and left exception, with files flipped. For example:

```
FROM A RIGHT OUTER JOIN B ON A.C1=B.C1
```

is implemented as B LEFT OUTER JOIN A ON B.C1=A.C1. The only allowed join order is 2-1.

Related information

[Open Query File \(OPNQRYF\) command](#)

[Change Query Attributes \(CHGQRYA\) command](#)

Full outer join

Full outer join is supported by the SQE optimizer. Just as right outer and right exception join are rewritten to the supported join types of inner, left outer or left exception, a full outer join is also rewritten.

A full outer join of A FULL OUTER JOIN B is equivalent to a (A LEFT OUTER JOIN B) UNION ALL (B LEFT EXCEPTION JOIN A). The following example illustrates the rewrite.

```
SELECT EMPNO, LASTNAME, DEPTNAME
FROM CORPDATA.EMPLOYEE XXX
FULL OUTER JOIN CORPDATA.DEPARTMENT YYY
ON XXX.WORKDEPT = YYY.DEPTNO
```

This query is rewritten as the following:

```
SELECT EMPNO, LASTNAME, DEPTNAME
FROM CORPDATA.EMPLOYEE XXX
LEFT OUTER JOIN CORPDATA.DEPARTMENT YYY
ON XXX.WORKDEPT = YYY.DEPTNO
UNION ALL
SELECT EMPNO, LASTNAME, DEPTNAME
FROM CORPDATA.DEPARTMENT YYY
LEFT EXCEPTION JOIN CORPDATA.EMPLOYEE XXX
ON XXX.WORKDEPT = YYY.DEPTNO
```

A query with multiple FULL OUTER JOIN requests, such as A FULL OUTER JOIN B FULL OUTER JOIN C can quickly become complicated in this rewritten state. This complication is illustrated in the following example.

If not running in live data mode, the optimizer could facilitate performance both during optimization and runtime by encapsulating intermediate results in a temporary data object. This object can be optimized once and plugged into both the scanned and probed side of the rewrite. These shared temporary objects eliminate the need to make multiple passes through the specific tables to satisfy the request.

In this example, the result of the (A FULL OUTER JOIN B) is a candidate for encapsulation during its FULL OUTER join with C.


```
A FULL OUTER JOIN B FULL OUTER JOIN C
```

This query is rewritten as the following:

```
((A LEFT OUTER JOIN B) UNION ALL (B LEFT EXCEPTION JOIN A)) LEFT OUTER JOIN C )  
UNION ALL  
(C LEFT EXCEPTION JOIN ((A LEFT OUTER JOIN B) UNION ALL (B LEFT EXCEPTION JOIN A)))
```

FULL OUTER implies that both sides of the join request can generate NULL values in the resulting answer set. Local selection in the WHERE clause of the query could result in the appropriate downgrade of the FULL OUTER to a LEFT OUTER or INNER JOIN.

If you want FULL OUTER JOIN behavior and local selection applied, specify the local selection in the ON clause of the FULL OUTER JOIN, or use common table expressions. For example:

```
WITH TEMPEMP AS (SELECT * FROM CORPDATA.EMPLOYEE XXX WHERE SALARY > 10000)  
SELECT EMPNO, LASTNAME, DEPTNAME  
FROM TEMPEMP XXX  
FULL OUTER JOIN CORPDATA.DEPARTMENT YYY  
ON XXX.WORKDEPT = YYY.DEPTNO
```

Join cost estimation and index selection

As the query optimizer compares the various possible access choices, it must assign a numeric cost value to each candidate. The optimizer uses that value to determine the implementation which consumes the least amount of processing time. This costing value is a combination of CPU and I/O time

In steps 3 and 5 in “Join order optimization” on page 69, the optimizer estimates cost and chooses an access method for a given dial combination. The choices made are like the choices for row selection, except that a plan using a probe must be chosen.

The costing value is based on the following assumptions:

- Table pages and index pages must be retrieved from auxiliary storage. For example, the query optimizer is not aware that an entire table might be loaded into active memory as the result of a **Set Object Access (SETOBJACC)** CL command. Use of this command could significantly improve the performance of a query. However, the optimizer does not change the query implementation to take advantage of the memory resident state of the table.
- The query is the only process running on the system. No allowance is given for system CPU utilization or I/O waits which occur because of other processes using the same resources. CPU-related costs are scaled to the relative processing speed of the system running the query.
- The values in a column are uniformly distributed across the table. For example, if 10% of the table rows have the same value, then on average, every 10th row in the table contains that value.
- The column values are independent from any other column values in a row, unless there is an index available whose key definition is (A, B). Multi-key field indexes allow the optimizer to detect when the values between columns are correlated.

For example, a column named A has a value of 1 in 50% of the rows in a table. A column named B has a value of 2 in 50% of the rows. It is expected that a query which selects rows where $A = 1$, and $B = 2$ selects 25% of the rows in the table.

The main factors in the join cost calculation for secondary dials are:

- the number of rows selected in all previous dials
- the number of rows which match, on average, each of the rows selected from previous dials.

Both of these factors can be derived by estimating the number of matching rows for a given dial.

When the join operator is something other than equal, the expected number of matching rows is based on the following default filter factors:

- 33% for less-than, greater-than, less-than-equal-to, or greater-than-equal-to
- 90% for not equal

- 25% for BETWEEN range (OPNQRYPF %RANGE)
- 10% for each IN list value (OPNQRYPF %VALUES)

For example, when the join operator is less-than, the expected number of matching rows is 0.33 * (number of rows in the dial). If no join specifications are active for the current dial, the Cartesian product is assumed to be the operator. For Cartesian products, the number of matching rows is every row in the dial, unless local row selection can be applied to the index.

When the join operator is equal, the expected number of rows is the average number of duplicate rows for a given value.

Related information

[Set Object Access \(SETOBJACC\) command](#)

Transitive closure predicates

For join queries, the query optimizer could do some special processing to generate additional selection. When the set of predicates that belong to a query logically infer extra predicates, the query optimizer generates additional predicates. The purpose is to provide more information during join optimization.

See the following examples:

```
SELECT * FROM EMPLOYEE, EMP_ACT
WHERE EMPLOYEE.EMPNO = EMP_ACT.EMPNO
AND EMPLOYEE.EMPNO = '000010'
```

The optimizer modifies the query to:

```
SELECT * FROM EMPLOYEE, EMP_ACT
WHERE EMPLOYEE.EMPNO = EMP_ACT.EMPNO
AND EMPLOYEE.EMPNO = '000010'
AND EMP_ACT.EMPNO = '000010'
```

The following rules determine which predicates are added to other join dials:

- The dials affected must have join operators of equal.
- The predicate is **isolatable**, which means that a false condition from this predicate omits the row.
- One operand of the predicate is an equal join column and the other is a constant or host variable.
- The predicate operator is not LIKE (OPNQRYPF %WLDCRD, or *CT).
- The predicate is not connected to other predicates by OR.

The query optimizer generates a new predicate, whether a predicate exists in the WHERE clause (OPNQRYPF QRYSLT parameter).

Some predicates are redundant. Redundant predicates occur when a previous evaluation of other predicates in the query already determines the result that predicate provides. Redundant predicates can be specified by you or generated by the query optimizer during predicate manipulation. Redundant predicates with operators of =, >, >=, <, <=, or BETWEEN (OPNQRYPF *EQ, *GT, *GE, *LT, *LE, or %RANGE) are merged into a single predicate to reflect the most selective range.

Look ahead predicate generation (LPG)

A special type of transitive closure called look ahead predicate generation (LPG) might be costed for joins. In this case, the optimizer tries to minimize the random I/O of a join by pre-applying the query results to a large fact table. LPG is typically used with a class of queries referred to as star join queries. However, it can possibly be used with any join query.

Look at the following query:

```
SELECT * FROM EMPLOYEE, EMP_ACT
WHERE EMPLOYEE.EMPNO = EMP_ACT.EMPNO
AND EMPLOYEE.EMPNO = '000010'
```

The optimizer could decide to internally modify the query to be:

```
WITH HT AS (SELECT *
             FROM EMPLOYEE
             WHERE EMPLOYEE.EMPNO='000010')

SELECT *
FROM HT, EMP_ACT
WHERE HT.EMPNO = EMP_ACT.EMPNO
AND EMP_ACT.EMPNO IN (SELECT DISTINCT EMPNO
                     FROM HT)
```

The optimizer places the results of the "subquery" into a temporary hash table. The hash table of the subquery can be applied in one of two methods against the EMP_ACT (fact) table:

- The distinct values of the hash tables are retrieved. For each distinct value, an index over EMP_ACT is probed to determine which records are returned for that value. Those record identifiers are normally then stored and sorted (sometimes the sorting is omitted, depending on the total number of record ids expected). Once the ids are determined, the subset of EMP_ACT records can be accessed more efficiently than in a traditional nested loop join processing.
- EMP_ACT can be scanned. For each record, the hash table is probed to see if the record joins at all to EMPLOYEE. This method allows for efficient access to EMP_ACT with a more efficient record rejection method than in a traditional nested loop join process.

Note: LPG processing is part of the normal processing in the SQL Query Engine. CQE only considers the first method, requires that the index in question be an EVI and also requires use of the STAR_JOIN and FORCE_JOIN_ORDER QAQQINI options.

Tips for improving performance when selecting data from more than two tables

The following suggestion is only applicable to CQE and is directed specifically to select-statements that access several tables. For joins that involve more than two tables, you might want to provide redundant information about the join columns. The CQE optimizer does not generate transitive closure predicates between two columns. If you give the optimizer extra information to work with when requesting a join, it can determine the best way to do the join. The additional information might seem redundant, but is helpful to the optimizer.

If the select-statement you are considering accesses two or more tables, all the recommendations suggested in [“Creating an index strategy”](#) on page 252 apply. For example, instead of coding:

```
EXEC SQL
  DECLARE EMPACTDATA CURSOR FOR
  SELECT LASTNAME, DEPTNAME, PROJNO, ACTNO
     FROM CORPDATA.DEPARTMENT, CORPDATA.EMPLOYEE,
          CORPDATA.EMP_ACT
     WHERE DEPARTMENT.MGRNO = EMPLOYEE.EMPNO
          AND EMPLOYEE.EMPNO = EMP_ACT.EMPNO
END-EXEC.
```

Provide the optimizer with a little more data and code:

```
EXEC SQL
  DECLARE EMPACTDATA CURSOR FOR
  SELECT LASTNAME, DEPTNAME, PROJNO, ACTNO
     FROM CORPDATA.DEPARTMENT, CORPDATA.EMPLOYEE,
          CORPDATA.EMP_ACT
     WHERE DEPARTMENT.MGRNO = EMPLOYEE.EMPNO
          AND EMPLOYEE.EMPNO = EMP_ACT.EMPNO
          AND DEPARTMENT.MGRNO = EMP_ACT.EMPNO
END-EXEC.
```

Multiple join types for a query

Multiple join types (inner, left outer, right outer, left exception, and right exception) can be specified in the query using the JOIN syntax. However, the Db2 for i can only support one join type of inner, left outer, or

left exception join for the entire query. The optimizer determines the overall join type for the query and reorders the files to achieve the correct semantics.

Note: This section does not apply to SQE or OPNQRYF.

The optimizer evaluates the join criteria, along with any row selection, to determine the join type for each dial and the entire query. Then the optimizer generates additional selection using the relative row number of the tables to simulate the different types of joins that occur within the query.

Null values are returned for any unmatched rows in either a left outer or an exception join. Any isolatable selection specified for that dial, including any additional join criteria specified in the WHERE clause, causes all the unmatched rows to be eliminated. (The exception is when the selection is for an IS NULL predicate.) This elimination causes the dial join type to change to an inner join (or an exception join) if the IS NULL predicate was specified.

In the following example, a left outer join is specified between the tables EMPLOYEE and DEPARTMENT. In the WHERE clause, there are two selection predicates that also apply to the DEPARTMENT table.

```
SELECT EMPNO, LASTNAME, DEPTNAME, PROJNO
FROM CORPDATA.EMPLOYEE XXX LEFT OUTER JOIN CORPDATA.DEPARTMENT YYY
    ON XXX.WORKDEPT = YYY.DEPTNO
LEFT OUTER JOIN CORPDATA.PROJECT ZZZ
    ON XXX.EMPNO = ZZZ.RESPEMP
WHERE XXX.EMPNO = YYY.MGRNO AND
      YYY.DEPTNO IN ('A00', 'D01', 'D11', 'D21', 'E11')
```

The first selection predicate, XXX.EMPNO = YYY.MGRNO, is an additional join condition that is evaluated as an "inner join" condition. The second is an isolatable selection predicate that eliminates any unmatched rows. Either of these predicates can cause the join type for the DEPARTMENT table to change from a left outer join to an inner join.

Even though the join between the EMPLOYEE and DEPARTMENT tables was changed to an inner join, the entire query remains a left outer join to satisfy the join condition for the PROJECT table.

Note: Care must be taken when specifying multiple join types since they are supported by appending selection to the query for any unmatched rows. The number of rows satisfying the join criteria can become large before selection that either selects or omits the unmatched rows based on that individual dial join type is applied.

Sources of join query performance problems

The optimization algorithms described earlier benefit most join queries, but the performance of a few queries might be degraded.

This occurs when:

- An index is not available which provides average number of duplicate values statistics for the potential join columns.
- The optimizer uses default filter factors to estimate the number of rows when applying local selection to the table when indexes or column statistics do not exist over the selection columns.

Creating indexes over the selection columns allows the optimizer to make a more accurate filtering estimate by using key range estimates.

- The particular values selected for the join columns yield a greater number of matching rows than the average number of duplicate values for all values of the join columns in the table. For example, the data is not uniformly distributed.

Join performance tips

If you have a join query performing poorly, or you are creating an application which uses join queries, these tips could be useful.

<i>Table 31. Checklist for Creating an Application that Uses Join Queries</i>	
What to Do	How It Helps
<p>Check the database design. Make sure that there are indexes available over all the join columns and row selection columns or both. The optimizer provides index advice in several places to aid in this process:</p> <ul style="list-style-type: none"> • the index advisor under System i Navigator - Database • the advised information under Visual Explain • the advised information in the 3020 record in the database monitor 	<p>The query optimizer can select an efficient access method because it can determine the average number of duplicate values. Many queries could use the existing index and avoid the cost of creating a temporary index or hash table.</p>
<p>Check the query to see whether some complex predicates could be added to other dials to allow the optimizer to get better selectivity for each dial.</p>	<p>The query optimizer does not add predicates for predicates connected by OR or non-isolatable predicates, or predicate operator LIKE. Modify the query by adding additional predicates to help.</p>
<p>Specify ALWCPYDTA(*OPTIMIZE) or ALWCPYDTA(*YES)</p>	<p>The query is creating a temporary index or hash table, and the processing time could be better if the existing index or hash table was used. Specify ALWCPYDTA(*YES).</p> <p>The query is not creating a temporary index or hash table, and the processing time could be better if a temporary index was created. Specify ALWCPYDTA(*OPTIMIZE).</p> <p>Alternatively, specify OPTIMIZE FOR n ROWS to inform the optimizer that the application reads every resulting row. Set n to a large number. You can also set n to a small number before ending the query.</p>
<p>For OPNQRYF, specify OPTIMIZE(*FIRSTIO) or OPTIMIZE(*ALLIO)</p>	<p>Specify the OPTIMIZE(*FIRSTIO) or OPTIMIZE(*ALLIO) option to accurately reflect your application. Use *FIRSTIO, if you want the optimizer to optimize the query to retrieve the first block of rows most efficiently. This biases the optimizer toward using existing objects. If you want to optimize the retrieval time for the entire answer set, use *ALLIO. This option could cause the optimizer to create temporary indexes or hash tables to minimize I/O.</p>

Table 31. Checklist for Creating an Application that Uses Join Queries (continued)

What to Do	How It Helps
Star join queries	<p>A join in which one table is joined with all secondary tables consecutively is sometimes called a star join. If all secondary join predicates contain a column reference to a particular table, place that table in join position one. In Example A, all tables are joined to table EMPLOYEE. The query optimizer can freely determine the join order. For SQE, the optimizer uses Look Ahead Predicate generation to determine the optimal join order. For CQE, the query could be changed to force EMPLOYEE into join position one by using the query options file (QAQQINI) FORCE_JOIN_ORDER parameter of *YES. In these examples, the join type is a join with no default values returned (an inner join.). The reason for forcing the table into the first position is to avoid random I/O processing. If EMPLOYEE is not in join position one, every row in EMPLOYEE can be examined repeatedly during the join process. If EMPLOYEE is fairly large, considerable random I/O processing occurs resulting in poor performance. By forcing EMPLOYEE to the first position, random I/O processing is minimized.</p> <p>Example A: Star join query</p> <pre> DECLARE C1 CURSOR FOR SELECT * FROM DEPARTMENT, EMP_ACT, EMPLOYEE, PROJECT WHERE DEPARTMENT.DEPTNO=EMPLOYEE.WORKDEPT AND EMP_ACT.EMPNO=EMPLOYEE.EMPNO AND EMPLOYEE.WORKDEPT=PROJECT.DEPTNO </pre> <p>Example B: Star join query with order forced using FORCE_JOIN_ORDER</p> <pre> DECLARE C1 CURSOR FOR SELECT * FROM EMPLOYEE, DEPARTMENT, EMP_ACT, PROJECT WHERE DEPARTMENT.DEPTNO=EMPLOYEE.WORKDEPT AND EMP_ACT.EMPNO=EMPLOYEE.EMPNO AND EMPLOYEE.WORKDEPT=PROJECT.DEPTNO </pre>
Specify ALWCOPYDTA(*OPTIMIZE) to allow the query optimizer to use a sort routine.	Ordering is specified and all key columns are from a single dial. The optimizer can consider all possible join orders with ALWCOPYDTA(*OPTIMIZE).
Specify join predicates to prevent all the rows from one table from being joined to every row in the other table.	Improves performance by reducing the join fan-out. It is best if every secondary table has at least one join predicate that references one of its columns as a 'join-to' column.

Distinct optimization

Distinct is used to compare a value with another value.

There are two methods to write a query that returns distinct values in SQL. One method uses the DISTINCT keyword:

```

SELECT DISTINCT COL1, COL2
FROM TABLE1

```

The second method uses GROUP BY:

```

SELECT COL1, COL2
FROM TABLE1
GROUP BY COL1, COL2

```

All queries that contain a DISTINCT, and are run using SQE, rewritten into queries using GROUP BY. This rewrite enables queries using DISTINCT to take advantage of the many grouping techniques available to the optimizer.

Distinct to Grouping implementation

The following example query has a DISTINCT:

```
SELECT DISTINCT COL1, COL2
FROM T1
WHERE COL2 > 5 AND COL3 = 2
```

The optimizer rewrites it into this query:

```
SELECT COL1, COL2
FROM T1
WHERE COL2 > 5 AND COL3 = 2
GROUP BY COL1, COL2
```

Distinct removal

A query containing a DISTINCT over whole-file aggregation (no grouping or selection) allows the DISTINCT to be removed. For example, look at this query with DISTINCT:

```
SELECT DISTINCT COUNT(C1), SUM(C1)
FROM TABLE1
```

The optimizer rewrites this query as the following:

```
SELECT COUNT(C1), SUM(C1)
FROM TABLE1
```

If the DISTINCT and the GROUP BY fields are identical, the DISTINCT can be removed. If the DISTINCT fields are a subset of the GROUP BY fields (and there are no aggregates), the DISTINCTs can be removed.

Grouping optimization

Db2 for i has certain techniques to use when the optimizer encounters grouping. The query optimizer chooses its methods for optimizing your query.

Hash grouping implementation

This technique uses the base hash access method to perform grouping or summarization of the selected table rows. For each selected row, the specified grouping value is run through the hash function. The computed hash value and grouping value are used to quickly find the entry in the hash table corresponding to the grouping value.

If the current grouping value already has a row in the hash table, the hash table entry is retrieved and summarized (updated) with the current table row values based on the requested grouping column operations (such as SUM or COUNT). If a hash table entry is not found for the current grouping value, a new entry is inserted into the hash table and initialized with the current grouping value.

The time required to receive the first group result for this implementation is most likely longer than other grouping implementations because the hash table must be built and populated first. Once the hash table is populated, the database manager uses the table to start returning the grouping results. Before returning any results, the database manager must apply any specified grouping selection criteria or ordering to the summary entries in the hash table.

Where the hash grouping method is most effective

The hash grouping method is most effective when the consolidation ratio is high. The **consolidation ratio** is the ratio of the selected table rows to the computed grouping results. If every database table row has

its own unique grouping value, then the hash table becomes too large. The size in turn slows down the hashing access method.

The optimizer estimates the consolidation ratio by first determining the number of unique values in the specified grouping columns (that is, the expected number of groups in the database table). The optimizer then examines the total number of rows in the table and the specified selection criteria and uses the result of this examination to estimate the consolidation ratio.

Indexes over the grouping columns can help make the ratio estimate of the optimizer more accurate. Indexes improve the accuracy because they contain statistics that include the average number of duplicate values for the key columns.

The optimizer also uses the expected number of groups estimate to compute the number of partitions in the hash table. As mentioned earlier, the hashing access method is more effective when the hash table is well-balanced. The number of hash table partitions directly affects how entries are distributed across the hash table and the uniformity of this distribution.

The hash function performs better when the grouping values consist of columns that have non-numeric data types, except for the integer (binary) data type. In addition, specifying grouping value columns that are not associated with the variable length and null column attributes allows the hash function to perform more effectively.

Index grouping implementation

There are two primary ways to implement grouping using an index: Ordered grouping and pre-summarized processing.

Ordered grouping

This implementation uses the Radix Index Scan or the Radix Index Probe access methods to perform the grouping. An index is required that contains all the grouping columns as contiguous leftmost key columns. The database manager accesses the individual groups through the index and performs the requested summary functions.

Since the index, by definition, already has all the key values grouped, the first group result can be returned in less time than the hashing method. This index performance is faster because the hashing method requires a temporary result. This implementation can be beneficial if an application does not need to retrieve all the group results, or if an index exists that matches the grouping columns.

When the grouping is implemented with an index and a permanent index does not exist that satisfies grouping columns, a temporary index is created. The grouping columns specified within the query are used as the key columns for this index.

Pre-summarized processing

This SQE-only implementation uses an Encoded Vector Index to extract the summary information already in the symbol table of the index. The EVI symbol table contains the unique key values and a count of the number of table records that have that unique value. The grouping for the columns of the index key is already performed. If the query references a single table and performs simple aggregation, the EVI might be used for quick access to the grouping results. For example, consider the following query:

```
SELECT COUNT(*), col1
FROM t1
GROUP BY col1
```

If an EVI exists over t1 with a key of col1, the optimizer can rewrite the query to access the precomputed grouping answer in the EVI symbol table.

This rewrite can result in dramatic improvements when the number of table records is large and the number of resulting groups is small, relative to the size of the table.

This method is also possible with selection (WHERE clause), as long as the reference columns are in the key definition of the EVI.

For example, consider the following query:

```
SELECT COUNT(*), col1
FROM t1
WHERE col1 > 100
GROUP BY col1
```

This query can be rewritten by the optimizer to use the EVI. This pre-summarized processing works for DISTINCT processing, GROUP BY and for column function COUNT. All columns of the table referenced in the query must also be in the key definition of the EVI.

So, for example, the following query can be made to use the EVI:

```
SELECT DISTINCT col1
FROM t1
```

However, this query cannot:

```
SELECT DISTINCT col1
FROM t1
WHERE col2 > 1
```

This query cannot use the EVI because it references col2 of the table, which is not in the key definition of the EVI. If multiple columns are defined in the EVI key, for example, col1 and col2, it is important to use the left-most columns of the key. For example, if an EVI existed with a key definition of (col1, col2), but the query referenced only col2, it is unlikely the EVI is used.

EVI INCLUDE aggregates

A more powerful example of pre-summarized processing can be facilitated by the use of the INCLUDE keyword on the index create. By default, COUNT(*) is implied on the creation of an EVI. Additional numeric aggregates specified over non-key data can further facilitate pre-determined or ready-made aggregate results during query optimization.

For example, suppose the following query is a frequently requested result set, queried in whole or as part of a subquery comparison.

```
SELECT AVG(col2)
FROM t1
GROUP BY col1
```

Create the following EVI to predetermine the value of AVG(col2).

```
CREATE ENCODED VECTOR INDEX eviT1 ON t1(col1) INCLUDE(AVG(col2))
```

eviT1 delivers distinct values for col1 and COUNT(*) specific to the group by of col1. eviT1 can be used to generate an asynchronous bitmap or RRN list for accessing the table rows for specific col1 values. In addition, eviT1 computes an additional aggregate, AVG(col2), over the same group by column (col1) by specifying the INCLUDE aggregate.

INCLUDE aggregates are limited to those aggregates that result in numeric values: SUM, COUNT, AVG, STDDEV, and so on. These values can be readily maintained as records are inserted, deleted, or updated in the base table.

MIN or MAX are two aggregates that are not supported as INCLUDE aggregates. Deleting the current row contributing to the MIN or MAX value would result in the need to recalculate, potentially accessing many rows, and reducing performance.

INCLUDE values can also contain aggregates over derivations. For example, if you have a couple of columns that contribute to an aggregate, that derivation can be specified, for example, as SUM(col1+col2+col3).

It is recommended that EVIs with INCLUDE aggregates only contain references to columns or column-specific derivations, for example, SUM(salary+bonus).

In many environments, queries that contain derivations using constants convert those constants to parameter markers. This conversion allows a much higher degree of ODP reuse. However, it can be more difficult to match the parameter value to a literal in the index definition.

The optimizer does attempt to match constants in the EVI with parameter markers or host variable values in the query. However, in some complex cases this support is limited and could result in the EVI not matching the query.

Pre-summarized processing can also take advantage of EVIs with INCLUDE in a JOIN situation.

For example, see the following aggregate query over the join of two tables.

EVI INCLUDE aggregate example

```
SELECT deptname, sum(salary)
FROM DEPARTMENT, EMPLOYEE
WHERE deptno=workdept
GROUP BY deptname
```

By providing an EVI with INCLUDE index, as follows, and with optimizer support to push down aggregates to the table level when possible, the resulting implementation takes advantage of the ready-made aggregates already supplied by EVI employeeSumByDept. The implementation never needs to touch or aggregate rows in the Employee table.

```
CREATE ENCODED VECTOR INDEX employeeSumByDept ON employee(workdept)
INCLUDE(sum(salary))
```

Aggregate pushdown results in a rewrite with EVI INCLUDE implementation, conceptually like the following query.

```
SELECT deptname, sum(sum(salary))
FROM department,
  (SELECT workdept, sum(salary) FROM employee group by workdept) employee_2
WHERE deptno=workdept
```

Instead of department joining to all the rows in the employee table, it now has the opportunity to join to the predetermined aggregates, the sum of salary by department number, in the EVI symbol table. This results in significant reduction in processing and IO.

Related concepts

[How the EVI works](#)

EVIIs work in different ways for costing and implementation.

Related reference

[Encoded vector index symbol table scan](#)

An encoded vector index symbol table scan operation is used to retrieve the entries from the symbol table portion of the index.

Related information

[SQL INCLUDE statement](#)

Optimizing grouping by eliminating grouping columns

All the grouping columns are evaluated to determine if they can be removed from the list of grouping columns. Only those grouping columns that have isolatable selection predicates with an equal operator specified can be considered. This guarantees that the column can only match a single value and does not help determine a unique group.

This processing allows the optimizer to consider more indexes to implement the query. It also reduces the number of columns that are added as key columns to a temporary index or hash table.

The following example illustrates a query where the optimizer might eliminate a grouping column.

```
DECLARE DEPTEMP CURSOR FOR
SELECT EMPNO, LASTNAME, WORKDEPT
```

```
FROM CORPDATA.EMPLOYEE
WHERE EMPNO = '000190'
GROUP BY EMPNO, LASTNAME, WORKDEPT
```

OPNQRYF example:

```
OPNQRYF FILE(EMPLOYEE) FORMAT(FORMAT1)
QRYSLT('EMPNO *EQ '000190''')
GRPFLLD(EMPNO LASTNAME WORKDEPT)
```

In this example, the optimizer can remove EMPNO from the list of grouping columns because of the EMPNO = '000190' selection predicate. An index that only has LASTNAME and WORKDEPT specified as key columns could implement the query. If a temporary index or hash is required then EMPNO is not used.

Note: Even though EMPNO can be removed from the list of grouping columns, the optimizer might use a permanent index that exists with all three grouping columns.

Optimizing grouping by adding additional grouping columns

The same logic that is applied to removing grouping columns can also be used to add additional grouping columns to the query. Additional grouping columns are added only when you are trying to determine if an index can be used to implement the grouping.

The following example illustrates a query where the optimizer might add an additional grouping column.

```
CREATE INDEX X1 ON EMPLOYEE
(LASTNAME, EMPNO, WORKDEPT)

DECLARE DEPTEMP CURSOR FOR
SELECT LASTNAME, WORKDEPT
FROM CORPDATA.EMPLOYEE
WHERE EMPNO = '000190'
GROUP BY LASTNAME, WORKDEPT
```

For this query request, the optimizer can add EMPNO as an additional grouping column when considering X1 for the query.

Optimizing grouping by using index skip key processing

Index Skip Key processing can be used when grouping with the keyed sequence implementation algorithm which uses an existing index. It is a specialized version of ordered grouping that processes few records in each group rather than all records in each group.

The index skip key processing algorithm:

1. Uses the index to position to a group and
2. finds the first row matching the selection criteria for the group, and if specified the first non-null MIN or MAX value in the group
3. Returns the group to the user
4. "Skip" to the next group and repeat processing

This algorithm improves performance by potentially not processing all index key values for a group.

Index skip key processing can be used:

- For single table queries using the keyed sequence grouping implementation when:
 - There are no column functions in the query, or
 - There is only a single MIN or MAX column function and the MIN or MAX operand is the next index key column after the grouping columns. There can be no other grouping functions in the query. For the MIN function, the key column must be an ascending key; for the MAX function, the key column must be a descending key. If the query is whole table grouping, the operand of the MIN or MAX must be the first key column.

Example 1, using SQL:

```
CREATE INDEX IX1 ON EMPLOYEE (SALARY DESC)

DECLARE C1 CURSOR FOR
  SELECT MAX(SALARY) FROM EMPLOYEE;
```

The query optimizer chooses to use the index IX1. The SLIC runtime code scans the index until it finds the first non-null value for SALARY. Assuming that SALARY is not null, the runtime code positions to the first index key and return that key value as the MAX of salary. No more index keys are processed.

Example 2, using SQL:

```
CREATE INDEX IX2 ON EMPLOYEE (WORKDEPT, JOB, SALARY)

DECLARE C1 CURSOR FOR
  SELECT WORKDEPT, MIN(SALARY)
  FROM EMPLOYEE
  WHERE JOB='CLERK'
  GROUP BY WORKDEPT
```

The query optimizer chooses to use Index IX2. The database manager positions to the first group for DEPT where JOB equals 'CLERK' and returns the SALARY. The code then skips to the next DEPT group where JOB equals 'CLERK'.

- For join queries:
 - All grouping columns must be from a single table.
 - For each dial, there can be at most one MIN or MAX column function operand that references the dial. No other column functions can exist in the query.
 - If the MIN or MAX function operand is from the same dial as the grouping columns, then it uses the same rules as single table queries.
 - If the MIN or MAX function operand is from a different dial, then the join column for that dial must join to one of the grouping columns. The index for that dial must contain the join columns followed by the MIN or MAX operand.

Example 1, using SQL:

```
CREATE INDEX IX1 ON DEPARTMENT (DEPTNAME)

CREATE INDEX IX2 ON EMPLOYEE (WORKDEPT, SALARY)

DECLARE C1 CURSOR FOR
  SELECT DEPARTMENT.DEPTNO, MIN(SALARY)
  FROM DEPARTMENT, EMPLOYEE
  WHERE DEPARTMENT.DEPTNO=EMPLOYEE.WORKDEPT
  GROUP BY DEPARTMENT.DEPTNO;
```

Optimizing grouping by removing read triggers

For queries involving physical files or tables with read triggers, group by triggers always involve a temporary file before the group by processing. Therefore, these queries slow down.

Note: Read triggers are added when the **Add Physical File Trigger (ADDPFTRG)** command has been used on the table with TRGTIME (*AFTER) and TRGEVENT (*READ).

The query runs faster if the read trigger is removed (RMVPFTRG TRGTIME (*AFTER) TRGEVENT (*READ)).

Related information

[Add Physical File Trigger \(ADDPFTRG\) command](#)

Grouping set optimization

The optimizer uses all the previously mentioned grouping optimizations for individual grouping sets specified in the query.

If multiple temporary result sets are needed to implement all the grouping sets, they can all be populated using one pass through the data. This one-pass population occurs even if different types of temporary result sets are used to implement various grouping sets.

A temporary result type called sorted distinct list is used specifically for ROLLUP implementations. This temporary result set is used to compute the aggregate rows: the grouping set that includes all expressions listed in the ROLLUP clause. Hash grouping is used internally to quickly find the current grouping value. The entries in the temporary result sets are also sorted. This sorting allows the aggregate results to be used to compute the super-aggregate rows in the rollup result set without creating additional temporary result sets.

ROLLUPS can also be implemented using a radix index over the columns in the rollup without creating a temporary result set.

The optimizer can compute all the grouping sets in a given ROLLUP using at most one temporary result set. Therefore, it is advantageous for the optimizer to look for the rollup pattern in grouping set queries.

The optimizer tries to find the ROLLUP pattern in a list of individual grouping sets. For example, the following GROUP BY clause:

```
GROUP BY GROUPING SETS ((A, B, C), (B, D), (A, B), (A), ())
```

is rewritten to:

```
GROUP BY GROUPING SETS ((ROLLUP(A, B, C)), (B, D))
```

This rewrite allows the query to be implemented using at most two temporary results sets rather than 4.

Queries containing a CUBE clause is broken down into a union of ROLLUPS and grouping sets. For example:

```
CUBE(A, B, C)
```

is equivalent to:

```
(ROLLUP(A, B, C)), (ROLLUP'(B, C)), (ROLLUP'(C, A))
```

The ROLLUP' notation is an internal representation of a ROLLUP operation that does not include a grand total row in its result set. So, ROLLUP'(B, C) is equivalent to GROUP BY GROUPING SETS ((B,C), (B)). This CUBE rewrite implements at most three temporary result sets, rather than the 8 that might be needed had the query not been rewritten.

Ordering optimization

This section describes how Db2 for i implements ordering techniques, and how optimization choices are made by the query optimizer. The query optimizer can use either index ordering or a sort to implement ordering.

Sort Ordering implementation

The sort algorithm reads the rows into a sort space and sorts the rows based on the specified ordering keys. The rows are then returned to the user from the ordered sort space.

Index Ordering implementation

The index ordering implementation requires an index that contains all the ordering columns as contiguous leftmost key columns. The database manager accesses the individual rows through the index in index order, which results in the rows being returned in order to the requester.

This implementation can be beneficial if an application does not need to retrieve all the ordered results, or if an index exists that matches the ordering columns. When the ordering is implemented with an index, and a permanent index does not exist that satisfies ordering columns, a temporary index is created. The ordering columns specified within the query are used as the key columns for this index.

Index Merge Ordering Implementation

Index Merge Ordering takes some of the features of both index and sorted ordering by using an index to apply the selection and then using an indexed list to sort the rows. This optimization is primarily aimed at queries that have an optimization goal of *FIRSTIO and the WHERE clause selection contains non-equal predicates.

Optimizing ordering by eliminating ordering columns

All the ordering columns are evaluated to determine if they can be removed from the list of ordering columns. Only those ordering columns that have isolatable selection predicates with an equal operator specified can be considered. This guarantees that the column can match only a single value, and does not help determine in the order.

In general, for an index to be eligible for Index Merge Ordering, the index should be created as follows:

- Any columns with equal predicates that are not in the ORDER BY list should be first keys in the index.
- Followed by any columns for non-equal predicates that are not in the ORDER BY list. The columns should be arranged so the most selective predicate column is the first column following the equal predicate columns.
- Followed by the ORDER BY columns in the same order as the ORDER BY clause of the query and with the same ASC or DESC attribute.
- Optionally, at the end you may include other selection columns or other selected columns for Index Only Access.

For example:

```
SELECT * FROM SALES
WHERE SALES_DATE BETWEEN '1996-03-29' AND '1996-04-29'
AND REGION IN ('Quebec','Manitoba')
ORDER BY SALES_PERSON
```

The Index to create would be with keys:

```
SALES_DATE, REGION, SALES_PERSON
```

In order to allow a more general indexing scheme, it is possible to use an index for index merge ordering even if there are leading keys not used in the query. The above index could also be used for the following query:

```
SELECT SALES_DATE, REGION, SALES_PERSON, SALES
FROM SALES
WHERE REGION IN ('Quebec','Manitoba')
ORDER BY SALES_PERSON
```

If SALES was added as a trailing key, then Index Only Access would be used for both queries.

```
CREATE INDEX SALES_IMOX ON SALES (SALES_DATE, REGION, SALES_PERSON, SALES)
```

The optimizer can now consider more indexes as it implements the query. The number of columns that are added as key columns to a temporary index is also reduced. The following SQL example illustrates a query where the optimizer might eliminate an ordering column.

```
DECLARE DEPTEMP CURSOR FOR
SELECT EMPNO, LASTNAME, WORKDEPT
FROM CORPDATA.EMPLOYEE
WHERE EMPNO = '000190'
ORDER BY EMPNO, LASTNAME, WORKDEPT
```

Optimizing ordering by adding additional ordering columns

The same logic that is applied to removing ordering columns can also be used to add additional grouping columns to the query. This logic is done only when you are trying to determine if an index can be used to implement the ordering.

The following example illustrates a query where the optimizer might add an additional ordering column.

```
CREATE INDEX X1 ON EMPLOYEE (LASTNAME, EMPNO, WORKDEPT)

DECLARE DEPTEMP CURSOR FOR
SELECT LASTNAME, WORKDEPT
FROM CORPDATA.EMPLOYEE
WHERE EMPNO = '000190'
ORDER BY LASTNAME, WORKDEPT
```

For this query request, the optimizer can add EMPNO as an additional ordering column when considering X1 for the query.

Index Ordering Implementation Using Reverse Ordered Indexes

It is also possible to use reverse ordered indexes to provide ordering. In that case, the index keys for all ORDER BY columns must be reversed. The database manager will process the index in reverse order by starting at the end and reading through the index backwards.

The following index and query illustrates an example where the optimizer could use a reversed ordered index.

```
CREATE INDEX CORPDATA.INDEX1 ON CORPDATA.EMPLOYEE (SALARY ASC, LASTNAME DESC)

SELECT EMPNO, LASTNAME, WORKDEPT, SALARY
FROM CORPDATA.EMPLOYEE
ORDER BY SALARY DESC, LASTNAME ASC
```

View implementation

Views, derived tables (nested table expressions or NTEs), and common table expressions (CTEs) are implemented by the query optimizer using one of two methods.

These methods are:

- The optimizer combines the query select statement with the select statement of the view.
- The optimizer places the results of the view in a temporary table and then replaces the view reference in the query with the temporary table.

View composite implementation

The view composite implementation takes the query select statement and combines it with the select statement of the view to generate a new query. The new, combined select statement query is then run directly against the underlying base tables.

This single, composite statement is the preferred implementation for queries containing views, since it requires only a single pass of the data.

See the following examples:

```
CREATE VIEW D21EMPL AS
SELECT * FROM CORPDATA.EMPLOYEE
WHERE WORKDEPT='D21'
```

Using SQL:

```
SELECT LASTNAME, FIRSTNME, SALARY
FROM D21EMPL
WHERE JOB='CLERK'
```

The query optimizer generates a new query that looks like the following example:

```
SELECT LASTNAME, FIRSTNME, SALARY
FROM CORPDATA.EMPLOYEE
WHERE WORKDEPT='D21' AND JOB='CLERK'
```

The query contains the columns selected by the user query, the base tables referenced in the query, and the selection from both the view and the user query.

Note: The new composite query that the query optimizer generates is not visible to users. Only the original query against the view is seen by users and database performance tools.

View materialization implementation

The view materialization implementation runs the query of the view and places the results in a temporary result. The view reference in the user query is then replaced with the temporary, and the query is run against the temporary result.

View materialization is done whenever it is not possible to create a view composite. For SQE, view materialization is optional. The following types of queries require view materialization:

- The outermost view select contains grouping, the query contains grouping, and refers to a column derived from a column function in the view HAVING or select-list.
- The query is a join and the outermost select of the view contains grouping or DISTINCT.
- The outermost select of the view contains DISTINCT, and the query has UNION, grouping, or DISTINCT and one of the following:
 - Only the query has a shared weight NLSS table
 - Only the view has a shared weight NLSS table
 - Both the query and the view have a shared weight NLSS table, but the tables are different.
- The query contains a column function and the outermost select of the view contains a DISTINCT
- The view does not contain an access plan. Occurs when a view references a view, and a view composite cannot be created because of one of the previous listed reasons. Does not apply to nested table expressions and common table expressions.
- The Common table expression (CTE) is referenced more than once in the query FROM clause. Also, the CTE SELECT clause references a MODIFIES or EXTERNAL ACTION UDF.

When a temporary result table is created, access methods that are allowed with ALWCOPYDTA(*OPTIMIZE) could be used to implement the query. These methods include hash grouping, hash join, and bitmaps.

See the following examples:

```
CREATE VIEW AVGSALVW AS
SELECT WORKDEPT, AVG(SALARY) AS AVGSAL
FROM CORPDATA.EMPLOYEE
GROUP BY WORKDEPT
```

SQL example:

```
SELECT D.DEPTNAME, A.AVGSAL
FROM CORPDATA.DEPARTMENT D, AVGSALVW A
WHERE D.DEPTNO=A.WORKDEPT
```


In this case, a view composite cannot be created since a join query references a grouping view. The results of AVGSALVW are placed in a temporary result table (*QUERY0001). The view reference AVGSALVW is replaced with the temporary result table. The new query is then run. The generated query looks like the following:

```
SELECT D.DEPTNAME, A.AVGSAL
FROM CORPDATA.DEPARTMENT D, *QUERY0001 A
WHERE D.DEPTNO=A.WORKDEPT
```

Note: The new query that the query optimizer generates is not visible to users. Only the original query against the view is seen by users and database performance tools.

Whenever possible, isolatable selection from the query, except subquery predicates, is added to the view materialization process. This results in smaller temporary result tables and allows existing indexes to be used when materializing the view. This process is not done if there is more than one reference to the same view or common table expression in the query. The following is an example where isolatable selection is added to the view materialization:

```
SELECT D.DEPTNAME, A.AVGSAL
FROM CORPDATA.DEPARTMENT D, AVGSALVW A
WHERE D.DEPTNO=A.WORKDEPT AND
A.WORKDEPT LIKE 'D%' AND AVGSAL>10000
```

The isolatable selection from the query is added to the view resulting in a new query to generate the temporary result table:

```
SELECT WORKDEPT, AVG(SALARY) AS AVGSAL
FROM CORPDATA.EMPLOYEE
WHERE WORKDEPT LIKE 'D%'
GROUP BY WORKDEPT
HAVING AVG(SALARY)>10000
```

Materialized query table optimization

Materialized query tables (MQTs) (also referred to as automatic summary tables or materialized views) can provide performance enhancements for queries.

This performance enhancement is done by precomputing and storing results of a query in the materialized query table. The database engine can use these results instead of recomputing them for a user specified query. The query optimizer looks for any applicable MQTs. The optimizer can implement the query using a given MQT, provided it is a faster implementation choice.

Materialized Query Tables are created using the SQL CREATE TABLE statement. Alternatively, the ALTER TABLE statement could be used to convert an existing table into a materialized query table. The REFRESH TABLE statement is used to recompute the results stored in the MQT. For user-maintained MQTs, the MQTs could also be maintained by the user using INSERT, UPDATE, and DELETE statements.

Related information

[Create Table statement](#)

MQT supported function

Although an MQT can contain almost any query, the optimizer only supports a limited set of query functions when matching MQTs to user specified queries. The user-specified query and the MQT query must both be supported by the SQE optimizer.

The supported function in the MQT query by the MQT matching algorithm includes:

- Single table and join queries
- WHERE clause
- GROUP BY and optional HAVING clauses
- ORDER BY

- FETCH FIRST n ROWS
- Views, common table expressions, and nested table expressions
- UNIONS
- Partitioned tables

There is limited support in the MQT matching algorithm for the following:

- Scalar subselects
- User Defined Functions (UDFs) and user-defined table functions
- Recursive Common Table Expressions (RCTE)
- The following scalar functions:
 - ATAN2
 - DAYNAME
 - DBPARTITIONNAME
 - DECRYPT_BIT
 - DECRYPT_BINARY
 - DECRYPT_CHAR
 - DECRYPT_DB
 - DIFFERENCE
 - DLVALUE
 - DLURLPATH
 - DLURLPATHONLY
 - DLURLSEVER
 - DLURLSCHEME
 - DLURLCOMPLETE
 - ENCRYPT_AES
 - ENCRYPT_RC2
 - ENCRYPT_TDES
 - GENERATE_UNIQUE
 - GETHINT
 - IDENTITY_VAL_LOCAL
 - INSERT
 - MONTHNAME
 - MONTHS_BETWEEN
 - NEXT_DAY
 - RAND
 - RAISE_ERROR
 - REPEAT
 - REPLACE
 - ROUND_TIMESTAMP
 - SOUNDEX
 - TIMESTAMP_FORMAT
 - TIMESTAMPDIFF
 - TRUNC_TIMESTAMP
 - VARCHAR_FORMAT

- WEEK_ISO

It is recommended that the MQT only contain references to columns and column functions. In many environments, queries that contain constants have the constants converted to parameter markers. This conversion allows a much higher degree of ODP reuse. The MQT matching algorithm attempts to match constants in the MQT with parameter markers or host variable values in the query. However, in some complex cases this support is limited and could result in the MQT not matching the query.

Related concepts

Query dispatcher

The function of the dispatcher is to route the query request to either CQE or SQE, depending on the attributes of the query. All queries are processed by the dispatcher. It cannot be bypassed.

Related reference

Details on the MQT matching algorithm

What follows is a generalized discussion of how the MQT matching algorithm works.

Using MQTs during query optimization

Before using MQTs, you need to consider your environment attributes.

To even consider using MQTs during optimization the following environmental attributes must be true:

- The query must specify ALWCPYDTA(*OPTIMIZE) or INSENSITIVE cursor.
- The query must not be a SENSITIVE cursor.
- The table to be replaced with an MQT must not be update or delete capable for this query.
- The MQT currently has the ENABLE QUERY OPTIMIZATION attribute active
- The MATERIALIZED_QUERY_TABLE_USAGE QAQQINI option must be set to *ALL or *USER to enable use of MQTs. The default setting of MATERIALIZED_QUERY_TABLE_USAGE does not allow usage of MQTs.
- The timestamp of the last REFRESH TABLE for an MQT is within the duration specified by the MATERIALIZED_QUERY_TABLE_REFRESH_AGE QAQQINI option. Or *ANY is specified, which allows MQTs to be considered regardless of the last REFRESH TABLE. The default setting of MATERIALIZED_QUERY_TABLE_REFRESH_AGE does not allow usage of MQTs.
- The query must be run through SQE.
- The following QAQQINI options must match: IGNORE_LIKE_REDUNDANT_SHIFTS, NORMALIZE_DATA, and VARIABLE_LENGTH_OPTIMIZATION. These options are stored at CREATE materialized query table time and must match the options specified at query run time.
- The commit level of the MQT must be greater than or equal to the query commit level. The commit level of the MQT is either specified in the MQT query using the WITH clause. Or it is defaulted to the commit level that the MQT was run under when it was created.
- The field procedure encoded comparison (QAQQINI FIELDPROC_ENCODED_COMPARISON option) level of the MQT must be greater than or equal to the query specified field procedure encoded comparison level.

MQT examples

The following are examples of using MQTs.

Example 1

The first example is a query that returns information about employees whose job is DESIGNER. The original query:

```
SELECT D.deptname, D.location, E.firstnme, E.lastname, E.salary+E.comm+E.bonus as total_sal
FROM Department D, Employee E
WHERE D.deptno=E.workdept
AND E.job = 'DESIGNER'
```

Create a table, MQT1, that uses this query:

```

CREATE TABLE MQT1
  AS (SELECT D.deptname, D.location, E.firstnme, E.lastname, E.salary, E.comm, E.bonus,
E.job
FROM Department D, Employee E
WHERE D.deptno=E.workdept)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER

```

Resulting new query after replacing the specified tables with the MQT.

```

SELECT M.deptname, M.location, M.firstnme, M.lastname, M.salary+M.comm+M.bonus as total_sal
FROM MQT1 M
WHERE M.job = 'DESIGNER'

```

In this query, the MQT matches part of the user query. The MQT is placed in the FROM clause and replaces tables DEPARTMENT and EMPLOYEE. Any remaining selection not done by the MQT query (M.job= 'DESIGNER') is done to remove the extra rows. The result expression, M.salary+M.comm+M.bonus, is calculated. JOB must be in the select-list of the MQT so that the additional selection can be performed.

Visual Explain diagram of the query when using the MQT:

Example 2

Get the total salary for all departments that are located in 'NY'. The original query:

```

SELECT D.deptname, sum(E.salary)
FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept AND D.location = 'NY'
GROUP BY D.deptname

```

Create a table, MQT2, that uses this query:

```

CREATE TABLE MQT2
  AS (SELECT D.deptname, D.location, sum(E.salary) as sum_sal

```

```

FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept
GROUP BY D.Deptname, D.location)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER

```

Resulting new query after replacing the specified tables with the MQT:

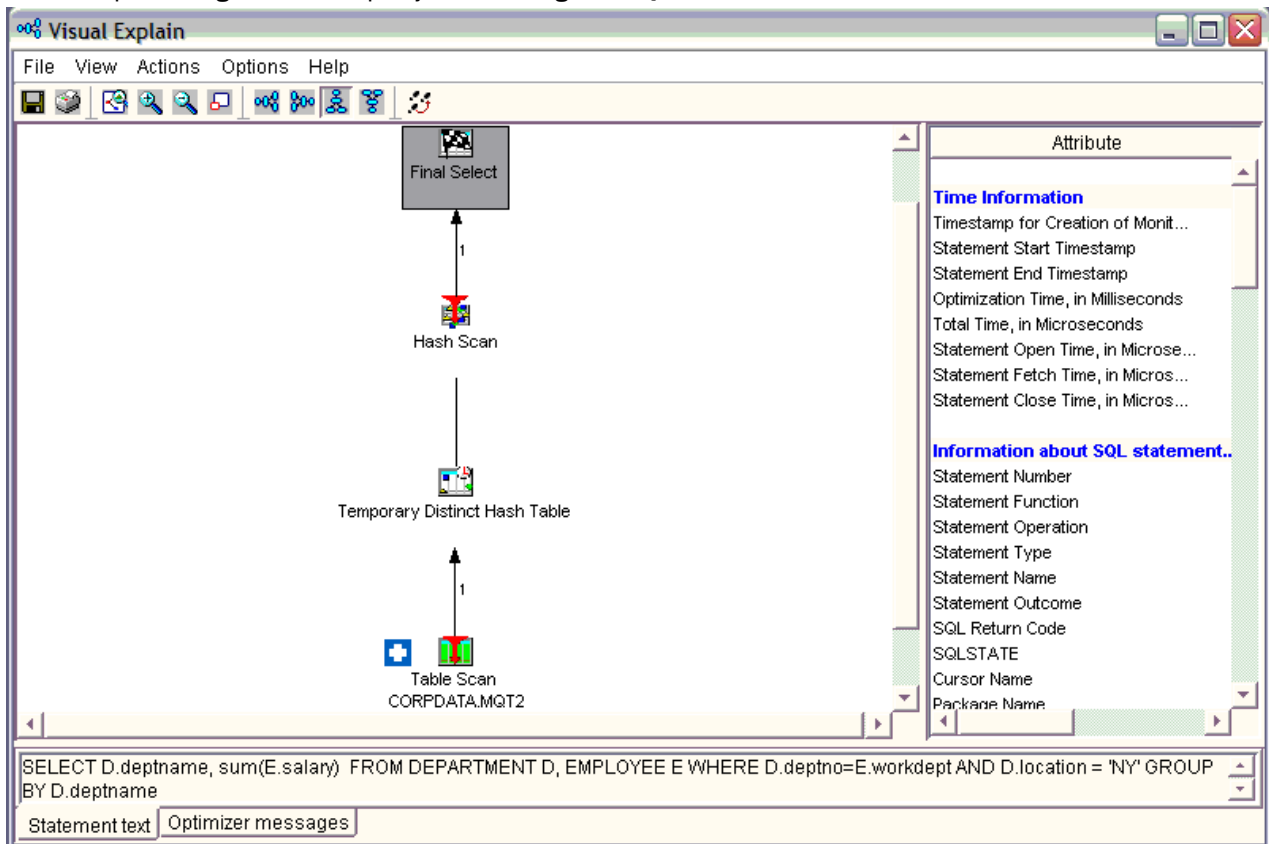
```

SELECT M.deptname, sum(M.sum_sal)
FROM MQT2 M
WHERE M.location = 'NY'
GROUP BY M.deptname

```

Since the MQT could potentially produce more groups than the original query, the final resulting query must group again and SUM the results to return the correct answer. Also, the selection M.location='NY' must be part of the new query.

Visual Explain diagram of the query when using the MQT:



Details on the MQT matching algorithm

What follows is a generalized discussion of how the MQT matching algorithm works.

The tables specified in the query and the MQT are examined. If the MQT and the query specify the same tables, then the MQT can potentially be used and matching continues. If the MQT references tables not referenced in the query, then the unreferenced table is examined to determine if it is a parent table in referential integrity constraint. If the foreign key is non-nullable and the two tables are joined using a primary key or foreign key equal predicate, then the MQT can still be potentially used.

Example 3

The MQT contains fewer tables than the query:

```

SELECT D.deptname, p.projname, sum(E.salary)
FROM DEPARTMENT D, EMPLOYEE E, EMPPROJECT EP, PROJECT P

```

```
WHERE D.deptno=E.workdept AND E.empno=ep.empno
AND ep.projno=p.projno
GROUP BY D.DEPTNAME, p.projname
```

Create an MQT based on the preceding query:

```
CREATE TABLE MQT3
AS (SELECT D.deptname, sum(E.salary) as sum_sal, e.workdept, e.empno
FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept
GROUP BY D.Deptname, e.workdept, e.empno)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER
```

The rewritten query:

```
SELECT M.deptname, p.projname, SUM(M.sum_sal)
FROM MQT3 M, EMPPROJECT EP, PROJECT P
WHERE M.empno=ep.empno AND ep.projno=p.projno
GROUP BY M.deptname, p.projname
```

All predicates specified in the MQT, must also be specified in the query. The query could contain additional predicates. Predicates specified in the MQT must match exactly the predicates in the query. Any additional predicates specified in the query, but not in the MQT must be able to be derived from columns projected from the MQT. See previous example 1.

Example 4

Set the total salary for all departments that are located in 'NY'.

```
SELECT D.deptname, sum(E.salary)
FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept AND D.location = ?
GROUP BY D.Deptname
```

Create an MQT based on the preceding query:

```
CREATE TABLE MQT4
AS (SELECT D.deptname, D.location, sum(E.salary) as sum_sal
FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept AND D.location = 'NY'
GROUP BY D.deptname, D.location)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER
```

In this example, the constant 'NY' was replaced by a parameter marker and the MQT also had the local selection of location='NY' applied to it when the MQT was populated. The MQT matching algorithm matches the parameter marker and to the constant 'NY' in the predicate D.Location=?. It verifies that the values of the parameter marker are the same as the constant in the MQT; therefore the MQT can be used.

The MQT matching algorithm also attempts to match where the predicates between the MQT and the query are not the same. For example, if the MQT has a predicate SALARY > 50000, and the query has the predicate SALARY > 70000, the MQT contains the rows necessary to run the query. The MQT is used in the query, but the predicate SALARY > 70000 is left as selection in the query, so SALARY must be a column of the MQT.

Example 5

```
SELECT D.deptname, sum(E.salary)
FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept AND D.location = 'NY'
GROUP BY D.deptname
```

Create an MQT based on the preceding query:

```
CREATE TABLE MQT5
  AS (SELECT D.deptname, E.salary
      FROM DEPARTMENT D, EMPLOYEE E
      WHERE D.deptno=E.workdept)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER
```

In this example, since D.Location is not a column of the MQT, the user query local selection predicate Location='NY' cannot be determined, so the MQT cannot be used.

Example 6

```
SELECT D.deptname, sum(E.salary)
FROM DEPARTMENT D, EMPLOYEE E
WHERE D.deptno=E.workdept
GROUP BY D.deptname
```

Create an MQT based on the preceding query:

```
CREATE TABLE MQT6(workdept, sumSalary)
AS (SELECT workdept, sum(salary)
    FROM EMPLOYEE
    GROUP BY workdept )
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER
```

In this example, the SUM(salary) aggregation is pushed down through the join to the EMPLOYEE table, allowing for a match and substitution of MQT6. A regrouping to (sum(sum(salary))) is defined at the top of the query to compensate for the grouping pushdown.

Instead of department joining to all the rows in the employee table, it now has the opportunity to join to the predetermined aggregates in MQT6. This type of MQT substitution can result in significant reduction of processing and IO.

If the MQT contains grouping, then the query must be a grouping query. The simplest case is where the MQT and the query specify the same list of grouping columns and column functions.

In some cases, if the MQT specifies group by columns that are a superset of query group by columns, the query can be rewritten to do regrouping. This regrouping reaggregates the groups of the MQT into the groups required by the query. When regrouping is required, the column functions need to be recomputed. The following table shows the supported regroup expressions.

The regrouping expression/aggregation rules are:

<i>Table 32. Expression/aggregation rules for MQTs</i>		
Query	MQT	Final query
COUNT(*)	COUNT(*) as cnt	SUM(cnt)
COUNT(*)	COUNT(C2) as cnt2 (where c2 is non-nullable)	SUM(cnt2)
COUNT(c1)	COUNT(c1) as cnt	SUM(cnt)
COUNT(C1) (where C1 is non-nullable)	COUNT(C2) as cnt2 (where C2 is non-nullable)	SUM(cnt2)
COUNT(distinct C1)	C1 as group_c1 (where C1 is a grouping column)	COUNT(group_C1)
COUNT(distinct C1)	where C1 is not a grouping column	MQT not usable

Table 32. Expression/aggregation rules for MQTs (continued)

Query	MQT	Final query
COUNT(C2) where C2 is from a table not in the MQT	COUNT(*) as cnt	cnt*COUNT(C2)
COUNT(distinct C2) where C2 is from a table not in the MQT	Not applicable	COUNT(distinct C2)
SUM(C1)	SUM(C1) as sm	SUM(sm)
SUM(C1)	C1 as group_c1, COUNT(*) as cnt (where C1 is a grouping column)	SUM(group_c1 * cnt)
SUM(C2) where C2 is from a table not in the MQT	COUNT(*) as cnt	cnt*SUM(C2)
SUM(distinct C1)	C1 as group_c1 (where C1 is a grouping column)	SUM(group_C1)
SUM(distinct C1)	where C1 is not a grouping column	MQT not usable
SUM(distinct C2) where C2 is from a table not in the MQT	Not applicable	SUM(distinct C2)
MAX(C1)	MAX(C1) as mx	MAX(mx)
MAX(C1)	C1 as group_C1 (where C1 is a grouping column)	MAX(group_c1)
MAX(C2) where C2 is from a table not in the MQT	Not applicable	MAX(C2)
MIN(C1)	MIN(C1) as mn	MIN(mn)
MIN(C1)	C1 as group_C1 (where C1 is a grouping column)	MIN(group_c1)
MIN(C2) where C2 is from a table not in the MQT	Not applicable	MIN(C2)
GROUPING(C1)	GROUPING(C1) as grp	grp
GROUPING(C2) where C2 is from a table not in the MQT	Not applicable	GROUPING(C2)

MQT matching does not support ARRAY_AGG, XMLAGG, and XMLGROUP grouping functions. AVG, STDDEV, STDDEV_SAMP, VARIANCE_SAMP and VAR_POP are calculated using combinations of COUNT and SUM. If AVG, STDDEV, or VAR_POP are included in the MQT and regroup requires recalculation of these functions, the MQT cannot be used. It is recommended that the MQT only use COUNT, SUM, MIN, and MAX. If the query contains AVG, STDDEV, or VAR_POP, it can be recalculated using COUNT and SUM.

If FETCH FIRST N ROWS is specified in the MQT, then FETCH FIRST N ROWS must also be specified in the query. Also, the number of rows specified for the MQT must be greater than or equal to the number of rows specified in the query. It is not recommended that an MQT contain the FETCH FIRST N ROWS clause.

The ORDER BY clause on the MQT can be used to order the data in the MQT if a REFRESH TABLE is run. It is ignored during MQT matching and if the query contains an ORDER BY clause, it is part of the rewritten query.

Related reference

[MQT supported function](#)

Although an MQT can contain almost any query, the optimizer only supports a limited set of query functions when matching MQTs to user specified queries. The user-specified query and the MQT query must both be supported by the SQE optimizer.

Determining unnecessary MQTs

You can easily determine which MQTs are being used for query optimization. However, you can now easily find all MQTs and retrieve statistics on MQT usage as a result of System i Navigator and IBM i functionality.

To assist you in tuning your performance, this function produces statistics on MQT usage in a query. To access through the System i Navigator, navigate to: **Database > Schemas > Tables**. Right-click your table and select **Show Materialized Query Tables**. You can also view MQT usage information by right-click on Tables or Views folder and select **Show Materialized Query Tables**. This action displays usage information for MQTs created over all the tables or view in that schema.

Note: You can also view the statistics through an application programming interface (API).

In addition to all existing attributes of an MQT, two fields can help you determine unnecessary MQTs.

These fields are:

Last Query Use States the timestamp when the MQT was last used by the optimizer to replace user specified tables in a query.

Query Use Count Lists the number of instances the MQT was used by the optimizer to replace user specified tables in a query.

The fields start and stop counting based on your situation, or the actions you are currently performing on your system. A save and restore procedure does not reset the statistics counter if the MQT is restored over an existing MQT. If an MQT is restored that does not exist on the system, the statistics are reset.

Related information

[Retrieve member description \(QUSRMBRD\) command](#)

Summary of MQT query recommendations

Follow these recommendations when using MQT queries.

- Do not include local selection or constants in the MQT because that limits the number of user-specified queries where the optimizer can use the MQT.
- For grouping MQTs, only use the SUM, COUNT, MIN, and MAX grouping functions. The query optimizer can recalculate AVG, STDDEV, and VAR_POP in user specified queries.
- Specifying FETCH FIRST N ROWS in the MQT limits the number of user-specified queries where the query optimizer can use the MQT. Not recommended.
- If the MQT is created with DATA INITIALLY DEFERRED, consider specifying DISABLE QUERY OPTIMIZATION to prevent the optimizer from using the MQT until it has been populated. When the MQT is populated and ready for use, the ALTER TABLE statement with ENABLE QUERY OPTIMIZATION enables the MQT.

In addition, consider using a sparse index or EVI INCLUDE additional aggregates rather than an MQT if you are concerned with stale data.

MQT tables need to be optimized just like non-MQT tables. It is recommended that indexes are created over the MQT columns used for selection, join, and grouping, as appropriate. Column statistics are collected for MQT tables.

The database monitor shows the list of MQTs considered during optimization. This information is in the 3030 record. If MQTs have been enabled through the QAQQINI file, and an MQT exists over at least one of the tables in the query, there is a 3030 record for the query. Each MQT has a reason code indicating that it was used or if it was not used, why it was not used.

Related concepts

[How the EVI works](#)

EVIIs work in different ways for costing and implementation.

Related reference

Sparse index optimization

An SQL sparse index is like a select/omit access path. Both the sparse index and the select/omit logical file contain only keys that meet the selection specified. For a sparse index, the selection is specified with a WHERE clause. For a select/omit logical file, the selection is specified in the DDS using the COMP operation.

Recursive query optimization

Certain applications and data are recursive by nature. Examples of such applications are a bill-of-material, reservation, trip planner, or networking planning system. Data in one results row has a natural relationship (call it a parent, child relationship) with data in another row or rows. The kinds of recursion implemented in these systems can be performed by using SQL Stored Procedures and temporary results tables. However, the use of a recursive query to facilitate the access of this hierarchical data can lead to a more elegant and better performing application.

Recursive queries can be implemented by defining either a Recursive Common Table Expression (RCTE) or a Recursive View.

Recursive query example

A recursive query is one that is defined by a Union All with an initialization fullselect that seeds the recursion. The iterative fullselect contains a direct reference to itself in the FROM clause.

There are additional restrictions as to what can be specified in the definition of a recursive query. Those restrictions can be found in SQL Programming topic.

Functions like grouping, aggregation, or distinct require a materialization of all the qualifying records before performing the function. These functions cannot be allowed within the iterative fullselect itself. The functions must be placed in the main query, allowing the recursion to complete.

The following is an example of a recursive query over a table called flights, that contains information about departure and arrival cities. The query returns all the flight destinations available by recursion from the two specified cities (New York and Chicago). It also returns the number of connections and total cost to arrive at that final destination.

This example uses the recursion process to also accumulate information like the running cost and number of connections. Four values are put in the queue entry. These values are:

- The originating departure city (either Chicago or New York) because it remains fixed from the start of the recursion
- The arrival city which is used for subsequent joins
- The incrementing connection count
- The accumulating total cost to reach each destination

Typically the data needed for the queue entry is less than the full record (sometimes much less) although that is not the case for this example.

```
CREATE TABLE flights
(
  departure CHAR (10) NOT NULL WITH DEFAULT,
  arrival CHAR (10) NOT NULL WITH DEFAULT,
  carrier CHAR (15) NOT NULL WITH DEFAULT,
  flight_num CHAR (5) NOT NULL WITH DEFAULT,
  ticket INT NOT NULL WITH DEFAULT)

WITH destinations (departure, arrival, connects, cost ) AS
(
  SELECT f.departure,f.arrival, 0, ticket
  FROM flights f
  WHERE f.departure = 'Chicago' OR
        f.departure = 'New York'
  UNION ALL
  SELECT
```

```

        r.departure, b.arrival, r.connects + 1,
        r.cost + b.ticket
    FROM destinations r, flights b
    WHERE r.arrival = b.departure
)
SELECT DISTINCT departure, arrival, connects, cost
FROM destinations

```

The following is the initialization fullselect of the preceding query. It seeds the rows that start the recursion process. It provides the initial destinations (arrival cities) that are a direct flight from Chicago or New York.

```

SELECT f.departure,f.arrival, 0, ticket
FROM flights f
WHERE f.departure='Chicago' OR
      f.departure='New York'

```

The following is the iterative fullselect of the preceding query. It contains a single reference in the FROM clause to the destination recursive common table expression. It also sources further recursive joins to the same flights table. The arrival values of the parent row (initially direct flights from New York or Chicago) are joined with the departure value of the subsequent child rows. It is important to identify the correct parent/child relationship on the recursive join predicate or infinite recursion can occur. Other local predicates can also be used to limit the recursion. For example, for a limit of at most 3 connecting flights, a local predicate using the accumulating connection count, `r.connects<=3`, can be specified.

```

SELECT
    r.departure, b.arrival, r.connects + 1 ,
    r.cost + b.ticket
FROM destinations r, flights b
WHERE r.arrival=b.departure

```

The main query is the query that references the recursive common table expression or view. It is in the main query where requests like grouping, ordering, and distinct are specified.

```


SELECT DISTINCT departure, arrival, connects, cost
FROM destinations

```

Implementation considerations

To implement a source for the recursion, a new temporary data object is provided called a queue. As rows meet the requirements of either the initialization fullselect or the iterative fullselect, they are pulled up through the union all. Values necessary to feed the continuing recursion process are captured and placed in an entry on the queue: an enqueue operation.

At query runtime, the queue data source then takes the place of the recursive reference in the common table expression or view. The iterative fullselect processing ends when the queue is exhausted of entries or a fetch N rows limitation has been met. The recursive queue feeds the recursion process and holds transient data. The join between dequeuing of these queue entries and the rest of the fullselect tables is always a constrained join, with the queue on the left.



Multiple initialization and iterative fullselects

The use of multiple initialization and iterative fullselects specified in the recursive query definition allows for a multitude of data sources and separate selection requirements to feed the recursion process.

For example, the following query allows for final destinations accessible from Chicago by both flight and train travel.

```
WITH destinations (departure, arrival, connects, cost ) AS
(
  SELECT f.departure, f.arrival, 0 , ticket
  FROM flights f
  WHERE f.departure='Chicago'
  UNION ALL
  SELECT t.departure, t.arrival, 0 , ticket
  FROM trains t
  WHERE t.departure='Chicago'
```

```

UNION ALL
SELECT
    r.departure,b.arrival, r.connects + 1 ,
    r.cost + b.ticket
FROM destinations r, flights b
WHERE r.arrival=b.departure
UNION ALL
SELECT
    r.departure,b.arrival, r.connects+1 ,
    r.cost+b.ticket
FROM destinations r, trains b
WHERE r.arrival=b.departure)
SELECT departure, arrival, connects,cost
FROM destinations;

```

All rows coming out of the RCTE/View are part of the recursion process and need to be fed back in. When there are multiple fullselects referencing the common table expression, the query is rewritten by the optimizer to process all non-recursive initialization fullselects first. Then, using a single queue feed, those same rows and all other row results are sent equally to the remaining iterative fullselects. No matter how you order the initialization and iterative fullselects in the definition of the RCTE/view, the initialization fullselects run first. The iterative fullselects share equal access to the contents of the queue.

The Visual Explain window displays the execution plan for the provided SQL query. The plan starts with two Table Probe operations at the bottom, which feed into a Union all operation. This Union all feeds into a Nested Loop Join, which then feeds into another Union all. This second Union all feeds into an Enqueue operation, which finally feeds into the Final Select operation at the top. The diagram includes various icons representing different types of operations and their relationships.

Attribute	V:
Time Information	
Timestamp for Creation of Monit...	20
Statement Start Timestamp	20
Statement End Timestamp	20
Optimization Time, in Milliseconds	29
Total Time, in Microseconds	15
Statement Open Time, in Micros...	15
Statement Fetch Time, in Micros...	No
Statement Close Time, in Micros...	No
Information about SQL stateme...	
Statement Number	40
Statement Function	Se
Statement Operation	Op
Statement Type	Dy
Statement Name	ST
Statement Outcome	Su
SQL Return Code	0
SQLSTATE	00
Cursor Name	CF
Package Name	
Package Library	
Statement Text	WI
Host Variable Values	0, 1
Rows Fetched	No
Additional information about SQ...	
CLOSQLCSR Value	

Statement text Optimizer messages

Predicate pushing

When processing most queries with non-recursive common table expressions or views, local predicates specified on the main query are pushed down so fewer records need to be materialized. Pushing local predicates from the main query into the defined recursive part of the query (through the Union ALL), however, could considerably alter the process of recursion itself. So as a rule, the Union All specified in a recursive query is currently a predicate fence. Predicates are not pushed down or up, through this fence.

The following is an example of how pushing a predicate in to the recursion limits the recursive results and alter the intent of the query.

The intent of the query is to find all destinations accessible from 'Chicago', not including the final destination of 'Dallas'. Pushing the "arrival<>'Dallas'" predicate into the recursive query alters the output of the intended results. It prevents the output of final destinations where 'Dallas' was an intermediate stop.

```
WITH destinations (departure, arrival, connects, cost ) AS
(
  SELECT f.departure,f.arrival, 0, ticket
  FROM flights f
  WHERE f.departure='Chicago'
  UNION ALL
  SELECT
    r.departure, b.arrival, r.connects + 1 ,
    r.cost + b.ticket
  FROM destinations r, flights b
  WHERE r.arrival=b.departure
)
SELECT departure, arrival, connects, cost
FROM destinations
WHERE arrival != 'Dallas'
```

Conversely, the following is an example where a local predicate applied to all the recursive results is a good predicate to put in the body of the recursive definition because it could greatly decrease the number of rows materialized from the RCTE/View. The better query request here is to specify the r.connects <=3 local predicate with in the RCTE definition, in the iterative fullselect.

```
WITH destinations (departure, arrival, connects, cost ) AS
(
  SELECT f.departure,f.arrival, 0, ticket
  FROM flights f
  WHERE f.departure='Chicago' OR
    f.departure='New York'
  UNION ALL
  SELECT
    r.departure, b.arrival, r.connects + 1 ,
    r.cost + b.ticket
  FROM destinations r, flights b
  WHERE r.arrival=b.departure
)
SELECT departure, arrival, connects, cost
FROM destinations
WHERE r.connects<=3
```

Placement of local predicates is key in recursive queries. They can incorrectly alter the recursive results if pushed into a recursive definition. Or they can cause unnecessary rows to be materialized and then rejected, when a local predicate could legitimately help limit the recursion.

Specifying SEARCH consideration

Certain applications dealing with hierarchical, recursive data could have a requirement in how data is processed: by depth or by breadth.

Using a queuing (First In First Out) mechanism to track the recursive join key values implies the results are retrieved in breadth first order. Breadth first means retrieving all the direct children of a parent row before retrieving any of the grandchildren of that same row. This retrieval is an implementation distinction, however, and not a guarantee.

Applications might want to guarantee how the data is retrieved. Some applications might want to retrieve the hierarchical data in depth first order. Depth first means that all the descendents of each immediate child row are retrieved before the descendents of the next child are retrieved.

The SQL architecture allows for the guaranteed specification of how the application retrieves the resulting data by the use of the SEARCH DEPTH FIRST or BREADTH FIRST keyword. When this option is specified, name the recursive join value, identify a set sequence column, and provide the sequence column in an outer ORDER BY clause. The results are output in depth or breadth first order. Note this ordering is ultimately a relationship sort and not a value-based sort.

Here is the preceding example output in depth first order.

```
WITH destinations (departure, arrival, connects, cost ) AS
(
  SELECT f.departure, f.arrival, 0 , ticket
  FROM flights f
  WHERE f.departure='Chicago' OR f.departure='New York'
  UNION ALL
  SELECT
    r.departure,b.arrival, r.connects+1 ,
    r.cost+b.ticket
  FROM destinations r, flights b
  WHERE r.arrival=b.departure)

SEARCH DEPTH FIRST BY arrival SET depth_sequence

SELECT *
FROM destinations
ORDER BY depth_sequence
```

If the ORDER BY clause is not specified in the main query, the sequencing option is ignored. To facilitate the correct sort there is additional information put on the queue entry during recursion. With BREADTH FIRST, it is the recursion level number and the immediate ancestor join value, so sibling rows can be sorted together. A depth first search is a little more data intensive. With DEPTH FIRST, the query engine needs to represent the entire ancestry of join values leading up to the current row and put that information in a queue entry. Also, because these sort values are not coming from an external data source, the sort implementation is always a temporary sorted list (no indexes possible).

Do not use the SEARCH option if you do not need your data materialized in a depth or breadth first manner. There is additional CPU and memory overhead to manage the sequencing information.

Specifying CYCLE considerations

Recognizing that data in the tables used in a recursive query might be cyclic in nature is important to preventing infinite loops.

The SQL architecture allows for the optional checking for cyclic data and discontinuing the repeating cycles at that point. This additional checking is done by the use of the CYCLE option. The correct join recursion value must be specified on the CYCLE request and a cyclic indicator must be specified. The cyclic indicator could be optionally output in the main query and can be used to help determine and correct errant cyclic data.

```
WITH destinations (departure, arrival, connects, cost , itinerary) AS
(
  SELECT f.departure, f.arrival, 1 , ticket, CAST(f.departure||f.arrival AS VARCHAR(2000))
  FROM flights f
  WHERE f.departure='New York'
  UNION ALL
  SELECT r.departure,b.arrival, r.connects+1 ,
    r.cost+b.ticket, cast(r.itinerary||b.arrival AS varchar(2000))
  FROM destinations r, flights b
  WHERE r.arrival = b.departure)
CYCLE arrival SET cyclic TO '1' DEFAULT '0' USING Cycle_Path

SELECT departure, arrival, itinerary, cyclic
FROM destinations
```

When a cycle is determined to be repeating, the output of that cyclic sequence of rows is stopped. To check for a 'repeated' value however, the query engine needs to represent the entire ancestry of the join

values leading up to the current row in order to look for the repeating join value. This ancestral history is information that is appended to with each recursive cycle and put in a field on the queue entry.

To implement this history field, the query engine uses a compressed representation of the recursion values on the ancestry chain. The query engine can then do a fixed length, quicker scan through the accumulating ancestry to determine if the value has been seen before. This compressed representation is determined by the use of a distinct node in the query tree.

Do not use the CYCLE option unless you know your data is cyclic, or you want to use it specifically to help find the cycles for correction or verification purposes. There is additional CPU and memory overhead to manage and check for repeating cycles before a given row is materialized.

The screenshot shows the Visual Explain interface. The main area displays a query execution plan as a tree diagram. The root node is 'Final Select', which connects to 'ListScan', then 'Temporary List', 'Eqseteq', 'Distinct', and 'Union all'. The 'Union all' node branches into 'Table Probe' and 'Nested Loop Join'. 'Table Probe' further branches into 'Index Probe' and 'Deqseteq'. 'Nested Loop Join' connects to 'Table Probe'. Arrows with 'z' or 'c 1' indicate the flow and cardinality between nodes.

On the right side, there is a list of attributes under the heading 'Attribute'. The list is divided into three sections:

- Time Information:**
 - Timestamp for Creation of Monit...
 - Statement Start Timestamp
 - Statement End Timestamp
 - Optimization Time, in Milliseconds
 - Total Time, in Microseconds
 - Statement Open Time, in Micros...
 - Statement Fetch Time, in Micros...
 - Statement Close Time, in Micros...
- Information about SQL stateme...:**
 - Statement Number
 - Statement Function
 - Statement Operation
 - Statement Type
 - Statement Name
 - Statement Outcome
 - SQL Return Code
 - SQLSTATE
 - Cursor Name
 - Package Name
 - Package Library
 - Statement Text
 - Host Variable Values
 - Rows Fetched
- Additional information about SQ...:**
 - CLOSESQLCSR Value
 - ALWCOPYDTA Value
 - Pseudo Open
 - Pseudo Close

At the bottom, there is a text area containing the SQL statement:

```
WITH destinations (departure, arrival, connects, cost, itinerary) AS (
  SELECT f.departure, f.arrival, 1, ticket,
  cast(f.departure||f.arrival as varchar(2000))FROM flights f WHERE f.departure='New York' UNION ALL
  SELECT f.departure, f.arrival, connects+1, cost+ticket, cost||itinerary||f.arrival as varchar(2000))
```

Below the text area are two tabs: 'Statement text' and 'Optimizer messages'.

SMP and recursive queries

Recursive queries can benefit as much from symmetric multiprocessing (SMP) as do other queries on the system.

Recursive queries and parallelism, however, present some unique requirements. The initialization fullselect of a recursive query is the fullselect that seeds the initial values of the recursion. It is likely to produce only a small fraction of the ultimate results that cycle through the recursion process. The

query optimizer does not want each of the threads running in parallel to have a unique queue object that feeds only itself. This results in some threads having way too much work to do and others threads quickly depleting their work.

The best way to handle this work is to have all the threads share the same queue. This method allows a thread to enqueue a new recursive key value just as a waiting thread is there to dequeue that request. A shared queue allows all threads to actively contribute to the overall depletion of the queue entries until no thread is able to contribute more results.

Having multiple threads share the same queue, however, requires some management by the Query runtime so that threads do not prematurely end. Some buffering of the initial seed values might be necessary. This buffering is illustrated in the following query, where there are two fullselects that seed the recursion. A buffer is provided so that no thread hits a dequeue state and terminates before the query has seeded enough recursive values to get things going.

The following Visual Explain diagram shows the plan for the following query run with CHGQRYA DEGREE(*NBRTASKS 4). It shows how the results of the multiple initialization fullselects are buffered up. The multiple threads, illustrated by the multiple arrow lines, are acting on the enqueue and dequeue request nodes. As with all SMP queries, the multiple threads, in this case 4, put their results into a Temporary List object which becomes the output for the main query.

```
cl:chgqrya degree(*nbrtasks 4);

WITH destinations (departure, arrival, connects, cost )AS
(
  SELECT f.departure, f.arrival, 0 , ticket
  FROM flights f WHERE f.departure='Chicago'
  UNION ALL
  SELECT t.departure, t.arrival, 0 , ticket
  FROM trains t WHERE t.departure='Chicago'
  UNION ALL
  SELECT
    r.departure,b.arrival, r.connects+1 ,
    r.cost+b.ticket
  FROM destinations r, flights b
  WHERE r.arrival=b.departure
  UNION ALL
  SELECT
    r.departure,b.arrival, r.connects+1 ,
    r.cost+b.ticket
  FROM destinations r, trains b
  WHERE r.arrival=b.departure)
SELECT departure, arrival, connects,cost
FROM destinations;
```

Visual Explain

File View Actions Options Help

Attribute	Value
Time Information	
Timestamp for Creation of Monit...	2005
Statement Start Timestamp	2005
Statement End Timestamp	2005
Optimization Time, in Milliseconds	111
Total Time, in Microseconds	2497
Statement Open Time, in Micros...	2497
Statement Fetch Time, in Micros...	Not A
Statement Close Time, in Micros...	Not A
Information about SQL stateme...	
Statement Number	17
Statement Function	Sele
Statement Operation	Oper
Statement Type	Dyna
Statement Name	STM
Statement Outcome	Succ
SQL Return Code	0
SQLSTATE	000C
Cursor Name	CRS
Package Name	
Package Library	
Statement Text	WITH
Host Variable Values	0, C,
Rows Fetched	Not A
Additional information about SQ...	
CLOSESQLCSR Value	
ALWCOPYDTA Value	Any T
Pseudo Open	No
Pseudo Close	No
Hard Close Reason Code	Not A
ODP Implementation	Reus
Dynamic Replan Reason Code	Acce
Timestamp When Plan Was Cre...	0001
Data Conversion Reason Code	Not a
Blocking Enabled	ALW
Delay Prep	Yes
Statement is Explainable	Yes
Naming Convention	SQL
Type of Dynamic Processing	Loca
SQL Path	"QSY
Information common to most m...	
System Name	Y045
Job Name	QZD.
Job User	QUS
Job Number	1889

WITH destinations (departure, arrival, connects, cost)AS (SELECT f.departure, f.arrival, 0 , ticket FROM

Statement text Optimizer messages

System-period temporal tables

Querying a system-period temporal table can return results for a specified point or period in time. These results can include both current values and previous historic values. The following sample queries request policy information from a system-period temporal table (`policy_info`), which also implicitly get information from the associated history table (`hist_policy_info`). See the “Querying system-period temporal data” topic in the Database Administration book for the layout of these tables and more information on how to specify time criteria for a system-period temporal table in a query.

Query with FOR SYSTEM_TIME AS OF specified.

```
SELECT policy_id, coverage
FROM policy_info
FOR SYSTEM_TIME AS OF '2011-02-28-09.10.12.649592000000'
```

For this query, the begin column of the period is inclusive, while the end column is exclusive. The history table row(s) with the begin column value less than or equal to '2011-02-28-09.10.12.649592000000' and the end column value greater than '2011-02-28-09.10.12.649592000000' will be included in the result.

As a result Db2 rewrites the query as follows:

```
SELECT policy_id, coverage FROM
  (SELECT policy_id, coverage
   FROM policy_info
   WHERE sys_start <= '2011-02-28-09.10.12.649592000000'
  UNION ALL
   SELECT policy_id, coverage
   FROM hist_policy_info
   WHERE sys_start <= '2011-02-28-09.10.12.649592000000'
   AND sys_end > '2011-02-28-09.10.12.649592000000')
```

Query with FOR SYSTEM_TIME FROM...TO specified.

```
SELECT policy_id, coverage, sys_start, sys_end
FROM policy_info
FOR SYSTEM_TIME FROM '0001-01-01-00.00.00.000000'
                    TO '9999-12-30-00.00.00.000000000000'
WHERE policy_id = 'C567'
```

For this query, the begin column and end column of the period are exclusive. The history table row(s) with the begin column value less than '9999-12-30-00.00.00.000000000000' and the end column value greater than '0001-01-01-00.00.00.000000' will be included in the result.

As a result, Db2 rewrites the query as follows:

```
SELECT policy_id, coverage, sys_start, sys_end FROM
  (SELECT policy_id, coverage, sys_start, sys_end
   FROM policy_info
   WHERE sys_start < '9999-12-30-00.00.00.000000000000'
   AND TIMESTAMP('0001-01-01-00.00.00.000000') <
      TIMESTAMP('9999-12-30-00.00.00.000000000000')
  UNION ALL
   SELECT policy_id, coverage, sys_start, sys_end
   FROM hist_policy_info
   WHERE sys_start < '9999-12-30-00.00.00.000000000000'
   AND sys_end > '0001-01-01-00.00.00.000000')
   AND TIMESTAMP('0001-01-01-00.00.00.000000') <
      TIMESTAMP('9999-12-30-00.00.00.000000000000')
WHERE policy_id = 'C567'
```

Query with FOR SYSTEM_TIME BETWEEN...AND specified.

```
SELECT policy_id, coverage
FROM policy_info
```

```
FOR SYSTEM_TIME BETWEEN '2011-02-28-09.10.12.649592000000'  
AND '9999-12-30-00.00.00.000000000000'
```

For this query, the begin column of the period is inclusive, while the end column is exclusive. The history table row(s) with the begin column value less than or equal to '9999-12-30-00.00.00.000000000000' and the end column value greater than '2011-02-28-09.10.12.649592000000' will be included in the result.

As a result, Db2 rewrites the query as follows:

```
SELECT policy_id, coverage, sys_start, sys_end FROM  
(SELECT policy_id, coverage, sys_start, sys_end  
FROM policy_info  
WHERE sys_start <= '9999-12-30-00.00.00.000000000000'  
AND TIMESTAMP('2011-02-28-09.10.12.649592000000') <=  
TIMESTAMP('9999-12-30-00.00.00.000000000000')  
UNION ALL  
SELECT policy_id, coverage, sys_start, sys_end  
FROM hist_policy_info  
WHERE sys_start <= '9999-12-30-00.00.00.000000000000'  
AND sys_end > '2011-02-28-09.10.12.649592000000')  
AND TIMESTAMP('2011-02-28-09.10.12.649592000000') <=  
TIMESTAMP('9999-12-30-00.00.00.000000000000')
```

Query with time criteria specified via CURRENT TEMPORAL SYSTEM_TIME special register.

The advantage of using this method is that you can change the time criteria later and not have to modify the SQL. For example, assume that you want to retrieve data from policy_info for a given policy_id of C567 that is from one year ago. If the SYSTIME option is set to YES, you can set the CURRENT TEMPORAL SYSTEM_TIME special register and issue the SELECT statement as follows:

```
SET CURRENT TEMPORAL SYSTEM_TIME = CURRENT_TIMESTAMP - 1 YEAR;  
  
SELECT policy_id, coverage FROM policy_info  
WHERE policy_id = 'C567';
```

Db2 interprets the SELECT statement as follows:

```
SELECT policy_id, coverage FROM policy_info  
FOR SYSTEM_TIME AS OF CURRENT TEMPORAL SYSTEM_TIME  
WHERE policy_id = 'C567';
```

For this query, the begin column of the period is inclusive, while the end column is exclusive. The history table row(s) with the begin column value less than or equal to CURRENT TEMPORAL SYSTEM_TIME and the end column value greater than CURRENT TEMPORAL SYSTEM_TIME will be included in the result.

As a result, Db2 rewrites the query as follows:

```
SELECT policy_id, coverage FROM  
(SELECT policy_id, coverage  
FROM policy_info  
WHERE sys_start <= CURRENT TEMPORAL SYSTEM_TIME  
UNION ALL  
SELECT policy_id, coverage  
FROM hist_policy_info  
WHERE sys_start <= CURRENT TEMPORAL SYSTEM_TIME  
AND sys_end > CURRENT TEMPORAL SYSTEM_TIME)  
WHERE policy_id = 'C567';
```

Adaptive Query Processing

Adaptive Query Processing analyzes actual query run time statistics and uses that information for subsequent optimizations.

With rapidly increasing amounts of data, the price of miscalculating complex plans can result in dramatic performance problems. These problems might be measured in minutes or hours instead of seconds or minutes. Traditionally, optimizer architecture has attempted to overcome potential plan problems in

several ways. The most common technique is to increase the amount of time spent optimizing a query, searching for safe alternatives. While additional time reduces the likelihood of a failed plan, it does not fundamentally avoid the problem.

The Db2 optimizer relies on statistical estimates to optimize a query. These estimates can be inaccurate for a number of reasons. The reasons include a lack of statistical metadata for the query tables, complex join conditions, skewed or rapidly changing data within the tables, and others.

The SQE query engine uses a technique called Adaptive Query Processing (AQP). AQP analyzes actual query run time statistics and uses that information to correct previous estimates. These updated estimates can provide better information for subsequent optimizations.

Related reference

[Adaptive Query Processing in Visual Explain](#)

You can use Visual Explain to request a new plan.

How AQP works

There are three main parts to AQP support.

- **Global Statistics Cache (GSC):** The [“Global Statistics Cache”](#) on page 14 is a system-side repository of statistical information gathered from actual query runs. When the SQE query engine observes a discrepancy between record count estimates and actual observed values, an entry might be made in the GSC. This entry provides the optimizer with more accurate statistical information for subsequent optimizations.
- **AQP Request Support:** This support runs after a query completes. The processing is done in a system task so it does not affect the performance of user applications. Estimated record counts are compared to the actual values. If significant discrepancies are noted, the AQP Request Support stores the observed statistic in the GSC. The AQP Request Support might also make specific recommendations for improving the query plan the next time the query runs.
- **AQP Handler:** The AQP Handler runs in a thread parallel to a running query and observes its progress. The AQP handler wakes up after a query runs for at least 2 seconds without returning any rows. Its job is to analyze the actual statistics from the partial query run, diagnose, and possibly recover from join order problems. These join order problems are due to inaccurate statistical estimates.

The query can be reoptimized using partial observed statistics or specific join order recommendations or both. If this optimization results in a new plan, the old plan is terminated and the query restarted with the new plan, provided the query has not returned any results.

AQP looks for an unexpected starvation join condition when it analyzes join performance. Starvation join is a condition where a table late in the join order eliminates many records from the result set. In general, the query would perform better if the table that eliminates the large number of rows is first in the join order. When AQP identifies a table that causes an unexpected starvation join condition, the table is noted as the 'forced primary table'. The forced primary table is saved for a subsequent optimization of the query.

That subsequent optimization with the forced primary recommendation can be used in two ways:

- The forced primary table is placed first in the join order, overriding the join order implied by the statistical estimates. The rest of the join order is defined using existing techniques.
- The forced primary table can be used for LPG preselection against a large fact table in the join.

Related reference

[Adaptive Query Processing in Visual Explain](#)

You can use Visual Explain to request a new plan.

AQP example

Here is an example query with an explanation of how AQP could work.

```
SELECT * from t1, t2, t3, t4
WHERE t1.c1=t2.c1 AND t1.c2=t3.c2
AND t1.c3 = CURRENT DATE - t4.c3
```

```
AND t1.c5 < 50 AND t2.c6 > 40  
AND t3.c7 < 100 AND t4.c8 - t4.c9 < 5
```

The WHERE clause of the preceding query contains a predicate, `t1.c3 = CURRENT DATE - t4.c3`, that is difficult to estimate. The estimation difficulty is due to the derivation applied to column `t4.c3` and the derivation involving columns `t4.c8` and `t4.c9`. For the purposes of this example, the predicate `t1.c3 = CURRENT DATE - t4.c3` actually eliminates all or nearly all records in the join.

Due to characteristics of the columns involved in that predicate, the statistical estimate has many rows returned from the join. The optimizer selects join order `t1, t3, t2, t4` based on the following record count estimates.

- Join `t1` to `t3` produces 33,000,000 rows.
- Join `t1, t3` result to `t2` produces 1,300,000 rows.
- Join `t1, t3, t2` result to `t4` (final result set) produces 5 million rows.

The join order is reasonable assuming that the final result set actually produces 5 million rows, but the estimate is incorrect. The query performs poorly since tables `t1, t3, t2` are joined first, producing 1,300,000 rows. These rows are all rejected by table `t4` and the `t1.c3 = CURRENT DATE - t4.c3` predicate (join starvation).

AQP identifies `t4` as the forced primary table. The optimizer would choose `t1` as the second table in the join order since there are no join conditions between `t4` and `t2` or `t3`. Since the join condition between tables `t4` and `t1` selects few rows, this plan is likely many orders of magnitude faster than the original plan.

Related reference

[Adaptive Query Processing in Visual Explain](#)

You can use Visual Explain to request a new plan.

AQP join order

Adaptive Query Processing analyzes actual query run time join statistics and uses that information for subsequent join optimizations.

The SQE engine implements AQP join order recommendations in the following ways:

Subsequent to run

When each query completes, a fast check is done on key points of the query execution to compare actual selected records with the estimates. If there is a significant discrepancy, then a stand-alone task is notified to do a deeper analysis of the query execution.

The query plan and the execution statistics are passed to the task. A separate task is used for the in-depth analysis so the user job is not impacted while the deep analysis is done. Each step of the join is analyzed, looking for characteristics of starvation join. Starvation join shows a significant reduction in the number of rows produced compared to the previous step. The definition of what is considered significant depends on a number of factors.

If the criteria for starvation join are met, the actual number of records selected at key points of the query are compared to estimates. If there is a significant discrepancy between the actual and estimated record counts, the table at that join position is identified as a 'forced primary table'. This table is saved with the query plan in the system plan cache. When the query runs in the future, the optimizer retrieves the original plan from the system plan cache. The optimizer sees the forced primary table recommendation, and optimizes the query using this recommendation.

The forced primary recommendation is used in two ways by the optimizer:

- The forced primary table is placed first in the join order by the join order optimization strategy.
- The forced primary table is used by the strategy for LPG optimization. The preceding example is a star join since table `T1` is joined to the other tables in the query. `t1.c3` is the column used to join `T1` to `T4`. If an index exists over this join column, then it might be advantageous to do preselection against table `T1`.

using the records selected from table T4. The forced primary table recommendation is used as a hint for the optimizer to consider this technique.

Concurrent to run

The preceding logic to identify starvation join can also run in a thread in parallel to the executing query. The AQP handler thread is created for longer running queries. The thread monitors the query execution and can run the same logic described earlier against partial data from the query execution.

If the partial results show starvation join and significant differences with the record count estimates, the query is reoptimized in the thread. When the new plan is ready, the execution of the original plan is stopped and the new plan started. This scheme for correcting join problems 'on the fly' can only be carried out before any records are selected for the final result set.

Note: AQP can help correct query performance problems, but it is not a substitute for a good database design coupled with a good indexing strategy.

Related reference

[Adaptive Query Processing in Visual Explain](#)

You can use Visual Explain to request a new plan.

Database Monitor additions for AQP

Additional information is logged in the database monitor when the AQP handler code replaces an executing plan.

A new set of 30xx records is written to the database monitor reflecting the replaced plan. The user needs to be able to distinguish between records produced for the first plan iteration and records produced for subsequent optimization. To distinguish these records, an unused short integer column of the database monitor record is used as a 'plan iteration counter'.

Column QQSMINTF is used for this purpose. For the original optimization of the query, the 30xx records have this field set to 1. Subsequent reoptimization done by AQP processing will increment the value by 1.

The following is an example of how DB monitor output might look like when *a* is replaced 'on the fly'. The example query is the following two-file join with an ORDER BY clause over one of the tables:

```
SELECT a.orderkey,b.orderkey
FROM rvdstar/item_fact3 a, rvdstar/item_fact b
WHERE a.quarter - 8 = b.quarter
ORDER BY b.orderkey
```

Assume that an *order by pushdown* plan is chosen, then replaced using AQP while the query is running. The following is an example of what the DB monitor records might look like. The columns shown for the purposes of explaining the changes are QQRID, QQUCNT, QQSMINTF, and QQRCD. The other fields in the monitor are not affected by AQP processing.

QQRID	QQUCNT	QQSMINTF	QQRCD
3010	14	-	-
3006	14	1	A0
3001	14	1	I2
3000	14	1	T1
3023	14	1	-
3007	14	1	-
3020	14	1	I1
3014	14	1	-

Table 33. Database monitor records for example query (continued)

QQRID	QQUCNT	QQSMINTF	QQRCD
5005	14	1	-
5002	14	1	-
5004	14	1	-
5007	14	1	-
3006	14	2	B6
3000	14	2	T1
3000	14	2	T3
3023	14	2	-
3003	14	2	F7
3007	14	2	-
3020	14	2	I1
3014	14	2	-
5005	14	2	-
5002	14	2	-
5004	14	2	-
1000	14	2	-
5007	14	2	-
3019	14	-	-
1000	14	-	-

Notes on the preceding table:

- There is a full set of optimizer-generated records that reflect the first choice of the optimizer: an *order by pushdown* plan. These records have the QQSMINTF column value set to 1. There is a 3001 record indicating an index was used to provide the ordering. There are 3000 and 3023 records indicating a Table Scan of the second table and a temporary hash table built to aid join performance. The remaining records, including the 3014 and the 500x records, have QQSMINTF set to 1 to reflect their association with the original *order by pushdown* plan.
- There is a second full set of optimizer-generated records that reflect the second choice of the optimizer: a *sorted temporary* plan to implement the ORDER BY. These records have the QQSMINTF column value set to 2. This time there are two 3000 records indicating table scan was used to access both tables. There is a 3023 record indicating a temporary hash table was built and a 3003 record indicating the results were sorted. The remaining records, including the 3014 and the 500x records, have QQSMINTF set to 2 to reflect their association with the replacement plan.
- Both sets of optimizer records have the same unique count (QQUCNT value).
- There is a 3006 (Access Plan Rebuilt) record generated for each replacement plan (QQSMINTF > 0). The QQRCD (reason code) value is set to a new value, 'B6'. The 'B6' value indicates the access plan was rebuilt due to AQP processing. In the example, there is a 3006 record with QQSMINTF = 1 and a QQRCD value of 'A0'. The 1 indicates that the original optimization built the plan for the first time. There might not be a 3006 record associated with the original optimization if the optimizer was able to reuse a plan from the plan cache.
- The 1000, 3010 and 3019 records are produced by XPF at open or close time. These records are not generated by the optimizer so there are no changes due to AQP. There are one set of the records, as in

previous releases, regardless of whether AQP replaced the plan. The QQSMINTF value is *NULL* for these records.

- The replacement plan is the plan that runs to completion and returns the results. To retrieve the DB monitor records from the plan that actually returns the records, it is necessary to query the DB monitor file using a subquery. Retrieve the records where the QQSMINTF value is equal to the maximum QQSMINTF value for a given QQUCNT.

Related concepts

[Database monitor formats](#)

This section contains the formats used to create the database monitor SQL tables and views.

Related reference

Monitoring your queries using the Database Monitor

Start Database Monitor (STRDBMON) command gathers information about a query in real time and stores this information in an output table. This information can help you determine whether your system and your queries are performing well, or whether they need fine-tuning. Database monitors can generate significant CPU and disk storage overhead when in use.

[Adaptive Query Processing in Visual Explain](#)

You can use Visual Explain to request a new plan.

[QAQQINI query options](#)

There are different options available for parameters in the QAQQINI file.

Row and column access control (RCAC)

Db2 for i introduces row and column access control (RCAC) as an additional layer of data security. RCAC controls access to a table at the row level, column level, or both. RCAC can be used to complement the existing table privileges model.

Indexing Strategy and RCAC

This section focuses on the consequence of RCAC to your SQL query performance when indexing is used.

Row and column access control (RCAC) places access control at the table level around the data itself. SQL rules, which are known as row permissions or column masks, created on rows and columns are the basis of the implementation of this capability.

You can use row and column access control to ensure that your users have access to only the data that is required for their work. For example, tellers in a bank can access customer rows in the CUSTOMER table only from their own branch. All tellers are members of the group user profile TELLER. Customer service representative or telemarketers are members of other groups and allowed to see all rows. A row permission is created by a user who is authorized to the QIBM_DB_SECADM function usage ID.

These SQL rules add additional predicates to any queries or data access requests over tables with defined and activated RCAC permissions. In this example, SQL rules are added to queries over the CUSTOMER table to enforce the following access rules. Depending on the nature of the rules, additional indexes might be advised or existing indexes might need to be enhanced or altered to accommodate the additional predicates enforcing the access. For example, when the TELLER_ROW_ACCESS permission is enabled, additional index advise might include the BRANCH_INFO table and key EMP_ID. In this particular example, index only access can be facilitated by creating an index over BRANCH_INFO that includes EMP_ID and HOME_BRANCH as key fields. The first to facilitate the probe, the second to prevent unnecessary access to the BRANCH_INFO table.

```
CREATE PERMISSION TELLER_ROW_ACCESS ON CUSTOMER
-----
-- Teller information:
-- Group TELLER is allowed to access customer data only
-- in their branch.
-----
FOR ROWS WHERE VERIFY_GROUP_FOR_USER(SESSION_USER, 'TELLER') = 1
AND
BRANCH = (SELECT HOME_BRANCH FROM BRANCH_INFO WHERE EMP_ID = SESSION_USER)
```

```
ENFORCED FOR ALL ACCESS
ENABLE;
```

```
ALTER TABLE CUSTOMER ACTIVATE ROW ACCESS CONTROL;
```

In the example below, not only are you verifying certain user groups for access to particular patient records but also masking certain data based on whether the patient has participated in a clinical trial. Extra security is that physicians can see only patient records for whom they are the primary care provider.

```
CREATE PERMISSION PCP ON patient
-----
-- Primary Care Physician Access
-- Group PCP is allowed to access patient data only
-- AND the Primary Care Physician must be assigned to patient
-- Group RESEARCH are allowed to access patient data for those patients
-- that opted in to a clinical trial
-----
FOR ROWS WHERE
(VERIFY_GROUP_FOR_USER(SESSION_USER, 'PCP') = 1
AND
PCPID = (SELECT PCPID FROM PHYSICIAN WHERE PCPUSER = SESSION_USER) )
OR
(VERIFY_GROUP_FOR_USER(SESSION_USER, 'RESEARCH') = 1
AND
(SELECT 1 FROM PATIENTCHOICE C
WHERE PATIENT.patientid = C.patientid
AND C.CHOICE = 'clinical trial'
AND C.VALUE = 'opt-in')=1
)
ENFORCED FOR ALL ACCESS
ENABLE;
```

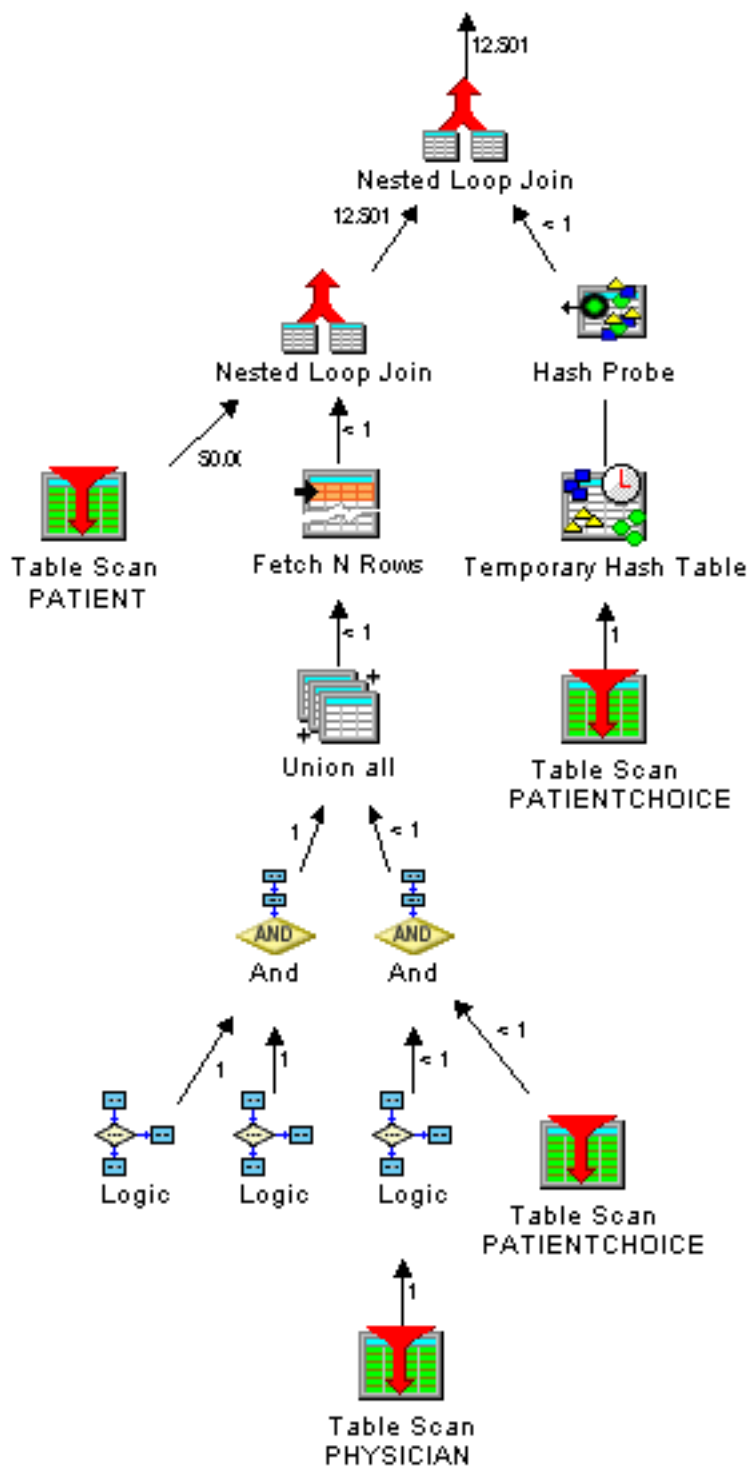
```
CREATE MASK PHARMACY_MASK ON PATIENT FOR
-----
-- Medical information:
-- Group PCP is allowed to access the full information in column PHARMACY.
-- For the purposes of drug research, Role DRUG_RESEARCH can
-- conditionally see a patient's medical information
-- provided that the patient has opted-in.
-- In all other cases, null values are rendered as column
-- values.
-----
COLUMN PHARMACY RETURN
CASE WHEN
    VERIFY_GROUP_FOR_USER(SESSION_USER, 'PCP') = 1 OR
    (VERIFY_GROUP_FOR_USER(SESSION_USER, 'DRUG_RSRCH')=1
    AND
    (SELECT 1 FROM PATIENTCHOICE C
    WHERE PATIENT.patientid = C.patientid
    AND C.CHOICE = 'drug-research'
    AND C.VALUE = 'opt-in')= 1
    )
THEN PHARMACY
ELSE NULL
END
ENABLE;
```

```
ALTER TABLE PATIENT ACTIVATE ROW ACCESS CONTROL ACTIVATE COLUMN ACCESS CONTROL;
```


The query in the next example, before the introduction of RCAC policies would have accessed only the PATIENT table. Now it accesses the PATIENT table and the supporting tables that are associated with the row and column permissions.

The next graphic is the Visual Explain for the next example query. As you can see, the PATIENT table is accessed along with any other tables mentioned in the ROW and COLUMN access control.

```
SELECT * FROM PATIENT WHERE PATIENTID = ?
```

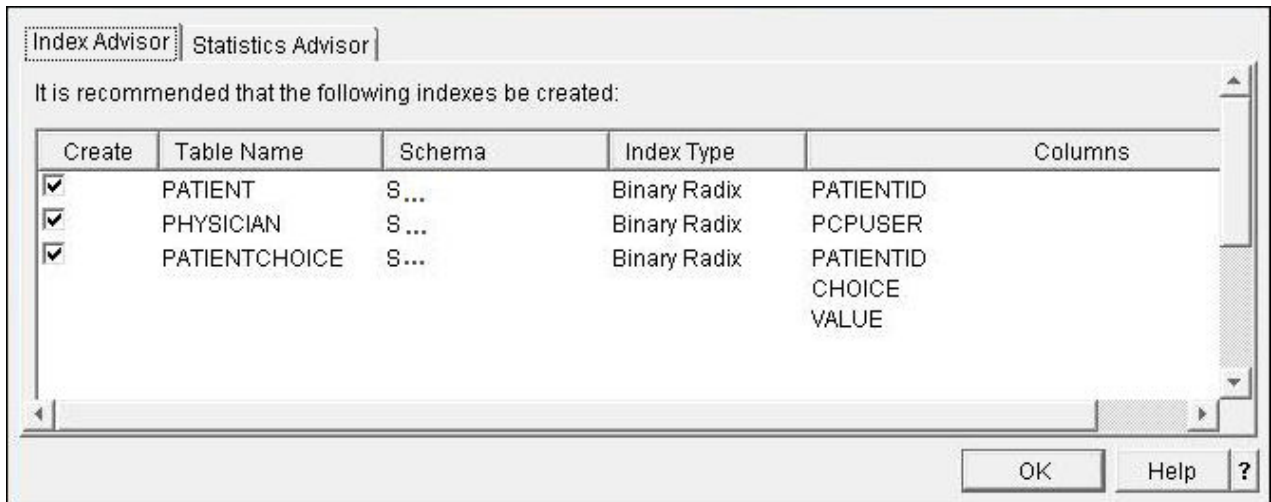


By clicking the index advised icon that is shown in the next graphic:



You get the resulting index advice depicted in the next graphic that shows that it is not only over the PATIENT table that is explicitly specified in the query, but also over the supporting RCAC tables.

Not considering additional advice per the introduction of RCAC SQL rules can affect query performance.



Materialized query tables and RCAC

This section focuses on the consequence of RCAC to your SQL query performance when MQTs are used.

Materialized Query Tables (MQTs) are heavily relied upon by data warehousing applications for better query performance. RCAC and MQTs coexist in harmony. This means:

1. MQTs must continue to provide their added performance benefit to data warehousing applications.
2. MQTs cannot become a means for gaining access to data protected through RCAC rules that are specified in the dependent base tables, either through direct access to the MQT or by MQT matching and substitution.

If a materialized query table that depends on the table (directly or indirectly through a view) for which access control is being activated and that materialized query table does not already have its own access control activated, row level access control is implicitly activated for the materialized query table. This restricts direct access to the contents of the materialized query table. A query that explicitly references the MQT table before such a row permission is defined returns Row Not Found as if there was no data in the table.

In this example MQT:

```
CREATE TABLE MQT1
AS (SELECT patientid, patientname,pcpid,pharmacy
FROM patient
WHERE diagnosis is not null)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER;
```

To provide access to this materialized query table, an appropriate row permission can be created, or an ALTER TABLE DEACTIVATE ROW ACCESS CONTROL on the materialized query table can be issued to remove the row level protection if that is appropriate. If the query optimizer substitutes one or more tables in a query with this materialized query table via MQT substitution, the row and column access controls on the replaced (base) tables remain in effect, and the access controls, if any, on the materialized query table do not apply.

```
SELECT * FROM MQT1
```

results in no rows because it does not have its own RCAC policy and therefore it cannot expose rows per the PATIENT table.

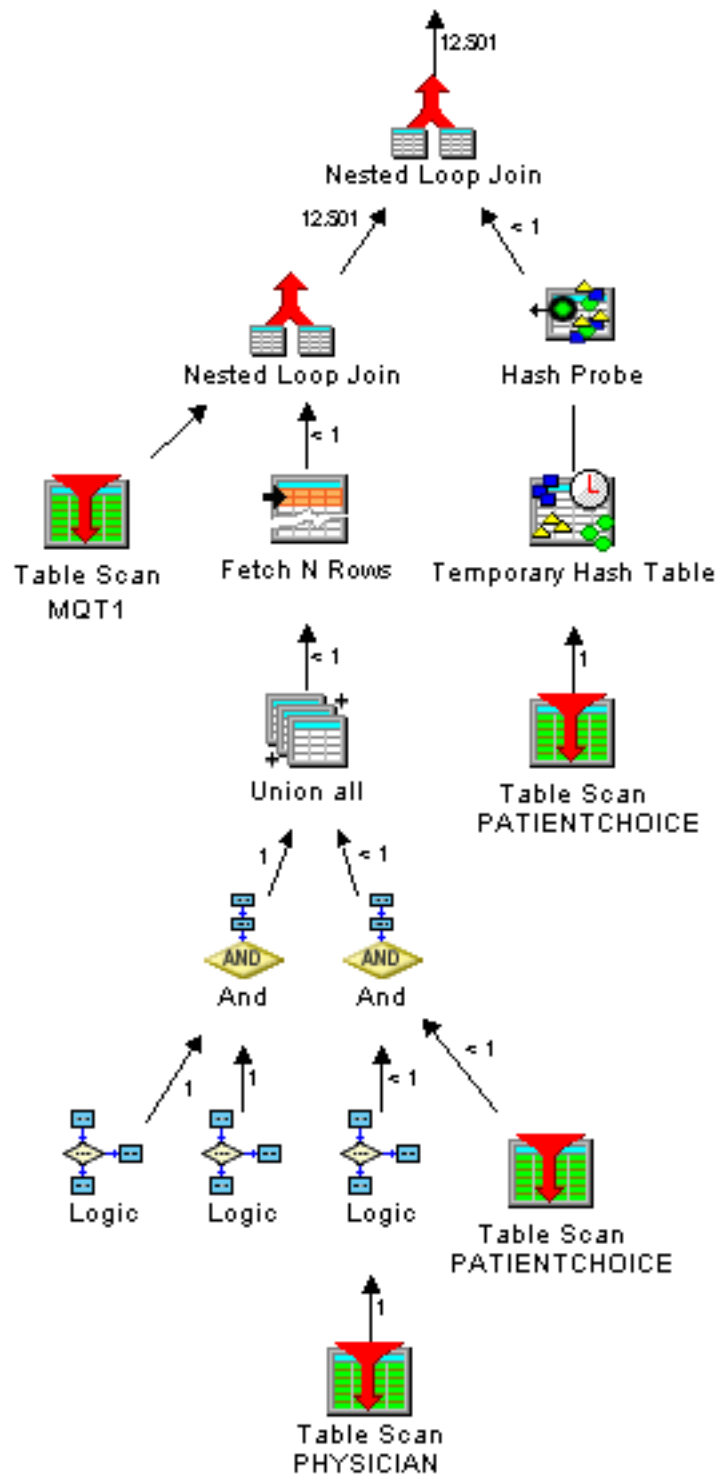
The following query however can be satisfied by the MQT1

```
SELECT patientid, patientname, pharmacy FROM patient WHERE patientid>4 and diagnosis is not null;
```

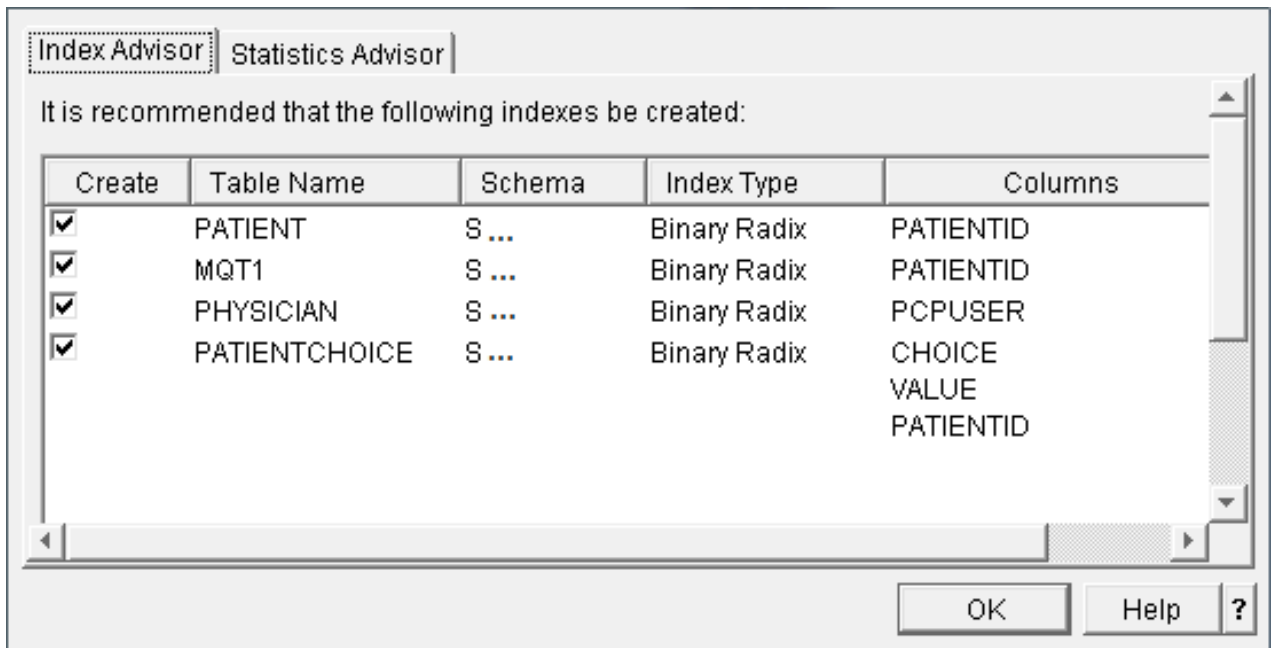
Row and column level access control does not affect the REFRESH TABLE statement. The table is refreshed as if row and column level access controls do not exist.

```
REFRESH TABLE mqt1;
```

The graphic below shows the Visual Explain that reflects the MQT match and substitution. Note that had the MQT1 not surfaced a required value of PCPID for the existing RCAC SQL Rules, it would not be able to satisfy the query request as an MQT match, even though that field is not in the required select list. In this example Visual Explain, you can see the MQT1 substituted but also inherited the RCAC rules of the base table PATIENT.



Index advice of the originating query, which is depicted in the graphic below, includes advice over the main query table, over the MQT and over the RCAC required tables.



As many MQTs, such as the one below, provide ready made aggregation values so aggregating queries in a data warehousing environment perform quickly, these MQTs are now likely not to match query requests with aggregated selection via MQT substitution.

The aggregation is based on the REFRESH TABLE with no RCAC applied and yet the matching is based on the underlying base table and all its RCAC requirements.

```
CREATE TABLE MQT_AGG
AS (SELECT pcpid, count(*) patientcnt
FROM patient group by pcpid
)
DATA INITIALLY IMMEDIATE REFRESH DEFERRED
ENABLE QUERY OPTIMIZATION
MAINTAINED BY USER;
```

The following query, although it appears to be a match for the above MQT_AGG, will not substitute the MQT per RCAC rules.

```
SELECT pcpid, count(*) FROM PATIENT WHERE pcpid in ( 1, ...) GROUP BY pcpid
```

All existing MQTs should be analyzed before deploying RCAC policy on base tables to make sure that performance does not unexpectedly start to suffer because MQTs are no longer available to facilitate the request.

Because most aggregating queries are not dealing with 'details' and so possibly less sensitive to the requirements of RCAC, aggregating MQT over base tables with RCAC might be best deployed by direct substitution in the query and restriction through table privileges and disabling the default RCAC rule, restricting all rows, as follows.

```
ALTER TABLE MQT_AGG DEACTIVATE ROW ACCESS CONTROL;
```

This deactivates the default RCAC applied due to base tables with RCAC and allows direct access to the MQT in a warehousing environment.

Optimizing query performance using query optimization tools

Query optimization is an iterative process. You can gather performance information about your queries and control the processing of your queries.

Db2 for IBM i – Health Center

Use the Db2 for IBM i Health Center to capture information about your database. You can view the total number of objects, the size limits of selected objects, the design limits of selected objects, environmental limits, and activity level.

Navigator view of Health Center

The System i Navigator provides a robust graphical interface to capture, view, and interact with the Health Center.

To start the health center, follow these steps:

1. In the System i Navigator window, expand the system that you want to use.
2. Expand **Databases**.
3. Right-click the database that you want to work with and select **Health Center**.

You can change your preferences by clicking **Change** and entering filter information. Click **Refresh** to update the information.

To save your health center history, do the following:

1. In the System i Navigator window, expand the system you want to use.
2. Expand **Databases**.
3. Right-click the database that you want to work with and select **Health Center**.
4. On the health center dialog, select the area that you want to save. For example, if you want to save the current overview, click **Save** on the Overview tab. Size limits and Design limits are not saved.
5. Specify a schema and table to save the information. You can view the contents of the selected table by clicking **View Contents**. If you select to save information to a table that does not exist, the system creates the table for you.

Health Center SQL procedures

The Health Center is implemented upon several Db2 for i SQL procedures.

IBM i users can call the Health Center SQL procedures directly.

QSYS2.Health_Database_Overview ()

The `QSYS2.Health_Database_Overview()` procedure returns counts of all the different types of Db2 for i objects within the target schema or schemas. The counts are broken down by object type and subtype.

Procedure definition:

```
CREATE PROCEDURE QSYS2.HEALTH_DATABASE_OVERVIEW(  
  IN ARCHIVE_OPTION INTEGER,  
  IN OBJECT_SCHEMA VARCHAR(258),  
  IN NUMBER_OF_ITEMS_ARCHIVE INTEGER,  
  IN OVERVIEW_SCHEMA VARCHAR(258),  
  IN OVERVIEW_TABLE VARCHAR(258))  
  DYNAMIC RESULT SETS 1  
  LANGUAGE C  
  SPECIFIC QSYS2.HEALTH_DATABASE_OVERVIEW  
  NOT DETERMINISTIC  
  MODIFIES SQL DATA  
  CALLED ON NULL INPUT  
  EXTERNAL NAME 'QSYS/QSQHEALTH(OVERVIEW)'  
  PARAMETER STYLE SQL;
```


Service Program Name: QSYS/QSQHEALTH

Default Public Authority: *USE

Threadsafe: Yes

IBM i release

This procedure was added to IBM i in V5R4M0.

Parameters

Archive_Option

(Input) The type of operation to perform for the Db2 for i Health Center overview detail.

The supported values are:

- 1 = Query only, no archive action is taken
- 2 = Archive only
- 3 = Create archive and archive
- 4 = Query the archive

Note: Option 1 produces a new result set. Options 2 and 3 simply use the results from the last Query option. Option 3 fails if the archive exists.

Object_Schema

(Input) The target schema or schemas for this operation. A single schema name can be entered. The '%' character can be used to direct the procedure to process all schemas with names that start with the same characters which appear before the '%'. When this parameter contains only the '%' character, the procedure processes all schemas within the database.

Number_Of_Items_Archive

(Input) The number of rows to archive.

The archive can be used to recognize trends over time. To have meaningful historical comparisons, choose the row count size carefully. This argument is ignored if the Archive_Option is 1.

Overview_Table

(Input) The table that contains the database overview archive.

This argument is ignored if the Archive_Option is 1.

Authorities

To query an existing archive, *USE object authority is required for the Overview_Schema and Overview_Table. To create an archive, *CHANGE object authority is required for the Overview_Schema. To add to an existing archive, *CHANGE object authority is required for the Overview_Table and *USE object authority is required for the Overview_Schema.

Result Set

When Archive_Option is 1 or 4, a single result set is returned.

The format of the result is as follows.

QSYS2.Health_Database_Overview () result set format:

```
"TIMESTAMP" TIMESTAMP NOT NULL DEFAULT CURRENT_TIMESTAMP ,  
SCHEMAS BIGINT NOT NULL ,  
GRP01 CHAR(1) DEFAULT NULL ,  
TABLES BIGINT NOT NULL ,  
PARTITIONED_TABLES FOR COLUMN TABLESRT BIGINT NOT NULL ,  
DISTRIBUTED_TABLES FOR COLUMN TABLES_DST BIGINT NOT NULL ,  
MATERIALIZED_QUERY_TABLES FOR COLUMN TABLES_MAT BIGINT NOT NULL ,
```

```

PHYSICAL_FILES FOR COLUMN TABLESHY BIGINT NOT NULL ,
SOURCE_FILES FOR COLUMN TABLES_SRC BIGINT NOT NULL ,
GRP02 CHAR(1) DEFAULT NULL ,
VIEWS BIGINT NOT NULL ,
LOGICAL_FILES FOR COLUMN VIEWS_LGL BIGINT NOT NULL ,
GRP03 CHAR(1) DEFAULT NULL ,
BINARY_RADIX_INDEXES FOR COLUMN INDEXES_BI BIGINT NOT NULL ,
EVI_INDEXES FOR COLUMN INDEXES_EV BIGINT NOT NULL ,
GRP04 CHAR(1) DEFAULT NULL ,
PRIMARY_KEY_CONSTRAINTS FOR COLUMN CSTSRI BIGINT NOT NULL ,
UNIQUE_CONSTRAINTS FOR COLUMN CSTS_UNQ BIGINT NOT NULL ,
CHECK_CONSTRAINTS FOR COLUMN CSTS_CHK BIGINT NOT NULL ,
REFERENTIAL_CONSTRAINTS FOR COLUMN CSTS_RI BIGINT NOT NULL ,
GRP05 CHAR(1) DEFAULT NULL ,
EXTERNAL_TRIGGERS FOR COLUMN TRGS_EXT BIGINT NOT NULL ,
SQL_TRIGGERS FOR COLUMN TRGS_SQL BIGINT NOT NULL ,
INSTEAD_OF_TRIGGERS FOR COLUMN TRGS_INSTD BIGINT NOT NULL ,
GRP06 CHAR(1) DEFAULT NULL ,
ALIASES BIGINT NOT NULL ,
DDM_FILES BIGINT NOT NULL ,
GRP07 CHAR(1) DEFAULT NULL ,
EXTERNALPROCEDURES FOR COLUMN PROCS_EXT BIGINT NOT NULL ,
SQLPROCEDURES FOR COLUMN PROCS_SQL BIGINT NOT NULL ,
GRP08 CHAR(1) DEFAULT NULL ,
EXTERNAL_SCALAR_FUNCTIONS FOR COLUMN FUNCS_EXTS BIGINT NOT NULL ,
EXTERNAL_TABLE_FUNCTIONS FOR COLUMN FUNCS_EXTT BIGINT NOT NULL ,
SOURCE_SCALAR_FUNCTIONS FOR COLUMN FUNCS_SRCS BIGINT NOT NULL ,
SOURCE_AGGREGATE_FUNCTIONS FOR COLUMN FUNCS_SRCA BIGINT NOT NULL ,
SQL_SCALAR_FUNCTIONS FOR COLUMN FUNCS_SQLS BIGINT NOT NULL ,
SQL_TABLE_FUNCTIONS FOR COLUMN FUNCS_SQLT BIGINT NOT NULL ,
GRP09 CHAR(1) DEFAULT NULL ,
SEQUENCES BIGINT NOT NULL ,
SQLACKAGES FOR COLUMN SQLPKGS BIGINT NOT NULL ,
USER_DEFINED_DISTINCT_TYPES FOR COLUMN UDTS BIGINT NOT NULL ,
JOURNALS BIGINT NOT NULL ,
JOURNAL_RECEIVERS FOR COLUMN JRNRCV BIGINT NOT NULL ,
"SCHEMA" VARCHAR(258) ALLOCATE(10) NOT NULL

```

```

LABEL ON COLUMN <result set>
( "TIMESTAMP" IS 'Timestamp' ,
  SCHEMAS IS 'Schemas' ,
  GRP01 IS 'Tables' ,
  TABLES IS 'Non-partitioned tables' ,
  PARTITIONED_TABLES IS 'Partitioned tables' ,
  DISTRIBUTED_TABLES IS 'Distributed tables' ,
  MATERIALIZED_QUERY_TABLES IS 'Materialized query tables' ,
  PHYSICAL_FILES IS 'Physical files' ,
  SOURCE_FILES IS 'Source files' ,
  GRP02 IS 'Views' ,
  VIEWS IS 'Views' ,
  LOGICAL_FILES IS 'Logical files' ,
  GRP03 IS 'Indexes' ,
  BINARY_RADIX_INDEXES IS 'Binary radix indexes' ,
  EVI_INDEXES IS 'Encoded vector indexes' ,
  GRP04 IS 'Constraints' ,
  PRIMARY_KEY_CONSTRAINTS IS 'PRIMARY KEY constraints' ,
  UNIQUE_CONSTRAINTS IS 'UNIQUE constraints' ,
  CHECK_CONSTRAINTS IS 'CHECK constraints' ,
  REFERENTIAL_CONSTRAINTS IS 'Referential constraints' ,
  GRP05 IS 'Triggers' ,
  EXTERNAL_TRIGGERS IS 'External triggers' ,
  SQL_TRIGGERS IS 'SQL triggers' ,
  INSTEAD_OF_TRIGGERS IS 'INSTEAD OF triggers' ,
  GRP06 IS 'Aliases' ,
  ALIASES IS 'Aliases' ,
  DDM_FILES IS 'DDM files' ,
  GRP07 IS 'Procedures' ,
  EXTERNALPROCEDURES IS 'External procedures' ,
  SQLPROCEDURES IS 'SQL procedures' ,
  GRP08 IS 'Functions' ,
  EXTERNAL_SCALAR_FUNCTIONS IS 'External scalar functions' ,
  EXTERNAL_TABLE_FUNCTIONS IS 'External table functions' ,
  SOURCE_SCALAR_FUNCTIONS IS 'Source scalar functions' ,
  SOURCE_AGGREGATE_FUNCTIONS IS 'Source aggregate functions' ,
  SQL_SCALAR_FUNCTIONS IS 'SQL scalar functions' ,
  SQL_TABLE_FUNCTIONS IS 'SQL table functions' ,
  GRP09 IS 'Miscellaneous' ,
  SEQUENCES IS 'Sequences' ,
  SQLACKAGES IS 'SQL packages' ,
  USER_DEFINED_DISTINCT_TYPES IS 'User-defined distinct types' ,
  JOURNALS IS 'Journals' ,

```

```
JOURNAL_RECEIVERS IS 'Journal          receivers' ,  
"SCHEMA" IS 'Schema          mask' ) ;
```

Error Messages

Table 34. Error messages	
Message ID	Error Message Text
SQL0462 W	This warning appears in the job log if the procedure encounters objects for which the user does not have *USE object authority. The warning is provided as an indication that the procedure was unable to process all available objects.

Usage Notes

None

Related Information

None

Examples

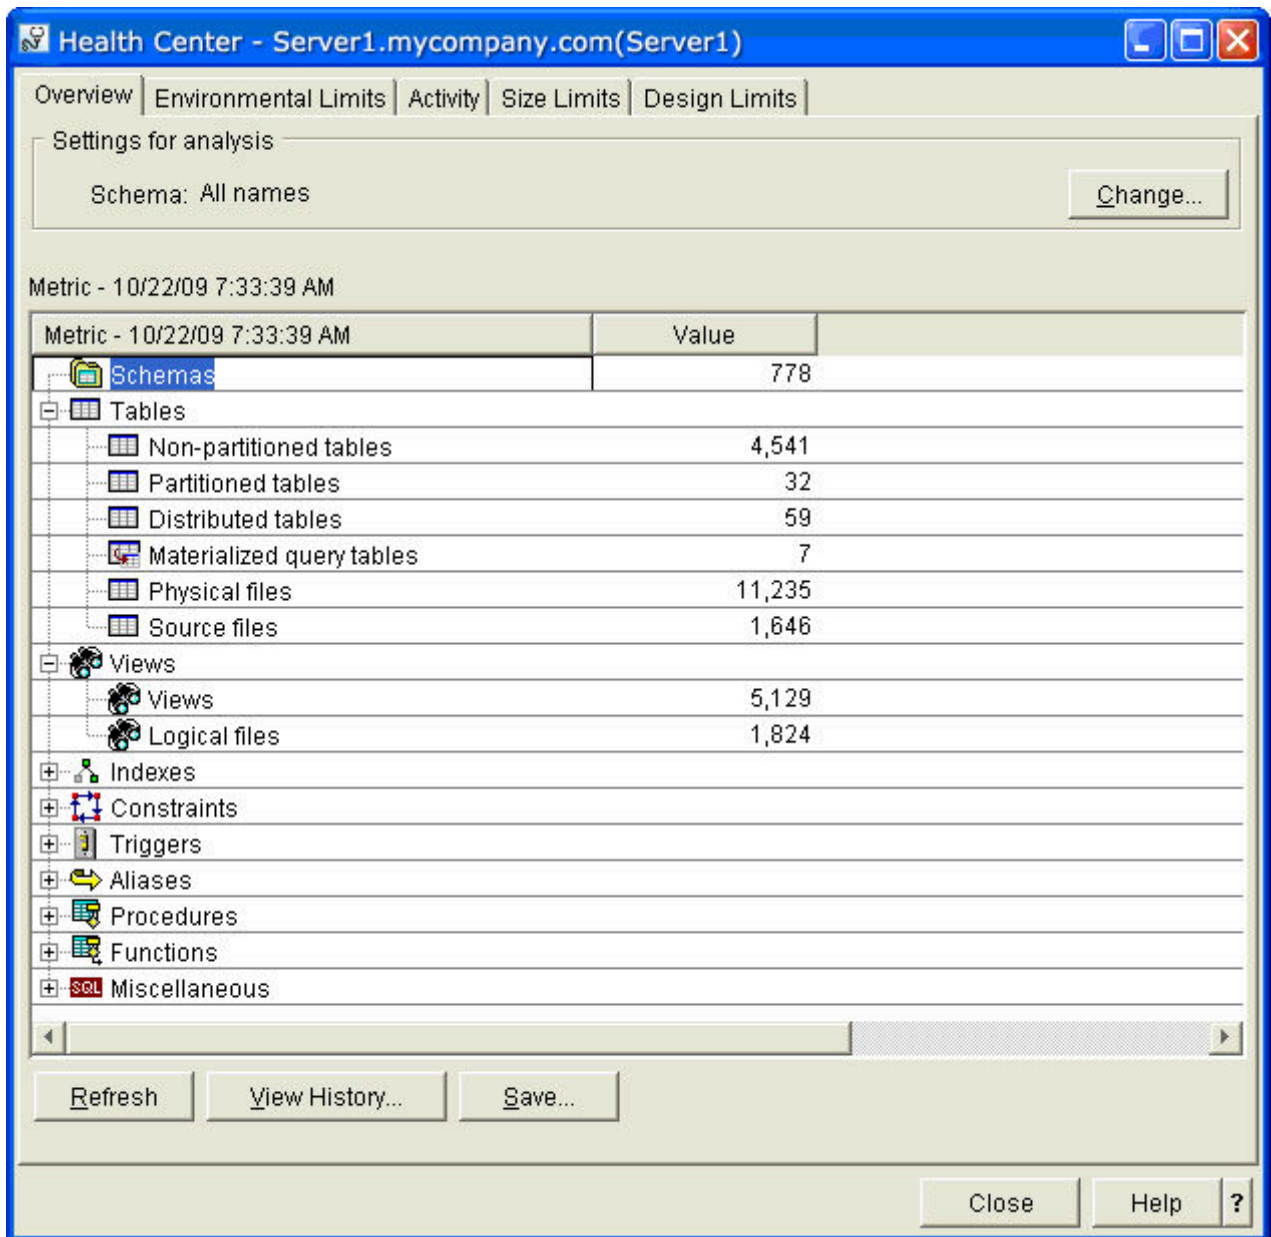
Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

Example 1

Retrieve the overview for the entire database.

```
CALL QSYS2.Health_Database_Overview(1, '%', NULL, NULL, NULL);
```

Example results in System i Navigator:



Example 2

Archive all rows in the overview to an SQL table named MYLIB/ARCHIVE1.

```
CALL QSYS2.Health_Database_Overview(3, '%', 2147483647, 'MYLIB', 'ARCHIVE1')
```

Example 3

Retrieve the overview from MYLIB/ARCHIVE1.

```
CALL QSYS2.Health_Database_Overview(4, '%', NULL, 'MYLIB', 'ARCHIVE1')
```

Example results in System i Navigator:

History table: MYLIB.ARCHIVE1

Metric

Metric	Value
Collections	10/22/09 7:58:39 AM
Schemas	778
Tables	
Non-partitioned tables	4,542
Partitioned tables	32
Distributed tables	59
Materialized query tables	7
Physical files	11,235
Source files	1,646
Views	
Views	5,129
Logical files	1,824
Indexes	

Refresh Select Collections to View... Delete Collections... Close Help ?

QSYS2.Health_Activity ()

The QSYS2.Health_Activity () procedure returns summary counts of database and SQL operations over a set of objects within one or more schemas.

Procedure definition:

```
CREATE PROCEDURE QSYS2.HEALTH_ACTIVITY(
  IN ARCHIVE_OPTION INTEGER,
  IN REFRESH_CURRENT_VALUES INTEGER,
  IN OBJECT_SCHEMA VARCHAR(258),
  IN OBJECT_NAME VARCHAR(258),
  IN NUMBER_OBJECTS_ACTIVITY_TO_ARCHIVE INTEGER,
  IN NUMBER_OF_ACTIVITY_ARCHIVE INTEGER,
  IN ACTIVITY_SCHEMA VARCHAR(258),
  IN ACTIVITY_TABLE VARCHAR(258))
  DYNAMIC RESULT SETS 1
  LANGUAGE C
  SPECIFIC QSYS2.HEALTH_ACTIVITY
  NOT DETERMINISTIC
  MODIFIES SQL DATA
  CALLED ON NULL INPUT
  EXTERNAL NAME 'QSYS/QSQHEALTH(ACTIVITY) '
  PARAMETER STYLE SQL;
```

Service Program Name: QSYS/QSQHEALTH

Default Public Authority: *USE

Threadsafe: Yes

IBM i release

This procedure was added to IBM i 6.1.

Parameters

Archive_Option

(Input) The type of operation to perform for the Db2 for i Health Center overview detail.

The supported values are:

- 1 = Query only, no archive action is taken
- 2 = Archive only
- 3 = Create archive and archive
- 4 = Query the archive

Note: Option 1 produces a new result set. Options 2 and 3 simply use the results from the last Query option. Option 3 fails if the archive exists.

Refresh_Current_Values

(Input) This option directs how the archive operation is done. This option is only valid with archive options 2 and 3.

The supported values are:

- 0 = No. Indicates that we capture the activity on the entire set of specified schemas and objects.
- 1 = Yes. Indicates that we only refresh the activity of the objects previously captured (based on the short names).
- 2 = None. Use the results from the prior call. A call must have been performed in this job before using this option

Object_Schema

(Input) The target schema or schemas for this operation. A single schema name can be entered. The '%' character can be used to direct the procedure to process all schemas with names that start with the same characters which appear before the '%'. When this parameter contains only the '%' character, the procedure processes all schemas within the database.

This name also affects the items refreshed if Refresh_Current_Values = 1.

Object_Name

(Input) The target object name for this operation. Only the '%' character is treated as a wildcard since an underscore is a valid character in a name. The name must be delimited, if necessary, and case sensitive.

This name also affects the items refreshed if Refresh_Current_Values = 1.

Number_Objects_Activity_to_Archive (Input) The number of objects to save for each activity.

Number_Of_Activity_Archive (Input) The number of rows to save per object activity.

The archive can be used to recognize trends over time. To have meaningful historical comparisons, choose the row count size carefully. This argument is ignored if the Archive_Option is 1 or 4.

Activity_Schema

(Input) The table that contains the database activity archive.

This argument is ignored if the Archive_Option is 1.

Activity_Table

The table that contains the database activity archive.

This argument is ignored if the Archive_Option is 1.

Authorities

To query an existing archive, *USE object authority is required for the Activity_Schema and Activity_Table. To create an archive, *CHANGE object authority is required for the Activity_Schema. To add to an existing archive, *CHANGE object authority is required for the Activity_Table and *USE object authority is required for the Activity_Schema.

When Archive_Option is 1 or 3, *USE object authority is required for the Object_Schema and for any objects which are indicated by Object_Name. When an object is encountered and the caller does not have *USE object authority, an SQL0462 warning is placed in the job log. The object is skipped and not included in the procedure result set.

Result Set

When Archive_Option is 1 or 4, a single result set is returned.

The format of the result is as follows. All these items were added for IBM i 6.1.

QSYS2.Health_Activity() result set format:

```
"TIMESTAMP" TIMESTAMP NOT NULL,
ACTIVITY VARCHAR(2000) ALLOCATE(20) DEFAULT NULL,
CURRENT_VALUE FOR COLUMN "VALUE" BIGINT DEFAULT NULL,
OBJECT_SCHEMA FOR COLUMN BSHEMA VARCHAR(128)ALLOCATE(10) DEFAULT NULL,
OBJECT_NAME FOR COLUMN BNAME VARCHAR(128) ALLOCATE(20) DEFAULT NULL,
OBJECT_TYPE FOR COLUMN BTYPE VARCHAR(24) ALLOCATE(10) DEFAULT NULL,
SYSTEM_OBJECT_SCHEMA FOR COLUMN SYS_DNAME VARCHAR(10) ALLOCATE(10)DEFAULT NULL,
SYSTEM_OBJECT_NAME FOR COLUMN SYS_ONAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
PARTITION_NAME FOR COLUMN MBRNAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
ACTIVITY_ID FOR COLUMN ACTIV00001 INTEGER DEFAULT NULL
```

```
LABEL ON COLUMN <result set>
("TIMESTAMP" IS 'Timestamp',
ACTIVITY IS 'Activity',
CURRENT_VALUE IS 'Current Value',
OBJECT_SCHEMA IS 'Object Schema',
OBJECT_NAME IS 'Object Name',
OBJECT_TYPE IS 'Object Type',
SYSTEM_OBJECT_SCHEMA IS 'System Object Schema',
SYSTEM_OBJECT_NAME IS 'System Object Name',
PARTITION_NAME IS 'Partition Name',
ACTIVITY_ID IS 'Activity ID');
```

Limit Detail

The supported Database Health Center Activity can be seen on any machine by executing this query. The supported value column contains zeros because this category of Health Center information is not tied to a limit.

```
SELECT * FROM QSYS2.SQL_SIZING WHERE SIZING_ID BETWEEN 18000 AND 18199;
```

Note: The **bold** rows were added in IBM i 7.1.

<i>Table 35. Summary counts of database and SQL operations within a schema.</i>		
SIZING_ID	SIZING_NAME	SUPPORTED_VALUE
18100	INSERT OPERATIONS	0
18101	UPDATE OPERATIONS	0
18102	DELETE OPERATIONS	0
18103	LOGICAL READS	0
18104	PHYSICAL READS	0
18105	CLEAR OPERATIONS	0

Table 35. Summary counts of database and SQL operations within a schema. (continued)

SIZING_ID	SIZING_NAME	SUPPORTED_VALUE
18106	INDEX BUILDS/REBUILDS	0
18107	DATA SPACE REORGANIZE OPERATIONS	0
18108	DATA SPACE COPY OPERATIONS	0
18109	FULL OPENS	0
18110	FULL CLOSES	0
18111	DAYS USED	0
18112	INDEX QUERY USE	0
18113	INDEX QUERY STATISTICS USE	0
18114	INDEX LOGICAL READS	0
18115	INDEX RANDOM READS	0
18116	SQL STATEMENT COMPRESSION COUNT	0
18117	SQL STATEMENT CONTENTION COUNT	0
18118	RANDOM READS	0
18119	SEQUENTIAL READS	0

Error Messages

Table 36. Error messages

Message ID	Error Message Text
SQL0462 W	This warning appears in the job log if the procedure encounters objects for which the user does not have *USE object authority. The warning is provided as an indication that the procedure was unable to process all available objects.

Usage Notes

None

Related Information

None

Example

Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

Retrieve the activity information for all objects within the QSYS2 schema, using a maximum of 10 objects per each activity.

```
CALL QSYS2.Health_Activity(1, 0, 'QSYS2', '%', 10, NULL, NULL, NULL);
```

Example results in System i Navigator:

Health Center - Server1(Server1)

Overview | Environmental Limits | Activity | Size Limits | Design Limits

Settings for analysis

Use the following filters

Schema: QSYS2

Object: All names Change...

Objects per activity: 10

Use the following history file

History file:

Activity - 10/26/09 10:18:18 AM

Activity - 10/26/09 10:18:18 AM	Value	Status
+	Insert operations	
+	Update operations	
+	Delete operations	
+	Logical reads	
+	Physical reads	
+	Clear operations	
+	Full opens	
+	Full closes	
+	Days used	
-	Index query use	
SQL	QSYS2.SYSTXTINDI (SYSTE00001)	21,155 ■ Normal
SQL	QSYS2.SYSTXTINDI (SYSTE00001)	14,756 ■ Normal
SQL	QSYS2.QASQRESL (SYSRO00001)	10,109 ■ Normal
SQL	QSYS2.SYSTXTCOLI (SYSTE00001)	9,872 ■ Normal
SQL	QSYS2.QASQSPDP (QASQDRDP)	9,818 ■ Normal
SQL	QSYS2.QASQDRDP (QASQDRDP)	9,818 ■ Normal
SQL	QSYS2.QASEQOBJ (SYSSEQOBJ)	9,724 ■ Normal
SQL	QSYS2.SYSTXTINDI (SYSTE00001)	7,796 ■ Normal
SQL	QSYS2.QASEQOBJ (QASEQOBJ)	814 ■ Normal
SQL	QSYS2.SYSTXTCOLI (SYSTXTCOLI)	439 ■ Normal
+	Index query statistics use	
+	Index logical reads	

Refresh | View History... | Save... | Change Status Threshold...

Close | Help | ?

QSYS2.Health_Design_Limits ()

The QSYS2.Health_Design_Limits () procedure returns detailed counts of design limits over a set of objects within one or more schemas. Design limits correspond to architectural constructs, such as 'Maximum number of columns in a table or view'.

Procedure definition:

```
CREATE PROCEDURE QSYS2.HEALTH_DESIGN_LIMITS(  
  ARCHIVE_OPTION INTEGER,  
  IN REFRESH_CURRENT_VALUES INTEGER,  
  IN OBJECT_SCHEMA VARCHAR(258),  
  IN OBJECT_NAME VARCHAR(258),  
  IN NUMBER_OBJECTS_LIMIT_TO_ARCHIVE INTEGER,  
  IN NUMBER_OF_LIMITS_ARCHIVE INTEGER,  
  IN LIMIT_SCHEMA VARCHAR(258),  
  IN LIMIT_TABLE VARCHAR(258),  
  DYNAMIC RESULT SETS 1  
  LANGUAGE C  
  SPECIFIC QSYS2.HEALTH_DESIGN_LIMITS  
  NOT DETERMINISTIC  
  MODIFIES SQL DATA  
  CALLED ON NULL INPUT  
  EXTERNAL NAME 'QSYS/QSQHEALTH(DESIGN)'  
  PARAMETER STYLE SQL;
```

Service Program Name: QSYS/QSQHEALTH

Default Public Authority: *USE

Threadsafe: Yes

IBM i release

This procedure was added to IBM i V5R4M0.

Parameters

Archive_Option

(Input) The type of operation to perform for the Db2 for i Health Center activity detail.

The supported values are:

- 1 = Query only, no archive action is taken
- 2 = Archive only
- 3 = Create archive and archive
- 4 = Query the archive

Note: Option 1 produces a new result set. Options 2 and 3 simply use the results from the last Query option. Option 3 fails if the archive exists.

Refresh_Current_Values

(Input) This option directs how the archive operation is done. This option is only valid with archive options 2 and 3.

The supported values are:

- 0 = No. Indicates that we capture the activity on the entire set of specified schemas and objects.
- 1 = Yes. Indicates that we only refresh the activity of the objects previously captured (based on the short names).
- 2 = None. Use the results from the prior call. A call must have been performed in this job before using this option

Object_Schema	(Input) The target schema or schemas for this operation. A single schema name can be entered. The '%' character can be used to direct the procedure to process all schemas with names that start with the same characters which appear before the '%'. When this parameter contains only the '%' character, the procedure processes all schemas within the database. This name also affects the items refreshed if Refresh_Current_Values = 1.
Object_Name	(Input) The target object name for this operation. Only the '%' character is treated as a wildcard since an underscore is a valid character in a name. The name must be delimited, if necessary, and case sensitive. This name also affects the items refreshed if Refresh_Current_Values = 1.
Number_Objects_Limit_to_Archive	(Input) The number of objects to save for each design limit.
Number_Of_Limits_Archive	(Input) The number of rows to save per object design limit. The archive can be used to recognize trends over time. To have meaningful historical comparisons, choose the row count size carefully. This argument is ignored if the Archive_Option is 1 or 4.
Limit_Schema	(Input) The schema that contains the database limit archive. This argument is ignored if the Archive_Option is 1.
Limit_Table	The table that contains the database limit archive. This argument is ignored if the Archive_Option is 1.

Authorities

To query an existing archive, *USE object authority is required for the Limit_Schema and Limit_Table. To create an archive, *CHANGE object authority is required for the Limit_Schema. To add to an archive, *CHANGE object authority is required for the Limit_Table.

When Archive_Option is 1 or 3, *USE object authority is required for the Object_Schema and for any objects which are indicated by Object_Name. When an object is encountered and the caller does not have *USE object authority, an SQL0462 warning is placed in the job log. The object is skipped and not included in the procedure result set.

Result Set

When Archive_Option is 1 or 4, a single result set is returned.

The format of the result is as follows. All these items were added for IBM i V5R4M0.

QSYS2.Health_Design_Limits() result set format:

```
"TIMESTAMP" TIMESTAMP NOT NULL,
LIMIT VARCHAR(2000) ALLOCATE(20) DEFAULT NULL,
CURRENT_VALUE FOR COLUMN "VALUE" BIGINT DEFAULT NULL,
PERCENT_DECIMAL(5, 2) DEFAULT NULL,
OBJECT_SCHEMA FOR COLUMN BSHEMA VARCHAR(128) ALLOCATE(10) DEFAULT NULL,
OBJECT_NAME FOR COLUMN BNAME VARCHAR(128) ALLOCATE(20) DEFAULT NULL,
OBJECT_TYPE FOR COLUMN BTYPE VARCHAR(24) ALLOCATE(10) DEFAULT NULL,
SYSTEM_OBJECT_SCHEMA FOR COLUMN SYS_DNAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
SYSTEM_OBJECT_NAME FOR COLUMN SYS_ONAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
PARTITION_NAME FOR COLUMN MBRNAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL
```

```

MAXIMUM_VALUE FOR COLUMN "MAXVALUE" BIGINT DEFAULT NULL
LIMIT_ID INTEGER DEFAULT NULL

```

```

LABEL ON COLUMN <result set>
( "TIMESTAMP" IS 'Timestamp',
  LIMIT IS 'Limit',
  CURRENT_VALUE IS 'Current Value',
  PERCENT IS 'Percent',
  OBJECT_SCHEMA IS 'Object Schema',
  OBJECT_NAME IS 'Object Name',
  OBJECT_TYPE IS 'Object Type',
  SYSTEM_OBJECT_SCHEMA IS 'System Object Schema',
  SYSTEM_OBJECT_NAME IS 'System Object Name',
  PARTITION_NAME IS 'Partition Name',
  MAXIMUM_VALUE IS 'Maximum Value',
  LIMIT_ID IS 'Limit ID');

```

Limit Detail

The supported Database Health Center Design limits can be seen on any machine by executing this query:

```
SELECT * FROM QSYS2.SQL_SIZING WHERE SIZING_ID BETWEEN 16000 AND 16999;
```

Table 37. Design limits over objects within a schema.

SIZING_ID	SIZING_NAME	SUPPORTED_VALUE
16100	MAXIMUM NUMBER OF MEMBERS	327670
16101	MAXIMUM NUMBER OF RECORD FORMATS	32
16800	MAXIMUM JOURNAL RECEIVER SIZE	1.09951E+12 (~1 TB)
16801	TOTAL SQL STATEMENTS	0
16802	TOTAL ACTIVE SQL STATEMENTS	0
16803	MAXIMUM SQL PACKAGE SIZE	520093696 (~500 MB)
16804	MAXIMUM LARGE SQL PACKAGE SIZE	1056964608 (~1 GB)
16805	MAXIMUM SQL PROGRAM ASSOCIATED SPACE SIZE	16777216

Error Messages

Table 38. Error messages

Message ID	Error Message Text
SQL0462 W	This warning appears in the job log if the procedure encounters objects for which the user does not have *USE object authority. The warning is provided as an indication that the procedure was unable to process all available objects.

Usage Notes

None

Related Information

None

Example

Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

Retrieve the design limit information for all object names which start with the letter R, within the SYSIBM schema, using a maximum of 20 objects per each design limit.

```
CALL QSYS2.Health_Design_Limits(1, 0, 'SYSIBM', 'R%', 20, NULL, NULL, NULL);
```

Example results in System i Navigator:

The screenshot shows the 'Design Limits' tab in the Health Center application. The settings for analysis are: Schema: SYSIBM, Object: R%, Objects per design limit: 20. The table below shows the results for various design limits.

Design Limit - 10/22/09 10:37:44 AM	Value	Percent of Limit	Status
Maximum number of columns in a table or view (8,000)			
SYSIBM.ROUTINES_S	82	1.02	Normal
SYSIBM.ROUTINES	82	1.02	Normal
SYSIBM.REFERENTIAL_CONSTRAINTS	9	0.11	Normal
SYSIBM.REF_CONSTRAINTS	9	0.11	Normal
Maximum length of a row (32 KB)			
SYSIBM.ROUTINES_S	9.68 KB	30.24	Normal
SYSIBM.ROUTINES	9.68 KB	30.24	Normal
SYSIBM.REFERENTIAL_CONSTRAINTS	815 bytes	2.48	Normal
SYSIBM.REF_CONSTRAINTS	815 bytes	2.48	Normal
Maximum row length with LOBs (3.5 GB)			
Maximum number of tables referenced in a view or logical (256)			
Maximum number of members (32,767)			

QSYS2.Health_Size_Limits ()

The QSYS2.Health_Size_Limits () procedure returns detailed size information for database objects within one or more schemas. Size limits help you understand trends towards reaching a database limit such as 'Maximum size of the data in a table partition'.

Procedure definition:

```
CREATE PROCEDURE QSYS2.HEALTH_SIZE_LIMITS(
  IN ARCHIVE_OPTION INTEGER,
  IN REFRESH_CURRENT_VALUES INTEGER,
  IN OBJECT_SCHEMA VARCHAR(258),
  IN OBJECT_NAME VARCHAR(258),
  IN NUMBER_OBJECTS_LIMIT_TO_ARCHIVE INTEGER,
  IN NUMBER_OF_LIMITS_ARCHIVE INTEGER,
  IN LIMIT_SCHEMA VARCHAR(258),
  IN LIMIT_TABLE VARCHAR(258))
  DYNAMIC RESULT SETS 1
  LANGUAGE C
  SPECIFIC QSYS2.HEALTH_SIZE_LIMITS
```

```
NOT DETERMINISTIC
MODIFIES SQL DATA
CALLED ON NULL INPUT
EXTERNAL NAME 'QSYS/QSQHEALTH(SIZE) '
PARAMETER STYLE SQL;
```

Service Program Name: QSYS/QSQHEALTH

Default Public Authority: *USE

Threadsafe: Yes

IBM i release

This procedure was added to IBM i V5R4M0.

Parameters

Archive_Option

(Input) The type of operation to perform for the Db2 for i Health Center activity detail.

The supported values are:

- 1 = Query only, no archive action is taken
- 2 = Archive only
- 3 = Create archive and archive
- 4 = Query the archive

Note: Option 1 produces a new result set. Options 2 and 3 simply use the results from the last Query option. Option 3 fails if the archive exists.

Refresh_Current_Values

(Input) This option directs how the archive operation is done. This option is only valid with archive options 2 and 3.

The supported values are:

- 0 = No. Indicates that we capture the activity on the entire set of specified schemas and objects.
- 1 = Yes. Indicates that we only refresh the activity of the objects previously captured (based on the short names).
- 2 = None. Use the results from the prior call. A call must have been performed in this job before using this option

Object_Schema

(Input) The target schema or schemas for this operation. A single schema name can be entered. The '%' character can be used to direct the procedure to process all schemas with names that start with the same characters which appear before the '%'. When this parameter contains only the '%' character, the procedure processes all schemas within the database.

This name also affects the items refreshed if Refresh_Current_Values = 1.

Object_Name

(Input) The target object name for this operation. Only the '%' character is treated as a wildcard since an underscore is a valid character in a name. The name must be delimited, if necessary, and case sensitive.

This name also affects the items refreshed if Refresh_Current_Values = 1.

Number_Objects_Limit_to_Archive	(Input) The number of objects to save for each size limit.
Number_Of_Limits_Archive	(Input) The number of rows to save per object size limit. The archive can be used to recognize trends over time. To have meaningful historical comparisons, choose the row count size carefully. This argument is ignored if the Archive_Option is 1 or 4.
Limit_Schema	(Input) The schema that contains the database activity archive. This argument is ignored if the Archive_Option is 1.
Limit_Table	The table that contains the database activity archive. This argument is ignored if the Archive_Option is 1.

Authorities

To query an existing archive, *USE object authority is required for the Limit_Schema and Limit_Table. To create an archive, *CHANGE object authority is required for the Limit_Schema. To add to an archive, *CHANGE object authority is required for the Limit_Table.

When Archive_Option is 1 or 3, *USE object authority is required for the Object_Schema and for any objects which are indicated by Object_Name. When an object is encountered and the caller does not have *USE object authority, an SQL0462 warning is placed in the job log. The object is skipped and not included in the procedure result set.

Result Set

When Archive_Option is 1 or 4, a single result set is returned.

The format of the result is as follows.

QSYS2.Health_Size_Limits() result set format:

```
"TIMESTAMP" TIMESTAMP NOT NULL,
LIMIT VARCHAR(2000) ALLOCATE(20) DEFAULT NULL,
CURRENT_VALUE FOR COLUMN "VALUE" BIGINT DEFAULT NULL,
PERCENT_DECIMAL(5, 2) DEFAULT NULL, OBJECT_SCHEMA FOR COLUMN BSHEMA VARCHAR(128) ALLOCATE(10) DEFAULT
NULL,
OBJECT_NAME FOR COLUMN BNAME VARCHAR(128) ALLOCATE(20) DEFAULT NULL,
OBJECT_TYPE FOR COLUMN BTYPE VARCHAR(24) ALLOCATE(10) DEFAULT NULL,
SYSTEM_OBJECT_SCHEMA FOR COLUMN SYS_DNAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
SYSTEM_OBJECT_NAME FOR COLUMN SYS_ONAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
MAXIMUM_VALUE FOR COLUMN "MAXVALUE" BIGINT DEFAULT NULL,
LIMIT_ID INTEGER DEFAULT NULL,
PARTITION_NAME FOR COLUMN MBRNAME VARCHAR(10) ALLOCATE(10) DEFAULT NULL,
"SCHEMA" VARCHAR(258) ALLOCATE(10) DEFAULT NULL,
OBJECT VARCHAR(258) ALLOCATE(10) DEFAULT NULL,
"REFRESH" INTEGER DEFAULT NULL
```

```
LABEL ON COLUMN <result set>
( "TIMESTAMP" IS 'Timestamp',
  LIMIT IS 'Limit',
  CURRENT_VALUE IS 'Current Value',
  PERCENT IS 'Percent',
  OBJECT_SCHEMA IS 'Object Schema',
  OBJECT_NAME IS 'Object Name',
  OBJECT_TYPE IS 'Object Type',
  SYSTEM_OBJECT_SCHEMA IS 'System Object Schema',
  SYSTEM_OBJECT_NAME IS 'System Object Name',
  MAXIMUM_VALUE IS 'Maximum Value',
  LIMIT_ID IS 'Limit ID',
  PARTITION_NAME IS 'Partition Name',
  "SCHEMA" IS 'Schema Mask',
  OBJECT IS 'Object Mask',
  "REFRESH" IS 'Refresh');
```

Limit Detail

The supported Database Health Center Size limits can be seen on any machine by executing this query:

```
SELECT * FROM QSYS2.SQL_SIZING WHERE SIZING_ID BETWEEN 15000 AND 15999;
```

Table 39. Size limit information for database objects within a schema.

SIZING_ID	SIZING_NAME	SUPPORTED_VALUE
15000	MAXIMUM NUMBER OF ALL ROWS	4.29E+09
15001	MAXIMUM NUMBER OF VALID ROWS	4.29E+09
15002	MAXIMUM NUMBER OF DELETED ROWS	4.29E+09
15003	MAXIMUM TABLE PARTITION SIZE	1.7E+12
15004	MAXIMUM NUMBER OF OVERFLOW ROWS	4.29E+09
15101	MAXIMUM ROW LENGTH	32766
15102	MAXIMUM ROW LENGTH WITH LOBS	3.76E+09
15103	MAXIMUM NUMBER OF PARTITIONS	256
15150	MAXIMUM NUMBER OF REFERENCED TABLES	256
15300	MAXIMUM NUMBER OF TRIGGERS	300
15301	MAXIMUM NUMBER OF CONSTRAINTS	300
15302	MAXIMUM LENGTH OF CHECK CONSTRAINT	2097151
15400	MAXIMUM *MAX4GB INDEX SIZE	4.29E+09
15401	MAXIMUM *MAX1TB INDEX SIZE	1.7E+12
15402	MAXIMUM NUMBER OF INDEX ENTRIES	0
15500	MAXIMUM KEY COLUMNS	120
15501	MAXIMUM KEY LENGTH	32767
15502	MAXIMUM NUMBER OF PARTITIONING KEYS	120
15700	MAXIMUM NUMBER OF FUNCTION PARAMETERS	2000
15701	MAXIMUM NUMBER OF PROCEDURE PARAMETERS	2000

Error Messages

Table 40. Error messages

Message ID	Error Message Text
SQL0462 W	This warning appears in the job log if the procedure encounters objects for which the user does not have *USE object authority. The warning is provided as an indication that the procedure was unable to process all available objects.

Usage Notes

None

Related Information

None

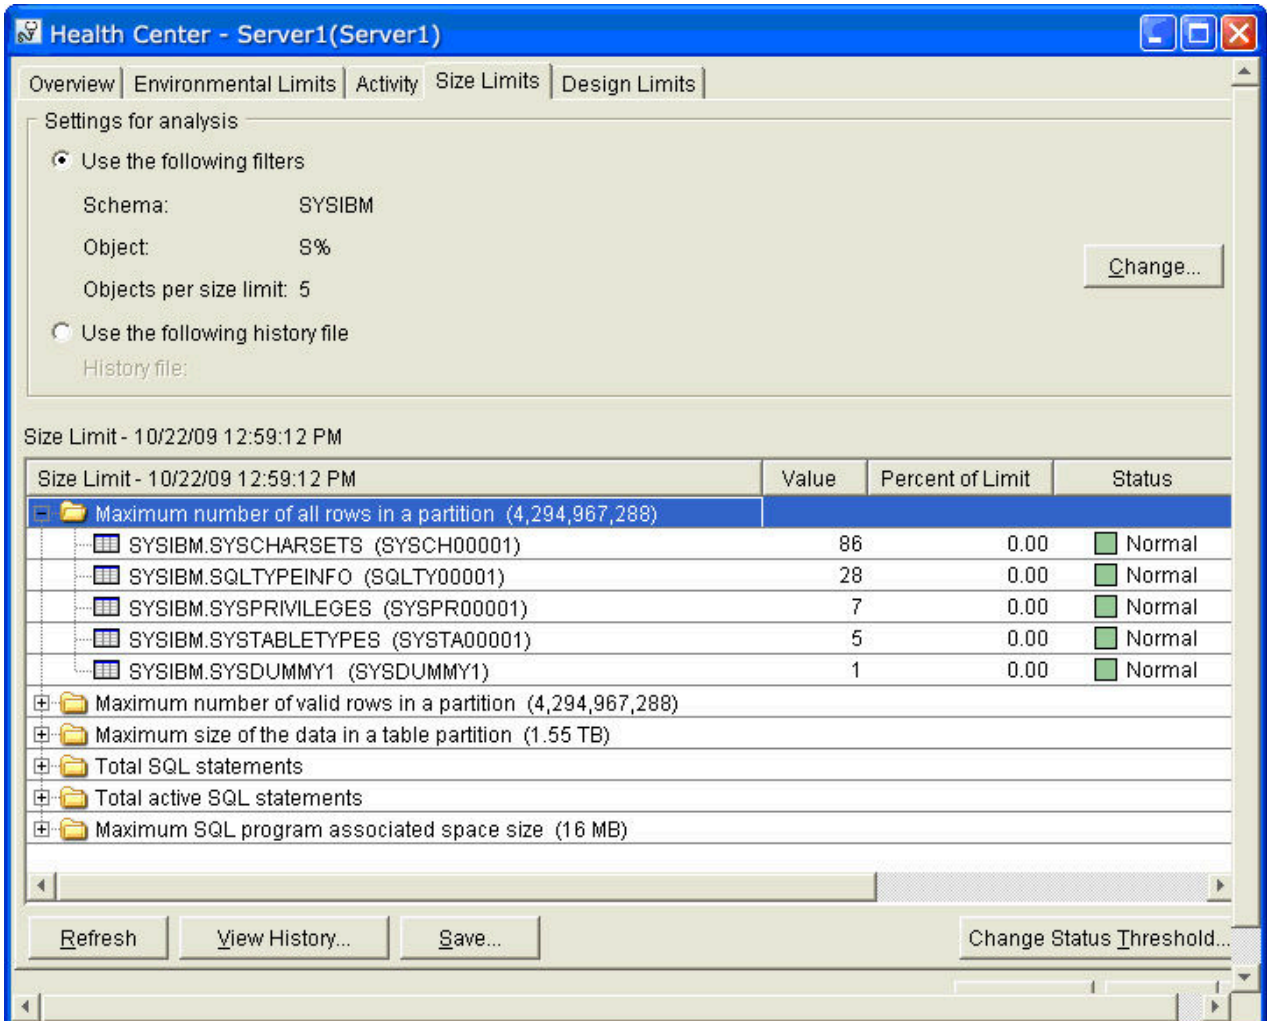
Example

Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

Retrieve the size limit information for all object names which start with the letter S, within the SYSIBM schema, using a maximum of five objects per each design limit.

```
CALL QSYS2.Health_Size_Limits(1, 0, 'SYSIBM', 'S%', 5, NULL, NULL, NULL);
```

Example results in System i Navigator:



Size Limit - 10/22/09 12:59:12 PM	Value	Percent of Limit	Status
Maximum number of all rows in a partition (4,294,967,288)			
SYSIBM.SYSCHARSETS (SYSCH00001)	86	0.00	Normal
SYSIBM.SQLTYPINFO (SQLTY00001)	28	0.00	Normal
SYSIBM.SYSPRIVILEGES (SYSPR00001)	7	0.00	Normal
SYSIBM.SYSTABLETYPES (SYSTA00001)	5	0.00	Normal
SYSIBM.SYSDUMMY1 (SYSDUMMY1)	1	0.00	Normal
Maximum number of valid rows in a partition (4,294,967,288)			
Maximum size of the data in a table partition (1.55 TB)			
Total SQL statements			
Total active SQL statements			
Maximum SQL program associated space size (16 MB)			

QSYS2.Health_Environmental_Limits ()

The QSYS2.Health_Environmental_Limits() procedure returns detail on the top 10 jobs on the system, for different SQL or application limits. The jobs do not have to be in existence. The top 10 information is maintained within Db2 for i and gets reset when the machine is IPLed, the IASP is varied ON, or when the QSYS2.Reset_Environmental_Limits() procedure is called.

Procedure definition:

```
CREATE PROCEDURE QSYS2.HEALTH_ENVIRONMENTAL_LIMITS(  
  IN ARCHIVE_OPTION INTEGER,  
  IN NUMBER_OF_LIMITS_ARCHIVE INTEGER,  
  IN LIMIT_SCHEMA VARCHAR(258),  
  IN LIMIT_TABLE VARCHAR(258))  
  DYNAMIC RESULT SETS 1  
  LANGUAGE C  
  SPECIFIC QSYS2.HEALTH_ENVIRONMENTAL_LIMITS
```

```
NOT DETERMINISTIC
MODIFIES SQL DATA
CALLED ON NULL INPUT
EXTERNAL NAME 'QSYS/QSQHEALTH(ENVIRON)'
PARAMETER STYLE SQL;
```

Service Program Name: QSYS/QSQHEALTH

Default Public Authority: *USE

Threadsafe: Yes

IBM i release

This procedure was added to IBM i 6.1.

Parameters

Archive_Option

(Input) The type of operation to perform for the Db2 for i Health Center activity detail.

The supported values are:

- 1 = Query only, no archive action is taken
- 2 = Archive only
- 3 = Create archive and archive
- 4 = Query the archive

Note: Option 1 produces a new result set. Options 2 and 3 simply use the results from the last Query option. Option 3 fails if the archive exists.

Number_Of_Limits_Archive

(Input) The number of rows to save per object health limit.

The archive can be used to recognize trends over time. To have meaningful historical comparisons, choose the row count size carefully. This argument is ignored if the Archive_Option is 1 or 4.

Limit_Schema

(Input) The schema that contains the database activity archive.

This argument is ignored if the Archive_Option is 1.

Limit_Table

The table that contains the database activity archive.

This argument is ignored if the Archive_Option is 1.

Authorities

To query an existing archive, *USE object authority is required for the Limit_Schema and Limit_Table. To create an archive, *CHANGE object authority is required for the Limit_Schema. To add to an archive, *CHANGE object authority is required for the Limit_Table.

When Archive_Option is 1 or 3, *USE object authority is required for the Object_Schema and for any objects which are indicated by Object_Name. When an object is encountered and the caller does not have *USE object authority, an SQL0462 warning is placed in the job log. The object is skipped and not included in the procedure result set.

Result Set

When Archive_Option is 1 or 4, a single result set is returned.

The format of the result is as follows. All these items were added for IBM i 6.1.

QSYS2.Health_Environmental_Limits() result set format:

```
"TIMESTAMP" TIMESTAMP NOT NULL,
LIMIT VARCHAR(2000) ALLOCATE(20) DEFAULT NULL,
HIGHWATER_MARK_VALUE FOR COLUMN HIMARK BIGINT DEFAULT NULL,
WHEN_VALUE_WAS_RECORDED FOR COLUMN TIMEHIT TIMESTAMP NOT NULL,
PERCENT DECIMAL(5, 2) DEFAULT NULL,
JOB_NAME VARCHAR(28) ALLOCATE(20) DEFAULT NULL,
"CURRENT_USER" FOR COLUMN CUSER VARCHAR(128) ALLOCATE(10) DEFAULT NULL,
JOB_TYPE VARCHAR(26) ALLOCATE(20) DEFAULT NULL,
MAXIMUM_VALUE FOR COLUMN MAXVAL BIGINT DEFAULT NULL,
JOB_STATUS VARCHAR(13) DEFAULT NULL,
CLIENT_WRKSTNNAME FOR COLUMN "WRKSTNNAME" VARCHAR(255) DEFAULT NULL,
CLIENT_APPLNAME FOR COLUMN "APPLNAME" VARCHAR(255) DEFAULT NULL,
CLIENT_ACCTNG FOR COLUMN "ACCTNG" VARCHAR(255) DEFAULT NULL,
CLIENTROGRAMID FOR COLUMN "PROGRAMID" VARCHAR(255) DEFAULT NULL,
CLIENT_USERID FOR COLUMN "USERID" VARCHAR(255) DEFAULT NULL,
WHEN_LIMITS_ESTABLISHED FOR COLUMN TIMESET TIMESTAMP NOT NULL,
INTERFACE_NAME FOR COLUMN INTNAME VARCHAR(127) ALLOCATE(10) DEFAULT NULL,
INTERFACE_TYPE FOR COLUMN INTTYPE VARCHAR(63) ALLOCATE(10) DEFAULT NULL,
INTERFACE_LEVEL FOR COLUMN INTLEVEL VARCHAR(63) ALLOCATE(10) DEFAULT NULL,
LIMIT_ID INTEGER DEFAULT NULL
```

```
LABEL ON COLUMN <result set>
( "TIMESTAMP" IS 'Timestamp',
  LIMIT IS 'Limit',
  HIGHWATER_MARK_VALUE IS 'Largest Value',
  WHEN_VALUE_WAS_RECORDED IS 'Timestamp When Recorded',
  PERCENT IS 'Percent',
  JOB_NAME IS 'Job Name',
  "CURRENT_USER" IS 'Current User',
  JOB_TYPE IS 'Job Type',
  MAXIMUM_VALUE IS 'Maximum Value',
  JOB_STATUS IS 'Job Status',
  CLIENT_WRKSTNNAME IS 'Client Workstation Name',
  CLIENT_APPLNAME IS 'Client Application Name',
  CLIENT_ACCTNG IS 'Client Accounting Code',
  CLIENTROGRAMID IS 'Client Program Identifier',
  CLIENT_USERID IS 'Client User Identifier',
  WHEN_LIMITS_ESTABLISHED IS 'Timestamp Limits Established',
  INTERFACE_NAME IS 'Interface Name',
  INTERFACE_TYPE IS 'Interface Type',
  INTERFACE_LEVEL IS 'Interface Level',
  LIMIT_ID IS 'Limit ID' );
```

Limit Detail

The supported Database Health Center Environmental limits can be seen on any machine by executing this query:

```
SELECT * FROM QSYS2.SQL_SIZING WHERE SIZING_ID BETWEEN 18200 AND 18299;
```

Note: The **bold** row was added in IBM i 7.1.

<i>Table 41. SQL environmental limits.</i>		
SIZING_ID	SIZING_NAME	SUPPORTED_VALUE
18200	MAXIMUM NUMBER OF LOB or XML LOCATORS PER JOB	16000000
18201	MAXIMUM NUMBER OF LOB or XML LOCATORS PER SERVER JOB	209000
18202	MAXIMUM NUMBER OF ACTIVATION GROUPS	0
18203	MAXIMUM NUMBER OF DESCRIPTORS	0
18204	MAXIMUM NUMBER OF CLI HANDLES	160000
18205	MAXIMUM NUMBER OF SQL OPEN CURSORS	21754
18206	MAXIMUM NUMBER OF SQL PSEUDO OPEN CURSORS	0

Table 41. SQL environmental limits. (continued)

SIZING_ID	SIZING_NAME	SUPPORTED_VALUE
18207	MAXIMUM LENGTH OF SQL STATEMENT2097152	2097152

Error Messages

None

Usage Notes

None

Related Information

None


Example

Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

Retrieve the SQL environmental limits for the current database.

```
CALL QSYS2.Health_Environmental_Limits(1, 0, NULL, NULL);
```

Example results in System i Navigator:



QSYS2.Reset_Environmental_Limits ()

The QSYS2.Reset_Environmental_Limits () procedure clears out the environment limit cache for the database. If IASPs are being used, this procedure clears the environment limit cache for the IASP within which it is called.

Procedure definition:

```
CREATE PROCEDURE QSYS2.RESET_ENVIRONMENTAL_LIMITS(  
LANGUAGE C  
SPECIFIC QSYS2.RESET_ENVIRONMENTAL_LIMITS  
NOT DETERMINISTIC  
MODIFIES SQL DATA  
CALLED ON NULL INPUT  
EXTERNAL NAME 'QSYS/QSQSSUDF(RESETENV) '  
PARAMETER STYLE SQL;
```

Service Program Name: QSYS/QSQHEALTH

Default Public Authority: *USE

Threadsafe: Yes

IBM i release

This procedure was added to IBM i 6.1.

Parameters

None.

Authorities

This procedure requires the user to have *JOBCTL user special authority or be authorized to the QIBM_DB_SQLADM Function through Application Administration in System i Navigator. The Change Function Usage (CHGFCNUSG) command can also be used to allow or deny use of the function.

For example:

```
CHGFCNUSG FCNID(QIBM_DB_SQLADM) USER(XXXXX) USAGE(*ALLOWED)
```

Result Set

None.

Error Messages

<i>Table 42. Error messages</i>	
Message ID	Error Message Text
SQL0552	Not authorized to PROCEDURE.

Usage Notes

None

Related Information

None

Example

Note: By using the code examples, you agree to the terms of the [“Code license and disclaimer information”](#) on page 1058.

Reset the SQL environmental limits for the current database.

```
CALL QSYS2.RESET_ENVIRONMENTAL_LIMITS;
```

Monitoring your queries using the Database Monitor

Start Database Monitor (STRDBMON) command gathers information about a query in real time and stores this information in an output table. This information can help you determine whether your system and your queries are performing well, or whether they need fine-tuning. Database monitors can generate significant CPU and disk storage overhead when in use.

You can gather performance information for a specific query, for every query on the system, or for a group of queries on the system. When a job is monitored by multiple monitors, each monitor is logging rows to a different output table. You can identify rows in the output database table by its unique identification number.

When you start a monitor using the **Start Database Monitor (STRDBMON)** command, the monitor is automatically registered with System i Navigator and appears in the System i Navigator monitor list.

Note: Database monitors also contain the SQL statement text and variable values. If the variable values or SQL statements contain sensitive data you should create database monitors in a library that is not publicly authorized to prevent exposure to the sensitive data.

What kinds of statistics you can gather

The database monitor provides the same information that is provided with the query optimizer debug messages (**Start Debug (STRDBG)**) and the **Print SQL information (PRTSQLINF)** command. The following is a sampling of the additional information that is gathered by the database monitors:

- System and job name
- SQL statement and subselect number
- Start and end timestamp
- Estimated processing time
- Total rows in table queried
- Number of rows selected
- Estimated number of rows selected
- Estimated number of joined rows
- Key columns for advised index
- Total optimization time
- Join type and method
- ODP implementation

How you can use performance statistics

You can use these performance statistics to generate various reports. For instance, you can include reports that show queries that:

- Use an abundance of the system resources.
- Take a long time to execute.
- Did not run because of the query governor time limit.
- Create a temporary index during execution

- Use the query sort during execution
- Might perform faster with the creation of a keyed logical file containing keys suggested by the query optimizer.

Note: A query that is canceled by an end request generally does not generate a full set of performance statistics. However, it does contain all the information about how a query was optimized, except for runtime or multi-step query information.

Related information

[Start Debug \(STRDBG\) command](#)

[Print SQL Information \(PRTSQLINF\) command](#)

[Start Database Monitor \(STRDBMON\) command](#)

Start Database Monitor (STRDBMON) command

The **Start Database Monitor (STRDBMON)** command starts the collection of database performance statistics for a specified job, for all jobs on the system or for a selected set of jobs. The statistics are placed in a user-specified database table and member. If the table or member do not exist, one is created based on the QAQQDBMN table in library QSYS, with the public authority for the file the same as the create authority specified for the library in which the file is created. If the table and member do exist, the record format of the specified table is verified to ensure it is the same.

For each monitor started using the STRDBMON command, the system generates a monitor ID that can be used to uniquely identify each individual monitor. The monitor ID can be used on the ENDDBMON command to uniquely identify which monitor is to be ended. The monitor ID is returned in the informational message CPI436A which is generated for each occurrence of the STRDBMON command. The monitor ID can also be found in either the column QQC101 of the QQQ3018 database monitor record or in the QSYS2.DATABASE_MONITOR_INFO catalog.

Informally there are two types of monitors. A private monitor is a monitor over one, specific job (or the current job). Only one (1) monitor can be started on a specific job at a time. For example, STRDBMON JOB(*) followed by another STRDBMON JOB(*) within the same job is not allowed. A public monitor is a monitor which collects data across multiple jobs. There can be a maximum of 10 public monitors active at any one time. For example, STRDBMON JOB(*ALL) followed by another STRDBMON JOB(*ALL) is allowed providing the maximum number of public monitors does not exceed 10. You could have 10 public monitors and 1 private monitor active at the same time for any specific job.

If multiple monitors specify the same output file, only one copy of the database statistic records is written to the file for each job. For example, STRDBMON OUTFILE(LIB/TABLE1) JOB(*) and STRDBMON OUTFILE(LIB/TABLE1) JOB(*ALL) target the same output file. For the current job, there are not two copies of the database statistic records—one copy for the private monitor and one copy for the public monitor. There is only one copy of the database statistic records.

If the monitor is started on all jobs (a public monitor), any jobs waiting on queues or started during the monitoring period are included in the monitor data. If the monitor is started on a specific job (a private monitor) that job must be active in the system when the command is issued. Each job in the system can be monitored concurrently by one private monitor and a maximum of 10 public monitors.

The STRDBMON command allows you to collect statistic records for a specific set or subset of the queries running on any job. This filtering can be performed over the job name, user profile, query table names, query estimated run time, TCP/IP address, or any combination of these filters. Specifying a STRDBMON filter helps minimize the number of statistic records captured for any monitor.

Example 1: Starting Public Monitoring

```
STRDBMON  OUTFILE(QGPL/FILE1)  OUTMBR(MEMBER1 *ADD)
JOB(*ALL)  FRCRCD(10)
```

This command starts database monitoring for all jobs on the system. The performance statistics are added to the member named MEMBER1 in the file named FILE1 in the QGPL library. 10 records are held before being written to the file.

Example 2: Starting Private Monitoring

```
STRDBMON  OUTFILE(*LIBL/FILE3)  OUTMBR(MEMBER2)
JOB(134543/QPGMR/DSP01)  FRCRCD(20)
```

This command starts database monitoring for job number 134543. The job name is DSP01 and was started by the user named QPGMR. The performance statistics are added to the member named MEMBER2 in the file named FILE3. 20 records are held before being written to the file.

Example 3: Starting Private Monitoring to a File in a Library in an Independent ASP

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(134543/QPGMR/DSP01)
```

This command starts database monitoring for job number 134543. The job name is DSP01 and was started by the user named QPGMR. The performance statistics are added to the member name DBMONFILE (since OUTMBR was not specified) in the file named DBMONFILE in the library named LIB41. This library could exist in more than one independent auxiliary storage pool (ASP); the library in the name space of the originator's job is always used.

Example 4: Starting Public Monitoring For All Jobs That Begin With 'QZDA'

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL/*ALL/QZDA*)
```

This command starts database monitoring for all jobs that whose job name begins with 'QZDA'. The performance statistics (monitor records) are added to member DBMONFILE (since OUTMBR was not specified) in file DBMONFILE in library LIB41. This library could exist in more than one independent auxiliary storage pool (ASP); the library in the name space of the originator's job is always used.

Example 5: Starting Public Monitoring and Filtering SQL Statements That Run Over 10 Seconds

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)  RUNTHLD(10)
```

This command starts database monitoring for all jobs. Monitor records are created only for those SQL statements whose estimated run time meets or exceeds 10 seconds.

Example 6: Starting Public Monitoring and Filtering SQL Statements That Have an Estimated Temporary Storage Over 200 MB

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)  STGTHLD(200)
```

This command starts database monitoring for all jobs. Monitor records are created only for those SQL statements whose estimated temporary storage meets or exceeds 200 MB.

Example 7: Starting Private Monitoring and Filtering Over a Specific File

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*)
FTRFILE(LIB41/TABLE1)
```

This command starts database monitoring for the current job. Monitor records are created only for those SQL statements that use file LIB41/TABLE1.

Example 8: Starting Private Monitoring for the Current User

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*)  FTRUSER(*CURRENT)
```

This command starts database monitoring for the current job. Monitor records are created only for those SQL statements that are executed by the current user.

Example 9: Starting Public Monitoring For Jobs Beginning With 'QZDA' and Filtering Over Run Time and File

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL/*ALL/QZDA*)
RUNTHLD(10)  FTRUSER(DEVLPR1)  FTRFILE(LIB41/TTT*)
```

This command starts database monitoring for all jobs whose job name begins with 'QZDA'. Monitor records are created only for those SQL statements that meet all the following conditions:

- The estimated run time, as calculated by the query optimizer, meets, or exceeds 10 seconds
- Was executed by user 'DEVLPR1'.
- Use any file whose name begins with 'TTT' and resides in library LIB41.

Example 10: Starting Public Monitoring and Filtering SQL Statements That Have Internet Address '9.10.111.77'.

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)
FTRINTNETA('9.10.111.77')
```

This command starts database monitoring for all jobs. Monitor records are created only for TCP/IP database server jobs that are using the client IP version 4 address of '9.10.111.77'.

Example 11: Starting Public Monitoring and Filtering SQL Statements That Have a Port Number of 8471

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)  FTRLCLPORT(8471)
```

This command starts database monitoring for all jobs. Monitor records are created only for TCP/IP database server jobs that are using the local port number 8471.

Example 12: Starting Public Monitoring Based on Feedback from the Query Governor

```
CHGSYSVAL  QQRYSIMLMT(200)
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)  FTRQRYGOVR(*COND)
```

This command starts database monitoring for all jobs whose estimated run time is expected to exceed 200 seconds, based on the response to the query governor. In this example, data is collected only if the query is canceled or a return code of 2 is returned by a query governor exit program. The query can be canceled by a user response to the inquiry message CPA4259, issued because the query exceeded the query governor limits. It can also be canceled by the program logic inside the registered query governor exit program.

Example 13: Collecting database monitor for Interactive SQL use

```
STRDBMON  OUTFILE(QGPL/STRSQLMON1)  OUTMBR(*FIRST *REPLACE)
JOB(*ALL/*ALL/*ALL)  TYPE(*DETAIL)
FTRCLTPGM(STRSQL)
```

This command uses the database monitor pre-filter by Client Special Register Program ID to collect monitor records for all the SQL statements executed by Interactive SQL (STRSQL command) usage.

Example 14: Starting Public Monitoring and Filtering SQL Statements Run From IBM i Access Client Solutions Run SQL Scripts

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)
FTRCLTAPP('IBM i Access Client Solutions - RUN SQL Scripts')
```

This command starts database monitoring for all jobs. Monitor records are created only for those statements where the client special register CLIENT APPLNAME is IBM i Access Client Solutions - RUN SQL Scripts.

Example 15: Starting Public Monitoring and Filtering SQL Statements Run From IBM System i Access for Windows Run SQL Scripts

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)
          FTRCLTPGM('cwbunnav.exe')
```

This command starts database monitoring for all jobs. Monitor records are created only for those statements where the client special register CLIENT_PROGRAMID is cwbunnav.exe.

Example 16: Starting Public Monitoring and Filtering for the client user dbmusr1

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)
          FTRCLTUSR(dbmusr1)
```

This command starts database monitoring for all jobs. Monitor records are created only for those statements where the client special register CLIENT_USERID is dbmusr1.

Example 17: Starting Public Monitoring and Filtering SQL Statements Using FTRUSER

```
STRDBMON  OUTFILE(LIB41/DBMONFILE)  JOB(*ALL)
          FTRUSER((STC* *NE) (S*) (TS*) (TSA* *NE))
```

This command starts database monitoring for all jobs. Monitor records are created only for those statements where the user profile is logically equivalent to: User profile NOT LIKE 'STC%' AND user profile NOT LIKE 'TSA%' AND (user profile LIKE 'S%' OR user profile LIKE 'TS%')

Related information

[Start Database Monitor \(STRDBMON\) command](#)

End Database Monitor (ENDDBMON) command

The **End Database Monitor (ENDDBMON)** command ends the collection of database performance statistics for a specified job, all jobs on the system, or a selected set of jobs (for example, a generic job name).

To end a monitor, you can specify the job or the monitor ID or both. If only the JOB parameter is specified, the monitor that was started using the same exact JOB parameter is ended - if there is only one monitor which matches the specified JOB. If more than one monitor is active which matches the specified JOB, then the user uniquely identifies which monitor is to be ended by use of the MONID parameter.

When only the MONID parameter is specified, the specified MONID is compared to the monitor ID of the monitor for the current job and to the monitor ID of all active public monitors (monitors that are open across multiple jobs). The monitor matching the specified MONID is ended.

The monitor ID is returned in the informational message CPI436A. This message is generated for each occurrence of the STRDBMON command. Look in the job log for message CPI436A to find the system generated monitor ID, if needed. The monitor ID can also be found in column QQC101 of the QQQ3018 database monitor record.

Restrictions

- If a specific job name and number or JOB(*) was specified on the **Start Database Monitor (STRDBMON)** command, the monitor can only be ended by specifying the same job name and number or JOB(*) on the ENDDBMON command.
- If JOB(*ALL) was specified on the **Start Database Monitor (STRDBMON)** command, the monitor can only be ended by specifying ENDDBMON JOB(*ALL). The monitor cannot be ended by specifying ENDDBMON JOB(*).

When monitoring is ended for all jobs, all the jobs on the system are triggered to close the database monitor output table. However, the ENDDBMON command can complete before all the monitored jobs have written their final statistic records to the log. Use the **Work with Object Locks (WRKOBJLCK)**

command to determine that all the monitored jobs no longer hold locks on the database monitor output table before assuming that the monitoring is complete.

Example 1: End Monitoring for a Specific Job

```
ENDDBMON JOB(*)
```

This command ends database monitoring for the current job.

Example 2: End Monitoring for All Jobs

```
ENDDBMON JOB(*ALL)
```

This command ends the monitor open across all jobs on the system. If more than one monitor with JOB(*ALL) is active, then the MONID parameter must also be specified to uniquely identify which specific public monitor to end.

Example 3: End Monitoring for an Individual Public Monitor with MONID Parameter

```
ENDDBMON JOB(*ALL) MONID(061601001)
```

This command ends the monitor that was started with JOB(*ALL) and that has a monitor ID of 061601001. Because there were multiple monitors started with JOB(*ALL), the monitor ID must be specified to uniquely identify which monitor that was started with JOB(*ALL) is to be ended.

Example 4: End Monitoring for an Individual Public Monitor with MONID Parameter

```
ENDDBMON MONID(061601001)
```

This command performs the same function as the previous example. It ends the monitor that was started with JOB(*ALL) or JOB(*) and that has a monitor ID of 061601001.

Example 5: End Monitoring for All JOB(*ALL) Monitors

```
ENDDBMON JOB(*ALL/*ALL/*ALL) MONID(*ALL)
```

This command ends all monitors that are active across multiple jobs. It does not end any monitors open for a specific job or the current job.

Example 6: End Monitoring for a Generic Job

```
ENDDBMON JOB(QZDA*)
```

This command ends the monitor that was started with JOB(QZDA*). If more than one monitor with JOB(QZDA*) is active, then the MONID parameter must also be specified to uniquely identify which individual monitor to end.

Example 7: End Monitoring for an Individual Monitor with a Generic Job

```
ENDDBMON JOB(QZDA*) MONID(061601001)
```

This command ends the monitor that was started with JOB(QZDA*) and has a monitor ID of 061601001. Because there were multiple monitors started with JOB(QZDA*), the monitor ID must be specified to uniquely identify which JOB(QZDA*) monitor is to be ended.

Example 8: End Monitoring for a Group of Generic Jobs

```
ENDDBMON JOB(QZDA*) MONID(*ALL)
```

This command ends all monitors that were started with JOB(QZDA*).

Related information

[End Database Monitor \(ENDDBMON\) command](#)

Database monitor performance rows

The rows in the database table are uniquely identified by their row identification number. The information within the file-based monitor (**Start Database Monitor (STRDBMON)**) is written out based upon a set of logical formats which are defined in the database monitor formats. These views correlate closely to the debug messages and the **Print SQL Information (PRSQLINF)** messages.

The database monitor formats section also identifies which physical columns are used for each view and what information it contains. You can use the views to identify the information that can be extracted from the monitor. These rows are defined in several different views which are not shipped with the system and must be created by the user, if wanted. The views can be created with the SQL DDL. The column descriptions are explained in the tables following each figure.

Related concepts

[Database monitor formats](#)

This section contains the formats used to create the database monitor SQL tables and views.

Database monitor examples

The System i Navigator interface provides a powerful tool for gathering and analyzing performance monitor data using database monitor. However, you might want to do your own analysis of the database monitor files.

Suppose you have an application program with SQL statements and you want to analyze and performance tune these queries. The first step in analyzing the performance is collection of data. The following examples show how you might collect and analyze data using **Start Database Monitor (STRDBMON)** and **End Database Monitor (ENDDBMON)** commands. Performance data is collected in LIB/PERFDATA for an application running in your current job. The following sequence collects performance data and prepares to analyze it.

1. **STRDBMON** FILE(LIB/PERFDATA) TYPE(*DETAIL). If this table does not exist, the command creates one from the skeleton table in QSYS/QAQQDBMN.
2. Run your application
3. **ENDDBMON**
4. Create views over LIB/PERFDATA using the SQL DDL. Creating the views is not mandatory. All the information resides in the base table that was specified on the **STRDBMON** command. The views simply provide an easier way to view the data.

You are now ready to analyze the data. The following examples give you a few ideas on how to use this data. You must closely study the physical and logical view formats to understand all the data being collected. Then you can create queries that give the best information for your applications.

Related information

[Start Database Monitor \(STRDBMON\) command](#)

[End Database Monitor \(ENDDBMON\) command](#)

Application with table scans example

Determine which queries in your SQL application are implemented with table scans. The complete information can be obtained by joining two views: QQQ1000, which contains information about the SQL statements, and QQQ3000, which contains data about queries performing table scans.

The following SQL query can be used:

```
SELECT (B.End_Timestamp - B.Start_Timestamp) AS TOT_TIME, A.System_Table_Schema,  
A.System_Table_Name,  
A.Index_Advised, A.Table_Total_Rows, C.Number_Rows_Returned, A.Estimated_Rows_Selected,  
B.Statement_Text_Long
```

```

FROM LIB.QQ3000 A, LIB.QQ1000 B, LIB.QQ3019 C
WHERE A.Join_Column = B.Join_Column
AND A.Join_Column = C.Join_Column

```

Sample output of this query is shown in the following table. Key to this example are the join criteria:

```

WHERE A.Join_Column = B.Join_Column
AND A.Join_Column = C.Join_Column

```

Much data about many queries is contained in multiple rows in table LIB/PERFDATA. It is not uncommon for data about a single query to be contained in 10 or more rows within the table. The combination of defining the logical views and then joining the views together allows you to piece together all the data for a query or set of queries. Column QQJFLD uniquely identifies all queries within a job; column QQUCNT is unique at the query level. The combination of the two, when referenced in the context of the logical views, connects the query implementation to the query statement information.

<i>Table 43. Output for SQL Queries that Performed Table Scans</i>						
Lib Name	Table Name	Total Rows	Index Advised	Rows Returned	TOT_TIME	Statement Text
LIB1	TBL1	20000	Y	10	6.2	SELECT * FROM LIB1/TBL1 WHERE FLD1 = 'A'
LIB1	TBL2	100	N	100	0.9	SELECT * FROM LIB1/TBL2
LIB1	TBL1	20000	Y	32	7.1	SELECT * FROM LIB1/TBL1 WHERE FLD1 = 'B' AND FLD2 > 9000

If the query does not use SQL, the SQL information row (QQQ1000) is not created. Without the SQL information row, it is more difficult to determine which rows in LIB/PERFDATA pertain to which query. When using SQL, row QQQ1000 contains the actual SQL statement text that matches the monitor rows to the corresponding query. Only through SQL is the statement text captured. For queries executed using the OPNQRYF command, the OPNID parameter is captured and can be used to tie the rows to the query. The OPNID is contained in column Open_Id of row QQQ3014.

Queries with table scans example

Like the preceding example that showed which SQL applications were implemented with table scans, the following example shows all queries that are implemented with table scans.

```

SELECT (D.End_Timestamp - D.Start_Timestamp) AS TOT_TIME, A.System_Table_Schema,
A.System_Table_Name,
A.Table_Total_Rows, A.Index_Advised,
B.Open_Id, B.Open_Time,
C.Clock_Time_to_Return_All_Rows, C.Number_Rows_Returned,
D.Result_Rows, D.Statement_Text_Long
FROM LIB.QQ3000 A INNER JOIN LIB.QQ3014 B
ON (A.Join_Column = B.Join_Column)
LEFT OUTER JOIN LIB.QQ3019 C
ON (A.Join_Column = C.Join_Column)
LEFT OUTER JOIN LIB.QQ1000 D
ON (A.Join_Column = D.Join_Column)

```

In this example, the output for all queries that performed table scans are shown in the following table.

Note: The columns selected from table QQQ1000 do return NULL default values if the query was not executed using SQL. For this example assume that the default value for character data is blanks and the default value for numeric data is an asterisk (*).

Table 44. Output for All Queries that Performed Table Scans

Lib Name	Table Name	Total Rows	Index Advised	Query OPNID	ODP Open Time	Clock Time	Recs Rtned	Rows Rtned	TOT_TIME	Statement Text
LIB1	TBL1	20000	Y		1.1	4.7	10	10	6.2	SELECT * FROM LIB1/TBL1 WHERE FLD1 = 'A'
LIB1	TBL2	100	N		0.1	0.7	100	100	0.9	SELECT * FROM LIB1/TBL2
LIB1	TBL1	20000	Y		2.6	4.4	32	32	7.1	SELECT * FROM LIB1/TBL1 WHERE FLD1 = 'A' AND FLD2 > 9000
LIB1	TBL4	4000	N	QRY04	1.2	4.2	724	*	*	*

If the SQL statement text is not needed, joining to table QQQ1000 is not necessary. You can determine the total time and rows selected from data in the QQQ3014 and QQQ3019 rows.

Table scan detail example

Your next step could include further analysis of the table scan data. The previous examples contained a column titled Index Advised. A 'Y' (yes) in this column is a hint from the query optimizer that the query could perform better with an index to access the data. For the queries where an index is advised, the rows selected by the query are low in comparison to the total number of table rows. This selectivity is another indication that a table scan might not be optimal. Finally, a long execution time might highlight queries that could be improved by performance tuning.

The next logical step is to look into the index advised optimizer hint. The following query can be used:

```
SELECT A.System_Table_Schema, A.System_Table_Name,
       A.Index_Advised, A.Index_Advised_Columns,
       A.Index_Advised_Columns_Count, B.Open_Id,
       C.Statement_Text_Long
FROM LIB.QQQ3000 A INNER JOIN LIB.QQQ3014 B
ON (A.Join_Column = B.Join_Column)
LEFT OUTER JOIN LIB.QQQ1000 C
ON (A.Join_Column = C.Join_Column)
WHERE A.Index_Advised = 'Y'
```

There are two slight modifications from the first example. First, the selected columns have been changed. Most important is the selection of Index_Advised_Columns containing a list of possible key columns to use when creating the suggested index. Second, the query selection limits the output to those table scan queries where the optimizer advises that an index is created (A.Index_Advised = 'Y'). The following table shows what the results might look like.

Table 45. Output with Recommended Key Columns

Lib Name	Table Name	Index Advised	Advised Key columns	Advised Primary Key	Query OPNID	Statement Text
LIB1	TBL1	Y	FLD1	1		SELECT * FROM LIB1/TBL1 WHERE FLD1 = 'A'
LIB1	TBL1	Y	FLD1, FLD2	1		SELECT * FROM LIB1/TBL1 WHERE FLD1 = 'B' AND FLD2 > 9000

Table 45. Output with Recommended Key Columns (continued)

Lib Name	Table Name	Index Advised	Advised Key columns	Advised Primary Key	Query OPNID	Statement Text
LIB1	TBL4	Y	FLD1, FLD4	1	QRY04	

Determine whether it makes sense to create a permanent index as advised by the optimizer. In this example, creating one index over LIB1/TBL1 satisfies all three queries since each use a primary or left-most key column of FLD1. By creating one index over LIB1/TBL1 with key columns FLD1, FLD2, there is potential to improve the performance of the second query even more. Consider how often these queries are run and the overhead of maintaining an additional index over the table when deciding whether to create the suggested index.

If you create a permanent index over FLD1, FLD2 the next sequence of steps is as follows:

1. Start the performance monitor again
2. Rerun the application
3. End the performance monitor
4. Re-evaluate the data.

It is likely that the three index-advised queries are no longer performing table scans.

Additional database monitor examples

The following are additional ideas or examples on how to extract information from the performance monitor statistics. All the examples assume that data has been collected in LIB/PERFDATA and the documented views have been created.

1. How many queries are performing dynamic replans?

```
SELECT COUNT(*)
FROM LIB.QQ01000
WHERE Dynamic_Replan_Reason_Code <> 'NA'
```

2. What is the statement text and the reason for the dynamic replans?

```
SELECT Dynamic_Replan_Reason_Code, Statement_Text_Long
FROM LIB.QQ01000
WHERE Dynamic_Replan_Reason_Code <> 'NA'
```

Note: You need to refer to the description of column Dynamic_Replan_Reason_Code for definitions of the dynamic replan reason codes.

3. How many indexes have been created over LIB1/TBL1?

```
SELECT COUNT(*)
FROM LIB.QQ03002
WHERE System_Table_Schema = 'LIB1'
AND System_Table_Name = 'TBL1'
```

4. What key columns are used for all indexes created over LIB1/TBL1 and what is the associated SQL statement text?

```
SELECT A.System_Table_Schema, A.System_Table_Name,
A.Index_Advised_Columns, B.Statement_Text_Long
FROM LIB.QQ03002 A, LIB.QQ01000 B
WHERE A.Join_Column = B.Join_Column
AND A.System_Table_Schema = 'LIB1'
AND A.System_Table_Name = 'TBL1'
```

Note: This query shows key columns only from queries executed using SQL.

5. What key columns are used for all indexes created over LIB1/TBL1 and what was the associated SQL statement text or query open ID?

```

SELECT A.System_Table_Schema, A.System_Table_Name, A.Index_Advised_Columns,
       B.Open_Id, C.Statement_Text_Long
FROM LIB.QQ3002 A INNER JOIN LIB.QQ3014 B
     ON (A.Join_Column = B.Join_Column)
LEFT OUTER JOIN LIB.QQ1000 C
     ON (A.Join_Column = C.Join_Column)
WHERE A.System_Table_Schema LIKE '%'
AND A.System_Table_Name = '%'

```

Note: This query shows key columns from all queries on the system.

6. What types of SQL statements are being performed? Which are performed most frequently?

```

SELECT CASE Statement_Function
       WHEN 'O' THEN 'Other'
       WHEN 'S' THEN 'Select'
       WHEN 'L' THEN 'DDL'
       WHEN 'I' THEN 'Insert'
       WHEN 'U' THEN 'Update'
       ELSE 'Unknown'
END, COUNT(*)
FROM LIB.QQ1000
GROUP BY Statement_Function
ORDER BY 2 DESC

```

7. Which SQL queries are the most time consuming? Which user is running these queries?

```

SELECT (End_Timestamp - Start_Timestamp), Job_User,
       Current_User_Profile, Statement_Text_Long
FROM LIB.QQ1000
ORDER BY 1 DESC

```

8. Which queries are the most time consuming?

```

SELECT (A.Open_Time + B.Clock_Time_to_Return_All_Rows),
       A.Open_Id, C.Statement_Text_Long
FROM LIB.QQ3014 A LEFT OUTER JOIN LIB.QQ3019 B
     ON (A.Join_Column = B.Join_Column)
LEFT OUTER JOIN LIB.QQ1000 C
     ON (A.Join_Column = C.Join_Column)
ORDER BY 1 DESC

```

Note: This example assumes that detail data was collected (STRDBMON TYPE(*DETAIL)).

9. Show the data for all SQL queries with the data for each SQL query logically grouped.

```

SELECT A.*
FROM LIB.PERFDATA A, LIB.QQ1000 B
WHERE A.QQJFLD = B.Join_Column

```

Note: This might be used within a report that will format the interesting data into a more readable format. For example, all reason code columns can be expanded by the report to print the definition of the reason code. Physical column QQRCOD = 'T1' means that a table scan was performed because no indexes exist over the queried table.

10. How many queries are implemented with temporary tables because a key length greater than 2000 bytes, or more than 120 key columns was specified for ordering?

```

SELECT COUNT(*)
FROM LIB.QQ3004
WHERE Reason_Code = 'F6'

```

11. Which SQL queries were implemented with nonreusable ODPs?

```

SELECT B.Statement_Text_Long
FROM LIB.QQ3010 A, LIB.QQ1000 B
WHERE A.Join_Column = B.Join_Column
AND A.ODP_Implementation = 'N';

```

12. What is the estimated time for all queries stopped by the query governor?


```
SELECT Estimated_Processing_Time, Open_Id
FROM LIB.QQ3014
WHERE Stopped_By_Query_Governor = 'Y'
```

Note: This example assumes that detail data was collected (STRDBMON TYPE(*DETAIL)).

13. Which queries estimated time exceeds actual time?

```
SELECT A.Estimated_Processing_Time,
       (A.Open_Time + B.Clock_Time_to_Return_All_Rows),
       A.Open_Id, C.Statement_Text_Long
FROM LIB.QQ3014 A LEFT OUTER JOIN LIB.QQ3019 B
ON (A.Join_Column = B.Join_Column)
LEFT OUTER JOIN LIB.QQ1000 C
ON (A.Join_Column = C.Join_Column)
WHERE A.Estimated_Processing_Time/1000 >
      (A.Open_Time + B.Clock_Time_to_Return_All_Rows)
```

Note: This example assumes that detail data was collected (STRDBMON TYPE(*DETAIL)).

14. Should you apply a PTF for queries containing UNIONS? Yes, if any queries are performing UNIONS. Do any of the queries perform this function?

```
SELECT COUNT(*)
FROM QQ3014
WHERE Has_Union = 'Y'
```

Note: If the result is greater than 0, apply the PTF.

15. You are a system administrator and an upgrade to the next release is planned. You want to compare data from the two releases.

- Collect data from your application on the current release and save this data in LIB/CUR_DATA
- Move to the next release
- Collect data from your application on the new release and save this data in a different table: LIB/NEW_DATA
- Write a program to compare the results. You need to compare the statement text between the rows in the two tables to correlate the data.

Using System i Navigator with detailed monitors

You can work with detailed monitors from the System i Navigator interface. The detailed SQL performance monitor is the System i Navigator version of the STRDBMON database monitor, found on the native interface.

You can start this monitor by right-clicking SQL Performance Monitors under the Database portion of the System i Navigator tree and selecting **New > Monitor**. This monitor saves detailed data in real time to a hard disk. It does not need to be paused or ended in order to analyze the results. You can also choose to run a Visual Explain based on the data gathered by the monitor. Since this monitor saves data in real time, it might have a performance impact on your system.

Starting a detailed monitor

You can start a detailed monitor from the System i Navigator interface.

You can start this monitor by right-clicking **SQL Performance Monitors** under the Database portion of the System i Navigator tree and selecting **New > SQL Performance Monitor**.

When you create a detailed monitor, you can filter the information that you want to capture.

Initial number of records: Select to specify the number of records initially allocated for the monitor. The 'Initial number of records' option is used to pre-allocate storage to the database monitor out file. When collecting large amounts of monitor records, this option improves the collection performance by avoiding automatic storage extensions that occur as a file grows in size.

Minimum estimated query runtime:	Select to include queries that exceed a specified amount of time. Select a number and then a unit of time.
Minimum estimated temporary storage:	Select to include queries that exceed a certain amount of temporary storage. Specify a size in MB.
Job name:	Select to filter by a specific job name. Specify a job name in the field. You can specify the entire ID or use a wildcard. For example, 'QZDAS*' finds all jobs where the name starts with 'QZDAS.'
Job user:	Select to filter by a job user. Specify a user ID in the field. You can specify the entire ID or use a wildcard. For example, 'QUSER*' finds all user IDs where the name starts with 'QUSER.'
Current user:	Select to filter by the current user of the job. Specify a user ID in the field. You can specify the entire ID or use a wildcard. For example, 'QSYS*' finds all users where the name starts with 'QSYS.'
Client location:	Select to filter by Internet access. The input needs to be in IPv4 or IPv6 form. <ol style="list-style-type: none"> 1. IP version 4 address in dotted decimal form. Specify an Internet Protocol version 4 address in the form nnn.nnn.nnn.nnn where each nnn is a number in the range 0 through 255. 2. IP version 6 address in colon hexadecimal form. Specify an internet protocol version 6 address in the form xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx, where each xxxx is a hex number in the range 0 through FFFF. IP version 6 includes the IPv4-mapped IPv6 address form (for example, ::FFFF:1.2.3.4). For IP version 6, the compressed form of the address is allowed. 3. IP host domain name. Specify an internet host domain name of up to 254 characters in length.
Local port:	Select to filter by port number. You can select a port from the list or else enter your own port number. Ports in the list include: <ul style="list-style-type: none"> • 446 - DRDA/DDM • 447 - DRDA/DDM • 448 - Secure DRDA/DDM (SSL) • 4402 - QXDAEDRSQL server • 8471 - Database server • 9471 - Secure database server (SSL)
Query Governor limits:	Select to search for queries that have exceeded or are expected to exceed the query governor limits set for the system. Choose from the following options: <ul style="list-style-type: none"> • Always collect information when exceeded • Conditional collection of information when exceeded
Client registers:	Select to filter by the client register information.
Statements that access these objects:	Select to filter by only queries that use certain tables. Click Browse to select tables to include. To remove a table from the list, select the table and click Remove . A maximum of 10 table names can be specified.
Activity to monitor:	Select to collect monitor output for user-generated queries or for both user-generated and system-generated queries.

You can choose which jobs you want to monitor or choose to monitor all jobs. You can have multiple instances of monitors running on your system at one time. You can create up to 10 detailed monitors to monitor all jobs. When collecting information for all jobs, the monitor will collect on previously started jobs or new jobs that are started after the monitor is created. You can edit this list by selecting and removing jobs from the **Selected jobs** list.

Analyzing detailed monitor data

SQL performance monitors provides several predefined reports that you can use to analyze your monitor data.

To view these reports, right-click a monitor and select **Analyze**. The monitor does not need to be ended in order to view this information.

	Value	Summary Available	Statements Available
7/21/05 12:22:58 PM to 7/21/05 12:24:46 PT			
Overview			
• SQL Statements	51	✓	✓
• Users	1	✓	
• Jobs	1	✓	
• Threads	1		
• Average Table Rows	14		
• Average Rows Returned	0		
• Average Runtime	0.691		
• Average Parallel Degree Used	1		
• Maximum Parallel Degree	1		
• SQE	1		
• COE	0	✓	✓
• System Naming	39	✓	✓
• SQL Naming	11	✓	✓
• Full Opens	1	✓	✓
• Pseudo Opens	0	✓	✓
• Average MQTs Used	0	✓	✓
• Average Indexes Used	1	✓	✓
• Full Indexes Created	0	✓	✓
• Sparse Indexes Created	0	✓	✓
• Index From Index Created	0	✓	✓
• Index Creates Advised	2	✓	✓
• Advised Statistics	0	✓	✓
• Temporary Tables	0	✓	✓
• Sorts	1	✓	✓
• Access Plans Rebuilt	2	✓	✓
• Sort Sequence	0	✓	✓
• Call Statements	3	✓	✓
• Error	0	✓	✓
How much work was requested?			
What options were provided to the optimizer?			

On the Analysis Overview dialog, you can view overview information or else choose one of the following categories:

- How much work was requested?
- What options were provided to the optimizer?
- What implementations did the optimizer use?
- What types of SQL statements were requested?
- Miscellaneous information
- I/O information

From the Actions menu, you can choose one of the following summary predefined reports:

User summary Contains a row of summary information for each user. Each row summarizes all SQL activity for that user.

Job summary	Contains a row of information for each job. Each row summarizes all SQL activity for that job. This information can be used to tell which jobs on the system are the heaviest users of SQL. These jobs are perhaps candidates for performance tuning. You could then start a separate detailed performance monitor on an individual job to get more detailed information without having to monitor the entire system.
Operation summary	Contains a row of summary information for each type of SQL operation. Each row summarizes all SQL activity for that type of SQL operation. This information provides the user with a high-level indication of the type of SQL statements used. For example, are the applications mostly read-only, or is there a large amount of update, delete, or insert activity. This information can then be used to try specific performance tuning techniques. For example, if many INSERTs are occurring, you might use an OVRDBF command to increase the blocking factor or the QDBENCWT API.
Program summary	Contains a row of information for each program that performed SQL operations. Each row summarizes all SQL activity for that program. This information can be used to identify which programs use the most or most expensive SQL statements. Those programs are then potential candidates for performance tuning. A program name is only available if the SQL statements are embedded inside a compiled program. SQL statements that are issued through ODBC, JDBC, or OLE DB have a blank program name unless they result from a procedure, function, or trigger.

In addition, when a green check is displayed under Summary column, you can select that row and click **Summary** to view information about that row type. Click **Help** for more information about the summary report. To view information organized by statements, click **Statements**.

Comparing monitor data

You can use System i Navigator to compare data sets in two or more monitors.

System i Navigator provides you with two types of comparison. The first is a simple compare that provides a high-level comparison of the monitors or snapshots. The second is a detailed comparison. The simple compare provides you with enough data about the monitors or snapshots to help you determine if a detailed comparison is helpful.

To launch a simple compare, go to **System i Navigator > system name > SQL performance monitors**. Right-click one or more monitors and select **Compare**.

To launch a detailed comparison, select the Detailed Comparison tab.

On the Detailed Comparison dialog, you can specify information about the data sets that you want to compare.

Name	The name of the monitors that you want to compare.
Schema mask	Select any names that you want the comparison to ignore. For example, consider the following scenario: You have an application running in a test schema and it is optimized. Now you move it to the production schema and you want to compare how it executes there. The statements in the comparison are identical except that the statements in the test schema use "TEST" and the statements in the production schema use "PROD". You can use the schema mask to ignore "TEST" in the first monitor and "PROD" in the second monitor. Then the statements in the two monitors appear identical.
Statements that ran longer than	The minimum runtime for statements to be compared.
Minimum percent difference	The minimum difference in key attributes of the two statements being compared that determines if the statements are considered equal or not. For example, if you select 25% as the minimum percent different, only matching statements whose key attributes differ by 25% or more are returned.

When you click **Compare**, both monitors are scanned for matching statements. Any matches found are displayed side-by-side for comparison of key attributes of each implementation.

On the Comparison output dialog, you view statements that are included in the monitor by clicking **Show Statements**. You can also run Visual Explain by selecting a statement and clicking **Visual Explain**.

Any matches found are displayed side-by-side for comparison of key attributes of each implementation.

Viewing statements in a monitor

You can view SQL statements that are included in a detailed monitor.

Right-click any detailed monitor in the SQL performance monitor window and select **Show statements**.

The filtering options provide a way to focus in on a particular area of interest:

Minimum runtime for the longest execution of the statement:	Select to include statements that exceed a certain amount of time. Select a number and then a unit of time.
Statements that ran on or after this date and time:	Select to include statements run at a specified date and time. Select a date and time.
Statements that reference the following objects:	Select to include statements that use or reference certain objects. Click Browse to select objects to include.
Statements that contain the following text:	Select to include only those statements that contain a specific type of SQL statement. For example, specify SELECT if you only want to include statements that are using SELECT. The search is case insensitive for ease of use. For example, the string 'SELECT' finds the same entries as the search string 'select'.

Multiple filter options can be specified. In a multi-filter case, the candidate entries for each filter are computed independently. Only those entries that are present in all the candidate lists are shown. For example, if you specified options **Minimum runtime for the longest execution of the statement** and **Statements that ran on or after this date and time**, you will be shown statements with the minimum runtime that ran on or after the specified date and time.

Related reference

[Index advisor](#)

The query optimizer analyzes the row selection in the query and determines, based on default values, if creation of a permanent index improves performance. If the optimizer determines that a permanent index might be beneficial, it returns the key columns necessary to create the suggested index.

Importing a monitor

You can import monitor data that has been collected using some other interface by using System i Navigator.

Monitors that are created using the **Start Database Monitor (STRDBMON)** command are automatically registered with System i Navigator. They are also included in the list of monitors displayed by System i Navigator.

To import monitor data, right-click **SQL Performance monitors** and select **Import**. Once you have imported a monitor, you can analyze the data.

Index advisor

The query optimizer analyzes the row selection in the query and determines, based on default values, if creation of a permanent index improves performance. If the optimizer determines that a permanent index might be beneficial, it returns the key columns necessary to create the suggested index.

The optimizer is able to perform radix index probe over any combination of the primary key columns, plus one additional secondary key column. Therefore it is important that the first secondary key column is

the most selective secondary key column. The optimizer uses radix index scan with any of the remaining secondary key columns. While radix index scan is not as fast as radix index probe, it can still reduce the number of keys selected. It is recommended that secondary key columns that are fairly selective are included.

Determine the true selectivity of any secondary key columns and whether you include those key columns in the index. When building the index, make the primary key columns the left-most key columns, followed by any of the secondary key columns chosen, prioritized by selectivity.

After creating the suggested index and executing the query again, it is possible that the query optimizer will choose not to use the suggested index. It does not include join, ordering, and grouping criteria. The SQE optimizer includes selection, join, ordering, and grouping criteria when suggesting indexes. Local selection advice can now factor in both AND and OR predicates with the qualifications mentioned below.

You can access index advisor information in many different ways. These ways include:

- The index advisor interface in System i Navigator
- SQL performance monitor Show statements
- Visual Explain interface
- Querying the Database monitor view 3020 - Index advised.

Related reference

[Overview of information available from Visual Explain](#)

You can use Visual Explain to view many types of information.

[Database monitor view 3020 - Index advised \(SQE\)](#)

Displays the SQL logical view format for database monitor QQQ3020.

[Viewing statements in a monitor](#)

You can view SQL statements that are included in a detailed monitor.

Index advice and OR predicates

Index advice generation to handle OR predicates

Index Advisor has been extended to include queries that OR together local selection (WHERE clause) columns over a single table. OR advice requires two or more indexes to be created as a dependent set.

If any of the OR'd indexes are missing, the optimizer won't be able to cost and choose these dependent indexes for implementation of the OR based query.

This relationship between OR based indexes in the SYSIXADV index advice table is with a new **DEPENDENT_ADVICE_COUNT** column.

Some restrictions with this support:

- OR'd predicate advice appears only if no other advice is generated
- Maximum of 5 predicates OR'd together
- Advised for files with OR'd local selection that get costed in the primary join dial when optimizing a join query

When Index Advisor shows highly dependent advice, use of the exact match capability from Show Statements to find the query in the plan cache is helpful. Once found, use Visual Explain to discover the dependent index advice specific to that query.

Index Or Advice example

- Should advise indexes over all OR'd predicate columns
- All 3 advised indexes will have **DEPENDENT_ADVICE_COUNT>0**
- Execution with indexes should produce bitmap implementation and register no new advice

```
SELECT orderkey, partkey, suppkey,  
linenumber, shipmode orderpriority
```

```
FROM item_fact
WHERE OrderKey <= 10 OR
SuppKey <= 10 OR
PartKey <= 10
OPTIMIZE FOR ALL ROWS
```

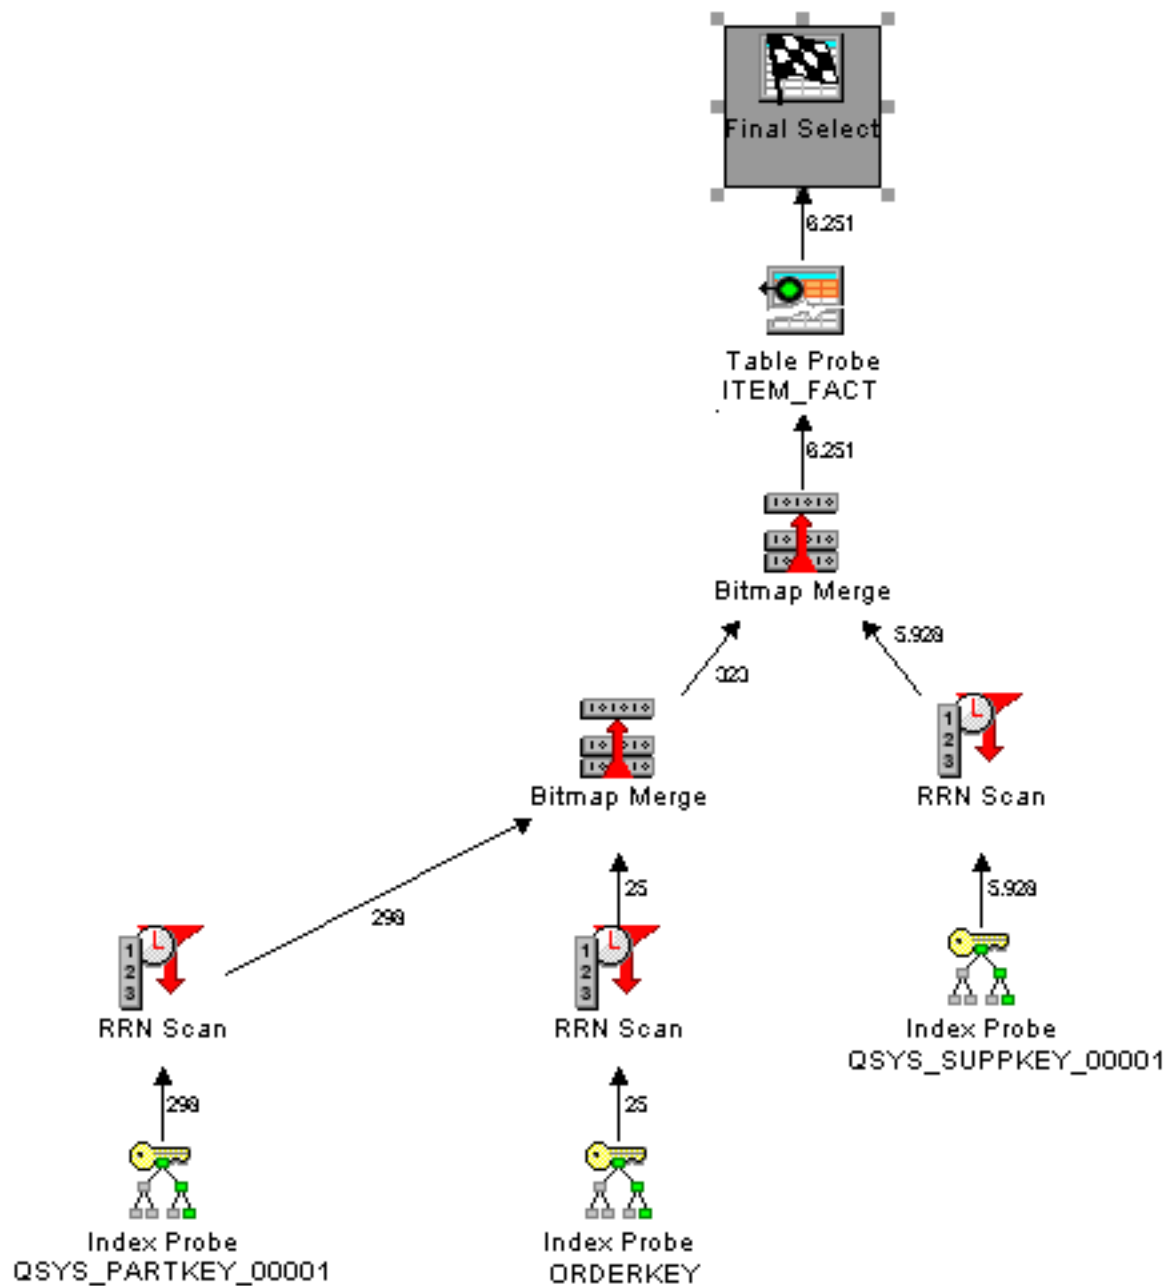
The graphic below shows the Index advice for the OR'd predicate columns:

The screenshot shows the 'Index and Statistics Advisor' window. The 'Index Advisor' tab is active. A message states: 'It is recommended that the following indexes be created:'. Below this is a table with the following data:

Create	Table Name	Schema	Index Type	Columns
<input checked="" type="checkbox"/>	ITEM_FACT	AAA	Binary Radix	SUPPKEY
<input checked="" type="checkbox"/>	ITEM_FACT	AAA	Binary Radix	PARTKEY
<input checked="" type="checkbox"/>	ITEM_FACT	AAA	Binary Radix	ORDERKEY

At the bottom right of the window is a 'Create ...' button.

The graphic below depicts the Visual Explain showing the implementation of merge bitmap representation using the OR'd advice indexes:



Index advice for Encoded Vector Index RRN Probe Plans

Index Advisor may advise EVI indexes for use with the EVI RRN Probe plan.

EVI RRN Probe Plan advice requires one or more indexes to be created as a dependent set. If any of the EVI indexes are missing, the optimizer won't be able to cost and choose these dependent indexes for implementation of the EVI RRN Probe.

In the SYSIXADV advice table, all the dependent EVI based indexes will have a value greater than 0 in the DEPENDENT_ADVICE_COUNT column. In addition, the REASON_ADVISED column will have reason code 'I8'.

The advice is on a per table basis for the query.

Indexes are only advised when the following is true:

- At least one of the columns in the select list of the query needs to have an existing single column EVI.
- The number of columns that do not match to an existing EVI must be less than 20.

- All columns must be eligible to be an index key.

When Index Advisor shows highly dependent advice, use of the exact match capability from Show Statements to find the query in the plan cache is helpful. Once found, use Visual Explain to discover the dependent index advice specific to that query.

EVI RRN Probe Plan Advice example


In the following example, assume the following indexes already exist.

- A radix index with a key of SUPPKEY
- An EVI Index with a key of PARTKEY


```
SELECT PARTKEY, QUANTITY, EXTENDEDPRICE
FROM ITEM_FACT
WHERE SUPPKEY = 50;
```

The index advisor will advise EVI indexes over all over missing columns: QUANTITY and EXTENDEDPRICE. The 2 advised indexes will have index type set to “EVI” and DEPENDENT_ADVICE_COUNT>0.

The following graphic shows the EVI advice generated for the query:



The following Visual Explain diagram shows the query implementation after the advised EVI indexes have been created:



Displaying index advisor information

You can display index advisor information from the optimizer using System i Navigator.

System i Navigator displays information found in the QSYS2/SYSIXADV system table.

To display index advisor information, follow these steps:

1. In the System i Navigator window, expand the system that you want to use.
2. Expand **Databases**.
3. Right-click the database that you want to work with and select **Index Advisor > Index Advisor**.

You can also find index advisor information for a specific schema or a specific table by right-clicking on a schema or table object.

Once you have displayed the information, you have several options. You can create an index from the list, remove the index advised from the list, or clear the list entirely. You can also right-click on an index and select **Show SQL**, launching a Run SQL Scripts session with the index creation statement. Finally, you can right-click on an advised index and select Show Statements. With additional information automatically provided in the advised index filter for the Plan Cache search, the resulting SQL statements shown will be a better match to the original queries that generated that specific index advice.

Depending on if you are viewing the index advice at the database level or the schema level your list could be large. Once you have the list displayed, follow these steps to subset your list:

1. Go to the **View** menu option, and select **Customize this view > Include ...**
2. Enter the information you would like to filter the list by.
3. Press the **OK** button to get the refreshed list of index advice.

Database manager indexes advised system table

This topic describes the indexes advised system table.

<i>Table 46. SYSIXADV system table</i>			
Column name	System column name	Data type	Description
TABLE_NAME	TBNAME	VARCHAR(258)	Table over which an index is advised
TABLE_SCHEMA	DBNAME	VARCHAR(128)	SQL schema containing the table
SYSTEM_TABLE_NAME	SYS_TNAME	CHAR(10)	System table name on which the index is advised
PARTITION_NAME	TBMEMBER	CHAR(10)	Partition detail for the index
KEY_COLUMNS_ADVISED	KEYSADV	VARCHAR(16000)	Column names for the advised index
LEADING_COLUMN_KEYS	LEADKEYS	VARCHAR(16000)	Leading, Order Independent keys. the keys at the beginning of the KEY_COLUMNS_ADVISED field which could be reordered and still satisfy the index being advised.
INDEX_TYPE	INDEX_TYPE	CHAR(14)	Radix (default) or EVI
LAST_ADVISED	LASTADV	TIMESTAMP	Last time this row was updated
TIMES_ADVISED	TIMESADV	BIGTINT	Number of times this index has been advised
ESTIMATED_CREATION_TIME	ESTTIME	INT	Estimated number of seconds for index creation
REASON_ADVISED	REASON	CHAR(2)	Coded reason why index was advised
LOGICAL_PAGE_SIZE	PAGESIZE	INT	Recommended page size for index
MOST_EXPENSIVE_QUERY	QUERYCOST	INT	Execution time in seconds of the query
AVERAGE_QUERY_ESTIMATE	QUERYEST	INT	Average execution time in seconds of the query
TABLE_SIZE	TABLE_SIZE	BIGINT	Number of rows in table when the index was advised
NLSS_TABLE_NAME	NLSSNAME	CHAR(10)	NLSS table to use for the index
NLSS_TABLE_SCHEMA	NLSSDBNAME	CHAR(10)	Schema name of the NLSS table
MTI_USED	MTIUSED	BIGINT	The number of times that this specific Maintained Temporary Index (MTI) has been used by the optimizer. The optimizer stops using a matching MTI once a permanent index is created.

Table 46. SYSIXADV system table (continued)

Column name	System column name	Data type	Description
MTI_CREATED	MTICREATED	INTEGER	The number of times that this specific Maintained Temporary Index (MTI) has been created by the optimizer. MTIs do not persist across system IPLs.
LAST_MTI_USED	LASTMTIUSE	TIMESTAMP	The timestamp representing the last time this specific Maintained Temporary Index (MTI) was used by the optimizer to improve the performance of a query. The MTI Last Used field can be blank. The blank field indicates that an MTI which exactly matches this advice has never been used by the queries which generated this index advice.
AVERAGE_QUERY_ESTIMATE_MICRO	QRYMICRO	BIGINT	Average execution time in microseconds of the query which drove the index advice
EVI_DISTINCT_VALUES	EVIVALS	INTEGER	Recommended value to use when creating the advised EVI index. This value is <i>n</i> within the WITH <i>n</i> DISTINCT VALUES clause on the CREATE INDEX SQL statement.
INCLUDE_COLUMNS	INCLCOL	CLOB(10000)	EVI INCLUDE expressions for index creation.
FIRST_ADVISED	FIRSTADV	TIMESTAMP	When this row was inserted.
SYSTEM_TABLE_SCHEMA	SYS_DNAME	CHAR(10)	System name of the table schema.
MTI_USED_FOR_STATS	MTISTATS	BIGINT	Number of times Maintained Temporary Index was used as a source for optimizer statistics.
LAST_MTI_USED_FOR_STATS	LASTMTISTA	TIMESTAMP	The timestamp representing the last time this specific Maintained Temporary Index was used as a source of statistics by the optimizer to improve the performance of a query.
DEPENDENT_ADVICE_COUNT	DEPCNT	BIGINT	The number of times this index advice was dependent upon other advice.

Index advisor column descriptions

Displays the columns that are used in the Index advisor window.

<i>Table 47. Columns used in Index advisor window</i>	
Column name	Description
Table for Which Index was Advised	The optimizer is advising creation of a permanent index over this table. This value is the long name for the table. The advice was generated because the table was queried and no existing permanent index could be used to improve the performance of the query.
Schema	Schema or library containing the table.
System Schema	System name of the schema.
System Name	System table name on which the index is advised
Partition	Partition detail for the index. Possible values: <ul style="list-style-type: none"> • <blank>, which means For all partitions • For Each Partition • specific name of the partition
Keys Advised	Column names for the advised index. The order of the column names is important. The names are listed in the same order as in the CREATE INDEX SQL statement. An exception is when the leading, order independent key information indicates that the ordering can be changed.
Leading Keys Order Independent	Leading, Order Independent keys. the keys at the beginning of the KEY_COLUMNS_ADVISED field which could be reordered and still satisfy the index being advised.
Index Type Advised	Radix (default) or EVI
Last Advised for Query Use	The timestamp representing the last time this index was advised for a query.
Times Advised for Query Use	The cumulative number of times this index has been advised. This count ceases to increase once a matching permanent index is created. The row of advice remains in this table until the user removes it
Estimated Index Creation Time	Estimated time in seconds to create this index.
Reason advised	Reason why index was advised. Possible values are: <ul style="list-style-type: none"> Row selection Ordering/Grouping Row selection and Ordering/Grouping
Logical Page Size Advised (KB)	Recommended page size to be used on the PAGESIZE keyword of the CREATE INDEX SQL statement when creating this index.
Most Expensive Query Estimate	Execution time in seconds of the longest running query which generated this index advice.
Average of Query Estimates	Average execution time in seconds of all queries that generated this index advice.

Table 47. Columns used in Index advisor window (continued)

Column name	Description
Rows in Table when Advised	Number of rows in table for the last time this index was advised.
NLSS Table Advised	The sort sequence table in use by the query which generated the index advice. For more detail on sort sequences:
NLSS Schema Advised	The schema of the sort sequence table.
MTI Used	The number of times that this specific Maintained Temporary Index (MTI) has been used by the optimizer.
MTI Created	The number of times that this specific Maintained Temporary Index (MTI) has been created by the optimizer. MTIs do not persist across system IPLs.
MTI Last Used	The timestamp representing the last time this specific Maintained Temporary Index (MTI) was used by the optimizer to improve the performance of a query. The MTI Last Used field can be blank. A blank field indicates that an MTI which exactly matches this advice has never been used by the queries which generated this index advice.
EVI Distinct Values	Recommended value to use when creating the advised EVI index. This value is n within the WITH n DISTINCT VALUES clause on the CREATE INDEX SQL statement.
First Advised	The date/time when a row is first added to the Index Advisor table for this advice.
MTI Used for Stats	The number of times that this specific Maintained Temporary Index (MTI) has been used by the optimizer.
MTI Last Used for Stats	The timestamp representing the last time this specific Maintained Temporary Index (MTI) was used as a source of statistics by the optimizer to improve the performance of a query. The MTI Last Used field can be blank.
Dependent Advice Count	<p>Dependent implies that this advised index is dependent on the creation of other dependent advised indexes and all of the other dependent indexes must be created in order for the index implementation to be costed and utilized.</p> <ul style="list-style-type: none"> • Zero - this advised index stands on its own. • Greater than Zero – Compare this column against the TIMES_ADVISED column to understand how often this advised index has been advised in conjunction with other indexes.

Querying database monitor view 3020 - Index advised

The index advisor information can be found in the Database Monitor view 3020 - Index advised (SQE).

The advisor information is stored in columns QQIDXA, QQIDXK, and QQIDXD. When the QQIDXA column contains a value of 'Y' the optimizer is advising you to create an index using the key columns shown in column QQIDXD. The intention of creating this index is to improve the performance of the query.

In the list of key columns contained in column QQIDXD, the optimizer has listed what it considers the suggested primary and secondary key columns. Primary key columns are columns that can significantly

reduce the number of keys selected based on the corresponding query selection. Secondary key columns are columns that might or might not significantly reduce the number of keys selected.

Column QQIDXK contains the number of suggested primary key columns that are listed in column QQIDXJ. These primary key columns are the left-most suggested key columns. The remaining key columns are considered secondary key columns and are listed in order of expected selectivity based on the query. For example, assuming QQIDXK contains the value of four and QQIDXJ specifies seven key columns, then the first four key columns are the primary key columns. The remaining three key columns are the suggested secondary key columns.

Condensing index advice

Many times, the index advisor advises several different indexes for the same table. You can condense these advised indexes into the best matches for your queries.

1. In the **System i Navigator** window, expand the system you want to use.
2. Expand **Databases**.
3. Right-click the database that you want to work with and select **Index Advisor > Condense Advised Indexes**.

Depending on if you are viewing the condensed index advice at the database level or the schema level your list could be large. Once you have the list displayed, follow these steps to subset your list:

1. Go to the **View** menu option, and select **Customize this view > Include ...**
2. Enter the information you would like to filter the list by.
3. Select **OK** to get the refreshed list of condensed index advice.

Table for Which Index was Advised	Schema	System Schema	System Name	Partition	Keys Advised	Advised Index Type	Last Advised for Query Use	Times Advised for Query Use	Estimated Index Creation	Logical Page Size Advised	Most Expensive Query Estimate	Average of Query Estimates	Rows in Table when Advised	NLSS Table Advised	MTI Used	MTI Creat	MTI Last Used
MQT1	SAMPLE	SAMPLE	MQT1		JOB	Binary Radix	9/1/09 3:57:14 PM	10	00:00:01	64	1	1	42	*HEX	4	3	9/1/09 2:58:10 PM
EMPLOYEE	SAMPLE	SAMPLE	EMPLOYEE		JOB, WORKDEPT	Binary Radix	9/1/09 4:06:54 PM	8	00:00:01	64	1	1	42	*HEX	0	0	
DEPARTMENT	SAMPLE	SAMPLE	DEPARTMENT		LOCATION, DEPTNAME	Binary Radix	9/1/09 3:53:05 PM	8	00:00:01	64	1	1	14	*HEX	0	0	
DEPARTMENT	SAMPLE	SAMPLE	DEPARTMENT		LOCATION, DEPTNO	Binary Radix	9/1/09 3:53:05 PM	8	00:00:01	64	1	1	14	*HEX	0	0	
DEPARTMENT	SAMPLE	SAMPLE	DEPARTMENT		DEPTNAME, LOCATION	Binary Radix	9/1/09 3:53:12 PM	6	00:00:01	64	1	1	14	*HEX	0	0	
MQT2	SAMPLE	SAMPLE	MQT2		LOCATION, DEPTNAME	Binary Radix	9/1/09 3:53:15 PM	4	00:00:01	64	1	1	8	*HEX	0	0	
MQT2	SAMPLE	SAMPLE	MQT2		DEPTNAME, SUM_SAL	Encoded vector (not unique)	9/1/09 3:53:15 PM	4	00:00:01	64	1	1	8	*HEX	0	0	

Viewing your queries with Visual Explain

You can use the **Visual Explain** tool with System i Navigator to create a query graph that graphically displays the implementation of an SQL statement. You can use this tool to see information about both static and dynamic SQL statements. Visual Explain supports the following types of SQL statements: SELECT, INSERT, UPDATE, and DELETE.

Queries are displayed using a graph with a series of icons that represent different operations that occur during implementation. This graph is displayed in the main window. In the lower portion of the pane, the SQL statement that the graph is based on is displayed. If Visual Explain is started from Run SQL Scripts, you can view the debug messages issued by the optimizer by clicking the **Optimizer messages** tab. The query attributes are displayed in the right pane.

Visual Explain can be used to graphically display the implementations of queries stored in the detailed SQL performance monitor. However, it does not work with tables resulting from the memory-resident monitor.

Starting Visual Explain

There are two ways to invoke the Visual Explain tool. The first, and most common, is through System i Navigator. The second is through the Visual Explain (QQQVEXPL) API.

You can start Visual Explain from any of the following windows in System i Navigator:

- Enter an SQL statement in the **Run SQL Scripts** window. Select the statement and choose **Explain** or **Run and Explain** from the **Visual Explain** menu.

- Expand the list of available SQL Performance Monitors. Right-click a detailed SQL Performance Monitor and choose the **Show Statements** option. Select filtering information and select the statement in the List of Statements window. Right-click and select **Visual Explain**. You can also start an SQL Performance Monitor from Run SQL Scripts. Select **Start SQL Performance monitor** from the **Monitor** menu.
- Start the SQL Details for Jobs function by right-clicking **Databases** and select **SQL Details for Jobs**. Click **Apply**. Select a job from the list and right-click and select **Show Details**. When the SQL is displayed in the lower pane, you can start Visual Explain by right-clicking on **Statement** and selecting **Visual Explain**.
- Right-click SQL Plan Cache and select **Show Statements**. Select filtering information and select the statement in the List of Statements window. Right-click and select **Visual Explain**.
- Expand the list of available SQL Plan Cache Snapshots. Right-click a snapshot and select **Show Statements**. Select filtering information and select the statement in the List of Statements window. Right-click and select **Visual Explain**.
- Expand the list of SQL Plan Cache Event Monitors. Right-click an event monitor and select **Show Statements**. Select filtering information and select the statement in the List of Statements window. Right-click and select **Visual Explain**.

You have three options when running Visual Explain from Run SQL Scripts.

Visual Explain only	This option allows you to explain the query without actually running it. The data displayed represents the estimate of the query optimizer. Note: Some queries might receive an error code 93 stating that they are too complex for displaying in Visual Explain. You can circumvent this error by selecting the "Run and Explain" option.
Run and Explain	If you select Run and Explain, the query is run by the system before the diagram is displayed. This option might take a significant amount of time, but the information displayed is more complete and accurate.
Explain while running	For long running queries, you can choose to start Visual Explain while the query is running. By refreshing the Visual Explain diagram, you can view the progress of the query.

In addition, a database monitor table that was not created as a result of using System i Navigator can be explained through System i Navigator. First you must import the database monitor table into System i Navigator. To import, right-click the SQL Performance Monitors and choose the **Import** option. Specify a name for the performance monitor (name it is known by within System i Navigator) and the qualified name of the database monitor table. Be sure to select Detailed as the type of monitor. Detailed represents the file-based (STRDBMON) monitor while Summary represents the memory-resident monitor (which is not supported by Visual Explain). Once the monitor has been imported, follow the steps to start Visual Explain from within System i Navigator.

You can save your Visual Explain information as an SQL Performance monitor. This monitor can be useful if you started the query from Run SQL Scripts and want to save the information for later comparison. Select **Save as Performance monitor** from the **File** menu.

Related information

[Visual Explain \(QQQVEXPL\) API](#)

Overview of information available from Visual Explain

You can use Visual Explain to view many types of information.

The information includes:

- Information about each operation (icon) in the query graph
- Highlight expensive icons
- The statistics and index advisor

- The predicate implementation of the query
- Basic and detailed information in the graph

Information about each operation (icon) in the query graph

As stated before, the icons in the graph represent operations that occur during the implementation of the query. The order of operations is shown by the arrows connecting the icons. If parallelism was used to process an operation, the arrows are doubled. Occasionally, the optimizer "shares" hash tables with different operations in a query, causing the lines of the query to cross.

You can view information about an operation by selecting the icon. Information is displayed in the **Attributes** table in the right pane. To view information about the environment, click an icon and then select **Display query environment** from the **Action** menu. Finally, you can view more information about the icon by right-clicking the icon and selecting **Help**.

Highlight expensive icons

You can highlight problem areas (expensive icons) in your query using Visual Explain. Visual Explain offers you two types of expensive icons to highlight: by processing time or number of rows. You can highlight icons by selecting **Highlight expensive icons** from the **View** menu.

The statistics and index advisor

During the query implementation, the optimizer can determine if statistics need to be created or refreshed, or if an index might make the query run faster. You can view these recommendations using the Statistics and Index Advisor from Visual Explain. Start the advisor by selecting **Advisor** from the **Action** menu. Additionally, you can begin collecting statistics or create an index directly from the advisor.

The predicate implementation of the query

Visual explain allows you to view the implementation of query predicates. Predicate implementation is represented by a blue plus sign next to an icon. You can expand this view by right-clicking the icon and selecting **Expand**, or open it into another window. Click an icon to view attributes about the operation. To collapse the view, right-click anywhere in the window and select **Collapse**. This function is only available on V5R3 or later systems.

The optimizer can also use the Look Ahead Predicate Generation to minimize the random the I/O costs of a join. To highlight predicates that used this method, select **Highlight LPG** from the **View** menu.

Basic and full information in the graph

Visual Explain also presents information in two different views: basic and full. The basic view only shows those icons that are necessary to understand the implementation of the SQL statement. It excludes some preliminary, or intermediate operations that are not essential for understanding the main flow of query implementation. The full view might show more icons that further depict the flow of the execution tree. You can change the graph detail by select **Graph Detail** from the **Options** menu and selecting either **Basic** or **Full**. The default view is **Basic**. In order to see all the detail for a **Full** view, change the Graph Detail to **Full**, close out Visual Explain, and run the query again. The setting for Graph Detail persists.

For more information about Visual Explain and the different options that are available, see the Visual Explain online help.

Refresh the Visual Explain diagram

For long running queries, you can refresh the visual explain graph with runtime statistical information before the query is complete. Refresh also updates the appropriate information in the attributes section of the icon shown on the right of the screen. In order to use the **Refresh** option, you need to select **Explain while Running** from the Run SQL Scripts window.

To refresh the diagram, select **Refresh** from the **View** menu. Or click the **Refresh** button in the toolbar.

Related reference

Index advisor

The query optimizer analyzes the row selection in the query and determines, based on default values, if creation of a permanent index improves performance. If the optimizer determines that a permanent index might be beneficial, it returns the key columns necessary to create the suggested index.

Adaptive Query Processing in Visual Explain

You can use Visual Explain to request a new plan.

There might be times when you are asked to performance tune a query while the query is still running. For instance, a query might be taking a long time to finish. After viewing the plan in Visual Explain, you decide to create the recommended index to improve the speed of the query. So you create the index and then want to signal the database optimizer to consider a new plan based on the new index.

Here are the steps to request the database engine to consider a new plan while running in Visual Explain:

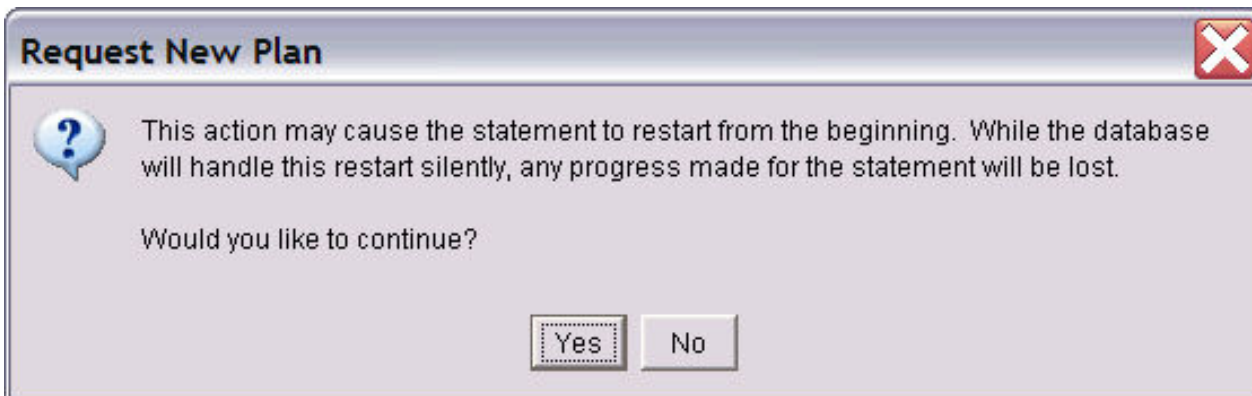
1. Open Run SQL Scripts.
2. Type in a query.
3. Go to the **Visual Explain** menu and select **Explain While Running**.
4. The Visual Explain window is displayed.
5. Next, go to the **Actions** menu and select **Request New Plan**.

The screenshot shows the Visual Explain window for a query on the 'Server1(Server1)' instance. The query is: `select * from qsys2.syscolumns where dbname = 'QSYS2'`. The query plan diagram shows a 'Final Select' operation connected to a 'Nested Loop Join' (2,335 rows), which is connected to another 'Nested Loop Join' (2,335 rows). This second join is connected to a 'Fetch N Rows' operation (2 rows), which is connected to a 'Union all' operation (<1 rows). The 'Union all' is connected to a 'Logic' operation (<1 rows). The 'Logic' is connected to an 'Index Probe' (<1 rows), which is connected to a 'Table Probe' (<1 rows). The 'Table Probe' is connected to another 'Index Probe' (2,635 rows), which is connected to a 'Table Probe' (2,088 rows). The 'Table Probe' is connected to another 'Index Probe' (2,635 rows). A red box highlights a message box that says: 'Request that the database engine consider a new plan for this statement'. The message box is located in the toolbar area of the window.

Attribute	Value
Time Information	
Timestamp for Creation of Monit...	2009
Statement Start Timestamp	2009
Statement End Timestamp	2009
Total Estimated Run Time (ms)	186.
Actual Runtime Information	
Optimization Time (ms)	88
Run Time (ms)	105
Statement Open Time (ms)	Not.
Statement Fetch Time (ms)	105
Statement Close Time (ms)	Not.
Rows Fetched	51
Total Times Query Was Run	2
Total Time For All Runs (ms)	211
Synchronous Database Reads	4
Asynchronous Database Reads	17
Page Faults	41
Information about SQL stateme...	
Statement Number	34,2
Statement Function	Sele

A message box appears.

Select **Yes** to restart the query.



The database optimizer considers any changes to the query environment, and determines whether it is appropriate to generate a new plan. It might be possible that the database optimizer decides it is better to continue using the existing plan.

Note: This capability could also be available when selecting Visual Explain of a statement in the SQL Details for a Job window, or the SQL Plan Cache Show Statements window.

Related reference

[Adaptive Query Processing](#)

Adaptive Query Processing analyzes actual query run time statistics and uses that information for subsequent optimizations.

Optimizing performance using the Plan Cache

The SQL Plan Cache contains a wealth of information about the SQE queries being run through the database. Its contents are viewable through the System i Navigator GUI interface. Certain portions of the plan cache can also be modified.

In addition, procedures are provided to allow users to programmatically work with the plan cache. These procedures can be invoked using the SQL CALL statement.

The Plan Cache interface provides a window into the database query operations on the system. The interface to the Plan Cache resides under the **System i Navigator > system name > Database**.

Within the SQL Plan Cache folder are two folders, SQL Plan Cache Snapshots and SQL Plan Cache Event Monitors.

Clicking the SQL Plan Cache Snapshots folder shows a list of any snapshots gathered so far. A snapshot is a database monitor file generated from the plan cache at the time a 'New Snapshot' is requested. It can be treated much the same as the SQL Performance Monitors list. The same analysis capability exists for snapshots as exists for traditional SQL performance monitors.

Clicking the SQL Plan Cache Event Monitors shows a list of any events that have been defined. Plan Cache event monitors, when defined, generate database monitor information from plans as they are being removed from the cache. The list includes currently active events as well as ones that have completed. Like a snapshot, the event monitor is a database monitor file. Consequently, the same analysis capability available to SQL performance monitors and snapshots can be used on the event file.

The plan cache is an actively changing cache. Therefore, it is important to realize that it contains timely information. If information over long periods of time is of interest, an event monitor could be defined to ensure that information is captured on any plans that are removed from the cache over time. Alternatively, you could consider implementing a method of performing periodic snapshots of the plan cache to capture trends and heavy usage periods. See the discussion on IBM supplied, callable SQL procedures later in this section on plan cache.

Related concepts

[Plan cache](#)

The plan cache is a repository that contains the access plans for queries that were optimized by SQE.

SQL Plan Cache - Show Statements

By right-clicking the SQL Plan Cache icon, a series of options are shown which allow different views of current plan cache of the database. The **SQL Plan Cache > Show Statements** option opens a screen with filtering capability. This screen provides a direct view of the current plan cache on the system.

Filters to apply:

- Minimum runtime for the longest execution of the statement: Seconds
- Statements that ran on or after this date and time:
- Top 'n' most frequently run statements:
- Top 'n' statements with the largest total accumulated runtime:
- Statements the following user has ever run:
- Statements that are currently active
- Statements for which an index has been advised
- Statements for which statistics have been advised
- Include statements initiated by the operating system
- Statements that reference the following objects:

Schema	Name
- Statements that contain the following text:

Statements:

Last Time Run	Most Expensive Time (sec)	T
8/27/09 10:12:55 PM	1372.3884	
8/27/09 3:09:42 PM	531.2134	
8/25/09 4:05:10 PM	324.2397	
8/26/09 6:05:33 PM	387.3384	
8/27/09 10:12:55 PM	87.8831	
8/27/09 3:17:07 PM	235.4072	
8/27/09 2:03:42 PM	98.5024	
8/27/09 2:09:28 PM	159.4982	
8/25/09 5:03:47 PM	119.6275	
8/25/09 5:03:47 PM	58.1911	
8/27/09 3:22:37 PM	106.4988	
8/25/09 5:03:47 PM	47.7330	
8/28/09 10:01:19 AM	62.9925	
8/28/09 10:04:53 AM	62.9820	
8/25/09 1:00:30 AM	62.9747	
8/28/09 5:38:33 AM	62.9747	
8/26/09 4:07:55 PM	44.6644	
8/25/09 2:15:53 AM	38.0097	
8/27/09 2:29:20 PM	17.4664	
8/25/09 1:05:35 PM	20.0067	
8/27/09 1:37:01 AM	3.9000	
8/25/09 1:01:40 AM	4.1502	
8/27/09 2:59:33 AM	3.0125	
8/26/09 9:23:47 PM	4.9244	
8/27/09 3:10:25 PM	9.8281	
8/27/09 2:24:32 PM	7.9627	
8/27/09 4:22:10 AM	2.7389	
8/26/09 4:13:25 PM	4.7626	
8/25/09 1:01:43 AM	0.6606	
8/25/09 12:47:53 AM	1.0081	
8/25/09 7:16:18 PM	7.4941	
8/27/09 3:01:23 AM	1.0579	
8/27/09 3:01:00 AM	0.8244	

Status: Complete

Buttons: Reset All Filters, Apply, Columns..., Refresh, Close, Help, ?

Press the **Apply** or **Refresh** button to display the current Plan Cache statements. The information shown includes the SQL query text, last time the query ran, most expensive single instance run, total processing time consumed, total number of times run, and information about the user and job that first created the plan entry.

The information also includes several per run averages, including average runtime, average result set size and average temporary storage usage. There is an adjusted average processing time which is the average discounting any anomalous runs.

The display also shows how many times, if any, that the database engine reused the results of a prior run, avoiding rerunning the entire statement. There is also a Save Results button (not shown) that allows you to save the statement list, for example, to a .csv file or spreadsheet.

Finally, the numeric identifier and plan score are also displayed. For more detail on the columns displayed, see [“SQL Plan Cache column descriptions” on page 173](#)

Statement Options

By highlighting one or more plans and right clicking, a menu with several possible actions appears.

Visual Explain Shows a visual depiction of the access plan and provides more detailed performance analysis. Note only one statement can be highlighted when performing this action.

Show Longest Runs Shows details of up to 10 of the longest running instances of that statement. Within the Longest Runs list, you can right click a statement and select **Visual Explain, Work With SQL Statement, Work With SQL Statement and Variables, Save to New...** snapshot or **Remove**. Snapshots are useful for capturing the information for that specific run in Visual Explain. Removing old or superfluous runs makes room to capture future runs. Only one statement can be highlighted when performing these actions. Any runs removed only affect which runs are shown in the list. The total time, total number of runs, and other information for the statement are still calculated including the runs removed from the list.

Show Active Jobs Displays a list of jobs on the system that are currently using that statement or statements.

Show User History Shows a list of all user IDs that have run that statement along with the last time they ran it.

Work with SQL Statement Displays a scripting window containing the SQL statement. The scripting window is useful for working with and tuning the statement directly, or for just viewing the statement in its own window. Only one statement can be highlighted when performing this action.

Work with SQL Statements and Variables Displays a scripting window containing the SQL Statement and any parameter markers entered with their specific values for that run of the SQL statement.

Save to New... Allows you to create a snapshot of the selected statements.

Plan Right-click to show options for modifying the plan:

Change Plan Score allows you to set the score to a specific value. The plan score is used to determine when a plan might be removed from the cache. A lower score plan is removed before a higher score plan. By setting the plan score high, the plan remains in the cache for a longer time. Setting the plan score to a low value causes the plan to be pruned sooner than might otherwise have occurred.

Delete allows you to remove the plan immediately from the cache. Note under normal circumstances there might not be a need to modify the attributes of a plan. Normal database processing ages and prunes plans appropriately. These modifying options are provided mostly as tools for minute analysis and for general interest.

The User and Job Name for each statement on the Statements screen is the user and job that created the initial plan with full optimization. This user is not necessarily the same as the last user to run that statement. The Longest Runs screen, however, does show the particular user and job for that individual run.

Filtering Options

The screen provides filtering options which allow the user to more quickly isolate specific criteria of interest. No filters are required to be specified (the default), though adding filtering shortens the time it takes to show the results. The list of statements that is returned is ordered so that the statement consuming the most total processing time is shown at the top. You can reorder the results by clicking the column heading for which you want the list ordered. Repeated clicking toggles the order from ascending to descending.

The filtering options provide a way to focus in on a particular area of interest:

Minimum runtime for the longest execution of the statement:	Show statements with at least one long individual statement instance runtime.
Statements that ran on or after this date and time:	Show statements that have been run recently.
Top 'n' most frequently run statements:	Show statements run most often.
Top 'n' statements with the largest total accumulated runtime:	Show the top resource consumers. Shows the first 'n' top statements by default when no filtering is given. Specifying a value for 'n' improves the performance of getting the first screen of statements, though the total statements displayed is limited to 'n'.
Statements the following user has ever run:	Show statements a particular user has run. The user and job name shown reflect the originator of the cached statement. This user is not necessarily the same as the user specified on the filter (there could be multiple users running the statement).
Statements that are currently active	Show statements that are still running or are in pseudo-close mode. The user and job name shown reflect the originator of the cached statement. This user is not necessarily the same as the user specified on the filter (there could be multiple users running the statement). Note: An alternative for viewing the active statement for a job is to right-click the Database icon and select SQL Details for Jobs...
Statements for which an index has been advised	Show only those statements where the optimizer advised an index to improve performance.
Statements for which statistics have been advised	Show only those statements where a statistic not yet collected might have been useful to the optimizer. The optimizer automatically collects these statistics in the background. This option is normally not that interesting unless, for whatever reason, you want to control the statistics collection yourself.
Include statements initiated by the operating system	Show the 'hidden' statements initiated by the database to process a request. By default the list only includes user-initiated statements.
Statements that reference the following objects:	Show statements that reference the tables or indexes specified.
Statements that contain the following text:	Show statements that include the text specified. This option is useful for finding particular types of statements. For example, statements with a FETCH FIRST clause can be found by specifying 'fetch'. The search is not case sensitive for ease of use. For example, the string 'FETCH' finds the same statements as the search string 'fetch'. This option provides a wildcard search capability on the SQL text itself.

Multiple filter options can be specified. The candidate statements for each filter are computed independently. Only those statements that are present in all the candidate lists are shown. For example,

you could specify options **Top 'n' most frequently run statements** and **Statements the following user has ever run**. The display shows those most frequently run statements in the cache that have been run by the specified user. It does not show the most frequently run statements by the user (unless those statements are also the most frequently run statements in the entire cache).

SQL Plan Cache column descriptions

Displays the columns that are used in the SQL Plan Cache Statements window.

<i>Table 48. Columns used in SQL Plan Cache Statements window</i>	
Column name	Description
Last Time Run	Displays the last time that this statement was run.
Most Expensive Time (sec)	The time taken for the longest run of this statement.
Total Processing Time (sec)	The sum total time that all runs of this statement took to process in seconds.
Total Times Run	The total number of times that this statement ran.
Average Processing Time (sec)	The average time per run that this statement took to process in seconds.
Statement	The statement text.
Plan Creation User Name	The name of the user id that created the plan.
Job Name	The name of the job that created the plan.
Job User	The name of the user id that owned the job that created the plan.
Job Number	The job number of the job that created the plan.
Adjusted Average Processing Time (sec)	The average time per run that this statement took to process in seconds where anomalous runs are removed from the average calculation. This time provides a realistic average for a statement by ignoring a single (or few) run that was atypical to the normal condition of the statement.
Average Result Set Rows	The average number of result set rows that are returned when this statement is run.
Average Temp Storage Used (MB)	The average amount of temporary storage used when this statement is run.
Plan Score	The rating of this plan relative to other plans in the cache. A plan with a higher rating relative to other plans remains in the cache for a longer time. A plan with a lower rating relative to other plans is removed from the cache sooner than the other plans.
Plan Identifier	A unique numeric identifier of the plan.
Total Cached Results Used	The number of times a result set from a prior run of the statement was reused on a subsequent run of the statement.
Optimization Time (sec)	The amount of time that it took to optimize this statement.
System Name	The system name.
Relational Database name	Relational database name

SQL plan cache properties

The Plan Cache tab of the SQL Performance Center in IBM i Access Client Solutions (ACS) shows high-level information about the SQL plan cache and overall query activity. This information includes cache size, number of plans, number of full opens and pseudo-opens that have occurred.

This information can be used to view overall database activity. If tracked over time, it provides trends to help you better understand the database utilization peaks and valleys throughout the day and week.

Several Plan Cache Properties can be changed by right-clicking a property and selecting Edit Value.

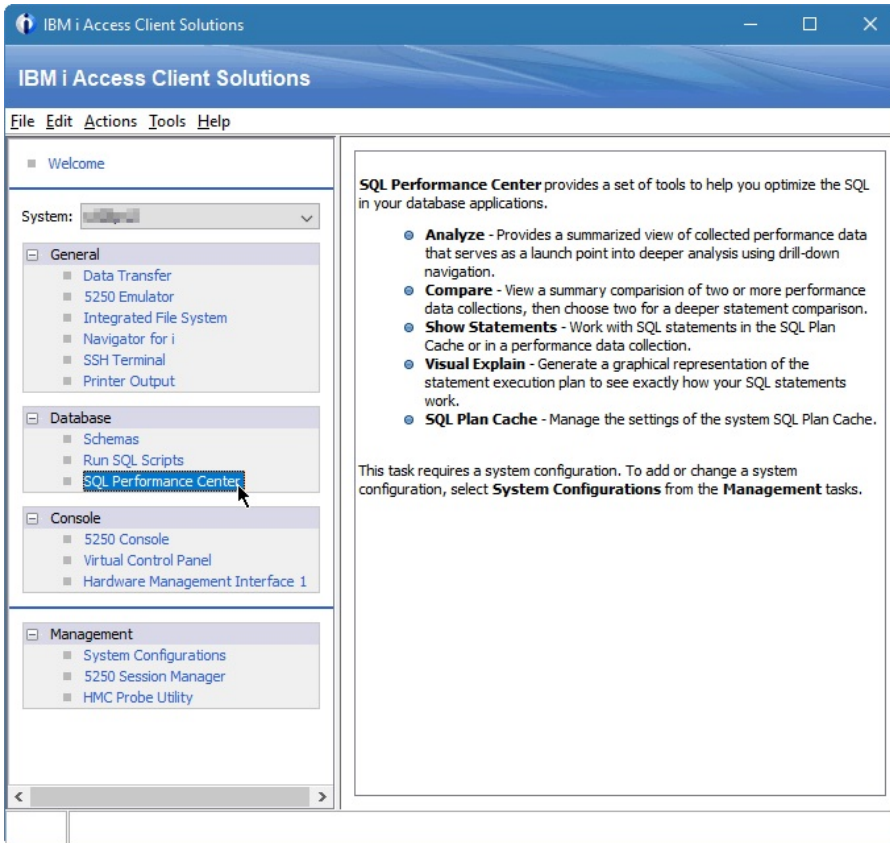
The following terms are used in regards to plan cache properties:

- Active Queries – queries that are currently open or in pseudo close mode.
- Full Open – describes a query run that requires a cursor to be built before the query executes and return rows. A query that reuses a cursor is called a Pseudo Open.
- Unique Queries – unique SQL query statements. For the plan cache, this is uniqueness once tables are resolved to their schemas.
- Hit Ratio – the percentage of time that the query optimizer, when searching the plan cache for a plan, finds an existing plan for the query.
- Target Hit Ratio – the hit ratio percentage that the database tries to achieve by adjusting the plan cache's size. A larger size means plans stay in the cache longer and are therefore more likely to be found and used for future runs of the query.
- Job Scoped – queries that reference tables or indexes that reside in the job's QTEMP library. This includes, for example, global temporary tables. By definition these job scoped plans and runtime objects cannot be reused across jobs.
- Temporary Runtime Objects – the actual runtime executable objects used to process a query. These include the execution tree (ROQ), hash tables, sorted lists, lists, and buffers. A certain number of these objects are cached with the query plan in the cache so they can be reused.
- Longest Runs - information kept about the longest running instances of a query.
- Activity Thresholds – highest or largest values that are tracked by the database.
- ...Since Start – the amount of activity since the cache was created (IPL).

Plan Cache Properties are divided into three main types:

- Summary and Usage – Information about plan usage, query usage and current conditions for the plan cache and queries.
- Configuration – properties that can be adjusted by the user.
- Activity Threshold – tracked 'high water mark' of several key indicators.

To view the Plan Cache properties, select the SQL Performance Center from the main ACS window or from the Tools menu of any ACS window.



The first tab in the SQL Performance Center window that appears will display all of the Plan Cache properties. Changes to configurable properties may be made by clicking the **Change Configuration...** button.

Description	Value	Value Unit
Plan Usage Summary		
Current Number of Plans in Cache	111372	
Total Number of Plans Built Since Start	1404089	
Total Number of SMP Plans Built Since Start	56242	
Total Number of Unique Queries Since Start	241083	
Current Plan Cache Size	23361	MB
Current Plan Cache Size Threshold	*AUTO	MB
Maximum Plan Cache Size For AutoSizing	*DEFAULT (24576)	MB
Current Plan Cache Hit Ratio	54	%
Target Plan Cache AutoSize Hit Ratio	*DEFAULT (90)	%
Total Number of Plan Cache Autosizing Adjustments	47	
Last Plan Cache AutoSizing Adjustment	2021-08-30-20.15.53.499818	
Last Autosizing Limited Due to Temporary Storage	0000-00-00-00.00.00.000000	
Current Number of Job Scoped (QTEMP) Plans	8035	
Total Number of Job Scoped (QTEMP) Plans Built Since Start	107758	
Total Number of Unique Queries With Job Scoped (QTEMP) Referen...	22668	
Total Times Plans Used from Cache	1659506	
Total Plans Removed	353754	
Total Plans Pruned	634311	
Number of Times Plan Cache Pruned	14	
Time Plan Cache was Last Pruned	2021-09-01-06.37.26.780238	
Current Number of Temporary Runtime Objects Stored in Cache	100291	
Current Total Size of Temporary Runtime Objects stored in Cache	50035	MB
Maximum Number of Temporary Runtime Objects Stored Per Plan	*DEFAULT (5)	
Total Number of Temporary Indexes Created	35353	
Current Number of Temporary Indexes	2949	
Current Total Size of Temporary Indexes	10104	MB
Number of Plans Rebuilt due to AQP	793	
Number of Query Mapping Errors Since Start	12498	

Done: 53 rows retrieved.

Table 49. Plan Cache Properties

Plan Cache Property	Description
Time Of Summary	Timestamp of when the properties were collected.
Plan Cache Creation Time	Timestamp of when the plan cache was created (IPL)
Active Query Summary	
Number of Currently Active Queries	Number of queries that are currently open or in pseudo close mode.
Number of Queries Run Since Start	Shows the total number of SQL queries run since IPL Note: This value includes job scoped queries.
Number of Query Full Opens Since Start	Number of queries run since IPL that required a cursor to be built before the query executed and returned rows.

Table 49. Plan Cache Properties (continued)

Plan Cache Property	Description
Plan Usage Summary	
Current Number of Plans in Cache	Current total number of plans in the plan cache
Total Number of Plans Built Since Start	Number of plans built since IPL. This includes the plans that have been pruned from the plan cache.
Total Number of SMP Plans Built Since Start	The number of plans built since IPL that run with SMP parallel processing.
Total Number of Unique Queries Since Start	This value reflects the total number of unique statements (SQL queries) run since IPL Note: This value includes unique job scoped queries
Current Plan Cache Size	Current size in MB of the plan cache. This does not include the size of cached temporary runtime objects.
Current Plan Cache Size Threshold	<p>The current maximum allowed size of the plan cache.</p> <p>This property is configurable.</p> <ul style="list-style-type: none"> • *AUTO indicates that the database manager will manage the maximum size of the plan cache. • A user specified value between 50 and 51200. Size is specified in MB.
Maximum Plan Cache Size For AutoSizing	<p>If AutoSizing is active, the maximum plan cache size.</p> <p>This property is configurable if the Current Plan Cache Size Threshold is *AUTO.</p> <ul style="list-style-type: none"> • *DEFAULT(nn) - The database manager determines, at IPL, the maximum size that the plan cache can grow to under autosizing. Only applicable if Size Threshold is set to *AUTO. The database determined size is shown in parentheses. • nn – A user specified size between 50 and 51200. Size is specified in MB. While supported, it should rarely be changed from *DEFAULT. • *DISABLED - Indicates that plan cache auto sizing has been disabled.
Current Plan Cache Hit Ratio	<p>The percentage of time the query optimizer found a matching plan in the plan cache.</p> <p>This value indicates the efficiency of the plan cache. The higher the percentage the better.</p>

Table 49. Plan Cache Properties (continued)

Plan Cache Property	Description
Target Plan Cache AutoSize Hit Ratio	<p>The target hit ratio percentage that the database manager tries to meet by adjusting the plan cache size.</p> <p>This property is configurable if the Current Plan Cache Size Threshold is *AUTO.</p> <ul style="list-style-type: none"> • DEFAULT(nn) – The database manager sets the target hit ratio. The database determined ratio is shown in parentheses. • nn – percentage from 1 to 99. While supported, it should rarely be changed from *DEFAULT. • *DISABLED - Indicates that plan cache auto sizing has been disabled.
Current Number of Job Scoped (QTEMP) Plans	<p>Current number of plans in the plan cache that for queries that reference tables or indexes that reside in the job's QTEMP library. This includes, for example, global temporary tables. By definition these job scoped plans and runtime objects cannot be reused across jobs.</p>
Total Number of Job Scoped (QTEMP) Plans Built Since Start	<p>Total number of plans built for queries that reference tables or indexes that reside in the job's QTEMP library.</p>
Total Number of Unique Queries With Job Scoped (QTEMP) References Since Start	<p>This value reflects the total number of unique statements (SQL queries) referencing temporary files that have been run since IPL.</p>
Total Times Plans Used from Cache	<p>Total number of plans that were reused from the plan cache. (i.e. Plans that did not require a reoptimization).</p>
Total Plans Pruned	<p>Total number of plans removed from the plan cache.</p>
Current Number of Temporary Runtime Objects Stored in Cache	<p>Current Number of Temporary Runtime Objects Stored in Cache</p>
Current Total Size of Temporary Runtime Objects stored in Cache	<p>Current Total Size of Temporary Runtime Objects stored in Cache</p>

Table 49. Plan Cache Properties (continued)

Plan Cache Property	Description
Maximum Number of Temporary Runtime Objects Stored Per Plan	<p>Maximum number of Temporary Runtime Objects stored per plan.</p> <p>This property is configurable.</p> <ul style="list-style-type: none"> • *DEFAULT(nn) - database determines the maximum number of runtime objects (ROQs) to keep per plan. The database determined number is shown in parentheses. • nn – A user specified value between 1 and 50 that is the maximum number of runtime objects to keep per plan. A runtime ‘object’ is all the runtime constructs (except the cursor) used to execute the query. It includes the ROQ, hash tables, sorts, etc... The database will automatically clear big hash tables and sorts of data contents (leaving only their shell) before storing them with the plan. However, smaller hashes and sorts will retain their contents. Setting this value smaller will lessen the temporary storage usage on the machine. Setting this value higher can improve performance by having more runtime objects ready to go during full opens rather than having to build them.
Total Number of Temporary Indexes Created	Total number of SQE created temporary indexes (MTIs) since IPL.
Current Number of Temporary Indexes	Current number of SQE created temporary indexes (MTIs) on the system.
Current Total Size of Temporary Indexes	The total size of all SQE created temporary indexes (MTIs) currently on the system.
Number of Plans Rebuilt due to AQP	Number of plans rebuilt to AQP.
Number of Query Mapping Errors Since Start	The number of SQE query mapping errors that have occurred since the last IPL. While a number of mapping errors greater than 0 does not indicate a problem, a significant number of mapping errors can negatively affect performance and may require further investigation.

Table 49. Plan Cache Properties (continued)

Plan Cache Property	Description
Maximum Number of Longest Runs Allowed Per Plan	The Maximum Number of Longest Runs Allowed Per Plan determines how many longest runs are maintained per plan. This property is configurable. <ul style="list-style-type: none"> • *DEFAULT(nn) - The database manager determines the maximum number of longest run entries to keep per plan. The database determined number is shown in parentheses. • nn – A user specified value between 1 and 50 indicating the maximum number of longest runs information to keep per plan Note: These can be seen by bringing up a show statements of the live cache, right clicking, and selecting Longest Runs
Plan Cache Activity Thresholds The Activity Thresholds group of properties track the upper thresholds of activity for both plans and query activity. The values represent the maximum values achieved since the Threshold Start Time. Each threshold shows both the maximum value and the timestamp of when that maximum was reached.	
Activity Thresholds Start Time	The time from which the tracking started. This value can be reset (to zero) to reset all the thresholds and restart tracking. Note: All thresholds are reset at one time, thresholds cannot be reset individually. Start Time restarts each IPL.
Highest Number of Active Queries at One Time	The highest number of open plus pseudo closed cursors (queries) at a given point in time
Highest Number of Plans in Cache	The largest number of plans in the plan cache at a given point in time
Highest Number of Temporary Runtime Objects Stored in Cache	The highest number of inactive runtime executable objects stored (cached) away (for future reuse) in the plan cache at a given point in time
Largest Total Size of Temporary Runtime Objects Stored in Cache	The largest amount of temporary storage, in MB, consumed by the inactive runtime executable objects stored (cached) away in the plan cache at a given point in time
Query Supervisor The Query Supervisor group of properties provide historical information about the interaction of the Query Supervisor with queries running on the system.	

Table 49. Plan Cache Properties (continued)

Plan Cache Property	Description
Number of thresholds reached	The number of Query Supervisor thresholds that have been reached since IPL.
Most recent job to reach a threshold	The name and timestamp of the most recent job to reach a Query Supervisor threshold.
Number of queries terminated	The number of queries that have been terminated since IPL at the request of a Query Supervisor exit program.
Most recent job to have a query terminated	The name and timestamp of the most recent job to have a query terminated at the request of a Query Supervisor exit program.
Most recent exit point program failure	The name and timestamp of the most recent job to encounter an error when using a Query Supervisor exit program.

Creating SQL plan cache snapshots

The **New > Snapshot** option allows for the creation of a snapshot from the plan cache.

Unlike the snapshot option under Show Statements, it allows you to create a snapshot without having to first view the queries.

New SQL Plan Cache Snapshot - Serv...

Name:

Schema:

Include all plan cache entries

Include plan cache entries that meet the following criteria

Filters to apply:

Minimum runtime for the longest execution of the statement:
 Seconds

Statements that ran on or after this date and time:

Top 'n' most frequently run statements:

Top 'n' statements with the largest total accumulated runtime:

Statements the following user has ever run:

Statements that are currently active

Statements for which an index has been advised

Statements for which statistics have been advised

Include statements initiated by the operating system

Statements that reference the following objects:

Schema	Name	
	Schema	

Statements that contain the following text:

The same filtering options are provided here as on the Show Statements screen.

The stored procedure, `qsys2.dump_plan_cache`, provides the simplest, programmatic way to create a database monitor file output (snapshot) from the plan cache. The `dump_plan_cache` procedure takes two parameters, library name and file name, for identifying the resulting database monitor file. If the file does

not exist, it is created. For example, to dump the plan cache to a database performance monitor file in library QGPL:

```
CALL qsys2.dump_plan_cache('QGPL', 'SNAPSHOT1');
```

SQL plan cache event monitor

The SQL plan cache event monitor records changes in your plan cache.

You can access the SQL plan cache event monitor through the System i Navigator interface or by calling the procedures directly.

The SQL plan cache event monitor captures monitor records of plans as they are removed from the plan cache. The event monitor is useful for ensuring that all performance information potentially available in the cache is captured even if plans are removed from the cache. Combining the event monitor output with a plan cache snapshot provides a composite view of the cache from when the event monitor was started until the snapshot is taken.

The event monitor allows the same filtering options as described for **Show statements** and **NewSnapshot**. The exceptions are that the *Top 'n' most frequently run statements* and the *Top 'n' statements with largest total accumulated runtime* are not shown. Once a statement is removed from the cache, it is no longer compared to other plans. Therefore, these two 'Top n' filters do not make sense for pruned plans.

Accessing the SQL plan cache with SQL stored procedures

The System i Navigator provides a visual interface into the plan cache. However, the plan cache is also accessible through stored procedures which can be called using the SQL CALL statement.

See [“Plan Cache Services” on page 338](#) for the descriptions of these procedures.

Verifying the performance of SQL applications

You can verify the performance of an SQL application by using commands.

The commands that can help you verify performance:

- | | |
|--|--|
| Display Job (DSPJOB) | <p>You can use the Display Job (DSPJOB) command with the OPTION(*OPNF) parameter to show the indexes and tables used by an application running in a job.</p> <p>You can also use DSPJOB with the OPTION(*JOBLOCK) parameter to analyze object and row lock contention. It displays the objects and rows that are locked and the name of the job holding the lock.</p> <p>Specify the OPTION(*CMTCTL) parameter on the DSPJOB command to show the isolation level, the number of rows locked during a transaction, and the pending DDL functions. The isolation level displayed is the default isolation level. The actual isolation level, used for any SQL program, is specified on the COMMIT parameter of the CRTSQLxxx command.</p> |
| Print SQL Information (PRTSQLINF) | <p>The Print SQL Information (PRTSQLINF) command lets you print information about the embedded SQL statements in a program, SQL package, or service program. The information includes the SQL statements, access plans used, and the command parameters used to precompile the source member.</p> |
| Start Database Monitor (STRDBMON) | <p>You can use the Start Database Monitor (STRDBMON) command to capture to a file information about every SQL statement that runs.</p> |
| Change Query Attribute (CHGQRYA) | <p>You can use the Change Query Attribute (CHGQRYA) command to change the query attributes for the query optimizer. Among the attributes that can be changed by this command are the predictive query governor, parallelism, and the query options.</p> |

Start Debug (STRDBG)

You can use the **Start Debug (STRDBG)** command to put a job into debug mode, and optionally add as many as 20 programs, 20 class files, and 20 service programs to debug mode. It also specifies certain attributes of the debugging session. For example, it can specify whether database files in production libraries can be updated while in debug mode.

Related information

[Display Job \(DSPJOB\) command](#)

[Print SQL Information \(PRTSQLINF\) command](#)

[Start Database Monitor \(STRDBMON\) command](#)

[Change Query Attributes \(CHGQRYA\) command](#)

[Start Debug \(STRDBG\) command](#)

Examining query optimizer debug messages in the job log

Query optimizer debug messages issue informational messages to the job log about the implementation of a query. These messages explain what happened during the query optimization process.

For example, you can learn:

- Why an index was or was not used
- Why a temporary result was required
- Whether joins and blocking are used
- What type of index was advised by the optimizer
- Status of the job queries
- Indexes used
- Status of the cursor

The optimizer automatically logs messages for all queries it optimizes, including SQL, call level interface, ODBC, OPNQRYF, and SQL Query Manager.

Viewing debug messages using STRDBG command:

STRDBG command puts a job into debug mode. It also specifies certain attributes of the debugging session. For example, it can specify whether database files in production schemas can be updated while in debug mode. For example, use the following command:

```
STRDBG PGM(Schema/program) UPDPROD(*YES)
```

STRDBG places in the job log information about all SQL statements that run.

Viewing debug messages using QAQQINI table:

You can also set the QRYOPTLIB parameter on the **Change Query Attributes (CHGQRYA)** command to a user schema where the QAQQINI table exists. Set the parameter on the QAQQINI table to MESSAGES_DEBUG, and set the value to *YES. This option places query optimization information in the job log. Changes made to the QAQQINI table are effective immediately and affect all users and queries that use this table. Once you change the MESSAGES_DEBUG parameter, all queries that use this QAQQINI table write debug messages to their respective job logs. Pressing F10 from the command Entry panel displays the message text. To see the second-level text, press F1 (Help). The second-level text sometimes offers hints for improving query performance.

Viewing debug messages in Run SQL Scripts:

To view debug messages in Run SQL Scripts, from the **Options** menu, select **Include Debug Messages in Job Log**. Then from the **View** menu, select **Job Log**. To view detailed messages, double-click a message.

Viewing debug messages in Visual Explain:

In Visual Explain, debug messages are always available. You do not need to turn them on or off. Debug messages appear in the lower portion of the window. You can view detailed messages by double-clicking a message.

Print SQL Information

The **Print SQL Information (PRTSQLINF)** command returns information about the embedded SQL statements in a program, SQL package (used to store the access plan for a remote query), or service program. This information is then stored in a spooled file.

PRTSQLINF provides information about:

- The SQL statements being executed
- The type of access plan used during execution. How the query is implemented, indexes used, join order, whether a sort is used, whether a database scan is used, and whether an index is created.
- A list of the command parameters used to precompile the source member for the object.
- The CREATE PROCEDURE and CREATE FUNCTION statement text used to create external procedures or User Defined Functions.

This output is like the information that you can get from debug messages. However, while query debug messages work at runtime, **PRTSQLINF** works retroactively. You can also see this information in the second-level text of the query governor inquiry message CPA4259.

You can issue **PRTSQLINF** in a couple of ways. First, you can run the **PRTSQLINF** command against a saved access plan. You must execute or at least prepare the query (using the SQL PREPARE statement) before you use the command. It is best to execute the query because the index created as a result of PREPARE is relatively sparse. It could well change after the first run. **PRTSQLINF**'s requirement of a saved access plan means that the command cannot be used with **OPNQRYF**.

You can also run **PRTSQLINF** against functions, stored procedures, triggers, SQL packages, and programs from System i Navigator. This function is called Explain SQL. To view **PRTSQLINF** information, right-click an object and select **Explain SQL**.

Related information

[Print SQL Information \(PRTSQLINF\) command](#)

Query optimization tools: Comparison

Use this table to find the information each tool can provide, when it analyzes your queries, and the tasks it can do to improve your queries.

PRTSQLINF	STRDBG or CHGQRYA	File-based monitor (STRDBMON)	Memory-Based Monitor	Visual Explain
Available without running query (after access plan has been created)	Only available when the query is run	Only available when the query is run	Only available when the query is run	Only available when the query is explained
Displayed for all queries in SQL program, whether executed or not	Displayed only for those queries which are executed	Displayed only for those queries which are executed	Displayed only for those queries which are executed	Displayed only for those queries that are explained
Information about host variable implementation	Limited information about the implementation of host variables	All information about host variables, implementation, and values	All information about host variables, implementation, and values	All information about host variables, implementation, and values

Table 50. Optimization tool comparison (continued)

PRTSQLINF	STRDBG or CHGQRYA	File-based monitor (STRDBMON)	Memory-Based Monitor	Visual Explain
Available only to SQL users with programs, packages, or service programs	Available to all query users (OPNQRYF , SQL, QUERY/400)	Available to all query users (OPNQRYF , SQL, QUERY/400)	Available only to SQL interfaces	Available through System i Navigator Database and API interface
Messages are printed to spool file	Messages are displayed in job log	Performance rows are written to database table	Performance information is collected in memory and then written to database table	Information is displayed visually through System i Navigator
Easier to tie messages to query with subqueries or unions	Difficult to tie messages to query with subqueries or unions	Uniquely identifies every query, subquery, and materialized view	Repeated query requests are summarized	Easy to view implementation of the query and associated information

Changing the attributes of your queries

You can modify different types of query attributes for a job with the **Change Query Attributes (CHGQRYA)** CL command. You can also use the System i Navigator Change Query Attributes interface.

Related concepts

[Plan cache](#)

The plan cache is a repository that contains the access plans for queries that were optimized by SQE.

[Objects processed in parallel](#)

The Db2 Symmetric multiprocessing feature provides the optimizer with additional methods for retrieving data that include parallel processing. Symmetrical multiprocessing is a form of parallelism achieved on a single system where multiple CPU and I/O processors sharing memory and disk work simultaneously toward a single result.

Related information

[Change Query Attributes \(CHGQRYA\) command](#)

Controlling queries dynamically with the query options file QAQQINI

The query options file QAQQINI support provides the ability to dynamically modify or override the environment in which queries are executed. This modification is done through the **Change Query Attributes (CHGQRYA)** command and the QAQQINI file. The query options file QAQQINI is used to set some attributes used by the database manager.

For each query that is run the query option values are retrieved from the QAQQINI file in the schema specified on the QRYOPTLIB parameter of the CHGQRYA CL command and used to optimize or implement the query.

Environmental attributes that you can modify through the QAQQINI file include:

- ACTIVE_JOBS
- ALLOW_ADAPTIVE_QUERY_PROCESSING
- ALLOW_ARRAY_VALUE_CHANGES
- ALLOW_DDL_CHANGES_WHILE_OPEN
- ALLOW_EVI_ONLY_ACCESS
- ALLOW_TEMPORARY_INDEXES
- APPLY_REMOTE

- ASYNC_JOB_USAGE
- CACHE_RESULTS
- COLLATE_ERRORS
- COMMITMENT_CONTROL_LOCK_LIMIT
- CONCURRENT_ACCESS_BEHAVIOR
- DETERMINISTIC_UDF_SCOPE
- FIELDPROC_ENCODED_COMPARISON
- FORCE_JOIN_ORDER
- IGNORE_LIKE_REDUNDANT_SHIFTS
- KEY_RANGE_ESTIMATE_TIMEOUT
- LIMIT_PREDICATE_OPTIMIZATION
- LOB_LOCATOR_THRESHOLD
- MATERIALIZED_QUERY_TABLE_REFRESH_AGE
- MATERIALIZED_QUERY_TABLE_USAGE
- MEMORY_POOL_PREFERENCE
- MESSAGES_DEBUG
- NORMALIZE_DATA
- OPEN_CURSOR_CLOSE_COUNT
- OPEN_CURSOR_THRESHOLD
- OPTIMIZATION_GOAL
- OPTIMIZE_STATISTIC_LIMITATION
- PARALLEL_DEGREE
- PARAMETER_MARKER_CONVERSION
- PSEUDO_OPEN_CHECK_HOST_VARS
- QUERY_TIME_LIMIT
- REOPTIMIZE_ACCESS_PLAN
- SQE_NATIVE_ACCESS
- SQE_NATIVE_ACCESS_POSITION_BEHAVIOR
- SQLSTANDARDS_MIXED_CONSTANT
- SQL_CONCURRENT_ACCESS_RESOLUTION
- SQL_DECFLOAT_WARNINGS
- SQL_FAST_DELETE_ROW_COUNT
- SQL_GVAR_BUILD_RULE
- SQL_MODIFIES_SQL_DATA
- SQL_PSEUDO_CLOSE
- SQL_STMT_COMPRESS_MAX
- SQL_STMT_REUSE
- SQL_SUPPRESS_MASKED_DATA_DETECTION
- SQL_SUPPRESS_WARNINGS
- SQL_TRANSLATE_ASCII_TO_JOB
- SQL_XML_DATA_CCSID
- STAR_JOIN
- STORAGE_LIMIT

- SYSTEM_SQL_STATEMENT_CACHE
- SYSTIME_PERIOD_ADJ
- TEXT_SEARCH_DEFAULT_TIMEZONE
- UDF_TIME_OUT
- VARIABLE_LENGTH_OPTIMIZATION

Specifying the QAQQINI file with CHGQRYA

Use the **Change Query Attributes (CHGQRYA)** command with the QRYOPLIB (query options library) parameter to specify which schema currently contains or contains the query options file QAQQINI.

The query options file is retrieved from the schema specified on the QRYOPLIB parameter for each query. It remains in effect for the duration of the job or user session, or until the QRYOPLIB parameter is changed by the **Change Query Attributes (CHGQRYA)** command.

If the **Change Query Attributes (CHGQRYA)** command is not issued, or is issued without the QRYOPLIB parameter specified, QUSRSYS is searched for the QAQQINI file. If a query options file is not found, no attributes are modified. Since the system ships without an INI file in QUSRSYS, you might receive a message indicating that there is no INI file. This message is not an error but an indication that a QAQQINI file that contains all default values is being used. The initial value of the QRYOPLIB parameter for a job is QUSRSYS.

Related information

[Change Query Attributes \(CHGQRYA\) command](#)

Creating the QAQQINI query options file

Each system is shipped with a QAQQINI template file in schema QSYS. The QAQQINI file in QSYS is to be used as a template when creating all user specified QAQQINI files.

To create your own QAQQINI file, use the **Create Duplicate Object (CRTDUPOBJ)** command. Create a copy of the QAQQINI file in the schema specified on the **Change Query Attributes (CHGQRYA)** QRYOPLIB parameter. The file name must remain QAQQINI. For example:

```
CRTDUPOBJ OBJ(QAQQINI)
          FROMLIB(QSYS)
          OBJTYPE(*FILE)
          TOLIB(MYLIB)
          DATA(*YES)
          TRG(*YES)
```

System-supplied triggers are attached to the QAQQINI file in QSYS therefore it is imperative that the only means of copying the QAQQINI file is through the CRTDUPOBJ CL command. If another means is used, such as **CPYF**, then the triggers could be corrupted. An error is signaled that the options file cannot be retrieved or that the options file cannot be updated.

Because of the trigger programs attached to the QAQQINI file, the following CPI321A informational message is displayed six times in the job log when the **CRTDUPOBJ** CL is used to create the file. These messages are not an error; they are only informational messages.

CPI321A Information Message: Trigger QSYS_TRIG_&1__QAQQINI__00000&N in library &1 was added to file QAQQINI in library &1. The ampersand variables (&1, &N) are replacement variables that contain either the library name or a numeric value.

Note: It is highly recommended that the file QAQQINI, in QSYS, not be modified. This file is the original template that is duplicated into QUSRSYS or a user specified library for use.

Related information

[Change Query Attributes \(CHGQRYA\) command](#)

[Create Duplicate Object \(CRTDUPOBJ\) command](#)

QAQQINI file override support

If you find working with the QAQQINI query options file cumbersome, consider using the QSYS2.OVERRIDE_QAQQINI procedure. Instead of creating, managing, and using a QAQQINI *FILE object directly, this procedure can be called to work with a temporary version of the INI file. It uses user-specified options and values. The support relies upon the QTEMP library, so any changes affect only the job which calls the procedure.

See [OVERRIDE_QAQQINI procedure](#) for more information.

QAQQINI query options file format

The QAQQINI file is shipped in the schema QSYS. It has a predefined format and has been pre-populated with the default values for the rows.

Query Options File:

A				UNIQUE
A	R QAQQINI			TEXT('Query options + file')
A	QQPARM	256A		VARLEN(10) + TEXT('Query + option parameter') + COLHDG('Parameter')
A	QQVAL	256A		VARLEN(10) + TEXT('Query option + parameter value') + COLHDG('Parameter Value')
A	QQTEXT	1000G		VARLEN(100) + TEXT('Query + option text') + ALWNULL + COLHDG('Query Option' + 'Text') + CCSID(13488) + DFT(*NULL)
A	K QQPARM			

Setting the options within the query options file

The QAQQINI file query options can be modified with the INSERT, UPDATE, or DELETE SQL statements.

For the following examples, a QAQQINI file has already been created in library MyLib. To update an existing row in MyLib/QAQQINI use the UPDATE SQL statement. This example sets MESSAGES_DEBUG = *YES so that the query optimizer prints out the optimizer debug messages:

```
UPDATE MyLib/QAQQINI SET QQVAL='*YES'  
WHERE QQPARM='MESSAGES_DEBUG'
```

To delete an existing row in MyLib/QAQQINI use the DELETE SQL statement. This example removes the QUERY_TIME_LIMIT row from the QAQQINI file:

```
DELETE FROM MyLib/QAQQINI  
WHERE QQPARM='QUERY_TIME_LIMIT'
```

To insert a new row into MyLib/QAQQINI use the INSERT SQL statement. This example adds the QUERY_TIME_LIMIT row with a value of *NOMAX to the QAQQINI file:

```
INSERT INTO MyLib/QAQQINI  
VALUES('QUERY_TIME_LIMIT','*NOMAX','New time limit set by DBAdmin')
```

QAQQINI query options file authority requirements

QAQQINI is shipped with a *PUBLIC *USE authority. This authority allows users to view the query options file, but not change it. Changing the values of the QAQQINI file affects all queries run on the system. Allow only the system or database administrator to have *CHANGE authority to the QAQQINI query options file.

The query options file, which resides in the library specified on the **Change Query Attributes (CHGQRYA)** CL command QRYOPTLIB parameter, is always used by the query optimizer. It is used even

if the user has no authority to the query options library and file. This authority provides the system administrator with an additional security mechanism.

When the QAQQINI file resides in the library QUSRSYS the query options affects all the query users on the system. To prevent anyone from inserting, deleting, or updating the query options, the system administrator must remove update authority from *PUBLIC to the file. This update authority prevents users from changing the data in the file.

A copy of the QAQQINI file can also reside in a user library. If that library is specified on the QRYOPLIB parameter of the **Change Query Attributes (CHGQRYA)** command, the query options affect all the queries run for that user job. To prevent the query options from being retrieved from a particular library the system administrator can revoke authority to the **Change Query Attributes (CHGQRYA)** CL command.

QAQQINI file system-supplied triggers

The query options file QAQQINI file uses a system-supplied trigger program in order to process any changes made to the file. A trigger cannot be removed from or added to the file QAQQINI.

If an error occurs on the update of the QAQQINI file (an INSERT, DELETE, or UPDATE operation), the following SQL0443 diagnostic message is issued:

```
Trigger program or external routine detected an error.
```

QAQQINI query options

There are different options available for parameters in the QAQQINI file.

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

Parameter	Value	Description
ACTIVE_JOBS This option specifies what should be used for active jobs during optimization. The number of active jobs is used in determining the optimizer fair share of memory.	*DEFAULT	Average active jobs will be computed based on system information. This is the recommended setting.
	Integer value (1 :: 8192)	The number of active jobs that will be used by the query optimizer to determine fair share. Specify a value between 1 and 8192. Specifying a value other than default means that you may not get all of the performance benefits from the optimizer determining a plan using the memory fair share based on actual system usage. Usage of this setting is intended primarily for debug or modeling query plans based on an optimizer fair share.
ALLOW_ADAPTIVE_QUERY_PROCESSING Specifies whether Adaptive Query Processing (AQP) processing is done for a query. Adaptive query processing uses runtime statistics to look for poor performing queries and potentially replace the poor plan with an improved plan.	*DEFAULT	The default value is set to *YES.
	*YES	Allows Adaptive query processing to occur for this query. The existing QAQQINI options that affect AQP are the following: <ul style="list-style-type: none"> If the REOPTIMIZE_ACCESS_PLAN QAQQINI option is set to *ONLY_REQUIRED, AQP does not reoptimize the original plan. *ONLY_REQUIRED indicates the user does not want the query reoptimized unless there is a functional reason to do so. *ONLY_REQUIRED takes precedence over AQP. Join order requirements specified by the user in the FORCE_JOIN_ORDER QAQQINI option take precedence over AQP. If the user specifies the primary table in the join order, any AQP primary recommendations will be placed after the primary table if they are different.
	*NO	Adaptive query processing cannot be used for this query.

Table 51. Query Options Specified on QAAQINI Command (continued)

Parameter	Value	Description
<p>ALLOW_ARRAY_VALUE_CHANGES</p> <p>Specifies whether changes to the values of array elements are visible to the query while the query is running.</p>	*DEFAULT	The default value is set to *NO.
	*NO	<p>Do not allow changes to values in arrays referenced in the query to be visible after the query is opened.</p> <p>All values which could be referenced in a query are copied during query open processing. Any changes to values in arrays after the query is opened are not visible.</p> <p>Produces queries with predictable and reproducible results, but might have a performance penalty when working with large arrays or large array elements. The penalty is less if all the references to arrays are simple non-column values, for example, :ARRAY[1] or :ARRAY[:hv2].</p> <p>Use of column values from a table to index the ARRAY, or using the UNNEST() function results in copies of the entire array being made. These copies have the largest performance penalty.</p>
	*YES	<p>Allow changes to values in arrays to be visible to the query while the query is running. The arrays are not copied during the open processing of the query. If the array values are changed during the processing of queries, the results of the query might be unpredictable.</p> <p>Performance might be improved for queries which reference large arrays in complex array index lookup operations, such as :Array[column-name], or when using UNNEST. Large arrays include arrays that have thousands of elements, or elements with a large size. Array index lookups using simple index values, such as :ARRAY[1] or :ARRAY[:hv2], see minimal performance improvements.</p> <p>Performance of some queries might be negatively impacted. For example, later queries that could reuse the results if they were cached to avoid recalculation where the cached result is applicable.</p> <p>Procedures that can run with *YES and still expect predictable results have the following characteristics:</p> <ol style="list-style-type: none"> 1. Contain no cursor declarations. 2. Receive arrays as input parameters: <ul style="list-style-type: none"> • and do not contain SET statements which reference arrays on the left side of the SET, and • and have no SQL statements with INTO clauses referencing arrays. 3. Do not contain SET statements which reference arrays on the left side of the set: <ul style="list-style-type: none"> • and have no SQL statements with INTO clauses referencing arrays while a cursor is open for a query which references an array.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
ALLOW_DDL_CHANGES_WHILE_OPEN Specifies whether certain DDL operations against a database file are allowed while another user has the database file (or table) open or a logical file (or view) over a physical file open in another job.	*DEFAULT	The default is set to *NO.
	*YES	The following DDL operations are supported: <ul style="list-style-type: none"> • Trigger operations <ul style="list-style-type: none"> – SQL CREATE TRIGGER, ALTER TRIGGER, and DROP TRIGGER Any open cursors in concurrent jobs that have open cursors over the target table will continue to operate as if the operation was not performed until the cursor is closed. – SQL COMMENT ON TRIGGER and LABEL ON TRIGGER statements – CL ADDPFTRG, RMVPFTRG, and CHGPFTRG commands There is no effect on concurrent jobs. Any open cursors in concurrent jobs that have open cursors over the target table will continue to operate as if the operation was not performed until the cursor is closed. This option is ignored for INSTEAD OF triggers and READ triggers. <ul style="list-style-type: none"> • Grant and revoke operations <ul style="list-style-type: none"> – SQL GRANT and REVOKE for database files Any fully open cursors in concurrent jobs that have open cursors over the target database file will continue to operate as if the operation was not performed until the cursor is closed. If this is not a grant of UPDATE, DELETE, or INSERT, pseudo-closed cursors in other jobs will not be closed. Otherwise, all pseudo-closed cursors will be fully closed. – GRTOBJAUT and RVKOBJAUT CL commands for database files Any fully open cursors in concurrent jobs that have open cursors over the target database file will continue to operate as if the operation was not performed until the cursor is closed. If this is not a grant of *UPD, *DLT, *ADD, or *EXCLUDE, pseudo-closed cursors in other jobs will not be closed. Otherwise, all pseudo-closed cursors will be fully closed.
	*NO	SQL DDL operations can fail to complete, with an SQL0913 error, if an exclusive lock cannot be acquired for the target object.
ALLOW_EVI_ONLY_ACCESS Specifies whether encoded vector index RRN probe can be considered by the optimizer.	*DEFAULT	The default is set to *YES.
	*YES	Specifies whether encoded vector index RRN probe can be considered by the optimizer to replace table accesses. An EVI must exist for every column being accessed in the table.
	*NO	Encoded vector index RRN probes cannot replace table access.
ALLOW_TEMPORARY_INDEXES Specifies whether temporary indexes can be considered by the optimizer. If temporary indexes are not allowed, then any other viable plan is chosen regardless of cost to implement this query.	*DEFAULT	The default value is set to *YES.
	*YES	Allow temporary indexes to be considered.
	*ONLY_REQUIRED	Do not allow any temporary indexes to be considered for this access plan. Choose any other implementation regardless of cost to avoid the creation of a temporary index. Only if no viable plan can be found, is a temporary index allowed.
APPLY_REMOTE Specifies for database queries involving distributed files, whether the CHGQRYA query attributes are applied to the jobs on the remote systems associated with this job.	*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 use the attributes associated to them on their systems.
	*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 system 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.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
ASYNCH_JOB_USAGE Specifies the circumstances in which asynchronous (temp writer) jobs can be used to help process database queries in the job. The option determines which types of database queries can be used in asynchronous jobs (running in parallel) to help complete the query. An asynchronous job is a separate job that handles query requests from jobs running the database queries on the system. The asynchronous job processes each request and puts the results into a temporary file. This intermediate temporary file is then used by the main job to complete the database query. The advantage of an asynchronous job is that it processes its request at the same time (in parallel) that the main job processes another query step. The disadvantage of using an asynchronous job is that it might encounter a situation that it cannot handle in the same way as the main job. For example, the asynchronous job might receive an inquiry message from which it cancels, whereas the main job can choose to ignore the message and continue. There are two different types of database queries that can run asynchronous jobs: 1. Distributed queries. These are database queries that involve distributed files. Distributed files are provided through the system feature Db2 Multi-System for IBM i. 2. Local queries. There are database queries that involve only files local to the system where the database queries are being run.	*DEFAULT	The default value is set to *LOCAL.
	*LOCAL	Asynchronous jobs might be used for database queries that involve only tables local to the system where the database queries are being run. In addition, this option allows the communications required for queries involving distributed tables to be asynchronous. Each system involved in the query of the distributed tables can run its portion of the query at the same time (in parallel).
	*DIST	Asynchronous jobs might be used for database queries that involve distributed tables.
	*ANY	Asynchronous jobs might 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 intersystem parallel processing occurs.
CACHE_RESULTS Specifies a way for the user to control the cache results processing. For queries involving temporary results, for example, sorts or hashes, the database manager often saves the results across query pseudo-close or pseudo-open. The results are saved as long as they are not large, with the hope that they can be reused for the next run of the query. Beginning in V5R3, the database manager saves these temporary results even when a job is finished with them. The database manager assumes that another job can later reuse the results. The database manager automatically controls the caching of these results, removing cache results as storage usage becomes large. However, the amount of temporary storage used by the database can be noticeably more than in previous releases.	*DEFAULT	The default value is the same as *SYSTEM.
	*SYSTEM	The database manager might cache a query result set. A subsequent run of the query by the same job can reuse the cached result set. Or, if the ODP for the query has been deleted, any job can reuse the cached result set. In many cases, this option works well. However, you need to consider if the query is doing work outside of the database manager which could affect temporary results. In that case, *JOB or *NONE may be a more appropriate setting. For example, if field procedures that mask data are used or swapping of user profiles in a UDF can occur, this option should specify *NONE.
	*JOB	The database manager might cache a query result set from one run to the next for a job. Caching can occur as long as the query uses a reusable ODP. When the reusable ODP is deleted, the cached result set is destroyed. This value mimics V5R2 processing.
	*NONE	The database does not cache any query results.
COLLATE_ERRORS Specifies how data errors are handled on the GROUP BY and ORDER BY expression during hash or sort processing within queries.	*DEFAULT	The default value is *NO.
	*NO	A value of *NO causes the query to be ended with an error when a grouping or ordering expressions results in an error.
	*YES	A value of *YES indicates that the grouping or sort continues.
COMMITMENT_CONTROL_LOCK_LIMIT Specifies the maximum number of records that can be locked to a commit transaction initiated after setting the new value. The value specified for COMMITMENT_CONTROL_LOCK_LIMIT does not affect transactions running in jobs that have already started commitment control. For the value to be effective, it must be changed before starting commitment control.	*DEFAULT	*DEFAULT is equivalent to 500,000,000. If multiple journals are involved in the transaction, the COMMITMENT_CONTROL_LOCK_LIMIT applies to each journal, not to the transaction as a whole. For example, files F1 to F5 are journaled to journal J1, and files F6 to F10 are journaled to J2. The COMMITMENT_CONTROL_LOCK_LIMIT is set to 100,000. 100,000 record locks can be acquired for files F1 to F5. 100,000 more locks can be acquired for files F6 to F10.
	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.

Table 51. Query Options Specified on QAAQINI Command (continued)

Parameter	Value	Description
CONCURRENT_ACCESS_BEHAVIOR Controls how queries with an isolation level of Cursor Stability (CS) or Read Stability (RS) interact with uncommitted table changes.	*DEFAULT	The default value is *OPTIMIZE
	*OPTIMIZE	Uncommitted changes that delete or update records so that they are no longer selected by the query will not be considered as candidates for query synchronization.
	*STRICTSCAN	All records referenced by a table scan query access plan will synchronize with any changes that are not yet committed. Serialization behavior depends on the concurrent access resolution used by the query, for example, SKIP LOCKED DATA, USE CURRENTLY COMMITTED, or WAIT FOR OUTCOME (default). Since the table scan attempts to serialize with any pending transactions for deleted and non-selected records, query performance will be reduced, as compared to *OPTIMIZE. Queries may contain many access plan types, but this option is only supported for table scan access. All other plan types will use *OPTIMIZE behavior.
DETERMINISTIC_UDF_SCOPE Specifies the scope or lifetime of the deterministic setting for User Defined Functions (UDFs) and User Defined Table Functions (UDTFs). It is recommended that you specify STATEMENT DETERMINISTIC on any CREATE FUNCTION statement that should be considered deterministic for a single instance of a query open rather than using the *OPEN option. DETERMINISTIC_UDF_SCOPE applies to all deterministic UDFs and UDTFs in every query while this QAAQINI option is in effect.	*DEFAULT	The default value is *ALWAYS.
	*ALWAYS	The UDF is always considered deterministic. Query temporary objects might be shared across query opens and the UDF might not run for a particular query open.
	*OPEN	The UDF is considered deterministic only for a single instance of a query open. Query temporary objects are not shared across query open. The UDF is run at least once in the query for a given set of input parameters.
FIELDPROC_ENCODED_COMPARISON Specifies the amount of optimization that the optimizer might use when queried columns have attached field procedures	*DEFAULT	The default value is *ALLOW_EQUAL.
	*NONE	No optimization to remove field procedure decode option 4 or transformations to optimize field procedure invocations is allowed. For example, the optimizer cannot transform <code>fieldProc(4, column) = 'literal'</code> to <code>column = fieldProc(0, 'literal')</code> . This option is used when the field procedure is not deterministic.
	*ALLOW_EQUAL	Optimization allowed for equal and not equal predicates, GROUP BY, and DISTINCT processing. For example, the optimizer might choose to change the predicate <code>fieldProc(4, column) = 'literal'</code> to <code>column = fieldProc(0, 'literal')</code> in order to facilitate index matching. This option is useful when the field procedure is deterministic but no ordering can be determined based on the result of the field encoding.
	*ALLOW_RANGE	Transformation allowed for MIN, MAX grouping functions, ORDER BY, and all predicates except LIKE in addition to the transformations supported by *ALLOW_EQUAL. This option is useful when the field procedure is deterministic and the encoded value implies ordering
	*ALL	Transformation allowed for all predicates including LIKE, in addition to the transformations supported by *ALLOW_RANGE.
FORCE_JOIN_ORDER Specifies to the query optimizer that the join of files is to occur in the order specified in the query.	*DEFAULT	The default is set to *NO.
	*NO	Allow the optimizer to reorder join tables.
	*SQL	Only force the join order for those queries that use the SQL JOIN syntax. This option mimics the behavior for the optimizer before V4R4M0.
	*PRIMARY nnn	Only force the join position for the file listed by the numeric value nnn into the primary position (or dial) for the join. nnn is optional and defaults to 1. The optimizer then determines the join order for all the remaining files based upon cost.
	*YES	Do not allow the query optimizer to specify the order of join tables as part of its optimization process. The join occurs in the order in which the tables were specified in the query.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
IGNORE_LIKE_REDUNDANT_SHIFTS Specifies whether redundant shift characters are ignored for DBCS-Open operands when processing the SQL LIKE predicate or OPNQRYF command %WLDCRD built-in function.	*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. The optimizer cannot use an index to perform key row positioning for SQL LIKE or OPNQRYF %WLDCRD predicates involving DBCS-Open, DBCS-Either, or DBCS-Only operands.
	*OPTIMIZE	When processing the SQL LIKE predicate or the OPNQRYF command %WLDCRD built-in function, redundant shift characters might be ignored for DBCS-Open operands. These characters are ignored depending on whether an index is used to perform key row positioning for these predicates. This option enables the query optimizer to consider key row positioning for SQL LIKE or OPNQRYF %WLDCRD predicates involving DBCS-Open, DBCS-Either, or DBCS-Only operands.
KEY_RANGE_ESTIMATE_TIMEOUT Specifies the amount of time the query optimizer may use for any individual key range estimate operation. Key range estimates are used with indexes to approximate the number of rows for a given predicate. This option may help to reduce the time spent waiting for some queries to complete optimization.	*DEFAULT	The default value is *OPTIMIZE.
	*OPTIMIZE	The amount of time used for key range estimate operations is determined by the query optimizer.
	*NONE	No time limit is specified for key range estimate operations, and every estimate will run to completion regardless of the time required. This is the behavior of the query optimizer in releases before IBM i 7.3.
	Integer Value	The number of seconds a key range estimate operation may execute before returning an estimate to the query optimizer. At the end of this time interval, the optimizer will continue optimizing the query, and the estimate operation will continue in a background process. A smaller value may reduce optimization time but may also cause less accurate estimates to be used by the optimizer. Valid values are 1-86400.
LIMIT_PREDICATE_OPTIMIZATION Specifies that the query optimizer can only use simple isolatable predicates (OIF) when performing its index optimization. An OIF is a predicate that can eliminate a record without further evaluation. Any predicate that cannot be classified as an OIF is ignored by the optimizer and needs to be evaluated as a non-key selection predicate. A=10 and (A => 10 AND B=9) are OIFs. A=10 OR B=9 are not OIFs. Note: *YES impairs or limits index optimization.	*DEFAULT	Do not eliminate the predicates that are not simple isolatable predicates (OIF) when doing index optimization. Same as *NO.
	*NO	Do not eliminate the predicates that are not simple isolatable predicates (OIF) when doing index optimization.
	*YES	Eliminate the predicates that are not simple isolatable predicates (OIF) when doing index optimization.
LOB_LOCATOR_THRESHOLD Specifies either *DEFAULT or an Integer Value -- the threshold to free eligible LOB locators that exist within the job.	*DEFAULT	The default value is set to 0. This option indicates that the database does not free locators.
	Integer Value	If the value is 0, then the database does not free locators. For values 1 through 250,000, on a FETCH request, the database compares the SQL current LOB locator count for the job against the threshold value. If the locator count is greater than or equal to the threshold, the database frees 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 Specifies the ability to examine which materialized query tables are eligible to be used based on the last time a REFRESH TABLE statement was run.	*DEFAULT	The default value is set to 0.
	0	No materialized query tables can be used.
	*ANY	Any tables indicated by the MATERIALIZED_QUERY_TABLE_USAGE INI parameter can be used.
	Timestamp_duration	Only tables indicated by MATERIALIZED_QUERY_TABLE_USAGE INI option which have a REFRESH TABLE performed within the specified timestamp duration can be used.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
MATERIALIZED_QUERY_TABLE_USAGE Specifies the usage of materialized query tables in query optimization and runtime.	*DEFAULT	The default value is set to *NONE.
	*NONE	Materialized query tables cannot be used in query optimization and implementation.
	*ALL	User-maintained materialized query tables may be used.
	*USER	User-maintained materialized query tables can be used.
MEMORY_POOL_PREFERENCE Specifies the preferred memory pool that database operations uses. This option does not guarantee use of the specified pool, but directs database to perform its paging into this pool when supported by the database operation.	*DEFAULT	The default value is set to *JOB.
	*JOB	Paging is done in the pool of the job. This option is normal paging behavior.
	*BASE	Attempt to page storage into the base pool when paging is needed and a database operation that supports targeted paging occurs.
	nn	Attempt to page storage into pool nn when paging is needed and a database operation that supports targeted paging occurs.
	*NAME PoolName	Attempt to page storage into a named storage pool when paging is needed and a database operation that supports targeted paging occurs.
	*PRIVATE Library/Subsystem/PoolNumber	Attempt to page storage into a private storage pool in specified library and subsystem when paging is needed and a database operation that supports targeted paging occurs.
MESSAGES_DEBUG Specifies whether Query Optimizer debug messages are displayed to the job log. These messages are regularly issued when the job is in debug mode.	*DEFAULT	The default is set to *NO.
	*NO	No debug messages are to be displayed.
	*YES	Issue all debug messages that are generated for STRDBG .
NORMALIZE_DATA Specifies whether normalization is performed on Unicode constants, host variables, parameter markers, and expressions that combine strings.	*DEFAULT	The default is set to *NO.
	*NO	Unicode constants, host variables, parameter markers, and expressions that combine strings is not normalized.
	*YES	Unicode constants, host variables, parameter markers, and expressions that combine strings is normalized
OPEN_CURSOR_CLOSE_COUNT Specifies either *DEFAULT or an Integer Value: the number of cursors to full close when the threshold is encountered.	*DEFAULT	*DEFAULT is equivalent to 0. See Integer Value for details.
	Integer Value	This value determines the number of cursors to be closed. The valid values for this parameter are 1 - 65536. The value for this parameter is less than or equal to the number in the OPEN_CURSOR_THRESHOLD parameter. If the number of open cursors reaches the value specified by the OPEN_CURSOR_THRESHOLD, pseudo-closed cursors are hard (fully) closed. The least recently used cursors are closed first. This value is ignored if OPEN_CURSOR_THRESHOLD is *DEFAULT. If OPEN_CURSOR_THRESHOLD is specified and the value is *DEFAULT, the number of cursors closed is equal to OPEN_CURSOR_THRESHOLD multiplied by 10 percent. The result is rounded up to the next integer value. OPEN_CURSOR_CLOSE_COUNT is used with OPEN_CURSOR_THRESHOLD to manage the number of open cursors within a job. Open cursors include pseudo-closed cursors.
OPEN_CURSOR_THRESHOLD Specifies either *DEFAULT or an Integer Value -- the threshold to start full close of pseudo-closed cursors.	*DEFAULT	*DEFAULT is equivalent to 0. See Integer Value for details.
	Integer Value	This value determines the threshold to start full close of pseudo-closed cursors. When the number of open 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. A default value of 0 indicates that there is no threshold. Hard closes are not forced based on the number of open cursors within a job. OPEN_CURSOR_THRESHOLD is used with OPEN_CURSOR_CLOSE_COUNT to manage the number of open cursors within a job. Open cursors include pseudo-closed cursors.

Table 51. Query Options Specified on QAAQINI Command (continued)

Parameter	Value	Description
OPTIMIZATION_GOAL Specifies the goal that the query optimizer uses when making costing decisions.	*DEFAULT	Optimization goal is determined by the interface (ODBC, SQL precompiler options, OPTIMIZE FOR nnn ROWS clause).
	*FIRSTIO	All queries are optimized with the goal of returning the first page of output as fast as possible. This option works well when the output is controlled by a user likely to cancel the query after viewing the first page of data. Queries coded with OPTIMIZE FOR nnn ROWS honor the goal specified by the clause.
	*ALLIO	All queries are optimized with the goal of running the entire query to completion in the shortest amount of elapsed time. This option is better when the output of a query is written to a file or report, or the interface is queuing the output data. Queries coded with OPTIMIZE FOR nnn ROWS honor the goal specified by the clause.
OPTIMIZE_STATISTIC_LIMITATION Specifies limitations on the statistics gathering phase of the query optimizer. One of the most time consuming aspects of query optimization is in gathering statistics from indexes associated with the queried tables. Generally, the larger the size of the tables involved in the query, the longer the gathering phase of statistics takes. This option provides the ability to limit the amount of resources spend during this phase of optimization. The more resources spent on statistics gathering, the more accurate (optimal) the optimization plan is.	*DEFAULT	The amount of time spent in gathering index statistics is determined by the query optimizer.
	*NO	No index statistics are gathered by the query optimizer. Default statistics are used for optimization. (Use this option sparingly.)
	*PERCENTAGE integer value	Specifies the maximum percentage of the index that is searched while gathering statistics. Valid values for are 1 - 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 does not gather statistics and uses default values.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
PARALLEL_DEGREE Specifies the parallel processing option that can be used when running database queries and database file keyed access path builds, rebuilds, and maintenance in the job. The specified parallel processing option determines the types of parallel processing allowed. There are two types of parallel processing: 1. Input/Output (I/O) parallel processing. With I/O parallel processing, the database manager uses multiple tasks for each query to do the I/O processing. The central processor unit (CPU) processing is still done serially. 2. Symmetric Multiprocessing (SMP). SMP assigns both CPU and I/O processing to tasks that run the query in parallel. Actual CPU parallelism requires a system with multiple processors. SMP can only be used if the system feature, Db2 Symmetric Multiprocessing, is installed. Use of SMP parallelism can affect the order in which records are returned.	*DEFAULT	The default value is *SYSVAL.
	*SYSVAL	Set to the current system value QQRVDEGREE.
	*IO	Any number of tasks can be used. SMP parallel processing is not allowed. The SQE optimizer considers I/O parallelism with or without this setting.
	*OPTIMIZE	Any number of tasks for: <ul style="list-style-type: none"> I/O or SMP parallel processing of the query database file keyed access path build, rebuild, or maintenance. SMP parallel processing is used only if the system feature, Db2 Symmetric Multiprocessing for IBM i, is installed. Use of parallel processing and the number of tasks used is determined by: <ul style="list-style-type: none"> the number of processors available in the system the job share of the amount of active memory available in the pool in which the job is run 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 share of the memory in the pool.
	*OPTIMIZE nnn	Like *OPTIMIZE, with the value nnn indicating a percentage from 1 to 200, used to influence the number of tasks. If not specified, 100 is used. The query optimizer determines the parallel degree for the query using the same processing as is done for *OPTIMIZE. Once determined, the optimizer adjusts the actual parallel degree used for the query by the percentage given. Allows the user to override the parallel degree used without having to specify a particular parallel degree under *NUMBER_OF_TASKS.
nnn	The query optimizer chooses to use either I/O or SMP parallel processing to process the query. SMP parallel processing is used only if the system feature, Db2 Symmetric Multiprocessing for IBM i, is installed. nnn is a percentage from 1 to 200 and is used to influence the number of tasks. If not specified, 100 is used. The choices made by the query optimizer are like those choices made for parameter value *OPTIMIZE. The exception is the assumption that all pool active memory can be used for query processing, database file keyed access path build, rebuild, or maintenance.	
PARALLEL_DEGREE (continued)	*NONE	No parallel processing is allowed for database query processing or database table index build, rebuild, or maintenance.
	*NUMBER_OF_TASKS nnn	Indicates the maximum number of tasks that can be used for a single query. The number of tasks is limited to either this value or the number of disk arms associated with the table. Not recommended if running SQE. The SQE optimizer attempts to use this degree and override many of the normal costing mechanisms. For SQE, use *OPTIMIZE with a percentage.
	*MAX xxx	Like *MAX, with the value xxx indicating the ability to specify an integer percentage value 1 - 200. The query optimizer determines the parallel degree for the query using the same processing as is done for *MAX. Once determined, the optimizer adjusts the actual parallel degree used for the query by the percentage given. This option provides the user the ability to override the parallel degree used to some extent without having to specify a particular parallel degree under *NUMBER_OF_TASKS.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
PARAMETER_MARKER_CONVERSION Specifies whether to allow literals to be implemented as parameter markers in dynamic SQL queries.	*DEFAULT	The default value is set to *YES.
	*NO	Constants cannot be implemented as parameter markers.
	*YES	Constants can be implemented as parameter markers.
PREVENT_ADDITIONAL_CONFLICTING_LOCKS The following SQL DDL statements require an exclusive, no read lock on the target table. If the application activity cannot be quiesced, it can be hard to accomplish these operations. The PREVENT_ADDITIONAL_CONFLICTING_LOCKS QAQQINI option provides a control for customers to use to direct the operating system to favor a request for an exclusive, no read lock over new requests to lock the object for reading.	*DEFAULT	The default value is set to *NO
	*NO	When a job requests an exclusive lock on an object, do not prevent concurrent jobs from acquiring additional locks on the object.
	*YES	When *YES is chosen, any new requests for these lower-level read locks will be kept behind the exclusive lock request and could surface to applications as the table is unavailable for use for querying. <ul style="list-style-type: none"> • ALTER TABLE (Add, Alter or Drop Column) • CREATE TRIGGER • LOCK TABLE • RENAME TABLE
PSEUDO_OPEN_CHECK_HOST_VARS This option can be used to allow SQE to check the selectivity of the host variable values at pseudo open time. If the new set of host variable values require a different plan to perform well, SQE will re-optimize the query. This option is most appropriate when there is considerable variability in the selectivity of host variable in the queries predicates.	*DEFAULT	The default value is set to *NO
	*NO	The optimizer does not check host variables for selectivity changes once in pseudo-open.
	*OPTIMIZE	The optimizer will determine when a host variable selectivity should be checked. In general, the optimizer will monitor the query and if after a certain number of runs it determines that there is no advantage to checking host variable selectivity at pseudo open time, it will stop checking. Full opens do normal plan validation.
QUERY_TIME_LIMIT Specifies a time limit for database queries allowed to be started based on the estimated number of elapsed seconds that the query requires to process.	*DEFAULT	The default value is set to *SYSVAL.
	*SYSVAL	The query time limit for this job is obtained from the system value, QQRYTIMLMT.
	*NOMAX	There is no maximum number of estimated elapsed seconds.
REOPTIMIZE_ACCESS_PLAN Specifies whether the query optimizer reoptimizes a query with a saved access plan. Queries can have a saved access plan stored in the associated storage of an HLL program, or in the plan cache managed by the optimizer itself. Note: If you specify *NO the query could still be revalidated. Some of the reasons this option might be necessary are: <ul style="list-style-type: none"> • The queried file was deleted and recreated. • The query was restored to a different system than the one on which it was created. • An OVRDBF command was used. 	*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 is optimized.
	*YES	Force the existing query to be reoptimized.
	*FORCE	Force the existing query to be reoptimized.
SQE_NATIVE_ACCESS This option controls how native access will be implemented for an open or query. It does not affect a simple native open, such as an open done using the OPNDBF command, unless opening an SQL view, a partition table with a MBR(*ALL) override or a file dependent on row or column access control. It also does not affect the Query (QQQRY) API.	*DEFAULT	The default value could be either *YES or *NO as determined by the Query Optimizer.
	*NO	Attempt open using the Classic Query Engine (CQE). If CQE is not possible, attempt open using the SQL Query Engine (SQE).
	*YES	Attempt open using the SQL Query Engine (SQE). If SQE is not possible, attempt open using the Classic Query Engine (CQE).

Table 51. Query Options Specified on QAAQINI Command (continued)

Parameter	Value	Description
SQE_NATIVE_ACCESS_POSITION_BEHAVIOR This option controls the positioning behavior of native opens or queries implemented by SQE. By specifying an option other than *DEFAULT, performance benefits may be realized.	*DEFAULT	Normal positioning behavior is performed.
	*NO_ROLLBACK_HOLD	The current cursor position is unchanged by a rollback.
	*NO_KEY_FAILURE_HOLD	If an attempted key positioning operation fails, the cursor position prior to the attempted operation will not be restored. It is assumed that another absolute positioning operation, such as first, last, or key equal, will be attempted before any relative positioning operations, such as next or previous.
	*NO_HOLD	Behavior is the same as defined for *NO_ROLLBACK_HOLD and *NO_KEY_FAILURE_HOLD values.
SQLSTANDARDS_MIXED_CONSTANT Specifies whether to allow IGC constants to always be treated as IGC-OPEN in SQL queries. Note: When *NO is specified, Db2 for i is not compatible with the other Db2 platforms.	*DEFAULT	The default value is set to *YES.
	*YES	SQL IGC constants are treated as IGC-OPEN constants.
	*NO	If the data in the IGC constant only contains shift-out DBCS-data shift-in, then the constant are treated as IGC-ONLY, otherwise it is treated as IGC-OPEN.
SQL_CONCURRENT_ACCESS_RESOLUTION Specifies the concurrent access resolution to use for an SQL query.	*DEFAULT	The default value is set to *WAIT.
	*WAIT	The database manager must wait for the commit or rollback when encountering data in the process of being updated, deleted, or inserted. Rows encountered that are in the process of being inserted are not skipped. This option applies if possible when the isolation level in effect is Cursor Stability or Read Stability and is ignored otherwise.
	*CURCMT	The database manager can use the currently committed version of the data for read-only scans when it is in the process of being updated or deleted. Rows in the process of being inserted can be skipped. This option applies if possible when the isolation level in effect is Cursor Stability and is ignored otherwise.
SQL_DECFLOAT_WARNINGS Specifies the warnings returned for SQL DECFLOAT computations and conversions involving: <ul style="list-style-type: none"> • division by 0. • overflow. • underflow. • an invalid operand. • an inexact result. • 	*DEFAULT	The default value is set to *NO.
	*YES	A warning is returned to the caller for DECFLOAT computations and conversions involving division by 0, overflow, underflow, invalid operand, inexact result, or subnormal number.
	*NO	An error or a mapping error is returned to the caller for DECFLOAT computations and conversions involving division by 0, overflow, underflow, or an invalid operand. A warning or error is not returned for an inexact result or a subnormal number.
SQL_FAST_DELETE_ROW_COUNT Specifies how the delete is implemented by the database manager. This value is used when processing a DELETE FROM table-name SQL statement without a WHERE clause.	*DEFAULT	The default value is set to 0. 0 indicates that the database manager chooses how many rows to consider when determining whether fast delete could be used instead of traditional delete. When using the default value, the database manager will most likely use 1000 as a row count. This means that using the INI option with a value of 1000 results in no operational difference from using 0 for the option.
	*NONE	This value forces the database manager to never attempt to fast delete on the rows.
	*OPTIMIZE	This value is same as using *DEFAULT.
	Integer Value	Specifying a value for this option allows the user to tune the behavior of DELETE. The target table for the DELETE statement must match or exceed the number of rows specified on the option for fast delete to be attempted. A fast delete does not write individual rows into a journal. The valid values are 1 - 999,999,999,999,999. A value of -1 indicates that fast delete should not be used.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
SQL_GVAR_BUILD_RULE Determines whether global variables must exist or not when building SQL routines or executing SQL pre-compiles. This option has no affect on dynamic SQL statements.	*DEFAULT	The default value is set to *DEFER
	*DEFER	Global variables do not need to exist when an SQL routine is created or the SQL pre-compiler is run. Since global variables are not required to exist, the create will not fail when an incorrect column name or routine variable is encountered. Incorrect name usage will result in SQL0206 - "Column or global variable &1 not found." failures when the statement is executed.
	*EXIST	Global variables referenced by SQL must exist when the SQL routine is created or the SQL pre-compiler is run. Using this option, an SQL0206 will be issued at create time.
SQL_MODIFIES_SQL_DATA From the SQL Standard, no MODIFIES SQL DATA operations are allowed in an SQL BEFORE trigger. The Informix® database allows MODIFIES SQL DATA operations in SQL BEFORE triggers. Setting the option to *YES allows SQL BEFORE triggers to perform the SQL MODIFIES SQL DATA operations.	*DEFAULT	The default value is set to *NO.
	*NO	No MODIFIES SQL DATA operations are allowed in an SQL BEFORE trigger.
	*YES	MODIFIES SQL DATA operations are allowed in an SQL BEFORE trigger.
SQL_PSEUDO_CLOSE Before V6R1: SQL cursor open processing checks for the presence of a data area named QSQPSCLS1 in the library list of the job. If the data area is found, all reusable cursors are marked as candidates for reuse. They are pseudo-closed the first time rather than the second time the application closes the cursor. Without this data area, a cursor does not become reusable until the second close. Pseudo-closing the first time results in leaving some cursors open that might not be reused. These open cursors can increase the amount of auxiliary and main storage required for the application. The storage can be monitored using the WRKSYSSTS command. For the amount of auxiliary storage used, look at the "% system ASP used." For the amount of main storage, examine the faulting rates on the WRKSYSSTS display. The format and the contents of the data area are not important. The data area can be deleted using the following command: DLTDTAARA DTAARA(QGPL/QSQPSCLS1). The existence of the data area is checked during the first SQL open operation for each job. It is checked only once and the processing mode remains the same for the life of the job. Because the library list is used for the search, the change can be isolated to specific jobs. Create the data area in a library that is included only in the library lists for those jobs.	*DEFAULT	The default behavior depends upon whether the QSQPSCLS1 *DTAARA exists. If the QSQPSCLS1 *DTAARA was found on the first OPEN within the job, then SQL cursors are marked as candidates for reuse. The cursors are pseudo-closed on the first close. If the QSQPSCLS1 *DTAARA was not found on the first OPEN within the job, then SQL cursors are marked as candidates for reuse. The cursors are pseudo-closed on the second close.
	Integer Value	Specifies a value greater than zero that indicates when a cursor is pseudo-closed. The value of this option minus 1 indicates how many times the cursor is hard closed before being marked as candidate for pseudo-close. Valid values are 1 - 65535.
SQL_STMT_COMPRESS_MAX Specifies the compression maximum setting, which is used when statements are prepared into a package.	*DEFAULT	The default value is set to 2. The default indicates that the access plan associated with any statement will be removed after a statement has been compressed twice without being executed.
	Integer Value	The integer value represents the number of times that a statement is compressed before the access plan is removed to create more space in the package. Executing the SQL statement resets the count for that statement to 0. The valid Integer values are 1 - 255.
SQL_STMT_REUSE Specifies the number of times the statement must be prepared in the same connection before the statement is stored in the SQL extended dynamic package. If the number of times the statement has been prepared in the same connection is less than the specified INI option, a temporary copy of the statement is used. Any other job preparing the statement does a complete prepare.	*DEFAULT	The default value is 3. The statement is stored on the third prepare of the statement.
	1::255	The number of times the statement must be prepared in the same connection before the statement is stored in the SQL package.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
SQL_SUPPRESS_MASKED_DATA_DETECTION	*DEFAULT	The default value is set to *NO.
	*YES	If masked data is being used to insert into or update a table, detection of this masked data will not be done and a SQ20478 with reason code 30 will not be sent.
	*NO	If masked data is being used to insert into or update a table with activated column access control directly from an expression involving a column with an active column mask, detection of this masked data will be done and a SQ20478 with reason code 30 will be sent.
SQL_SUPPRESS_WARNINGS For SQL statements, this parameter provides the ability to suppress SQL warnings.	*DEFAULT	The default value is set to *NO.
	*YES	Examine the SQLCODE in the SQLCA after execution of a statement. If the SQLCODE is one of the listed warnings, then alter the SQLCA so that no warning is returned to the caller. Set the SQLCODE to 0, the SQLSTATE to '00000' and SQLWARN to ' '. Warnings: <ul style="list-style-type: none"> • SQL0030 • SQL0335 • SQL0387 • SQL7909 (on a DROP PROCEDURE/ROUTINE/FUNCTION)
	*NO	Specifies that SQL warnings are returned to the caller.
SQL_TRANSLATE_ASCII_TO_JOB Specifies whether to translate SQL statement text on the application server (AS) according to the CCSID of the job. This option applies when using DRDA to connect to an IBM i as the AS where the application requestor (AR) machine is an ASCII-based platform.	*DEFAULT	The default value is set to *NO.
	*YES	Translate ASCII SQL statement text to the CCSID of the IBM i job.
	*NO	Translate ASCII SQL statement text to the EBCDIC CCSID associated with the ASCII CCSID.
SQL_XML_DATA_CCSID Specifies the CCSID to be used for XML columns, host variables, parameter markers, and expressions, if not explicitly specified. See "SQL_XML_DATA_CCSID QAQQINI option" on page 204	*DEFAULT	The default value is set to 1208.
	*JOB	The job CCSID is used for XML columns, host variables, parameter markers, and expressions, if not explicitly specified. If the job CCSID is 65535, the default CCSID of 1208 is used.
	Integer Value	The CCSID used for XML columns, host variables, parameter markers, and expressions, if not explicitly specified. This value must be a valid single-byte or mixed EBCDIC CCSID or Unicode CCSID. The value cannot be 65535.
STAR_JOIN Note: Only modifies the environment for the Classic Query Engine. Specifies enhanced optimization for hash queries where both a hash join table and a Distinct List of values is constructed from the data. This Distinct List of values is appended to the selection against the primary table of the hash join Any EVI indexes built over these foreign key columns can be used to perform bitmap selection against the table before matching the join values. The use of this option does not guarantee that star join is chosen by the optimizer. It only allows the use of this technique if the optimizer has decided to implement the query by using a hash join.	*DEFAULT	The default value is set to *NO
	*NO	The EVI Star Join optimization support is not enabled.
	*COST	Allow query optimization to cost the usage of EVI Star Join support. The optimizer determines whether the Distinct List selection is used based on how much benefit can be derived from using that selection.
STORAGE_LIMIT Specifies a temporary storage limit for database queries. If a query is expected to use more than the specified amount of storage, the query is not allowed to run. The value specified is in megabytes.	*DEFAULT	The default value is set to *NOMAX.
	*NOMAX	Never stop a query from running because of storage concerns.
	Integer Value	The maximum amount of temporary storage in megabytes that can be used by a query. This value is checked against the estimated amount of temporary storage required to run the query as calculated by the query optimizer. If the estimated amount of temporary storage is greater than this value, the query is not started. Valid values range from 0 through 2147352578.

Table 51. Query Options Specified on QAQQINI Command (continued)

Parameter	Value	Description
SUPPRESS_INQUIRY_MESSAGES Specifies that certain inquiry messages will not be issued.	*DEFAULT	The default value is set to *NO.
	*YES	System inquiry message CPA32B2 will not be issued when altering a table.
	*NO	No system inquiry messages will be suppressed.
SYSTEM_SQL_STATEMENT_CACHE Specifies whether to disable the system-wide SQL Statement Cache for SQL queries.	*DEFAULT	The default value is set to *YES.
	*YES	Examine the system-wide SQL Statement Cache when an SQL prepare request is processed. If a matching statement exists in the cache, use the results of that prepare. This option allows the application to potentially have better performing prepares.
	*NO	Specifies that the system-wide SQL Statement Cache is not examined when processing an SQL prepare request.
SYSTIME_PERIOD_ADJ For update and delete operations involving system-period temporal tables, this option resolves conflict for a historical row when its row end value could be less than its row begin value.	*DEFAULT	The default value is set to *ERROR.
	*ERROR	An SQ20528 error will be signaled.
	*ADJUST	These three things will happen: 1. The row end value of the history row will be set to the row begin value plus a small delta of one microsecond. 2. The row begin value of the temporal table current row will also be set to the adjusted value. 3. An SQL warning SQ20528 will be signaled.
TEXT_SEARCH_DEFAULT_TIMEZONE Specifies the time zone to apply to any date or dateTime value specified in an XML text search using the CONTAINS or SCORE function. The time zone is the offset from UTC (Greenwich mean time). It is only applicable when a specific time zone is not given for the value.	*DEFAULT	Use the default as defined by database. This option is equivalent to UTC.
	sHH:MM	A time zone formatted value where • s is the sign, + or – • HH is the hour • MM is the minute The valid range for HH is 00 - 23. The valid range for MM is 00 - 59. The format is specific. All values are required, including sign. If HH or MM is less than 10, it must have a leading zero specified.
UDF_TIME_OUT Note: Only modifies the environment for the Classic Query Engine. Specifies the amount of time, in seconds, that the database waits for a User Defined Function (UDF) to finish processing.	*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 waits for the UDF to finish.
	integer value	Specify the number of seconds that the database waits for a UDF to finish. If the value given exceeds the database maximum wait time, the maximum wait time is used by the database. Minimum value is 1 and maximum value is system defined.
VARIABLE_LENGTH_OPTIMIZATION Specifies whether aggressive optimization techniques are used on variable length columns.	*DEFAULT	The default value is set to *YES.
	*YES	Enables aggressive optimization of variable-length columns, including index-only access. It also allows constant value substitution when an equal predicate is present against the columns. As a consequence, the length of the data returned for the variable-length column might not include any trailing blanks that existed in the original data. As a result, the application can receive the substituted value back instead of the original data. Function calls could operate on the substituted value instead of the original string value.
	*NO	Do not allow aggressive optimization of variable length columns.

Note: The following QAQQINI options will be ignored for SQE native query access. These options were previously honored for CQE native query access.

- LIMIT_PREDICATE_OPTIMIZATION
- STAR_JOIN
- UDF_TIME_OUT

Note: The following QAQQINI options will be honored for SQE native query access. These options were previously ignored for CQE native query access.

- DETERMINISTIC_UDF_SCOPE
- FIELDPROC_ENCODED_COMPARISON
- MATERIALIZED_QUERY_TABLE_USAGE
- MATERIALIZED_QUERY_TABLE_REFRESH_AGE
- MEMORY_POOL_PREFERENCE
- VARIABLE_LENGTH_OPTIMIZATION

SQL_XML_DATA_CCSID QAQQINI option

The SQL_XML_DATA_CCSID QAQQINI option has several settings that affect SQL processing.

The SQL_XML_DATA_CCSID QAQQINI setting is applied within SQL in the following SQL processing:

SQL Processing item	Description
Valid values for the QAQQINI option are CCSIDs allowed on an XML column.	Valid values are all EBCDIC SBCS and mixed CCSIDs, and Unicode 1208, 1200, and 13488 CCSIDs.
Does not affect the promotion of SQL data types.	Other SQL data types cannot be directly promoted to the SQL XML data type.
XMLPARSE untyped parameter markers.	The QAQQINI setting applies to untyped parameter markers passed as string-expression. The type is CLOB(2G) for SBCS, mixed, and UTF-8 values. The type is DBLOB(1G) for Unicode 1200 and 13488.
XMLCOMMENT, XMLTEXT, XMLPI untyped parameter markers.	The QAQQINI setting applies to untyped parameter markers passed as string-expression. The type is VARCHAR(32740) for SBCS, mixed, and UTF-8 values. The type is VARGRAPHIC(16370) for Unicode 1200 and 13488.
Applies to parameter marker casts to the XML type for XMLCONCAT, and XMLDOCUMENT.	Applies to an untyped parameter marker passed as an XML-expression. Unless an explicit CCSID clause is specified, the CCSID of the parameter marker is obtained from the QAQQINI setting.
The QAQQINI setting does not affect storage and retrieval assignment rules.	The CCSID of the host variables and table columns apply.
String to column assignment on SQL INSERT and UPDATE.	An implicit or explicit XMLPARSE is required on the column assignment.
String to host variable assignment.	An implicit or explicit XMLSERIALIZE is required on the host variable assignment.
Column to column assignment.	When the target column is XML, an implicit XMLPARSE is applied if the source column is not XML. The target XML column has a defined XML CCSID. When the source column is XML, an explicit XMLSERIALIZE is required if the target column is not XML.
Host variable to column assignment.	The target column has a defined CCSID.
UNION ALL (if XML publishing functions in query).	The XML result CCSID is obtained from the QAQQINI setting.
Does not apply to SQL constants.	UX constants are defined as UTF-16. FX constants are defined as UTF-8.

Table 52. SQL_XML_DATA_CCSID setting application within SQL (continued)

SQL Processing item	Description
Result type of XML data built-in functions.	If the first operand of XMLPARSE and XMLVALIDATE is an untyped parameter marker, the CCSID is set from the QAQQINI setting, which then affects the XML result CCSID. The QAQQINI setting is used for XMLSERIALIZE for CHAR, VARCHAR, and LOB AS data-type. UTF-16 is used for GRAPHIC, DBCLOB, and NCHAR.
Result type of XML publishing functions - XMLAGG, XMLGROUP, XMLATTRIBUTES, XMLCOMMENT, XMLCONCAT, XMLDOCUMENT, XMELEMENT, XMLFOREST, XMLNAMESPACES, XMLPI, XMLROW, and XMLTEXT.	The XML result CCSID for XML publishing functions is obtained from the QAQQINI setting.
Result type of XML publishing functions in a view.	The XML result CCSID is set when the view is created.
XML data type on external procedure XML AS parameters.	The XML parameter CCSID is set when the procedure is created.
XML data type on external user-defined functions.	The XML parameter and result CCSID are set when the function is created.
CREATE TABLE XML column.	The QAQQINI setting is used for dynamic SQL. The QAQQINI setting is set in *PGM, *SRVPGM, and *SQLPKG objects when created.
MQTs containing select-statement with XML publishing functions.	The CCSID is set when the MQT is created. The CCSID is maintained for an ALTER TABLE.
ALTER TABLE ADD MATERIALIZED QUERY definition.	The QAQQINI setting is used if the select-statement contains XML publishing functions.
XML AS CLOB CCSID	The QAQQINI setting is built into *PGM and *SRVPGM objects when the program is created. The CCSID defaults to UTF-8 for CLOB when QAQQINI setting is UTF-16 or UCS2.
XML AS DBCLOB CCSID	The default for DBCLOB is always UTF-16 for XML.
SQL GET and SET DESCRIPTOR XML data type.	QAQQINI setting applied to XML data type.
SQL Global variables.	QAQQINI setting applied to global variables with the XML data type.

Related information

[XML values](#)

[SQL statements and SQL/XML functions](#)

Query Supervisor

The Db2 for i SQL Query Engine (SQE) provides a Query Supervisor which enables real-time monitoring of resource consumption by SQL and native queries, for example the Open Query File (OPNQRYF) command. Query Supervisor threshold levels can be established for general or specific scenarios. When a threshold is met or exceeded, the query execution is interrupted, and user-supplied exit program(s) are called.

Using Query Supervisor to monitor query resource usage

The Query Supervisor enables monitoring and management of queries that reach pre-determined thresholds of resource consumption. Unlike the Predictive Query Governor, which intervenes on the basis of *estimated* resource usage *before* a query runs, the Query Supervisor reacts to *actual* resource usage *while* the query runs.

The Query Supervisor monitors any SQL queries that are run to implement user SQL statements, including native database queries that are processed by SQE. SQL statements that do not require query processing such as a simple INSERT (for example, INSERT INTO <table> VALUES(123)) or COMMIT are not monitored by the Query Supervisor. Query supervision does not occur for queries initiated by the IBM i operating system or for SQL that has been classified as specific to operating system processing.

The Query Supervisor allows a Database Engineer (DBE) to:

- Deploy real-time notification solutions when a query exceeds a specific resource threshold
- Take actions to terminate queries based on real-time query consumption of system resources
- Capture and log performance details for long-running queries

By configuring the Query Supervisor, a DBE can proactively manage excessive resource usage, monitor for unexpected workload variation, and automatically end run-away queries. The Query Supervisor provides the infrastructure for establishing and detecting thresholds. The specific action(s) taken when a threshold is reached is determined by one or more exit programs that the DBE provides.

Query Supervisor configuration and operation

Multiple thresholds can be defined across the system for various resource types and can apply to all queries or be restricted to specific users, jobs, and subsystems.

Query Supervisor thresholds are configured using the add and remove procedures: “[ADD_QUERY_THRESHOLD procedure](#)” on page 324 and “[REMOVE_QUERY_THRESHOLD procedure](#)” on page 336. The “[QUERY_SUPERVISOR view](#)” on page 334 shows the defined thresholds.

The four threshold types which can be monitored by the Query Supervisor are:

1. **CPU Time** – The total processing unit time used by the query, in seconds
2. **Elapsed Time** – The total clock time, in seconds
3. **Temporary storage** – The amount of storage, in megabytes (MB), that the query uses
4. **Total I/O count** – The total number of I/O operations

Before a query runs, a list of applicable thresholds is passed to SQE. As the query executes, SQE monitors the resource usage and determines if a threshold has been reached. When a threshold is met or exceeded, the query execution is interrupted and any exit programs registered with the [Query Supervisor Exit Point \(QIBM_QQQ_QRY_SUPER\)](#) are called in sequence of their exit program registration number, from lowest to highest. The exit program is called inline within the same thread of the query that reached the threshold. As a result, the query is interrupted and paused until the exit program processing has completed.

Information passed to the exit programs includes details such as the threshold reached, the job name, the SQL statement or native query request, the client special register values, and host variable or parameter marker values. The structure passed to the exit program includes enough detail to allow the DBE to take a wide variety of actions, such as logging the information, sending a message, or retrieving detailed information about the query from the plan cache.

The exit program also has the option to indicate that the query should be terminated. If query termination is requested, no further exit programs are called, and the Query Supervisor issues a CPF5209 escape message with reason code 3. SQL statements will fail with SQLCODE -666 and SQLSTATE 57005.

If query termination is not requested, execution of the query resumes after all the registered exit programs have completed. It is possible for multiple, distinct thresholds to be reached simultaneously, in which case multiple calls to the exit programs will be made before the query resumes.

Resource usage is measured only for processing that the query engine performs while producing the query result set. It does not apply to other database or operating system processing that happens in conjunction with the operation. Consider the following SQL DELETE statement: DELETE FROM <table> WHERE <column> <= 100. Query supervision only applies to the system resources used to determine the rows selected by the WHERE predicate. The resources consumed by the deletion of rows, journal processing, and any associated index maintenance for the DELETE are not supervised.

The Query Supervisor will also supervise queries executed as part of a user-defined function (UDF) or user-defined table function (UDTF) invocation when the reference to that function is not inlined. When a query uses a function, the Query Supervisor independently supervises the resources consumed by any queries run during execution of the function. Resources consumed by the invoking query prior to the function invocation are not counted as part of the resources used by the function's queries; the resource usage counters for each query within the function start at 0. However, resources consumed by the queries within the function are added to the resource usage totals for the invoking query.

Any threshold that is applied to the function's queries and that is met while the function is executing will cause the exit program(s) to be called at that point in time. However, when a threshold that is applied to the invoking query is met while a function is running, the exit program call will be delayed until the function completes and the invoking query continues executing.

Consider a TOTAL IO COUNT threshold of 500 I/Os and a query that invokes a function that executes another query. Assume the invoking query consumes 400 I/Os and then invokes the function, whose query consumes 200 I/Os. The exit program will be called for the invoking query but only after the function completes. However, if the function's query consumes 500 I/Os, it is possible (subject to the threshold's DETECTION_FREQUENCY) that the exit program could be called both within the function execution (with 500 I/Os) and for the invoking query (with 900 I/Os) once the function completes.

In general, resource usage is checked on sub-second intervals while SQE is processing the query. As a result, there may be a small delay between when the query reaches a threshold and when the query supervisor detects this and calls the exit program. The delay is most likely to affect TOTAL IO COUNT and TEMPORARY STORAGE thresholds and may result in threshold consumption values that exceed the defined threshold value. Certain operations within the query engine, such as queries containing UDF calls as described above, can produce a more measurable delay that affects all threshold types. In all cases, however, the exit program will eventually be called and will be provided with the accurate and current resource consumption detail.

A specific Query Supervisor threshold will apply at most one time for each query execution within a thread. For example, a threshold defined for ELAPSED TIME of 5 seconds is reached only at the first 5 seconds of a query's execution and not repeatedly at further intervals of 5 seconds while the query runs.

Managing Query Supervisor activity

To make it easier to identify any Query Supervisor exit program problems, the first occurrence of an exit program failure each day will be noted with a CPD43B0 message sent to QSYSOPR.

Any misconfiguration or exit program failures are tolerated and will not affect the supervised query. The CPD43B0 message allows the DBE to easily recognize that the exit program processing is failing.

In the following figure, the exit program was registered, but the *PGM did not exist with the specified library and program name. As shown below, the CPD43B0 second level text includes the threshold name, the exit program library and name, the job that encountered the exit program problem, and the message identifier related to the exit program failure.

```

Message ID . . . . . : CPD43B0          Severity . . . . . : 80
Message type . . . . . : Diagnostic
Date sent . . . . . : 03/22/21        Time sent . . . . . : 13:04:46

Message . . . . . : An exit point program registered to exit point
QIBM_QQQ_QRY_SUPER encountered an error.
Cause . . . . . : An error was encountered while resolving to or running one
or more exit programs registered to exit point QIBM_QQQ_QRY_SUPER. The first
exit program was SUPERTERM in library SCOTTFF. The current user was TIMMR.
Refer to message ID MCH3401 in the job log for job 212791/TIMMR/QDFTJOB0 for
more information about the error.

```

Figure 1. CPD43B0 message in QSYSOPR highlights an exit program setup problem

Summary information about the Query Supervisor’s behavior is available in the SQE Plan Cache Properties which are accessed with the IBM i Access Client Solutions (ACS) SQL Performance Center. The properties can also be retrieved using “[DUMP_PLAN_CACHE_PROPERTIES procedure](#)” on page 341. The properties show the number of times since the last IPL that a Query Supervisor threshold has been reached along with the most recent job and time that the threshold was reached. Similar information is provided for the most recent instance of a Query Supervisor exit program that directed SQE to terminate a query. Finally, if there is a problem with the configuration or execution of an exit program, the most recent job and time of the failure is noted.

These Query Supervisor metrics are provided to give the database engineer (DBE) a basic understanding of whether thresholds are being encountered, whether queries are being terminated, and whether any exit program failures exist. The following figure illustrates the query supervision summary metrics.

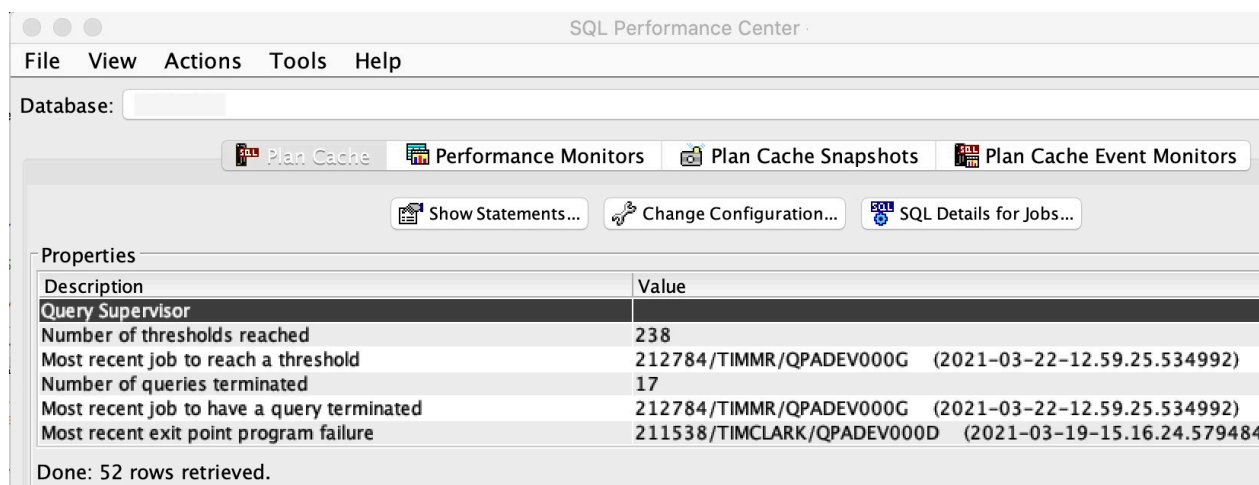


Figure 2. SQE Plan Cache properties include Query Supervisor detail

Using DETECTION_FREQUENCY to protect system resources

Because the Query Supervisor may be the first line of detection and defense against a critical performance problem, it is important to make the operation of the Query Supervisor as light-weight as possible. For this reason, each threshold has a DETECTION_FREQUENCY associated with it. This value, which defaults to 600 seconds (10 minutes), indicates how much time a given thread must wait after detecting the threshold before it may detect and signal the same threshold again. If the threshold is exceeded again in the thread before the defined interval is completed, the threshold will be ignored, and the exit program(s) will not be called.

Note that the DETECTION_FREQUENCY applies to multiple occurrences of a specific threshold within the same thread.

Consider a single threaded job, JOB1, which is executing queries. The following thresholds are defined:

- THRESHOLD_NAME => 'WARNING'
THRESHOLD_TYPE => 'ELAPSED TIME'

THRESHOLD_VALUE => 10
DETECTION_FREQUENCY => 600

- THRESHOLD_NAME => 'RUNAWAY'
THRESHOLD_TYPE => 'ELAPSED TIME'
THRESHOLD_VALUE => 120
DETECTION_FREQUENCY => 600

The following example timeline shows the impact of DETECTION_FREQUENCY:

00:00:00 JOB1 begins a query that will take 20 seconds to complete.

00:00:10 Query Supervisor calls the exit program in JOB1. The exit program input structure includes the name of the threshold: WARNING.

The WARNING threshold will be silenced for any future queries in JOB1 for 600 seconds, until 00:10:10.

00:00:30 JOB1 begins another query that runs for 25 seconds.

No Query Supervisor action is taken after 10 seconds of execution since the WARNING threshold's DETECTION_FREQUENCY time interval has not been met.

00:01:05 JOB1 begins a query that runs for 215 seconds.

00:03:05 Query Supervisor calls the exit program in JOB1. The exit program input structure includes the name of the threshold: RUNAWAY.

The RUNAWAY threshold will be silenced for any future queries in JOB1 for 600 seconds, until 00:13:05.

00:09:59 JOB1 begins a query that will run for 50 seconds.

No Query Supervisor action is taken after 10 seconds of execution since the WARNING threshold's DETECTION_FREQUENCY time interval has not been met. Since the threshold value is detected while reporting is silenced, this query will not call the exit program for the WARNING threshold.

A DETECTION_FREQUENCY of 0 is permitted, in which case every instance of the threshold is processed by a call to the exit program(s).

Writing a Query Supervisor exit program

A well-written exit program is the key to making the Query Supervisor effective. Each time a threshold is reached, the exit program receives information about the threshold, the job, and the query. The exit program can take action based on this information, including the option to terminate the query.

Because exit programs run inline and interrupt the query execution, it is recommended that the exit program should be as simple as possible. When more complex operations are required of the exit program, it is suggested that they be done asynchronously by passing information to another job for processing. By doing this, the exit program can complete faster, allowing the query to resume execution.

Some possible operations an exit program can perform are:

- Send an inquiry message, suspending the thread (and query) until a reply is received
- Send the query text, host variables, and parameter markers to another job to be logged to a database table for future review and analysis
- Terminate the query

SQL and native database I/O must not be run within an exit program. However, it is possible to run SQL native database I/O in a separate job. Some useful SQL operations include:

- Send a message to QSYSOPR using the QSYS2.SEND_MESSAGE procedure. See [“SEND_MESSAGE procedure”](#) on page 665 for details.

1. The exit program can send the message using the SQL7064 message identifier which has been defined for Query Supervisor notifications. The message is in the QSYS/QSQLMSG message file and has a severity level of 80.
 2. Alternatively, the exit program can send the message with a user created message identifier.
- Dump information about the plan for the query that was running when the exit program was called. The plan identifier that is part of the exit program data can be passed on the call to this procedure. See “[DUMP_PLAN_CACHE procedure](#)” on page 340 for details.
 - Use INSERT or MERGE statements to maintain a log of expensive queries.
 - Push details regarding the query supervision event to external systems management solutions.

Any messages generated by a failure to call the exit program or by an unhandled exception within the exit program will appear in the job log of the job that caused the threshold criteria to be met.

Query Supervisor example exit programs

The exit programs in this section demonstrate some of the ways an exit program can be used when a Query Supervisor threshold is reached. They are provided to enable quick and easy adoption of Query Supervisor.

The details of the exit program interface can be found here: [Query Supervisor Exit Program](#)

Exit program to send message to QSYSOPR

This Query Supervisor exit program sends an informational message to the QSYSOPR message queue.

This program sends an SQL7064 message to the QSYSOPR message queue when it is called for any threshold. Because SQL cannot be used directly in the exit program, the Run SQL (RUNSQL) CL command is submitted as a batch job.

The SQL7064 message ID is provided to be used with Query Supervisor. The caller provides a string containing the text for the message. In this example, the threshold name and the job name that reached the threshold are included in the message.

ILE RPG exit program to send message to QSYSOPR

This Query Supervisor exit program sends an informational message to the QSYSOPR message queue.

To compile this ILE RPG program, use the following command.

```
CRTBNDRPG PGM(SUPERVISOR/QSENDMSG) SRCFILE(SUPERVISOR/QRPGLESRC)
```

If CRTRPGMOD and CRTPGM are used to compile and build the program, be sure to include USRPRF(*OWNER) and ACTGRP(*CALLER) on the CRTPGM command.

The program can be registered as an exit program using this command:

```
ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QSENDMSG) THDSAFE(*YES) TEXT('Query Supervisor send message to QSYSOPR')
```

By including the THREAD, USRPRF, DFACTGRP, and ACTGRP options in the source, they can be omitted on the compiler command. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program uses THREAD(*CONCURRENT) to ensure thread safety. The program will run in the activation group of the caller.

```
**free
ctl-opt main(QS_sendmsg);
ctl-opt ccsid(*char : *jobrun);
ctl-opt thread(*concurrent);
/if defined(*crtbndrpg)
  ctl-opt usrprf(*owner);
  ctl-opt dfactgrp(*no);
  ctl-opt actgrp(*caller);
/endif
```

```

dcl-ds QQQ_QRYSV_QRYS0100_t qualified template;
Header_Size uns(10);
Format_Name char(8); // QRYS0100 - Query Supervisor
Job_Name char(10); // Name of the job running the query
Job_User char(10); // User profile used to initiate job
Job_Number char(6); // Number of the job running the query
Subsystem char(10); // Subsystem in which the job is running
User_Name char(10); // The effective user name of the
// thread that reached the threshold
Query_Identifier char(8); // Uniquely identifies the query text
// and its environment (QR0 hash)
Plan_Identifier uns(20); // Uniquely identifies the access plan
// implementing the query.
Threshold_Name ucs2(30) ccsid(1200); // User-defined name for threshold that
// was reached (in CCSID 1200)
Threshold_Type char(30); // Type of threshold reached
Threshold_Consumption_Value int(20); // Actual value reached by the
// query. Units vary with
// Threshold_Type.
Operation_Type int(5); // Type of query operation
Offset_to_SQL_Statement int(10); // 0, if not available
Length_of_SQL_Statement int(10); // 0, if not available
Offset_to_Host_Variable_List int(10); // 0, if no variables
Length_of_Host_Variable_List int(10); // 0, if no variables
Offset_to_CLIENT_ACCTNG int(10); // 0, if not defined
Length_of_CLIENT_ACCTNG int(10); // 0, if not defined
Offset_to_CLIENT_APPLNAME int(10); // 0, if not defined
Length_of_CLIENT_APPLNAME int(10); // 0, if not defined
Offset_to_CLIENT_PROGRAMID int(10); // 0, if not defined
Length_of_CLIENT_PROGRAMID int(10); // 0, if not defined
Offset_to_CLIENT_USERID int(10); // 0, if not defined
Length_of_CLIENT_USERID int(10); // 0, if not defined
Offset_to_CLIENT_WRKSTNNAME int(10); // 0, if not defined
Length_of_CLIENT_WRKSTNNAME int(10); // 0, if not defined
// char SQL_Statement_Text; /* Varying length, CCSID 1200 */
// char Host_Variable_List; /* Varying length, CCSID 1200 */
// char CLIENT_ACCTNG; /* Varying length */
// char CLIENT_APPLNAME; /* Varying length */
// char CLIENT_PROGRAMID; /* Varying length */
// char CLIENT_USERID; /* Varying length */
// char CLIENT_WRKSTNNAME; /* Varying length */
end-ds;

dcl-proc QS_sendmsg; // Main procedure
dcl-pi *n;
input likeds(QQQ_QRYSV_QRYS0100_t);
rc int(10);
end-pi;

dcl-pr qcmdexc extpgm;
cmd char(1000) const;
cmdlen packed(15:5) const;
end-pr;

dcl-s cmdstr char(1000);
dcl-s msg varchar(500);

//*****
// Init the return code to continue running the query
//*****
rc = 0;

msg = 'QUERY SUPERVISOR THRESHOLD ' + %trimr(input.Threshold_Name) +
      ' REACHED IN JOB ' + input.Job_Number + '/' +
      %trimr(input.Job_User) + '/' + input.Job_Name ;

cmdstr =
  'SBMJOB CMD(RUNSQL SQL(''call qsys2.send_message(''SQL7064'', '
    + %char(%len(msg)) + ' ,'''+ msg + '''))');

qcmdexc (cmdstr : %size(cmdstr));

return;
end-proc;

```

C exit program to send message to QSYSOPR

This Query Supervisor exit program sends an informational message to the QSYSOPR message queue.

To compile this C program, use the following commands. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program.

```
CRTCMOD MODULE(SUPERVISOR/SND_QS_MSG) SRCFILE(SUPERVISOR/QCSRC)
CRTPGM  PGM(SUPERVISOR/SND_QS_MSG) USRPRF(*OWNER) ACTGRP(*CALLER)
```

The program can be registered as an exit program using this command:

```
ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/SND_QS_MSG)
THDSAFE(*YES) TEXT('Query Supervisor send QSYSOPR message')
```

```
#include <stddef.h>
#include <string.h>
#include <stdlib.h>
#include <iconv.h>
#include <stdio.h>
#include <eqqqrsv.h>

static void convertThresholdNameToJobCCSID(const char* input, char* output)
{
    iconv_t converter;
    char from_code[32], to_code[32];
    size_t input_bytes, output_bytes;
    int iconv_rc;

    memset(from_code, 0, sizeof(from_code));
    memset(to_code, 0, sizeof(to_code));
    memcpy(from_code, "IBMCCSID012000000000", 20);
    memcpy(to_code, "IBMCCSID00000", 13);
    converter = iconv_open(to_code, from_code);

    if (converter.return_value == 0) {
        input_bytes = 60;
        output_bytes = 30;
        iconv_rc = iconv(converter,
                        &input, &input_bytes,
                        &output, &output_bytes);
        iconv_close(converter);

        if (iconv_rc >= 0)
            return; /* Conversion was successful. */
    }

    sprintf(output, "iconv_open() failed with: %d", converter.return_value);
}

int trimmed_length(const char* str, int len)
{
    const char* first_blank = strchr(str, ' ', len);
    if (first_blank)
        return first_blank - str;

    return len;
}

/*****
int main(int argc, char* argv[])
{
    char length_string[10];
    char cmd[600];
    char thresholdNameInJobCCSID[30];
    char msg[512];

    const QQQ_QRYSV_QRYS0100_t* input = (QQQ_QRYSV_QRYS0100_t*)argv[1];
    int* rc = (int*)argv[2];

    *rc = 0; /* don't terminate the query */

    memset(thresholdNameInJobCCSID, 0, sizeof(thresholdNameInJobCCSID));
    convertThresholdNameToJobCCSID(input->Threshold_Name,
```

```

        thresholdNameInJobCCSID);

memset(msg, 0, sizeof(msg));
strcat(msg, "QUERY SUPERVISOR THRESHOLD ");
strncat(msg, thresholdNameInJobCCSID, 30);
strcat(msg, " REACHED IN JOB ");
strncat(msg, input->Job_Number, trimmed_length(input->Job_Number,6));
strcat(msg, "/");
strncat(msg, input->Job_User, trimmed_length(input->Job_User,10));
strcat(msg, "/");
strncat(msg, input->Job_Name, trimmed_length(input->Job_Name,10));

memset(length_string, 0, sizeof(length_string));
sprintf(length_string, "%d", strlen(msg));

/*****
/* Use the SQL service send_message to send the message to QSYSOPR.      */
/* SQL cannot be run within the exit program, so submit to batch.      */
*****/
memset(cmd, 0, sizeof(cmd));
strcat(cmd, "SBMJOB CMD(RUNSQL SQL('call qsys2.send_message('SQL7064',"));
strcat(cmd, length_string);
strcat(cmd, ", '");
strcat(cmd, msg);
strcat(cmd, "'')'");
system(cmd);
}

```

CL exit program to send message to QSYSOPR

This Query Supervisor exit program sends an informational message to the QSYSOPR message queue.

To compile this ILE CL program, use the following command.

```

CRTBDCL PGM(SUPERVISOR/QS_SENDMSG) SRCFILE(SUPERVISOR/QCLSRC) DFTACTGRP(*NO) ACTGRP(*CALLER)
USRPRF(*OWNER)

```

The program can be registered as an exit program using this command:

```

ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QS_SENDMSG)
THDSAFE(*YES) TEXT('Query Supervisor send message to QSYSOPR')

```

By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program will run in the activation group of the caller.

```

        PGM          PARM(&QRYS0100 &RC)
DCL          VAR(&QRYS0100) TYPE(*CHAR) LEN(8192)
DCL          VAR(&RC) TYPE(*INT) LEN(4)
DCL          VAR(&MSGLEN) TYPE(*INT) LEN(2)
DCL          VAR(&MSGTXT) TYPE(*CHAR) LEN(1022)
DCL          VAR(&MSGDTA) TYPE(*CHAR) LEN(1024)
DCL          VAR(&BINLOW) TYPE(*UINT) LEN(4)
DCL          VAR(&BINHIGH) TYPE(*INT) LEN(4)
DCL          VAR(&MAXINT) TYPE(*DEC) LEN(15 0) +
            VALUE(4294967296)
DCL          VAR(&JOBNAME) TYPE(*CHAR) LEN(10)
DCL          VAR(&JOBUSER) TYPE(*CHAR) LEN(10)
DCL          VAR(&JOBNBR) TYPE(*CHAR) LEN(6)
DCL          VAR(&THRESHTYPE) TYPE(*CHAR) LEN(30)
DCL          VAR(&THRESHVAL) TYPE(*DEC) LEN(15 0)

        MONMSG      MSGID(CPF0000)
/* INIT RETURN CODE TO CONTINUE RUNNING QUERY */
        CHGVAR      VAR(&RC) VALUE(0)

/* GET VALUES FROM THE INPUT FORMAT */
        CHGVAR      VAR(&JOBNAME) VALUE(%SST(&QRYS0100 13 10))
        CHGVAR      VAR(&JOBUSER) VALUE(%SST(&QRYS0100 23 10))
        CHGVAR      VAR(&JOBNBR) VALUE(%SST(&QRYS0100 33 6))
        CHGVAR      VAR(&THRESHTYPE) VALUE(%SST(&QRYS0100 135 30))
        CHGVAR      VAR(&BINHIGH) VALUE(%BINARY(&QRYS0100 165 4))
        CHGVAR      VAR(&BINLOW) VALUE(%BINARY(&QRYS0100 169 4))
        CHGVAR      VAR(&THRESHVAL) VALUE(%DEC(&BINHIGH 15 0) * +
            &MAXINT + %DEC(&BINLOW 15 0))

/* GENERATE MESSAGE TEXT STRING */

```

```

CHGVAR      VAR(&MSGTXT) VALUE('QUERY SUPERVISOR +
      THRESHOLD TYPE ' *BCAT &THRESHTYPE *TCAT +
      ', THRESHOLD VALUE ' *BCAT
      %CHAR(&THRESHVAL) *BCAT 'REACHED IN JOB ' +
      *BCAT &JOBNBR *TCAT '/' *TCAT &JOBUSER +
      *TCAT '/' *TCAT &JOBNAME )
CHGVAR      VAR(&MSGLEN) VALUE(%LEN(&MSGTXT))
CHGVAR      VAR(%BINARY(&MSGDTA 1 2)) VALUE(&MSGLEN)
CHGVAR      VAR(&MSGDTA) VALUE(&MSGDTA *TCAT &MSGTXT)
SNDPGMMSG  MSGID(SQL7064) MSGF(QSQLMSG) MSGDTA(&MSGDTA) +
      TOMSGQ(*SYSOPR)

ENDPGM

```

Exit program to end query

This Query Supervisor exit program shows how to request that the query that was running when the threshold was reached is to be stopped.

When this exit program is called, it instructs the Query Supervisor to end the query if it is called for either:

1. An ELAPSED TIME threshold when the threshold consumption has reached at least 60 minutes
2. A threshold with a name that starts with the letters TERMINATE

The threshold definitions used by the C and ILE RPG sample programs are:

```

CALL QSYS2.ADD_QUERY_THRESHOLD(THRESHOLD_NAME=>'One hour limit',
      THRESHOLD_TYPE=>'ELAPSED TIME',
      THRESHOLD_VALUE=>'3600');
CALL QSYS2.ADD_QUERY_THRESHOLD(THRESHOLD_NAME=>'TERMINATE FOR STORAGE',
      THRESHOLD_TYPE=>'TEMPORARY STORAGE',
      THRESHOLD_VALUE=>'500');

```

ILE RPG exit program to end query

This Query Supervisor exit program shows how to request that the query should be ended.

To compile this ILE RPG program, use the following command.

```

CRTBNDRPG PGM(SUPERVISOR/QS_ENDQRY) SRCFILE(SUPERVISOR/QRPGLESRC)

```

If CRTRPGMOD and CRTPGM are used to compile and build the program, be sure to include USRPRF(*OWNER) and ACTGRP(*CALLER) on the CRTPGM command.

The program can be registered as an exit program using this command:

```

ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QS_ENDQRY)
THDSAFE(*YES) TEXT('End query after 60 minutes or TERMINATE threshold')

```

By including the THREAD, USRPRF, DFACTGRP, and ACTGRP options in the source, they can be omitted on the compiler command. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program uses THREAD(*CONCURRENT) to ensure thread safety. The program will run in the activation group of the caller.

```

**free
ctl-opt main(QS_endqry);
ctl-opt ccsid(*char : *jobrun);
ctl-opt thread(*concurrent);
/if defined(*crtbnrdpg)
  ctl-opt usrprf(*owner);
  ctl-opt dftactgrp(*no);
  ctl-opt actgrp(*caller);
/endif

dcl-ds QQQ_QRYSV_QRYS0100_t qualified template;
  Header_Size uns(10);
  Format_Name char(8); // QRYS0100 - Query Supervisor
  Job_Name char(10); // Name of the job running the query
  Job_User char(10); // User profile used to initiate job
  Job_Number char(6); // Number of the job running the query

```



```

Subsystem char(10); // Subsystem in which the job is running
User_Name char(10); // The effective user name of the
// thread that reached the threshold
Query_Identifier char(8); // Uniquely identifies the query text
// and its environment (QR0 hash)
Plan_Identifier uns(20); // Uniquely identifies the access plan
// implementing the query.
Threshold_Name ucs2(30) ccsid(1200); // User-defined name for threshold that
// was reached (in CCSID 1200)
Threshold_Type char(30); // Type of threshold reached
Threshold_Consumption_Value int(20); // Actual value reached by the
// query. Units vary with
// Threshold_Type.
Operation_Type int(5); // Type of query operation
Offset_to_SQL_Statement int(10); // 0, if not available
Length_of_SQL_Statement int(10); // 0, if not available
Offset_to_Host_Variable_List int(10); // 0, if no variables
Length_of_Host_Variable_List int(10); // 0, if no variables
Offset_to_CLIENT_ACCTNG int(10); // 0, if not defined
Length_of_CLIENT_ACCTNG int(10); // 0, if not defined
Offset_to_CLIENT_APPLNAME int(10); // 0, if not defined
Length_of_CLIENT_APPLNAME int(10); // 0, if not defined
Offset_to_CLIENT_PROGRAMID int(10); // 0, if not defined
Length_of_CLIENT_PROGRAMID int(10); // 0, if not defined
Offset_to_CLIENT_USERID int(10); // 0, if not defined
Length_of_CLIENT_USERID int(10); // 0, if not defined
Offset_to_CLIENT_WRKSTNNAME int(10); // 0, if not defined
Length_of_CLIENT_WRKSTNNAME int(10); // 0, if not defined
// char SQL_Statement_Text; /* Varying length, CCSID 1200 */
// char Host_Variable_List; /* Varying length, CCSID 1200 */
// char CLIENT_ACCTNG; /* Varying length */
// char CLIENT_APPLNAME; /* Varying length */
// char CLIENT_PROGRAMID; /* Varying length */
// char CLIENT_USERID; /* Varying length */
// char CLIENT_WRKSTNNAME; /* Varying length */
end-ds;

dcl-proc QS_endqry; // Main procedure
dcl-pi *n;
input likeds(QQQ_QRYSV_QRYS0100_t);
rc int(10);
end-pi;

//*****
// Init the return code to continue running the query
//*****
rc = 0;

/* Have we reached an ELAPSED TIME threshold running for 60 minutes? */
if input.Threshold_Type = 'ELAPSED TIME'
and input.Threshold_Consumption_Value >= 60 * 60;
rc = 1; /* Terminate the query */
endif;

/* Have we reached a threshold named TERMINATE...? */
if %subst(input.Threshold_Name:1:%len('TERMINATE')) = 'TERMINATE';
rc = 1;
endif;
return;
end-proc;

```

C exit program to end query

This Query Supervisor exit program shows how to request that the query should be ended.

To compile this C program, use the following commands:

```

CRTCMOD MODULE(SUPERVISOR/END_QUERY) SRCFILE(SUPERVISOR/QCSRC) LOCALETYPE(*LOCALEUCS2)
CRTPGM PGM(SUPERVISOR/END_QUERY) USRPRF(*OWNER) ACTGRP(*CALLER)

```

The program can be registered as an exit program using this command:

```

ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/END_QUERY)
THDSAFE(*YES) TEXT('End query after 60 minutes or TERMINATE threshold')

```

```
#include <string.h>
```

```

#include <wchar.h>
#include <eqqqrsv.h>

/* Threshold names are given in CCSID_1200 */
const wchar_t* TERMINATE_STRING_1200 = L"TERMINATE";

int main(int argc, char* argv[])
{
    const QQQ_QRYSV_QRYS0100_t* input = (QQQ_QRYSV_QRYS0100_t*)argv[1];
    int* rc = (int*)argv[2];

    *rc = 0; /* Default action is to allow query to continue running */

    /* Have we reached a threshold detecting that the query has been
       running for 60 minutes or more? */
    if (memcmp(input->Threshold_Type, "ELAPSED TIME", 12) == 0 &&
        input->Threshold_Consumption_Value >= 60 * 60)
        *rc = 1; /* Terminate the query */

    /* Have we reached a threshold named TERMINATE...? */
    if (memcmp(input->Threshold_Name,
              TERMINATE_STRING_1200,
              sizeof(TERMINATE_STRING_1200)-2) == 0)
        *rc = 1;

    return 0;
}

```

CL exit program to end query

This Query Supervisor exit program shows how to request that the query should be ended.

To compile this ILE CL program, use the following command.

```

CRTBNDCL PGM(SUPERVISOR/QS_ENDQRY) SRCFILE(SUPERVISOR/QCLSRC) DFTACTGRP(*NO) ACTGRP(*CALLER)
USRPRF(*OWNER)

```

The program can be registered as an exit program using this command:

```

ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QS_ENDQRY)
THDSAFE(*YES) TEXT('End query after 60 minutes or TERMINATE threshold')

```

By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program will run in the activation group of the caller.

```

PGM          PARM(&QRYS0100 &RC)
DCL          VAR(&QRYS0100) TYPE(*CHAR) LEN(8192)
DCL          VAR(&RC) TYPE(*INT) LEN(4)
DCL          VAR(&BINLOW) TYPE(*UINT) LEN(4)
DCL          VAR(&BINHIGH) TYPE(*INT) LEN(4)
DCL          VAR(&MAXINT) TYPE(*DEC) LEN(15 0) +
             VALUE(4294967296)
DCL          VAR(&THRESHNAME) TYPE(*CHAR) LEN(60)
DCL          VAR(&THRESHTYPE) TYPE(*CHAR) LEN(30)
DCL          VAR(&THRESHVAL) TYPE(*DEC) LEN(15 0)
DCL          VAR(&TERMINATE) TYPE(*CHAR) LEN(18)
DCL          VAR(&ONEHOUR) TYPE(*INT) LEN(4)

MONMSG      MSGID(CPF0000)
/* INIT RETURN CODE TO CONTINUE RUNNING QUERY */
CHGVAR     VAR(&RC) VALUE(0)
/* SET UNICODE HEX VALUE FOR THE THRESHOLD NAMED "TERMINATE" */
CHGVAR     VAR(&TERMINATE) +
             VALUE(X'005400450052004D0049004E004100540045')
/* VALUE FOR ONE HOUR OF ELAPSED TIME */
CHGVAR     VAR(&ONEHOUR) VALUE(60 * 60)

/* GET VALUES FROM THE INPUT FORMAT */
CHGVAR     VAR(&THRESHNAME) VALUE(%SST(&QRYS0100 75 60))
CHGVAR     VAR(&THRESHTYPE) VALUE(%SST(&QRYS0100 135 30))
CHGVAR     VAR(&BINHIGH) VALUE(%BINARY(&QRYS0100 165 4))
CHGVAR     VAR(&BINLOW) VALUE(%BINARY(&QRYS0100 169 4))
CHGVAR     VAR(&THRESHVAL) VALUE(%DEC(&BINHIGH 15 0) * +
                                   &MAXINT + %DEC(&BINLOW 15 0))

/* DECIDE WHETHER QUERY SHOULD BE ENDED */

```

```

IF          COND((&THRESHTYPE = 'ELAPSED TIME') *AND +
                (&THRESHVAL *GE &ONEHOUR)) +
                THEN(CHGVAR VAR(&RC) VALUE(1))
IF          COND(%SST(&THRESHNAME 1 18) *EQ &TERMINATE) +
                THEN(CHGVAR VAR(&RC) VALUE(1))

ENDPGM

```

Exit program to dump plan cache information for query

This Query Supervisor exit program dumps plan cache information about the query that reached a threshold.

The information is placed in a table named SUPERVISOR.PLAN_DUMPS. If the table does not exist, it will be created. If the table exists, the new plan cache snapshot will be appended to it. The SUPERVISOR schema must exist. See [“DUMP_PLAN_CACHE procedure”](#) on page 340 for details.

Because SQL cannot be used directly in the exit program, the Run SQL (RUNSQL) CL command is submitted as a batch job.

ILE RPG exit program to dump plan cache information for query

This Query Supervisor exit program dumps plan cache information about the query that reached a threshold.

To compile this ILE RPG program, use the following command. Before compiling, the sample source must be modified as indicated by the comment in the code. The USER parameter on the SBMJOB command must name a profile on the system that has the required authority to call QSYS2.DUMP_PLAN_CACHE(). This authority is *JOBCTL special authority or QIBM_DB_SQLADM function usage. In addition, the owner of the QS_DMPQRY program must have *USE authority to the profile specified on the USER parameter. One option is to make USER the same as the program owner.

```
CRTBNDRPG PGM(SUPERVISOR/QS_DMPQRY) SRCFILE(SUPERVISOR/QRPGLESRC)
```

If CRTRPGMOD and CRTPGM are used to compile and build the program, be sure to include USRPRF(*OWNER) and ACTGRP(*CALLER) on the CRTPGM command.

The program can be registered as an exit program using this command:

```
ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QS_DMPQRY)
THDSAFE(*YES) TEXT('Query supervisor dump plan cache')
```

By including the THREAD, USRPRF, DFACTGRP, and ACTGRP options in the source, they can be omitted on the compiler command. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program uses THREAD(*CONCURRENT) to ensure thread safety. The program will run in the activation group of the caller.

```

**free
ctl-opt main(QS_dmpqry);
ctl-opt ccsid(*char : *jobrun);
ctl-opt thread(*concurrent);
/if defined(*crtbndrpg)
  ctl-opt usrprf(*owner);
  ctl-opt dftactgrp(*no);
  ctl-opt actgrp(*caller);
/endif

dcl-ds QQQ_QRYSV_QRYS0100_t qualified template;
  Header_Size uns(10);
  Format_Name char(8); // QRYS0100 - Query Supervisor
  Job_Name char(10); // Name of the job running the query
  Job_User char(10); // User profile used to initiate job
  Job_Number char(6); // Number of the job running the query
  Subsystem char(10); // Subsystem in which the job is running
  User_Name char(10); // The effective user name of the
  // thread that reached the threshold
  Query_Identifier char(8); // Uniquely identifies the query text
  // and its environment (QR0 hash)

```

```

Plan_Identifier uns(20); // Uniquely identifies the access plan
                          // implementing the query.
Threshold_Name ucs2(30) ccsid(1200); // User-defined name for threshold that
                          // was reached (in CCSID 1200)
Threshold_Type char(30); // Type of threshold reached
Threshold_Consumption_Value int(20); // Actual value reached by the
                          // query. Units vary with
                          // Threshold_Type.
Operation_Type int(5); // Type of query operation
Offset_to_SQL_Statement int(10); // 0, if not available
Length_of_SQL_Statement int(10); // 0, if not available
Offset_to_Host_Variable_List int(10); // 0, if no variables
Length_of_Host_Variable_List int(10); // 0, if no variables
Offset_to_CLIENT_ACCTNG int(10); // 0, if not defined
Length_of_CLIENT_ACCTNG int(10); // 0, if not defined
Offset_to_CLIENT_APPLNAME int(10); // 0, if not defined
Length_of_CLIENT_APPLNAME int(10); // 0, if not defined
Offset_to_CLIENT_PROGRAMID int(10); // 0, if not defined
Length_of_CLIENT_PROGRAMID int(10); // 0, if not defined
Offset_to_CLIENT_USERID int(10); // 0, if not defined
Length_of_CLIENT_USERID int(10); // 0, if not defined
Offset_to_CLIENT_WRKSTNNAME int(10); // 0, if not defined
Length_of_CLIENT_WRKSTNNAME int(10); // 0, if not defined
// char SQL_Statement_Text; /* Varying length, CCSID 1200 */
// char Host_Variable_List; /* Varying length, CCSID 1200 */
// char CLIENT_ACCTNG; /* Varying length */
// char CLIENT_APPLNAME; /* Varying length */
// char CLIENT_PROGRAMID; /* Varying length */
// char CLIENT_USERID; /* Varying length */
// char CLIENT_WRKSTNNAME; /* Varying length */
end-ds;

dcl-proc QS_dmpqry; // Main procedure
  dcl-pi *n;
    input likeds(QQQ_QRYSV_QRY0100_t);
    rc int(10);
  end-pi;

dcl-pr qcmdexc extpgm;
  cmd char(1000) const;
  cmdlen packed(15:5) const;
end-pr;

dcl-s cmdstr char(1000);

//*****
// Init the return code to continue running the query
//*****
rc = 0;

cmdstr =
  'SBMJOB CMD(RUNSQL SQL(''call qsys2.dump_plan_cache('
  + 'FILESCHMA => ''SUPERVISOR''', FILENAME => ''PLANDUMPS'' '
  + ', PLAN_IDENTIFIER => ''
  + %char(input.Plan_Identifier) + ''')) '
  + 'USER('
  + 'MYAUTHUSER' // Replace MYAUTHUSER with a user profile authorized
  // to DUMP_PLAN_CACHE().
  + ')';
qcmdexc (cmdstr : %size(cmdstr));

return;
end-proc;

```

C exit program to dump plan cache information for query

This Query Supervisor exit program dumps plan cache information about the query that reached a threshold.

To compile this C program, use the following commands. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The macro definition for AUTHORIZED_USER should be customized, replacing MYAUTHUSER with a profile that has the required authority to call QSYS2.DUMP_PLAN_CACHE(). This authority is *JOBCTL special authority or QIBM_DB_SQLADM function usage. In addition, the owner of the DUMP_QUERY program must have *USE

authority to the profile identified by AUTHORIZED_USER. One option is to make AUTHORIZER_USER the same as the program owner.

```
CRTCMOD MODULE(SUPERVISOR/DUMP_QUERY) SRCFILE(SUPERVISOR/QCSRC)
DEFINE('AUTHORIZED_USER="MYAUTHUSER')

CRTPGM PGM(SUPERVISOR/DUMP_QUERY) USRPRF(*OWNER) ACTGRP(*CALLER)
```

The program can be registered as an exit program using this command:

```
ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/DUMP_QUERY)
THDSAFE(*YES) TEXT('Query Supervisor dump plan cache')
```

```
#include <stdio.h>
#include <stddef.h>
#include <string.h>
#include <stdlib.h>
#include <eqqqrsv.h>

int main(int argc, char* argv[])
{
    char plan_id_string[21];
    char cmd[600];

    const QQQ_QRYSV_QRYS0100_t* input = (QQQ_QRYSV_QRYS0100_t*)argv[1];
    int* rc = (int*)argv[2];

    *rc = 0; /* don't terminate the query */

    /*****
    /* SQL cannot be run within the exit program, so submit to batch. */
    *****/
    memset(cmd, 0, sizeof(cmd));
    strcat(cmd, "SBMJOB CMD(RUNSQL SQL('call qsys2.dump_plan_cache(");
    strcat(cmd, "FILESCHHEMA => 'SUPERVISOR', FILENAME => 'PLANDUMPS'"));
    strcat(cmd, ", PLAN_IDENTIFIER => ");
    sprintf(plan_id_string, "%llu", input->Plan_Identifier);
    strcat(cmd, plan_id_string);
    strcat(cmd, "')' ) USER(");
    strcat(cmd, AUTHORIZED_USER);
    strcat(cmd, ")");
    system(cmd);
}
```

CL exit program to dump plan cache information for query

This Query Supervisor exit program dumps plan cache information about the query that reached a threshold.

To compile this ILE CL program, use the following command. Before compiling, the sample source must be modified as indicated by the comment in the code. The USER parameter on the SBMJOB command must name a profile on the system that has the required authority to call QSYS2.DUMP_PLAN_CACHE(). This authority is *JOBCTL special authority or QIBM_DB_SQLADM function usage. In addition, the owner of the QS_DMPQRY program must have *USE authority to the profile specified on the USER parameter. One option is to make USER the same as the program owner.

```
CRTBNDCL PGM(SUPERVISOR/QS_DMPQRY) SRCFILE(SUPERVISOR/QCLSRC) DFTACTGRP(*NO) ACTGRP(*CALLER)
USRPRF(*OWNER)
```

The program can be registered as an exit program using this command:

```
ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QS_DMPQRY)
THDSAFE(*YES) TEXT('Query supervisor dump plan cache')
```

By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program will run in the activation group of the caller.

```
PGM          PARM(&QRYS0100 &RC)
DCL          VAR(&QRYS0100) TYPE(*CHAR) LEN(8192)
```

```

DCL      VAR(&RC) TYPE(*INT) LEN(4)
DCL      VAR(&BINLOW) TYPE(*UINT) LEN(4)
DCL      VAR(&BINHIGHU) TYPE(*UINT) LEN(4)
DCL      VAR(&MAXINT) TYPE(*DEC) LEN(15 0) +
          VALUE(4294967296)
DCL      VAR(&PLANID) TYPE(*DEC) LEN(15 0)

MONMSG   MSGID(CPF0000)
/* INIT RETURN CODE TO CONTINUE RUNNING QUERY */
CHGVAR   VAR(&RC) VALUE(0)

/* GET VALUES FROM THE INPUT FORMAT */
CHGVAR   VAR(&BINHIGHU) VALUE(%BINARY(&QRYS0100 67 4))
CHGVAR   VAR(&BINLOW) VALUE(%BINARY(&QRYS0100 71 4))
CHGVAR   VAR(&PLANID) VALUE(%DEC(&BINHIGHU 15 0) * +
          &MAXINT + %DEC(&BINLOW 15 0))

/* SUBMIT JOB TO DUMP THE PLAN CACHE FOR THIS QUERY */
SBMJOB   CMD(RUNSQL SQL('CALL QSYS2.DUMP_PLAN_CACHE(' +
          *BCAT 'FILESCHHEMA => 'SUPERVISOR', +
          FILENAME => 'PLANDUMPS', +
          PLAN_IDENTIFIER=> ' ' *TCAT +
          %CHAR(&PLANID) *TCAT ' '))) +
          USER(MYAUTHUSER)
          /* Replace MYAUTHUSER with a user profile */
          /* authorized to DUMP_PLAN_CACHE().
*/
ENDPGM

```

Exit program to log information using a data queue

This Query Supervisor exit program shows how to use an asynchronous job to log information about the query that reported a threshold.

This example is more complicated. It illustrates how a lightweight exit program can offload processing to a background job for more complex consumption. By doing this, the exit program can complete quicker, causing less of a disruption to the query that caused the threshold to be reported. The example has three pieces.

1. Configuration details
2. An exit program (written in ILE RPG or C)
3. An SQL procedure that performs the actual logging

The exit program produces a JSON object encoded in UTF-16 that contains all of the threshold information sent to the exit program. It sends this JSON object to a data queue named SUPERVISOR/SUPER_DQ. An SQL procedure running in a separate job consumes JSON objects from the data queue. In this example, the procedure logs the entry to a log table.

Setup configuration

There are several objects that need to be defined for this data queue processing solution.

The following table is used to log the Query Supervisor data received from the SEND_JSON exit program. It is populated by the process_QS_data_queue() procedure defined later.

```

create schema SQE_QUERY_SUPERVISOR for schema SUPERVISOR ;

create table SQE_QUERY_SUPERVISOR.SUPERVISOR_LOG for system name SUPERLOG (
  THRESHOLD_TIMESTAMP for column WHENHIT          timestamp,
  JOB_NAME              varchar(10)              default null,
  JOB_USER              varchar(10)              default null,
  JOB_NUMBER            for column JOBNUM          varchar(6)        default null,
  SUBSYSTEM             varchar(10)              default null,
  USER_NAME             varchar(10)              default null,
  THRESHOLD_NAME        for column THRESHNAME     varchar(30)       ccsid 1208 default null,
  THRESHOLD_TYPE        for column THRESHTYPE     varchar(30)       ccsid 1208 default null,
  THRESHOLD_CONSUMPTION_VALUE for column THRESHVAL bigint      default null,
  OPERATION_TYPE        for column OP             smallint         default null,
  SQL_STATEMENT_TEXT    for column STMTTEXT       varchar(10000)    ccsid 1208 default null,
  HOST_VARIABLE_LIST    for column HVLIST        varchar(10000)    ccsid 1208 default null,
  CLIENT_ACCTNG         for column "ACCTNG"       varchar(255)      ccsid 1208 default null,
  CLIENT_APPLNAME       for column "APPLNAME"     varchar(255)      ccsid 1208 default null,
  CLIENT_PROGRAMID      for column "PROGRAMID"    varchar(255)      ccsid 1208 default null,

```

CLIENT_USERID	for column "USERID"	varchar(255)	ccsid 1208 default null,
CLIENT_WRKSTNNAME	for column "WRKSTNNAME"	varchar(255)	ccsid 1208 default null,
QUERY_IDENTIFIER	for column QRO_HASH	varchar(8)	ccsid 1208 default null,
QUERY_PLAN_IDENTIFIER	for column plan_id	decimal(20,0)	default null);

The following data queue is the communication mechanism used by the exit program that got control because of a Query Supervisor threshold having been met or exceeded. The detail passed in the data queue is a JSON document where each key name is a piece of data passed to the Query Supervisor exit program.

```
c1: CRTDTAQ DTAQ(SUPERVISOR/SUPER_DQ) MAXLEN(45000) FORCE(*YES) SIZE(*MAX2GB 1000)
  AUTORCL(*YES) TEXT('Query Supervisor Threshold Processor');
```

Once the programs are in place, a job must be started to run the asynchronous SQL procedure. It will sit and wait for data queue entries to process.

This command can be added to the QSTRUP processing to allow it to be automatically restarted upon system IPL

```
c1: SBMJOB CMD(RUNSQL SQL('call supervisor.process_QS_data_queue()') COMMIT(*NONE))
  JOB(SUPERMGR)
  INLASPGRP(*NONE) LOG(4 00 *SECLVL) JOBMMSGQFL(*PRTWRAP) CCSID(37);
```

ILE RPG exit program to log information using a data queue

This Query Supervisor exit program shows how to use an asynchronous job to log information about the query that reported a threshold.

To compile this ILE RPG program, use the following command.

```
CRTBNDRPG PGM(SUPERVISOR/QS_DTAQ) SRCFILE(SUPERVISOR/QRPGLESRC)
```

If CRTRPGMOD and CRTPGM are used to compile and build the program, be sure to include USRPRF(*OWNER) and ACTGRP(*CALLER) on the CRTPGM command.

The program can be registered as an exit program using this command:

```
ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/QS_DTAQ)
  THDSAFE(*YES) TEXT('Query Supervisor logging with data queue')
```

By including the THREAD, USRPRF, DFACTGRP, and ACTGRP options in the source, they can be omitted on the compiler command. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program. The program uses THREAD(*CONCURRENT) to ensure thread safety. The program will run in the activation group of the caller.

```
**free
ctl-opt main(QS_dtaq);
ctl-opt ccsid(*char : *jobrun);
ctl-opt thread(*concurrent);
/if defined(*crtbnbrpg)
  ctl-opt usrprf(*owner);
  ctl-opt dftactgrp(*no);
  ctl-opt actgrp(*caller);
/endif

dcl-ds QQQ_QRYSV_QRYS0100_t qualified template;
  Header_Size uns(10);
  Format_Name char(8); // QRYS0100 - Query Supervisor
  Job_Name char(10); // Name of the job running the query
  Job_User char(10); // User profile used to initiate job
  Job_Number char(6); // Number of the job running the query
  Subsystem char(10); // Subsystem in which the job is running
  User_Name char(10); // The effective user name of the
  // thread that reached the threshold
  Query_Identifier char(8); // Uniquely identifies the query text
  // and its environment (QRO hash)
  Plan_Identifier uns(20); // Uniquely identifies the access plan
  // implementing the query.
  Threshold_Name ucs2(30) ccsid(1200); // User-defined name for threshold that
```

```

// was reached (in CCSID 1200)
Threshold_Type char(30); // Type of threshold reached
Threshold_Consumption_Value int(20); // Actual value reached by the
// query. Units vary with
// Threshold_Type.
Operation_Type int(5); // Type of query operation
Offset_to_SQL_Statement int(10); // 0, if not available
Length_of_SQL_Statement int(10); // 0, if not available
Offset_to_Host_Variable_List int(10); // 0, if no variables
Length_of_Host_Variable_List int(10); // 0, if no variables
Offset_to_CLIENT_ACCTNG int(10); // 0, if not defined
Length_of_CLIENT_ACCTNG int(10); // 0, if not defined
Offset_to_CLIENT_APPLNAME int(10); // 0, if not defined
Length_of_CLIENT_APPLNAME int(10); // 0, if not defined
Offset_to_CLIENT_PROGRAMID int(10); // 0, if not defined
Length_of_CLIENT_PROGRAMID int(10); // 0, if not defined
Offset_to_CLIENT_USERID int(10); // 0, if not defined
Length_of_CLIENT_USERID int(10); // 0, if not defined
Offset_to_CLIENT_WRKSTNNAME int(10); // 0, if not defined
Length_of_CLIENT_WRKSTNNAME int(10); // 0, if not defined
// char SQL_Statement_Text; /* Varying length, CCSID 1200 */
// char Host_Variable_List; /* Varying length, CCSID 1200 */
// char CLIENT_ACCTNG; /* Varying length */
// char CLIENT_APPLNAME; /* Varying length */
// char CLIENT_PROGRAMID; /* Varying length */
// char CLIENT_USERID; /* Varying length */
// char CLIENT_WRKSTNNAME; /* Varying length */
end-ds;

// Properties of the data queue that we will send the JSON to:
dcl-c DataQLib 'SUPERVISOR';
dcl-c DataQName 'SUPER_DQ';
dcl-c DataQMaxBytes 45000;

dcl-pr QSNDDTAQ extpgm;
DataQName char(10) const;
DataQLib char(10) const;
len packed(5) const;
data char(DataQMaxBytes) options(*varsize);
end-pr;

dcl-c MAX_SOURCE_LEN 11000;
dcl-c MAX_TARGET_LEN 22500;

// - Define a lookup table that indicates how to handle escapable JSON
// characters. Each entry represents one of the first 160 UTF-16
// codepoints.
// - The array JsonEscapeAction is initialized in subroutine
// initEscapeAction
dcl-c None 0;
dcl-c UTF16 1;
dcl-c Backspace 2;
dcl-c Tab 3;
dcl-c Newline 4;
dcl-c Formfeed 5;
dcl-c CR 6;
dcl-c Quote 7;
dcl-c Backslash 8;

dcl-s JsonEscapeAction int(10) dim(160);

// Defines properties for the buffer containing the JSON payload.
dcl-ds BufferInfo_t qualified;
baseptr pointer;
curptr pointer;
allocsize int(10);
truncated ind;
end-ds;

// Define constants used for generating JSON escape sequences.
dcl-c HexChars %ucs2('0123456789ABCDEF');
dcl-c EscapeChars %ucs2(' btnfr"\''); // "
dcl-c UnicodePrefix %ucs2('u00');

dcl-pr Qp0zLprintf extproc(*dclcase);
template pointer value options(*string);
parm pointer value options(*string : *nopass);
end-pr;
dcl-c NEWLINE_CH
x'15';

```



```

// SEND_JSON: main
procedure
                                                                    // SEND_JSON:
main procedure
dcl-proc send_json;

    dcl-pi *n;
        input likeds(QQQ_QRYSV_QRYV0100_t);
        rc int(10);
    end-pi;

    dcl-s based_char char(MAX_SOURCE_LEN) based(based_char_ptr);
    dcl-s based_ucs2 ucs2(MAX_SOURCE_LEN) based(based_ucs2_ptr);
    dcl-s based_buf_data char(1) based(buf.baseptr);
    dcl-s debug_based_buf_data_ucs2 ucs2(MAX_TARGET_LEN) based(buf.baseptr);
    dcl-s ts timestamp(0);
    dcl-s pInput pointer;

    dcl-ds buf
likeds(BufferInfo_t);

    // Don't terminate the query
    rc = 0;

    // Initialization
    pInput = %addr(input);
    exsr initEscapeAction;

    // Set up the buffer used for the final JSON payload.
    buf.baseptr = %alloc (DataQMaxBytes);
    buf.alloysize = DataQMaxBytes;
    buf.curptr = buf.baseptr;
    buf.truncated = *off;

    // JSON encoding starts here.
    json_append(buf : '{}');

    // Insert fixed length fields from the input.
    // All fields must be converted into UTF-16 strings, except
    // for the threshold name, which is already UTF-16.

    json_key(buf : 'threshold_timestamp');
    ts = %timestamp();
    json_append(buf : ''
        + %char(%date(ts))
        + ','
        + %char(%time(ts) : *hms)
        + '');

    json_append(buf : ',');
    json_key(buf : 'job_name');
    json_string_value(buf : %trimr(input.Job_Name));

    json_append(buf : ',');
    json_key(buf : 'job_user');
    json_string_value(buf : %trimr(input.Job_User));

    json_append(buf : ',');
    json_key(buf : 'job_number');
    json_string_value(buf : input.Job_Number);

    json_append(buf : ',');
    json_key(buf : 'subsystem');
    json_string_value(buf : %trimr(input.Subsystem));

    json_append(buf : ',');
    json_key(buf : 'user_name');
    json_string_value(buf : %trimr(input.User_Name));

    json_append(buf : ',');
    json_key(buf : 'query_identifier');
    json_string_value(buf : input.Query_Identifier);

    json_append(buf : ',');
    json_key(buf : 'query_plan_identifier');
    json_numeric_value(buf : input.Plan_Identifier);

    json_append(buf : ',');
    json_key(buf : 'threshold_name');
    json_escape_value(buf : %trimr(input.Threshold_Name));

```

```

json_append(buf : ',');
json_key(buf : 'threshold_type');
json_string_value(buf : %trimr(input.Threshold_Type));

json_append(buf : ',');
json_key(buf : 'threshold_consumption_value');
json_numeric_value(buf : input.Threshold_Consumption_Value);

json_append(buf : ',');
json_key(buf : 'operation_type');
json_numeric_value(buf : input.Operation_Type);

// Insert variable length fields.
// Client special registers need to be translated to UTF-16
// before they can be escaped and inserted.
// The SQL statement text and the host variable list can be
// copied over (with JSON escaping) as they are, since they
// are passed in as UTF-16 data.

if input.Length_of_CLIENT_ACCTNG > 0;
  json_append(buf : ',');
  json_key(buf : 'client_acctng');
  based_char_ptr = pInput + input.Offset_to_CLIENT_ACCTNG;
  json_escape_value(buf
    : %subst(based_char : 1
      : Input.Length_of_CLIENT_ACCTNG));
endif;

if input.Length_of_CLIENT_APPLNAME > 0;
  json_append(buf : ',');
  json_key(buf : 'client_applname');
  based_char_ptr = pInput + input.Offset_to_CLIENT_APPLNAME;
  json_escape_value(buf
    : %subst(based_char : 1
      : Input.Length_of_CLIENT_APPLNAME));
endif;

if input.Length_of_CLIENT_PROGRAMID > 0;
  json_append(buf : ',');
  json_key(buf : 'client_programid');
  based_char_ptr = pInput + input.Offset_to_CLIENT_PROGRAMID;
  json_escape_value(buf
    : %subst(based_char : 1
      : Input.Length_of_CLIENT_PROGRAMID));
endif;

if input.Length_of_CLIENT_USERID > 0;
  json_append(buf : ',');
  json_key(buf : 'client_userid');
  based_char_ptr = pInput + input.Offset_to_CLIENT_USERID;
  json_escape_value(buf
    : %subst(based_char : 1
      : Input.Length_of_CLIENT_USERID));
endif;

if input.Length_of_CLIENT_WRKSTNNAME > 0;
  json_append(buf : ',');
  json_key(buf : 'client_wrkstnname');
  based_char_ptr = pInput + input.Offset_to_CLIENT_WRKSTNNAME;
  json_escape_value(buf
    : %subst(based_char : 1
      : Input.Length_of_CLIENT_WRKSTNNAME));
endif;

if input.Length_of_SQL_Statement > 0;
  json_append(buf : ',');
  json_key(buf : 'sql_statement_text');
  based_ucs2_ptr = pInput + input.Offset_to_SQL_Statement;
  json_escape_value(buf
    : %subst(based_ucs2 : 1
      : %div(Input.Length_of_SQL_Statement : 2)));
endif;

if input.Length_of_Host_Variable_List > 0;
  json_append(buf : ',');
  json_key(buf : 'host_variable_list');
  based_ucs2_ptr = pInput + input.Offset_to_Host_Variable_List;
  json_escape_value(buf
    : %subst(based_ucs2 : 1
      : %div(Input.Length_of_Host_Variable_List : 2)));
endif;

```

```

json_append(buf : '{}');

if buf.truncated;
    Qp0zLprintf('SEND_JSON ran out of allocated space for '
        + 'constructing the Query Supervisor JSON data.' + NEWLINE_CH);
endif;

QSNDDTAQ(DataQName
    : DataQLib
    : buf.curptr - buf.baseptr
    : based_buf_data);

begsr initEscapeAction;

    jsonEscapeAction =
        %list(
            // x'00'
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            Backspace : Tab : Newline : UTF16 :
            Formfeed : CR : UTF16 : UTF16 :
            // x'10'
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            // x'20'
            None : None : Quote : None :
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            // x'30'
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            // x'40'
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            // x'50'
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            Backslash : None : None : None :
            // x'60'
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            // x'70'
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            None : None : None : None :
            // x'80'
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            // x'90'
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16 :
            UTF16 : UTF16 : UTF16 : UTF16);
        endsr;

on-exit;
    dealloc buf.baseptr;
end-proc;

// Copies a UTF-16 source stream to a UTF-16 target stream, escaping
// JSON characters as needed.
dcl-proc json_escape;
    dcl-pi *n varucs2(MAX_TARGET_LEN) rtnparm;
        source varucs2(MAX_SOURCE_LEN) const;
    end-pi;
    dcl-s i int(10);
    dcl-ds ucs2_to_num qualified;
        c ucs2(1);
        n uns(5) overlay(c);

```

```

end-ds;
dcl-s action int(10);
dcl-s target_pos int(10);
dcl-s target varucs2(40000) inz(*blanks);
dcl-s c ucs2(1);
dcl-s c_num packed(5);

// Loop through all the source UTF-16 characters, copying any that
// do not require escaping, and escaping those that are required.
target_pos = 1;
for i = 1 to %len(source);
  c = %subst(source : i : 1);
  ucs2_to_num.c = c;
  c_num = ucs2_to_num.n;

  // Is the current character in the defined escapable range?
  if c_num <= %elem(JsonEscapeAction);
    action = JsonEscapeAction(c_num + 1);
  else;
    action = None;
  endif;

  if action = None;
    // No escaping needed. Emit the source character.
    %subst(target : target_pos : 1) = c;
    target_pos += 1;
  else;
    %subst(target : target_pos : 1) =
      %subst(EscapeChars : Backslash + 1 : 1);
    target_pos += 1;
    if action = UTF16;
      // We escape this as a UTF16 codepoint: \u00XX
      %subst(target : target_pos : %len(UnicodePrefix)) = UnicodePrefix;
      target_pos += %len(UnicodePrefix);
      %subst(target : target_pos : 1) =
        %subst(HexChars : %div(c_num : 16) + 1 : 1);
      target_pos += 1;
      %subst(target : target_pos : 1) =
        %subst(HexChars : %rem(c_num : 16) + 1 : 1);
      target_pos += 1;
    else;
      // We escape this as a special character: e.g. \t
      %subst(target : target_pos : 1) =
        %subst(EscapeChars : Action + 1 : 1);
      target_pos += 1;
    endif;
  endif;
endfor;

return %subst(target : 1 : target_pos - 1);
end-proc json_escape;

// Returns *on if the buffer space can accommodate "len" Unicode chars
// Returns *off and sets the truncated flag if otherwise.
dcl-proc buffer_has_space;
  dcl-pi *n ind;
  buf likeds(BufferInfo_t);
  len int(10) value;
end-pi;
dcl-s required_buf_len int(10);

required_buf_len = (buf.curptr - buf.baseptr) + (len * 2);
if required_buf_len <= buf.alloclsize;
  return *on;
else;
  buf.truncated = *on;
  return *off;
endif;
end-proc;

// JSON_KEY takes a null-terminated UTF-16 string, delimits
// it with quotes, adds a colon (:), and outputs it to the
// output buffer.
dcl-proc json_key;
  dcl-pi *n;
  buf likeds(BufferInfo_t);
  key varucs2(MAX_SOURCE_LEN) const;
end-pi;

  dcl-s string ucs2(MAX_TARGET_LEN) based(buf.curptr);

```

```

    if buffer_has_space (buf : %len(key) + 3);
        %subst(string : 1 : 1) = '"';
        buf.curptr += 2;
        %subst(string : 1 : %len(key)) = key;
        buf.curptr += 2 * %len(key);
        %subst(string : 1 : 2) = ':';
        buf.curptr += 4;
    endif;
end-proc;

// json_string_value takes string,
// delimits it with quotes, and
// outputs it to the output buffer.
dcl-proc json_string_value;
    dcl-pi
*n;

        buf likeds(BufferInfo_t);
        string varucs2(MAX_SOURCE_LEN) const;
end-pi;

    dcl-s buf_string ucs2(MAX_TARGET_LEN) based(buf.curptr);

    if buffer_has_space (buf : %len(string) + 2);
        %subst(buf_string : 1 : 1) = '"';
        buf.curptr += 2;
        %subst(buf_string : 1 : %len(string)) = string;
        buf.curptr += 2 * %len(string);
        %subst(buf_string : 1 : 1) = '"';
        buf.curptr += 2;
    endif;
end-proc;

// json_numeric_value takes a numeric value
// converts it to UTF-16 and
// outputs it to the output buffer.
dcl-proc json_numeric_value;
    dcl-pi *n;
        buf likeds(BufferInfo_t);
        num int(20) value;
end-pi;

    dcl-s json_string varucs2(MAX_TARGET_LEN);
    dcl-s buf_string ucs2(MAX_TARGET_LEN) based(buf.curptr);

    json_string = %char(num);

    if buffer_has_space (buf : %len(json_string));
        %subst(buf_string : 1 : %len(json_string)) = json_string;
        buf.curptr += 2 * %len(json_string);
    endif;
end-proc;

// json_escape_value takes a UTF-16 string of a given size
// and outputs it to the output buffer after delimiting it
// and escaping any characters as defined by the JSON standard.
dcl-proc json_escape_value;
    dcl-pi *n;
        buf likeds(BufferInfo_t);
        string varucs2(MAX_SOURCE_LEN) const;
end-pi;

    dcl-s buf_string ucs2(MAX_TARGET_LEN) based(buf.curptr);
    dcl-s json_string varucs2(MAX_TARGET_LEN);

    json_string = json_escape (string);

    if buffer_has_space (buf : %len(json_string) + 2);
        %subst(buf_string : 1 : 1) = '"';
        buf.curptr += 2;
        %subst(buf_string : 1 : %len(json_string)) = json_string;
        buf.curptr += 2 * %len(json_string);
        %subst(buf_string : 1 : 1) = '"';
        buf.curptr += 2;
    endif;
end-proc;

// JSON_APPEND takes a UTF-16 string and
// outputs it to the output buffer.
dcl-proc json_append;
    dcl-pi *n;
        buf likeds(BufferInfo_t);

```

```

        string varucs2(MAX_SOURCE_LEN) const;
    end-pi;
    dcl-s buf_string ucs2(MAX_TARGET_LEN) based(buf.curptr);

    if buffer_has_space (buf : %len(string));
        %subst(buf_string : 1 : %len(string)) = string;
        buf.curptr += 2 * %len(string);
    endif;

end-proc;

```

C exit program to log information using a data queue

This Query Supervisor exit program shows how to use an asynchronous job to log information about the query that reported a threshold.

This is the registered exit program. It sends the exit program input data to a data queue named SUPERVISOR/SUPER_DQ for further processing by an SQL procedure. The data is sent as a JSON object, encoded as UTF-16 data.

To compile this C program, use the following commands. By using USRPRF(*OWNER), *PUBLIC does not need to be given access to the objects used by this program. The owning profile of the created program must have authority to the commands and objects used by the exit program.

```

CRTCMOD MODULE(SUPERVISOR/SEND_JSON) SRCFILE(SUPERVISOR/QCSRC) LOCALETYPE(*LOCALEUCS2)

CRTPGM  PGM(SUPERVISOR/SEND_JSON) USRPRF(*OWNER) ACTGRP(*CALLER)

```

The program can be registered as an exit program using this command:

```

ADDEXITPGM EXITPNT(QIBM_QQQ_QRY_SUPER) FORMAT(QRYS0100) PGMNBR(*LOW) PGM(SUPERVISOR/SEND_JSON)
THDSAFE(*YES) TEXT('Query Supervisor logging with data queue')

```

```

#include <except.h>
#include <decimal.h>
#include <qsnddtaq.h>
#include <stddef.h>
#include <string.h>
#include <stdlib.h>
#include <wchar.h>
#include <iconv.h>
#include <assert.h>
#include <time.h>
#include <qp0ztrc.h>
#include <eqqqrysv.h>

/* Properties of the data queue that we will send the JSON to: */
char* DataQLib = "SUPERVISOR";
char* DataQName = "SUPER_DQ ";
const int DataQMaxBytes = 45000;

/* Define a lookup table that indicates how to handle escapable JSON
   characters. Each entry represents one of the first 160 UTF-16
   codepoints. */
enum EscapeAction
    {None, UTF16, Backspace, Tab, Newline, Formfeed, CR, Quote, Backslash};
char JsonEscapeAction[160]={
/*x00*/ UTF16, UTF16, UTF16, UTF16,
        UTF16, UTF16, UTF16, UTF16,
        Backspace, Tab, Newline, UTF16,
        Formfeed, CR, UTF16, UTF16,
/*x10*/ UTF16, UTF16, UTF16, UTF16,
        UTF16, UTF16, UTF16, UTF16,
        UTF16, UTF16, UTF16, UTF16,
        UTF16, UTF16, UTF16, UTF16,
/*x20*/ None, None, Quote, None,
        None, None, None, None,
        None, None, None, None,
        None, None, None, None,
/*x30*/ None, None, None, None,
        None, None, None, None,
        None, None, None, None,
        None, None, None, None,
/*x40*/ None, None, None, None,
        None, None, None, None,
        None, None, None, None,

```

```

None,      None,      None,      None,
/*x50*/ None,      None,      None,      None,
None,      None,      None,      None,
None,      None,      None,      None,
Backslash, None,      None,      None,
/*x60*/ None,      None,      None,      None,
None,      None,      None,      None,
None,      None,      None,      None,
None,      None,      None,      None,
/*x70*/ None,      None,      None,      None,
None,      None,      None,      None,
None,      None,      None,      None,
None,      None,      None,      None,
/*x80*/ UTF16,     UTF16,     UTF16,     UTF16,
UTF16,     UTF16,     UTF16,     UTF16,
UTF16,     UTF16,     UTF16,     UTF16,
UTF16,     UTF16,     UTF16,     UTF16,
/*x90*/ UTF16,     UTF16,     UTF16,     UTF16,
UTF16,     UTF16,     UTF16,     UTF16,
UTF16,     UTF16,     UTF16,     UTF16,
UTF16,     UTF16,     UTF16,     UTF16
};

/* Define constants used for generating JSON escape sequences. */
const wchar_t HexChars[17]=L"0123456789ABCDEF";
const wchar_t EscapeChars[10]= L" bnf\\\\";
const wchar_t UnicodePrefix[4]= L"u00";

/* Copies a UTF-16 source stream to a UTF-16 target stream, escaping
   JSON characters as needed. */
void json_escape(wchar_t** target, int target_length,
                 const wchar_t* source, int source_length)
{
    const wchar_t** target_end = target + target_length;

    /* Loop through all the source UTF-16 characters, copying any that
       do not require escaping, and escaping those that are required. */
    while (source_length-- && target < target_end)
    {
        const wchar_t c = *source;
        enum EscapeAction action;

        /* Is the current character in the defined escapable range? */
        if(c<(sizeof(JsonEscapeAction)/sizeof(JsonEscapeAction[0])))
            action = (enum EscapeAction)JsonEscapeAction[c];
        else
            action = None;

        if (action == None)
        {
            /* No escaping needed. Emit the source character. */
            **target = c;
            ++*target;
        }
        else
        {
            **target = EscapeChars[Backslash];
            ++*target;
            if (action == UTF16)
            {
                if (target+6 >= target_end)
                    break;

                /* We escape this as a UTF16 codepoint: \u00XX */
                memcpy(*target, UnicodePrefix, sizeof(UnicodePrefix)-2);
                *target += sizeof(UnicodePrefix)/2-1;

                **target = HexChars[c >> 4];
                ++*target;

                **target = HexChars[c & 0x0f];
                ++*target;
            }
            else
            {
                if (target >= target_end)
                    break;

                /* We escape this as a special character: e.g. \t */
                **target = EscapeChars[action];
                ++*target;
            }
        }
    }
}

```

```

    }
    ++source;
}
}

/* Uses iconv to convert an input string in Job CCSID to UTF-16. */
size_t convert_to_utf16(const char* input, size_t input_bytes,
                       wchar_t** output, size_t output_bytes,
                       iconv_t converter)
{
    int iconv_rc;
    int original_output_bytes = output_bytes;
    iconv_rc = iconv(converter,
                    &input, &input_bytes,
                    (char**)output, &output_bytes);

    /* If conversion was successful, we'll return the number of bytes that
       were produced. */
    if (iconv_rc >= 0)
        return original_output_bytes - output_bytes;

    /* Conversion was unsuccessful. Put an ERROR value in the output if
       there is room. */
    if (original_output_bytes >= 10)
    {
        memcpy(*output, L"ERROR", 10);
        *output+=5;
        return 10;
    }

    return 0;
}

/* Returns the trimmed length of the given wide string after
   discarding any trailing blanks. */
size_t wtrimmed_length(const wchar_t* input, size_t input_len)
{
    while (input_len && input[input_len-1] == L' ')
        --input_len;

    return input_len;
}

/* Returns the trimmed length of the given string after
   discarding any trailing blanks. */
size_t trimmed_length(const char* input, size_t input_len)
{
    while (input_len && input[input_len-1] == ' ')
        --input_len;

    return input_len;
}

/* Defines properties for the buffer containing the JSON payload. */
typedef struct BufferInfo
{
    volatile wchar_t* baseptr;
    wchar_t* curptry;
    int allocsize;
    volatile iconv_t converter;
    int truncated;
} BufferInfo_t;

/* Calculates the bytes allocated but unused for a given BufferInfo
   object. */
#define BYTES_REMAINING(buf) \
    (buf.allocsize - ((char*)buf.curptry - (char*)buf.baseptr))

/* Returns 1 if the buffer space can accommodate "size" bytes.
   Returns 0 and sets the truncated flag if otherwise. */
inline int buffer_has_space(BufferInfo_t* buf, int size)
{
    if (!buf->truncated &&
        BYTES_REMAINING(*buf) >= size)
        return 1;

    buf->truncated = 1;
    return 0;
}

```



```

/* JSON_KEY takes a null-terminated UTF-16 string, delimits
   it with quotes, adds a colon (:), and outputs it to the
   output buffer. */
#define JSON_KEY(buf, key) \
    if (buffer_has_space(&buf, 6 + (sizeof(key)-2))) \
    { \
        wcsncpy(buf.curptry, L"\"); ++buf.curptry; \
        wcsncpy(buf.curptry, key); buf.curptry+=(sizeof(key)/2)-1; \
        wcsncpy(buf.curptry, L"\"); buf.curptry+=2; \
    }

/* JSON_STRING_VALUE takes a blank-padded string in the job CCSID,
   trims trailing blanks, converts the string to UTF-16,
   delimits it with quotes, and copies it to the output buffer. */
#define JSON_STRING_VALUE(buf, value) \
    if (buffer_has_space(&buf, 2)) \
    { \
        wcsncpy(buf.curptry, L"\"); ++buf.curptry; \
    } \
    if (!buf.truncated) \
    { \
        convert_to_utf16(value, \
            trimmed_length(value, sizeof(value)), \
            &buf.curptry, \
            BYTES_REMAINING(buf), \
            buf.converter); \
    } \
    if (buffer_has_space(&buf, 2)) \
    { \
        wcsncpy(buf.curptry, L"\"); ++buf.curptry; \
    }

/* JSON_NUMERIC_VALUE takes a UTF-16 string of a given size
   and outputs it to the output buffer. */
#define JSON_NUMERIC_VALUE(buf, value, count) \
    if (buffer_has_space(&buf, count * 2)) \
    { \
        wcsncpy(buf.curptry, value, count); buf.curptry+=count; \
    }

/* JSON_ESCAPE_VALUE takes a UTF-16 string of a given size
   and outputs it to the output buffer after delimiting it
   and escaping any characters as defined by the JSON standard. */
#define JSON_ESCAPE_VALUE(buf, value, count) \
    if (buffer_has_space(&buf, 2)) \
    { \
        wcsncpy(buf.curptry, L"\"); ++buf.curptry; \
    } \
    if (!buf.truncated) \
    { \
        json_escape(&buf.curptry, \
            BYTES_REMAINING(buf) / 2, \
            value, \
            count); \
    } \
    if (buffer_has_space(&buf, 2)) \
    { \
        wcsncpy(buf.curptry, L"\"); ++buf.curptry; \
    }

/* JSON_APPEND takes a null-terminated UTF-16 string and
   outputs it to the output buffer. */
#define JSON_APPEND(buf, string) \
    if (buffer_has_space(&buf, sizeof(string)-2)) \
    { \
        wcsncpy(buf.curptry, string); buf.curptry += (sizeof(string)/2)-1; \
    }

/* Cleanup handler function to ensure all resources are freed */
void cleanup(_CNL_Hndlr_Parms_T* cancel_info)
{
    Qp0zlprintf("SEND_JSON is cleaning up \n");
    BufferInfo_t* buf = (BufferInfo_t*)(cancel_info->Com_Area);
    free((void*)buf->baseptr);
    iconv_close(buf->converter);
}

int main(int argc, char* argv[])
{
    const QQQ_QRYSV_QRYS0100_t* input = (QQQ_QRYSV_QRYS0100_t*)argv[1];

```

```

int* rc = (int*)argv[2];

BufferInfo_t buf;
wchar_t translate_buffer[1024];
time_t current_time;
size_t char_count;
wchar_t* work_ptr;
size_t converted_bytes;
struct tm *timeptr;
char* inputAsCharData;
char from_code[32], to_code[32];

/* Don't terminate the query */
*rc = 0;

/* Set up the buffer used for the final JSON payload. */
buf.baseptr = (wchar_t*)malloc(DataQMaxBytes);
buf.alloysize = DataQMaxBytes;
buf.curptr = (wchar_t*) buf.baseptr;
buf.truncated = 0;

/* Prepare the CCSID converter. We convert from Job CCSID
   to CCSID 1200 (UTF-16). */
memset(from_code, 0, sizeof(from_code));
memset(to_code, 0, sizeof(to_code));
memcpy(from_code, "IBMCCSID000000000000", 20);
memcpy(to_code, "IBMCCSID01200", 13);
buf.converter = iconv_open(to_code, from_code);
assert(buf.converter.return_value == 0);

#pragma cancel_handler (cleanup, buf)

/* JSON encoding starts here. */
JSON_APPEND(buf, L"{}");

/* Insert fixed length fields from the input.
   All fields must be converted into UTF-16 strings, except
   for the threshold name, which is already UTF-16. */

JSON_KEY(buf, L"threshold_timestamp");
JSON_APPEND(buf, L"\");
current_time = time(NULL);
timeptr = localtime(&current_time);
char_count = wcsftime(buf.curptr,
                      sizeof(translate_buffer)/2,
                      L"%F %T",
                      timeptr);
buf.curptr += char_count;
JSON_APPEND(buf, L"\");

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"job_name");
JSON_STRING_VALUE(buf, input->Job_Name);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"job_user");
JSON_STRING_VALUE(buf, input->Job_User);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"job_number");
JSON_STRING_VALUE(buf, input->Job_Number);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"subsystem");
JSON_STRING_VALUE(buf, input->Subsystem);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"user_name");
JSON_STRING_VALUE(buf, input->User_Name);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"query_identifier");
JSON_STRING_VALUE(buf, input->Query_Identifier);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"query_plan_identifier");
char_count = swprintf(translate_buffer,
                      sizeof(translate_buffer)/2,
                      L"%llu",
                      input->Plan_Identifier);
JSON_NUMERIC_VALUE(buf, translate_buffer, char_count);

```

```

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"threshold_name");
JSON_ESCAPE_VALUE(buf,
    (wchar_t*)(input->Threshold_Name),
    wtrimmmed_length((wchar_t*)(input->Threshold_Name),
        sizeof(input->Threshold_Name)/2));

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"threshold_type");
JSON_STRING_VALUE(buf, input->Threshold_Type);

JSON_APPEND(buf, L",");
JSON_KEY(buf, L"threshold_consumption_value");
char_count = swprintf(translate_buffer,
    sizeof(translate_buffer)/2,
    L"%lld",
    input->Threshold_Consumption_Value);
JSON_NUMERIC_VALUE(buf, translate_buffer, char_count);

JSON_APPEND(buf, L",");
char_count = swprintf(translate_buffer,
    sizeof(translate_buffer)/2,
    L"%d",
    (int)input->Operation_Type);
JSON_KEY(buf, L"operation_type");
JSON_NUMERIC_VALUE(buf, translate_buffer, char_count);

/* Insert variable length fields.
   Client special registers need to be translated to UTF-16
   before they can be escaped and inserted.
   The SQL statement text and the host variable list can be
   copied over (with JSON escaping) as they are, since they
   are passed in as UTF-16 data. */

inputAsCharData = (char*)input;

if (input->Length_of_CLIENT_ACCTNG)
{
    JSON_APPEND(buf, L",");
    JSON_KEY(buf, L"client_acctng");
    work_ptr = translate_buffer;
    converted_bytes =
        convert_to_utf16(inputAsCharData+input->Offset_to_CLIENT_ACCTNG,
            input->Length_of_CLIENT_ACCTNG,
            &work_ptr,
            sizeof(translate_buffer),
            buf.converter);
    JSON_ESCAPE_VALUE(buf, translate_buffer, converted_bytes / 2);
}

if (input->Length_of_CLIENT_APPLNAME)
{
    JSON_APPEND(buf, L",");
    JSON_KEY(buf, L"client_applname");
    work_ptr = translate_buffer;
    converted_bytes =
        convert_to_utf16(inputAsCharData+input->Offset_to_CLIENT_APPLNAME,
            input->Length_of_CLIENT_APPLNAME,
            &work_ptr,
            sizeof(translate_buffer),
            buf.converter);
    JSON_ESCAPE_VALUE(buf, translate_buffer, converted_bytes / 2);
}

if (input->Length_of_CLIENT_PROGRAMID)
{
    JSON_APPEND(buf, L",");
    JSON_KEY(buf, L"client_programid");
    work_ptr = translate_buffer;
    converted_bytes =
        convert_to_utf16(inputAsCharData+input->Offset_to_CLIENT_PROGRAMID,
            input->Length_of_CLIENT_PROGRAMID,
            &work_ptr,
            sizeof(translate_buffer),
            buf.converter);
    JSON_ESCAPE_VALUE(buf, translate_buffer, converted_bytes / 2);
}

if (input->Length_of_CLIENT_USERID)
{
    JSON_APPEND(buf, L",");

```

```

JSON_KEY(buf, L"client_userid");
work_ptr = translate_buffer;
converted_bytes =
    convert_to_utf16(inputAsCharData+input->Offset_to_CLIENT_USERID,
                    input->Length_of_CLIENT_USERID,
                    &work_ptr,
                    sizeof(translate_buffer),
                    buf.converter);
JSON_ESCAPE_VALUE(buf, translate_buffer, converted_bytes / 2);
}

if (input->Length_of_CLIENT_WRKSTNNAME)
{
JSON_APPEND(buf, L",");
JSON_KEY(buf, L"client_wrkstnname");
work_ptr = translate_buffer;
converted_bytes =
    convert_to_utf16(inputAsCharData+input->Offset_to_CLIENT_WRKSTNNAME,
                    input->Length_of_CLIENT_WRKSTNNAME,
                    &work_ptr,
                    sizeof(translate_buffer),
                    buf.converter);
JSON_ESCAPE_VALUE(buf, translate_buffer, converted_bytes / 2);
}

if (input->Length_of_SQL_Statement)
{
JSON_APPEND(buf, L",");
JSON_KEY(buf, L"sql_statement_text");
JSON_ESCAPE_VALUE(buf,
    (wchar_t*)(inputAsCharData+input->Offset_to_SQL_Statement),
    input->Length_of_SQL_Statement/2);
}

if (input->Length_of_Host_Variable_List)
{
JSON_APPEND(buf, L",");
JSON_KEY(buf, L"host_variable_list");
JSON_ESCAPE_VALUE(buf,
    (wchar_t*)(inputAsCharData+input->Offset_to_Host_Variable_List),
    input->Length_of_Host_Variable_List/2);
}

JSON_APPEND(buf, L"}");

if (buf.truncated)
{
    Qp0zLprintf("SEND_JSON ran out of allocated space for "
               "constructing the Query Supervisor JSON data.\n");
}

QSNDTAQ(DataQName,
        DataQLib,
        (decimal(5,0))((char*)buf.curptr-(char*)buf.baseptr),
        (void *)buf.baseptr
        );

#pragma disable_handler

free((void*)buf.baseptr);
iconv_close(buf.converter);

return 0;
}

```

SQL procedure used to log information using a data queue

This SQL procedure, running in an asynchronous job, logs information about the query that reported a threshold.

The procedure receives entries from the SUPERVISOR/SUPER_DQ data queue. Each entry represents an instance of a Query Supervisor threshold having been met or exceeded. The Query Supervisor exit program that got control put the entry on the data queue so it could be handled asynchronously by this SQL procedure.

Using MONITOR = *SYSTEM, ensures that the procedure's activity is not supervised by Query Supervisor. This prevents possible recursion if the procedure itself were to reach a Query Supervisor threshold.

```

create or replace procedure supervisor.process_QS_data_queue ()
  modifies sql data
  set option commit = *none, datfmt = *usa, datsep = *slash, output = *print, monitor =
*system
begin
  declare local_sqlcode integer;
  declare local_sqlstate char(5) for sbcs data;
  declare v_message_text varchar(200) for sbcs data;
  declare v_message_data_utf8 clob(63k) ccsid 1208;
  declare v_message_data_binary blob(63k);
  declare v_message_data_binary_length integer;
  declare v_threshold_timestamp timestamp;
  declare v_job_name varchar(10) ccsid 1208;
  declare v_job_user varchar(10) ccsid 1208;
  declare v_job_number varchar(6) ccsid 1208;
  declare v_subsystem varchar(10) ccsid 1208;
  declare v_user_name varchar(10) ccsid 1208;
  declare v_threshold_name varchar(30) ccsid 1208;
  declare v_threshold_type varchar(30) ccsid 1208;
  declare v_threshold_consumption_value bigint;
  declare v_operation_type smallint;
  declare v_sql_statement_text varchar(10000) ccsid 1208;
  declare v_host_variable_list varchar(10000) ccsid 1208;
  declare v_client_acctng varchar(255) ccsid 1208;
  declare v_client_applname varchar(255) ccsid 1208;
  declare v_client_programid varchar(255) ccsid 1208;
  declare v_client_userid varchar(255) ccsid 1208;
  declare v_client_wrkstnname varchar(255) ccsid 1208;
  declare v_query_identifier varchar(8) ccsid 1208;
  declare v_query_plan_identifier decimal(20,0);
  declare continue handler for sqlexception
  begin
    get diagnostics condition 1
      local_sqlcode = db2_returned_sqlcode, local_sqlstate = returned_sqlstate,
      v_message_text = message_text;
    call systools.lprintf('process_QS_data_queue() - failed with SQLCODE:' concat
      error_sqlcode concat ' SQLSTATE:' concat error_sqlstate concat ' and MESSAGE:'
      concat v_message_text);
  end; /* end of continue handler */
  while (1 = 1) do
    select MESSAGE_DATA_BINARY
      into v_message_data_binary
      from
        table (
          qsys2.receive_data_queue(
            data_queue => 'SUPER_DQ', data_queue_library => 'SUPERVISOR',
            wait_time => -1) -- any negative means to wait forever for the next message
        );

    -- Convert the UTF16 data to UTF8
    set v_message_data_binary_length = length(v_message_data_binary) / 2;
    set v_message_data_utf8 = (
      select
        interpret(
          varbinary_format(hex(v_message_data_binary_length)) concat v_message_data_binary
          as dbclob(63k) ccsid 1200)
      from sysibm.sysdummy1);

    select THRESHOLD_TIMESTAMP, JOB_NAME, JOB_USER, JOB_NUMBER, SUBSYSTEM, USER_NAME,
      THRESHOLD_NAME, THRESHOLD_TYPE, THRESHOLD_CONSUMPTION_VALUE, OPERATION_TYPE,
      SQL_STATEMENT_TEXT, HOST_VARIABLE_LIST, CLIENT_ACCTNG, CLIENT_APPLNAME,
      CLIENT_PROGRAMID, CLIENT_USERID, CLIENT_WRKSTNNAME, QUERY_IDENTIFIER,
      QUERY_PLAN_IDENTIFIER
      into v_THRESHOLD_TIMESTAMP, v_JOB_NAME, v_JOB_USER, v_JOB_NUMBER, v_SUBSYSTEM,
      v_USER_NAME, v_THRESHOLD_NAME, v_THRESHOLD_TYPE, v_THRESHOLD_CONSUMPTION_VALUE,
      v_OPERATION_TYPE, v_SQL_STATEMENT_TEXT, v_HOST_VARIABLE_LIST,
      v_CLIENT_ACCTNG, v_CLIENT_APPLNAME, v_CLIENT_PROGRAMID, v_CLIENT_USERID,
      v_CLIENT_WRKSTNNAME, v_QUERY_IDENTIFIER, v_QUERY_PLAN_IDENTIFIER
      from
        json_table(
          v_message_data_utf8,
          '$'
          columns(
            threshold_timestamp timestamp path 'lax $.threshold_timestamp',
            job_name varchar(10) path 'lax $.job_name',
            job_user varchar(10) path 'lax $.job_user',
            job_number varchar(6) path 'lax $.job_number',

```

```

        subsystem varchar(10) path 'lax $.subsystem',
        user_name varchar(10) path 'lax $.user_name',
        threshold_name varchar(30) path 'lax $.threshold_name',
        threshold_type varchar(30) path 'lax $.threshold_type',
        threshold_consumption_value bigint path
        'lax $.threshold_consumption_value',
        operation_type smallint path 'lax $.operation_type',
        sql_statement_text varchar(10000) path 'lax $.sql_statement_text',
        host_variable_list varchar(10000) path 'lax $.host_variable_list',
        client_acctng varchar(255) path 'lax $.client_acctng',
        client_applname varchar(255) path 'lax $.client_applname',
        client_programid varchar(255) path 'lax $.client_programid',
        client_userid varchar(255) path 'lax $.client_userid',
        client_wrkstnname varchar(255) path 'lax $.client_wrkstnname',
        query_identifier varchar(8) path 'lax $.query_identifier',
        query_plan_identifier bigint path 'lax $.query_plan_identifier'
    )
);
insert into SQE_QUERY_SUPERVISOR.SUPERVISOR_LOG (
    THRESHOLD_TIMESTAMP, JOB_NAME, JOB_USER, JOB_NUMBER, SUBSYSTEM, USER_NAME,
    THRESHOLD_NAME, THRESHOLD_TYPE, THRESHOLD_CONSUMPTION_VALUE, OPERATION_TYPE,
    SQL_STATEMENT_TEXT, HOST_VARIABLE_LIST, CLIENT_ACCTNG, CLIENT_APPLNAME,
    CLIENT_PROGRAMID, CLIENT_USERID, CLIENT_WRKSTNNAME, QUERY_IDENTIFIER,
    QUERY_PLAN_IDENTIFIER)
values (v_THRESHOLD_TIMESTAMP, v_JOB_NAME, v_JOB_USER, v_JOB_NUMBER, v_SUBSYSTEM,
v_USER_NAME, v_THRESHOLD_NAME, v_THRESHOLD_TYPE, v_THRESHOLD_CONSUMPTION_VALUE,
v_OPERATION_TYPE, v_SQL_STATEMENT_TEXT, v_HOST_VARIABLE_LIST, v_CLIENT_ACCTNG,
v_CLIENT_APPLNAME, v_CLIENT_PROGRAMID, v_CLIENT_USERID, v_CLIENT_WRKSTNNAME,
v_QUERY_IDENTIFIER, v_QUERY_PLAN_IDENTIFIER);
end while;
end;

```

Setting resource limits with the Predictive Query Governor

The Db2 for i Predictive Query Governor can stop the initiation of a query if the estimated run time (elapsed execution time) or estimated temporary storage for the query is excessive. The governor acts *before* a query is run instead of while a query is run. The governor can be used in any interactive or batch job on the system. It can be used with all Db2 for i query interfaces and is not limited to use with SQL queries.

The ability of the governor to predict and stop queries before they are started is important because:

- Operating a long-running query and abnormally ending the query before obtaining any results wastes system resources.
- Some CQE operations within a query cannot be interrupted by the **End Request (ENDRQS)** CL command. The creation of a temporary index or a query using a column function without a GROUP BY clause are two examples of these types of queries. It is important to not start these operations if they take longer than the user wants to wait.

The governor in Db2 for i is based on two measurements:

- The estimated runtime for a query.
- The estimated temporary storage consumption for a query.

If the query estimated runtime or temporary storage usage exceed the user-defined limits, the initiation of the query can be stopped.

To define a time limit (in seconds) for the governor to use, do one of the following:

- Use the Query Time Limit (QRYTIMLMT) parameter on the **Change Query Attributes (CHGQRYA)** CL command. The command language used is the first place where the optimizer attempts to find the time limit.
- Set the Query Time Limit option in the query options file. The query options file is the second place where the query optimizer attempts to find the time limit.
- Set the QQRYTIMLMT system value. Allow each job to use the value *SYSVAL on the **Change Query Attributes (CHGQRYA)** CL command, and set the query options file to *DEFAULT. The system value is the third place where the query optimizer attempts to find the time limit.

To define a temporary storage limit (in megabytes) for the governor to use, do the following:

- Use the Query Storage Limit (QRYSTGLMT) parameter on the **Change Query Attributes (CHGQRYA)** CL command. The command language used is the first place where the query optimizer attempts to find the limit.
- Set the Query Storage Limit option STORAGE_LIMIT in the query options file. The query options file is the second place where the query optimizer attempts to find the time limit.

The time and temporary storage values generated by the optimizer are *only* estimates. The actual query runtime might be more or less than the estimate. In certain cases when the optimizer does not have full information about the data being queried, the estimate could vary considerably from the actual resource used. In those cases, you might need to artificially adjust your limits to correspond to an inaccurate estimate.

When setting the time limit for the entire system, set it to the maximum allowable time that any query must be allowed to run. By setting the limit too low you run the risk of preventing some queries from completing and thus preventing the application from successfully finishing. There are many functions that use the query component to internally perform query requests. These requests are also compared to the user-defined time limit.

You can check the inquiry message CPA4259 for the predicted runtime and storage. If the query is canceled, debug messages are still written to the job log.

You can also add the Query Governor Exit Program that is called when estimated runtime and temporary storage limits have exceeded the specified limits.

Related information

[Query Governor Exit Program](#)

[End Request \(ENDRQS\) command](#)

[Change Query Attributes \(CHGQRYA\) command](#)

Using the Query Governor

The resource governor works with the query optimizer.

When a user issues a request to the system to run a query, the following occurs:

1. The query access plan is created by the optimizer.

As part of the evaluation, the optimizer predicts or estimates the runtime for the query. This estimate helps determine the best way to access and retrieve the data for the query. In addition, as part of the estimating process, the optimizer also computes the estimated temporary storage usage for the query.
2. The estimated runtime and estimated temporary storage are compared against the user-defined query limit currently in effect for the job or user session.
3. If the estimates for the query are less than or equal to the specified limits, the query governor lets the query run without interruption. No message is sent to the user.
4. If the query limit is exceeded, inquiry message CPA4259 is sent to the user. The message states the estimates as well as the specified limits. Realize that only one limit needs to be exceeded; it is possible that you see that only one limit was exceeded. Also, if no limit was explicitly specified by the user, a large integer value is shown for that limit.

Note: A default reply can be established for this message so that the user does not have the option to reply. The query request is *always* ended.
5. If a default message reply is not used, the user chooses to do one of the following:
 - End the query request before it is run.
 - Continue and run the query even though the estimated value exceeds the associated governor limit.

Setting the resource limits for jobs other than the current job

You can set either or both resource limits for a job other than the current job. You set these limits by using the JOB parameter on the **Change Query Attributes (CHGQRYA)** command. Specify either a query options file library to search (QRYOPLIB) or a specific QRYTIMLMT, or QRYSTGLMT, or both for that job.

Using the resource limits to balance system resources

After the source job runs the **Change Query Attributes (CHGQRYA)** command, effects of the governor on the target job are not dependent upon the source job. The query resource limits remain in effect for the duration of the job or user session, or until a resource limit is changed by a **Change Query Attributes (CHGQRYA)** command.

Under program control, a user might be given different limits depending on the application function performed, time of day, or system resources available. These limits provide a significant amount of flexibility when trying to balance system resources with temporary query requirements.

Cancel a query with the Query Governor

When a query is expected to take more resources than the set limit, the governor issues inquiry message CPA4259.

You can respond to the message in one of the following ways:

- Enter a C to cancel the query. Escape message CPF427F is issued to the SQL runtime code. SQL returns SQLCODE -666.
- Enter an I to ignore the exceeded limit and let the query run to completion.

Control the default reply to the query governor inquiry message

The system administrator can control whether the interactive user has the option of ignoring the database query inquiry message by using the **Change Job (CHGJOB)** CL command.

Changes made include the following:

- If a value of *DFT is specified for the INQMSGRPY parameter of the **Change Job (CHGJOB)** CL command, the interactive user does not see the inquiry messages. The query is canceled immediately.
- If a value of *RQD is specified for the INQMSGRPY parameter of the **Change Job (CHGJOB)** CL command, the interactive user sees the inquiry. The user must reply to the inquiry.
- If a value of *SYSRPLY is specified for the INQMSGRPY parameter of the **Change Job (CHGJOB)** CL command, a system reply list is used to determine whether the interactive user sees the inquiry and whether a reply is necessary. The system reply list entries can be used to customize different default replies based on user profile name, user id, or process names. The fully qualified job name is available in the message data for inquiry message CPA4259. This algorithm allows the keyword CMPDTA to be used to select the system reply list entry that applies to the process or user profile. The user profile name is 10 characters long and starts at position 51. The process name is 10 character long and starts at position 27.
- The following example adds a reply list element that causes the default reply of C to cancel requests for jobs whose user profile is 'QPGMR'.

```
ADDRPYLE SEQNBR(56) MSGID(CPA4259) CMPDTA(QPGMR 51) RPY(C)
```

The following example adds a reply list element that causes the default reply of C to cancel requests for jobs whose process name is 'QPADEV0011'.

```
ADDRPYLE SEQNBR(57) MSGID(CPA4259) CMPDTA(QPADEV0011 27) RPY(C)
```

Related information

[Change Job \(CHGJOB\) command](#)

Testing performance with the query governor

You can use the query governor to test the performance of your queries.

To test the performance of a query with the query governor, do the following:

1. Set the query time limit to zero (QRYTIMLMT(0)) using the **Change Query Attributes (CHGQRYA)** command or in the INI file. This forces an inquiry message from the governor stating that the estimated time to run the query exceeds the query time limit.
2. Prompt for message help on the inquiry message and find the same information that you can find by running the **Print SQL Information (PRTSQLINF)** command.

The query governor lets you optimize performance without having to run through several iterations of the query.

Additionally, if the query is canceled, the query optimizer evaluates the access plan and sends the optimizer debug messages to the job log. This process occurs even if the job is *not* in debug mode. You can then review the optimizer tuning messages in the job log to see if additional tuning is needed to obtain optimal query performance.

This method allows you to try several permutations of the query with different attributes, indexes, and syntax, or both. You can then determine what performs better through the optimizer without actually running the query to completion. This process saves on system resources because the actual query of the data is never done. If the tables to be queried contain many rows, this method represents a significant savings in system resources.

Be careful when you use this technique for performance testing, because all query requests are stopped before they are run. This caution is especially important for a CQE query that cannot be implemented in a single query step. For these types of queries, separate multiple query requests are issued, and then their results are accumulated before returning the final results. Stopping the query in one of these intermediate steps gives you only the performance information for that intermediate step, and not for the entire query.

Related information

[Print SQL Information \(PRTSQLINF\) command](#)

[Change Query Attributes \(CHGQRYA\) command](#)

Examples of setting query time limits

You can set the query time limit for the current job or user session using query options file QAQQINI. Specify the QRYOPTLIB parameter on the **Change Query Attributes (CHGQRYA)** command. Use a user library where the QAQQINI file exists with the parameter set to QUERY_TIME_LIMIT, and the value set to a valid query time limit.

To set the query time limit for 45 seconds you can use the following **Change Query Attributes (CHGQRYA)** command:

```
CHGQRYA JOB(*) QRYTIMLMT(45)
```

This command sets the query time limit at 45 seconds. If the user runs a query with an estimated runtime equal to or less than 45 seconds, the query runs without interruption. The time limit remains in effect for the duration of the job or user session, or until the time limit is changed by the **Change Query Attributes (CHGQRYA)** command.

Assume that the query optimizer estimated the runtime for a query as 135 seconds. A message is sent to the user that stated that the estimated runtime of 135 seconds exceeds the query time limit of 45 seconds.

To set or change the query time limit for a job other than your current job, the **Change Query Attributes (CHGQRYA)** command is run using the JOB parameter. To set the query time limit to

45 seconds for job 123456/USERNAME/JOBNAME use the following **Change Query Attributes (CHGQRYA)** command:

```
CHGQRYA JOB(123456/USERNAME/JOBNAME) QRYTILMT(45)
```

This command sets the query time limit at 45 seconds for job 123456/USERNAME/JOBNAME. If job 123456/USERNAME/JOBNAME tries to run a query with an estimated runtime equal to or less than 45 seconds the query runs without interruption. If the estimated runtime for the query is greater than 45 seconds, for example, 50 seconds, a message is sent to the user. The message states that the estimated runtime of 50 seconds exceeds the query time limit of 45 seconds. The time limit remains in effect for the duration of job 123456/USERNAME/JOBNAME, or until the time limit for job 123456/USERNAME/JOBNAME is changed by the **Change Query Attributes (CHGQRYA)** command.

To set or change the query time limit to the QRYTILMT system value, use the following **Change Query Attributes (CHGQRYA)** command:

```
CHGQRYA QRYTILMT(*SYSVAL)
```

The QRYTILMT system value is used for duration of the job or user session, or until the time limit is changed by the **Change Query Attributes (CHGQRYA)** command. This use is the default behavior for the **Change Query Attributes (CHGQRYA)** command.

Note: The query time limit can also be set in the INI file, or by using the **Change System Value (CHGSYSVAL)** command.

Related information

[Change Query Attributes \(CHGQRYA\) command](#)

[Change System Value \(CHGSYSVAL\) command](#)

Test temporary storage usage with the query governor

The predictive storage governor specifies a temporary storage limit for database queries. You can use the query governor to test if a query uses any temporary object, such as a hash table, sort, or temporary index.

To test for usage of a temporary object, do the following:

- Set the query storage limit to zero (QRYSTGLMT(0)) using the **Change Query Attributes (CHGQRYA)** command or in the INI file. This forces an inquiry message from the governor anytime a temporary object is used for the query. The message is sent regardless of the estimated size of the temporary object.
- Prompt for message help on the inquiry message and find the same information that you can find by running the **Print SQL Information (PRTSQLINF)** command. This command allows you to see what temporary objects were involved.

Related information

[Print SQL Information \(PRTSQLINF\) command](#)

[Change Query Attributes \(CHGQRYA\) command](#)

Examples of setting query temporary storage limits

The temporary storage limit can be specified either in the QAQQINI file or on the **Change Query Attributes (CHGQRYA)** command.

You can set the query temporary storage limit for a job using query options file QAQQINI. Specify the QRYOPTLIB parameter on the **Change Query Attributes (CHGQRYA)** command. Use a user library where the QAQQINI file exists with a valid value set for parameter STORAGE_LIMIT.

To set the query temporary storage limit on the **Change Query Attributes (CHGQRYA)** command itself, specify a valid value for the QRYSTGLMT parameter.

If a value is specified both on the **Change Query Attributes (CHGQRYA)** command QRYSTGLMT parameter and in the QAQQINI file specified on the QRYOPTLIB parameter, the QRYSTGLMT value is used.

To set the temporary storage limit for 100 MB in the current job, you can use the following **Change Query Attributes (CHGQRYA)** command:

```
CHGQRYA JOB(*) QRYSTGLMT(100)
```

If the user runs any query with an estimated temporary storage consumption equal to or less than 100 MB, the query runs without interruption. If the estimate is more than 100 MB, the CPA4259 inquiry message is sent by the database. To set or change the query time limit for a job other than your current job, the CHGQRYA command is run using the JOB parameter. To set the same limit for job 123456/USERNAME/JOBNAME use the following CHGQRYA command:

```
CHGQRYA JOB(123456/USERNAME/JOBNAME) QRYSTGLMT(100)
```

This sets the query temporary storage limit to 100 MB for job 123456/USERNAME/JOBNAME.

Note: Unlike the query time limit, there is no system value for temporary storage limit. The default behavior is to let any queries run regardless of their temporary storage usage. The query temporary storage limit can be specified either in the INI file or on the **Change Query Attributes (CHGQRYA)** command.

Related information

[Change Query Attributes \(CHGQRYA\) command](#)

Controlling parallel processing for queries

There are two types of parallel processing available. The first is a parallel I/O that is available at no charge. The second is Db2 Symmetric Multiprocessing, a feature that you can purchase. You can turn parallel processing on and off.

Even if parallelism is enabled for a system or job, the individual queries that run in a job might not actually use a parallel method. This decision might be because of functional restrictions, or the optimizer might choose a non-parallel method because it runs faster.

Queries processed with parallel access methods aggressively use main storage, CPU, and disk resources. The number of queries that use parallel processing must be limited and controlled. Activating parallel processing system wide with the QQRYDEGREE system value is not recommended.

The parallel processing level can be controlled at a system or job level. For a job, either the CURRENT DEGREE special register or the PARALLEL_DEGREE QAQQINI option can be used.

Controlling system-wide parallel processing for queries

You can use the QQRYDEGREE system value to control parallel processing for a system.

The current value of the system value can be displayed or modified using the following CL commands:

- **WRKSYSVAL - Work with System Value**
- **CHGSYSVAL - Change System Value**
- **DSPSYSVAL - Display System Value**
- **RTVSYVAL - Retrieve System Value**

The special values for QQRYDEGREE control whether parallel processing is allowed by default for all jobs on the system. The possible values are:

- *NONE** No parallel processing is allowed for database query processing.
- *IO** I/O parallel processing is allowed for queries.
- *OPTIMIZE** The query optimizer can choose to use any number of tasks for either I/O or SMP parallel processing to process the queries. SMP parallel processing is used only if the Db2

Symmetric Multiprocessing feature is installed. The query optimizer chooses to use parallel processing to minimize elapsed time based on the job share of the memory in the pool.

***MAX** The query optimizer can choose to use either I/O or SMP parallel processing to process the query. SMP parallel processing can be used only if the Db2 Symmetric Multiprocessing feature is installed. The choices made by the query optimizer are like the choices made for parameter value *OPTIMIZE. The exception is that the optimizer assumes that all active memory in the pool can be used to process the query.

The default QORYDEGREE system value is *NONE. You must change the value if you want parallel query processing as the default for jobs run on the system.

Changing this system value affects all jobs that is run or are currently running on the system whose DEGREE query attribute is *SYSVAL. However, queries that have already been started or queries using reusable ODPs are not affected.

Controlling job level parallel processing for queries

You can also control query parallel processing at the job level using the DEGREE parameter of the **Change Query Attributes (CHGQRYA)** command or in the QAQQINI file. You can also use the SET_CURRENT_DEGREE SQL statement.

Using the Change Query Attributes (CHGQRYA) command

The parallel processing option allowed and, optionally, the number of tasks that can be used when running database queries in the job can be specified. You can prompt on the **Change Query Attributes (CHGQRYA)** command in an interactive job to display the current values of the DEGREE query attribute.

Changing the DEGREE query attribute does not affect queries that have already been started or queries using reusable ODPs.

The parameter values for the DEGREE keyword are:

- *SAME** The parallel degree query attribute does not change.
- *NONE** No parallel processing is allowed for database query processing.
- *IO** The CQE optimizer can use any number of tasks when it chooses to use I/O parallel processing for queries. SMP parallel processing is not allowed. The SQE optimizer considers I/O parallelism with or without this setting.
- *OPTIMIZE** The query optimizer can choose to use any number of tasks for either I/O or SMP parallel processing to process the query. SMP parallel processing can be used only if the Db2 Symmetric Multiprocessing feature is installed. Use of parallel processing and the number of tasks used is determined by:
- the number of system processors available
 - the job share of active memory available in the pool
 - whether the expected elapsed time is limited by CPU processing or I/O resources
- The query optimizer chooses an implementation that minimizes elapsed time based on the job share of the memory in the pool.
- *MAX** The query optimizer can choose to use either I/O or SMP parallel processing to process the query. SMP parallel processing can be used only if the Db2 Symmetric Multiprocessing feature is installed. The choices made by the query optimizer are like the choices made for parameter value *OPTIMIZE. The exception is that the optimizer assumes that all active memory in the pool can be used to process the query.
- *NBRTASKS**
number-of-tasks Specifies the number of tasks to be used when the query optimizer chooses to use SMP parallel processing to process a query. I/O parallelism is also allowed. SMP parallel processing can be used only if the Db2 Symmetric Multiprocessing feature is installed.

Using a number of tasks less than the number of system processors available restricts the number of processors used simultaneously for running a query. A larger number of tasks ensures that the query is allowed to use all the processors available on the system to run the query. Too many tasks can degrade performance because of the over commitment of active memory and the overhead cost of managing all the tasks.

***SYSVAL** Specifies that the processing option used is set to the current value of the QQRYDEGREE system value.

The initial value of the DEGREE attribute for a job is *SYSVAL.

Using the SET CURRENT DEGREE SQL statement

You can use the SET CURRENT DEGREE SQL statement to change the value of the CURRENT_DEGREE special register. The possible values for the CURRENT_DEGREE special register are:

- 1** No parallel processing is allowed.
- 2 through 32767** Specifies the degree of parallelism that is used.
- ANY** Specifies that the database manager can choose to use any number of tasks for either I/O or SMP parallel processing. Use of parallel processing and the number of tasks used is determined by:
- the number of system processors available
 - the job share of active memory available in the pool
 - whether the expected elapsed time is limited by CPU processing or I/O resources
- The database manager chooses an implementation that minimizes elapsed time based on the job share of the memory in the pool.
- NONE** No parallel processing is allowed.
- MAX** The database manager can choose to use any number of tasks for either I/O or SMP parallel processing. MAX is like ANY except the database manager assumes that all active memory in the pool can be used.
- IO** The CQE optimizer can use any number of tasks when it chooses to use I/O parallel processing for queries. SMP parallel processing is not allowed. The SQE optimizer considers I/O parallelism with or without this setting.

The value can be changed by invoking the SET CURRENT DEGREE statement.

The initial value of CURRENT DEGREE comes from the CHGQRYA CL command, PARALLEL_DEGREE parameter in the current query options file (QAQQINI), or the QQRYDEGREE system value.

Related information

[Set Current Degree statement](#)

[Change Query Attributes \(CHGQRYA\) command](#)

[DB2 Symmetric Multiprocessing](#)

Collecting statistics with the statistics manager

The collection of statistics is handled by a separate component called the statistics manager. Statistical information can be used by the query optimizer to determine the best access plan for a query. Since the query optimizer bases its choice of access plan on the statistical information found in the table, it is important that this information is current.

On many platforms, statistics collection is a manual process that is the responsibility of the database administrator. With IBM i products, the database statistics collection process is handled automatically, and only rarely is it necessary to update statistics manually.

The statistics manager does not actually run or optimize the query. It controls the access to the metadata and other information that is required to optimize the query. It uses this information to answer questions posed by the query optimizer. The answers can either be derived from table header information, from existing indexes, or from single-column statistics.

The use of indexes for estimates includes Maintained Temporary Indexes (MTIs), which the statistics manager is able to use in the same way as standard permanent indexes. When deciding which index to use for an estimate, the statistics engine will prioritize permanent indexes over equivalent MTIs.

The statistics manager must always provide an answer to the questions from the Optimizer. It uses the best method available to provide the answers. For example, it could use a single-column statistic or perform a key range estimate over an index. Along with the answer, the statistics manager returns a confidence level to the optimizer that the optimizer can use to provide greater latitude for sizing algorithms. If the statistics manager provides a low confidence in the number of groups estimated for a grouping request, the optimizer can increase the size of the temporary hash table allocated.

Related concepts

Statistics manager

In CQE, the retrieval of statistics is a function of the Optimizer. When the Optimizer needs to know information about a table, it looks at the table description to retrieve the row count and table size. If an index is available, the Optimizer might extract information about the data in the table. In SQE, the collection and management of statistics is handled by a separate component called the statistics manager. The statistics manager leverages all the same statistical sources as CQE, but adds more sources and capabilities.

Automatic statistics collection

When the statistics manager prepares its responses to the optimizer, it tracks the responses that were generated using default filter factors. Default filter factors are used when column statistics or indexes are not available. The statistics manager uses this information to automatically generate a statistic collection request for the columns. This request occurs while the access plan is written to the plan cache. If system resources allow, statistics collections occur in real time for direct use by the current query, avoiding a default answer to the optimizer.

Otherwise, as system resources become available, the requested column statistics are collected in the background. The next time the query is executed, the missing column statistics are available to the statistics manager. This process allows the statistics manager to provide more accurate information to the optimizer at that time. More statistics make it easier for the optimizer to generate a better performing access plan.

If a query is canceled before or during execution, the requests for column statistics are still processed. These requests occur if the execution reaches the point where the generated access plan is written to the Plan Cache.

To minimize the number of passes through a table during statistics collection, the statistics manager groups multiple requests for the same table. For example, two queries are executed against table T1. The first query has selection criteria on column C1 and the second over column C2. If no statistics are available for the table, the statistics manager identifies both of these columns as good candidates for column statistics. When the statistics manager reviews requests, it looks for multiple requests for the same table and groups them into one request. This grouping allows both column statistics to be created with only one pass through table T1.

One thing to note is that column statistics are usually automatically created when the statistics manager must answer questions from the optimizer using default filter factors. However, when an index is available that might be used to generate the answer, then column statistics are not automatically generated. In this scenario, there might be cases where optimization time benefits from column statistics. Using column statistics to answer questions from the optimizer is more efficient than using the index data. So if query performance seems extended, you might want to verify that there are indexes over the relevant columns in your query. If so, try manually generating column statistics for these columns.

As stated before, statistics collection occurs as system resources become available. If you have a low priority job permanently active on your system that is supposed to use all spare CPU cycles for processing, your statistics collection is never active.

Automatic statistics refresh

Column statistics are not maintained when the underlying table data changes. The statistics manager determines if columns statistics are still valid or if they no longer represent the column accurately (stale).

This validation is done each time one of the following occurs:

- A full open occurs for a query where column statistics were used to create the access plan
- A new plan is added to the plan cache, either because a new query was optimized or because an existing plan was reoptimized.
- A significant number of inserts / updates / deletes have been made to a table. The actual number of changes before checking the column statistics staleness depends on several factors, including the number of entries in the table.

To validate the statistics, the statistics manager checks to see if any of the following apply:

- Number of rows in the table has changed by more than 15% of the total table row count
- Number of rows changed in the table is more than 15% of the total table row count

If the statistics are stale, the statistics manager still uses them to answer the questions from the optimizer. However, the statistics manager marks the statistics as stale in the plan cache and generates a request to refresh them.

Viewing statistics requests

You can view the current statistics requests by using System i Navigator or by using Statistics APIs.

To view requests in System i Navigator, right-click **Database** and select **Statistic Requests**. This window shows all user requested statistics collections that are pending or active. The view also shows all system requested statistics collections that are being considered, are active, or have failed. You can change the status of the request, order the request to process immediately, or cancel the request.

Related reference

[Statistics manager APIs](#)

You can use APIs to implement the statistics function of System i Navigator.

Indexes and column statistics

While performing similar functions, indexes and column statistics are different.

If you are trying to decide whether to use statistics or indexes to provide information to the statistics manager, keep in mind the following differences.

One major difference between indexes and column statistics is that indexes are permanent objects that are updated when changes to the underlying table occur. Column statistics are not updated. If your data is constantly changing, the statistics manager might need to rely on stale column statistics. However, maintaining an index after each table change might use more system resources than refreshing stale column statistics after a group of changes have occurred.

Another difference is the effect that the existence of new indexes or column statistics has on the optimizer. When new indexes become available, the optimizer considers them for implementation. If they are candidates, the optimizer reoptimizes the query and tries to find a better implementation. However, this reoptimization is not true for column statistics. When new or refreshed column statistics are available, the statistics manager interrogates immediately. Reoptimization occurs only if the answers are different from the ones that were given before these refreshed statistics. It is possible to use statistics that are refreshed without causing a reoptimization of an access plan.

Also, column statistics and EVIs can be used to recognize data skew in joins, which standard radix indexes are not able to do. So, when join columns contain certain values that occur far more frequently than other values, column statistics or EVIs may be used to help provide better estimates.

Requests for selectivity information from indexes are obtained in the form of key range estimate operations. In order to get accurate information for a given predicate, the statistics manager may need to evaluate a significant portion of an index's data. This evaluation can take a long time for large indexes as index pages are loaded into active memory and then analyzed. The optimizer must wait for each estimate operation to complete before continuing with query optimization. Each query that is optimized will generally require multiple key range estimates. As a result, the time spent obtaining these estimates can significantly impact the overall time required to optimize a query.

In order to shorten the time required to optimize a query, the statistics manager uses several different strategies when faced with the need to perform a key range estimate. One strategy is to avoid getting the estimate when other information (e.g. file size or selectivity information from column statistics) indicate that the information returned by the estimate would not be worth the time required to obtain it. A second strategy is to cache prior estimates so that subsequent requests for the same key range can be returned quickly. Each index has a cache associated with it. Estimates are stored in this cache until the cache fills or becomes stale relative to the data in the index or until an IPL occurs. A third strategy, introduced in IBM i 7.3, is to quickly return a less accurate initial estimate while submitting a background request for a more accurate full estimate. With this strategy, the statistics manager adds a timeout value to each key range estimate operation. If the operation does not complete in the time indicated by the timeout value, the operation is interrupted. The statistics manager then looks at some other fast estimate sources (including similar ranges from the cache and a shallower evaluation of the index) and returns the best answer obtained thus far. This allows the query to resume optimizing with a reasonable estimate but without waiting for the full estimate to complete. In the meantime, the QDBFSTCCOL system job will take up the request to process the full key range estimate and will place an updated (and more accurate) final result in the estimate cache when it is ready. Once the full estimate is complete and in the cache, any query which used the initial estimate will have its plan invalidated and re-optimized on the next execution.

The amount of time that the statistics manager waits for a key range estimate to complete can be configured in the QAQQINI query options file with option KEY_RANGE_ESTIMATE_TIMEOUT. This control should be seen as a way to influence the optimization time and not as a hard limit. Although KEY_RANGE_ESTIMATE_TIMEOUT limits the amount of time that the optimizer may take for any individual estimate operation, it does not guarantee an upper bound on the overall optimization time. This is because the optimizer may request multiple key range estimates, each of which receives the full timeout duration.

When trying to determine the selectivity of predicates, the statistics manager considers column statistics and indexes as resources for its answers in the following order with some caveats and exceptions:

1. Try to recognize if an index-based estimate would be very long running, and if so circumvent the index estimate and use column statistics instead
2. Try to use a multi-column keyed index when ANDed or ORed predicates reference multiple columns
3. If there is no perfect index that contains all the columns in the predicates, it tries to find a combination of indexes that can be used.
4. For single column questions, it uses available column statistics
5. If the answer derived from the column statistics shows a selectivity of less than 2%, indexes are used to verify this answer

Accessing column statistics to answer questions is faster than trying to obtain these answers from indexes.

Column statistics can only be used by SQE. For CQE, all statistics are retrieved from indexes.

Finally, column statistics can be used only for query optimization. They cannot be used for the actual implementation of a query, whereas indexes can be used for both.

Monitoring background statistics collection

The system value QDBFSTCCOL controls who is allowed to create statistics in the background.

The following list provides the possible values:

- ***ALL** Allows all statistics to be collected in the background. *ALL is the default setting.
- ***NONE** Restricts everyone from creating statistics in the background. *NONE does not prevent immediate user-requested statistics from being collected, however.
- ***USER** Allows only user-requested statistics to be collected in the background.
- ***SYSTEM** Allows only system-requested statistics to be collected in the background.

When you switch the system value to something other than *ALL or *SYSTEM, the statistics manager continues to place statistics requests in the plan cache. When the system value is switched back to *ALL, for example, background processing analyzes the entire plan cache and looks for any existing column statistics requests. This background task also identifies column statistics that have been used by a plan in the plan cache. The task determines if these column statistics have become stale. Requests for the new column statistics as well as requests for refresh of the stale columns statistics are then executed.

All background statistic collections initiated by the system or submitted by a user are performed by the system job QDBFSTCCOL. User-initiated immediate requests are run within the user job. This job uses multiple threads to create the statistics. The number of threads is determined by the number of processors that the system has. Each thread is then associated with a request queue.

There are four types of request queues based on who submitted the request and how long the collection is estimated to take. The default priority assigned to each thread can determine to which queue the thread belongs:

- Priority 90 — short user requests
- Priority 93 — long user requests
- Priority 96 — short system requests
- Priority 99 — long system requests

Background statistics collections attempt to use as much parallelism as possible. This parallelism is independent of the SMP feature installed on the system. However, parallel processing is allowed only for immediate statistics collection if SMP is installed on the system. The job that requests the column statistics also must allow parallelism.

Related information

[Performance system values: Allow background database statistics collection](#)

Replication of column statistics with CRTDUPOBJ versus CPYF

You can replicate column statistics with the **Create Duplicate Object (CRTDUPOBJ)** or the **Copy File (CPYF)** commands.

Statistics are not copied to new tables when using the **Copy File (CPYF)** command. If statistics are needed immediately after using this command, then you must manually generate the statistics using System i Navigator or the statistics APIs. If statistics are not needed immediately, then they could be created automatically by the system after the first touch of a column by a query.

Statistics are copied when using **Create Duplicate Object (CRTDUPOBJ)** command with DATA(*YES). You can use this command as an alternative to creating statistics automatically after using a **Copy File (CPYF)** command.

Related information

[Create Duplicate Object \(CRTDUPOBJ\) command](#)

[Copy File \(CPYF\) command](#)

Determining what column statistics exist

You can determine what column statistics exist in a couple of ways.

The first is to view statistics by using System i Navigator. Right-click a table or alias and select **Statistic Data**. Another way is to create a user-defined table function and call that function from an SQL statement or stored procedure.

Manually collecting and refreshing statistics

You can manually collect and refresh statistics through System i Navigator or by using statistics APIs.

To collect statistics using System i Navigator, right-click a table or alias and select **Statistic Data**. On the **Statistic Data** dialog, click **New**. Then select the columns that you want to collect statistics for. Once you have selected the columns, you can collect the statistics immediately or collect them in the background.

To refresh a statistic using System i Navigator, right-click a table or alias and select **Statistic Data**. Click **Update**. Select the statistic that you want to refresh. You can collect the statistics immediately or collect them in the background.

There are several scenarios in which the manual management (create, remove, refresh, and so on) of column statistics could be beneficial and recommended.

High Availability (HA) solutions	High availability solutions replicate data to a secondary system by using journal entries. However, column statistics are not journaled. That means that, on your backup system, no column statistics are available when you first start using that system. To prevent the "warm up" effect, you might want to propagate the column statistics that were gathered on your production system. Recreate them on your backup system manually.
ISV (Independent Software Vendor) preparation	An ISV might want to deliver a customer solution that already includes column statistics frequently used in the application, rather than waiting for the automatic statistics collection to create them. Run the application on the development system for some time and examine which column statistics were created automatically. You can then generate a script file to execute on the customer system after the initial data load takes place. The script file can be shipped as part of the application
Business Intelligence environments	<p>In a large Business Intelligence environment, it is common for large data load and update operations to occur overnight. Column statistics are marked stale only when they are touched by the statistics manager, and then refreshed after first touch. You might want to consider refreshing the column statistics manually after loading the data.</p> <p>You can do this refresh easily by toggling the system value QDBFSTCCOL to *NONE and then back to *ALL. This process causes all stale column statistics to be refreshed. It also starts collection of any column statistics previously requested by the system but not yet available. Since this process relies on the access plans stored in the plan cache, avoid performing a system initial program load (IPL) before toggling QDBFSTCCOL. An IPL clears the plan cache.</p> <p>This procedure works only if you do not delete (drop) the tables and recreate them in the process of loading your data. When deleting a table, access plans in the plan cache that refer to this table are deleted. Information about column statistics on that table is also lost. The process in this environment is either to add data to your tables or to clear the tables instead of deleting them.</p>
Massive data updates	Updating rows in a column statistics-enabled table can significantly change the cardinality, add new ranges of values, or change the distribution of data values. These updates can affect query performance on the first query run against the new data. On the first run of such a query, the optimizer uses stale column statistics to determine the access plan. At that point, it starts a request to refresh the column statistics.

Prior to this data update, you might want to toggle the system value QDBFSTCCOL to *NONE and back to *ALL or *SYSTEM. This toggle causes an analysis of the plan cache. The analysis includes searching for column statistics used in access plan generation, analyzing them for staleness, and requesting updates for the stale statistics.

If you massively update or load data, and run queries against these tables at the same time, the automatic column statistics collection tries to refresh every time 15% of the data is changed. This processing can be redundant since you are still updating or loading the data. In this case, you might want to block automatic statistics collection for the tables and unblock it again after the data update or load finishes. An alternative is to turn off automatic statistics collection for the whole system before updating or loading the data. Switch it back on after the updating or loading has finished.

Backup and recovery

When thinking about backup and recovery strategies, keep in mind that creation of column statistics is not journaled. Column statistics that exist at the time a save operation occurs are saved as part of the table and restored with the table. Any column statistics created after the save took place are lost and cannot be recreated by using techniques such as applying journal entries. If you have a long interval between save operations and rely on journaling to restore your environment, consider tracking column statistics that are generated after the latest save operation.

Related information

[Performance system values: Allow background database statistics collection](#)

Statistics manager APIs

You can use APIs to implement the statistics function of System i Navigator.

- [Cancel Requested Statistics Collections \(QDBSTCRS, QdbstCancelRequestedStatistics\)](#) immediately cancels statistics collections that have been requested, but are not yet completed or not successfully completed.
- [Delete Statistics Collections \(QDBSTDS, QdbstDeleteStatistics\)](#) immediately deletes existing completed statistics collections.
- [List Requested Statistics Collections \(QDBSTLRS, QdbstListRequestedStatistics\)](#) lists all the columns and combination of columns and file members that have background statistic collections requested, but not yet completed.
- [List Statistics Collection Details \(QDBSTLDS, QdbstListDetailStatistics\)](#) lists additional statistics data for a single statistics collection.
- [List Statistics Collections \(QDBSTLS, QdbstListStatistics\)](#) lists all the columns and combination of columns for a given file member that have statistics available.
- [Request Statistics Collections \(QDBSTRS, QdbstRequestStatistics\)](#) allows you to request one or more statistics collections for a given set of columns of a specific file member.
- [Update Statistics Collection \(QDBSTUS, QdbstUpdateStatistics\)](#) allows you to update the attributes and to refresh the data of an existing single statistics collection

Related reference

[Viewing statistics requests](#)

You can view the current statistics requests by using System i Navigator or by using Statistics APIs.

Displaying materialized query table columns

You can display materialized query tables associated with another table using System i Navigator.

To display materialized query tables, follow these steps:

1. In the System i Navigator window, expand the system that you want to use.

2. Expand **Databases** and the database that you want to work with.
3. Expand **Schemas** and the schema that you want to work with.
4. Right-click a table and select **Show Materialized Query Tables**.

<i>Table 53. Columns used in Show materialized query table window</i>	
Column name	Description
Name	The SQL name for the materialized query table
Schema	Schema or library containing the materialized query table
Partition	Partition detail for the index. Possible values: <ul style="list-style-type: none"> • <blank>, which means For all partitions • For Each Partition • specific name of the partition
Owner	The user ID of the owner of the materialized query table.
System Name	System table name for the materialized query table
Enabled	Whether the materialized query table is enabled. Possible values are: <ul style="list-style-type: none"> • Yes • No <p>If the materialized query table is not enabled, it cannot be used for query optimization. It can, however, be queried directly.</p>
Creation Date	The timestamp of when the materialized table was created.
Last Refresh Date	The timestamp of the last time the materialized query table was refreshed.
Last Query Use	The timestamp when the materialized query table was last used by the optimizer to replace user specified tables in a query.
Last Query Statistics Use	The timestamp when the materialized query table was last used by the statistics manager to determine an access method.
Query Use Count	The number of instances the materialized query table was used by the optimizer to replace user specified tables in a query.
Query Statistics Use Count	The number of instances the materialized query table was used by the statistics manager to determine an access method.
Last Used Date	The timestamp when the materialized query table was last used.
Days Used Count	The number of days the materialized query table has been used.
Date Reset Days Used Count	The year and date when the days-used count was last set to 0.
Current Number of Rows	The total number of rows included in this materialized query table at this time.

<i>Table 53. Columns used in Show materialized query table window (continued)</i>	
Column name	Description
Current Size	The current size of the materialized query table.
Last Changed	The timestamp when the materialized query table was last changed.
Maintenance	The maintenance for the materialized query table. Possible values are: <ul style="list-style-type: none"> • User • System
Initial Data	Whether the initial data was inserted immediately or deferred. Possible values are <ul style="list-style-type: none"> • Deferred • Immediate
Refresh Mode	The refresh mode for the materialized query table. A materialized query table can be refreshed whenever a change is made to the table or deferred to a later time.
Isolation Level	The isolation level for the materialized query table.
Sort Sequence	The alternate character sorting sequence for National Language Support (NLS).
Language Identifier	The language code for the object.
SQL Statement	The SQL statement that is used to populate the table.
Text	The text description of the materialized query table.
Table	Schema and table name.
Table Partition	Table partition.
Table System Name	System name of the table.

Managing check pending constraints columns

You can view and change constraints that have been placed in a check pending state by the system. Check pending constraints refers to a state in which a mismatch exists between a parent and foreign key in a referential constraint. A mismatch can also occur between the column value and the check constraint definition in a check constraint.

To view constraints that have been placed in a check pending state, follow these steps:

1. Expand the system name and **Databases**.
2. Expand the database that you want to work with.
3. Expand the **Database Maintenance** folder.
4. Select **Check Pending Constraints**.
5. From this interface, you can view the definition of the constraint and the rows that are in violation of the constraint rules. Select the constraint that you want to work with and then select **Edit Check Pending Constraint** from the **File** menu.
6. You can either alter or delete the rows that are in violation.

<i>Table 54. Columns used in Check pending constraints window</i>	
Column name	Description
Name of Constraint in Check Pending	Displays the name of the constraint that is in a check pending state.
Schema	Schema containing the constraint that is in a check pending state.
Type	Displays the type of constraint that is in check pending. Possible values are: Check constraint Foreign key constraint
Table name	The name of the table associated with the constraint in check pending state.
Enabled	Displays whether the constraint is enabled. The constraint must be disabled or the relationship taken out of the check pending state before any input/output (I/O) operations can be performed.

Creating an index strategy

Db2 for i provides two basic means for accessing tables: a table scan and an index-based retrieval. Index-based retrieval is typically more efficient than table scan when less than 20% of the table rows are selected.

There are two kinds of persistent indexes: binary radix tree indexes, which have been available since 1988, and encoded vector indexes (EVIs), which became available in 1998 with V4R2. Both types of indexes are useful in improving performance for certain kinds of queries.

Binary radix indexes

A radix index is a multilevel, hybrid tree structure that allows many key values to be stored efficiently while minimizing access times. A key compression algorithm assists in this process. The lowest level of the tree contains the leaf nodes, which contain the base table row addresses associated with the key value. The key value is used to quickly navigate to the leaf node with a few simple binary search tests.

The binary radix tree structure is good for finding a few rows because it finds a given row with a minimal amount of processing. For example, create a binary radix index over a customer number column. Then create a typical OLTP request like "find the outstanding orders for a single customer". The binary index results in fast performance. An index created over the customer number column is considered the perfect index for this type of query. The index allows the database to find the rows it needs and perform a minimal number of I/Os.

In some situations, however, you do not always have the same level of predictability. Many users want on demand access to the detail data. For example, they might run a report every week to look at sales data. Then they want to "drill down" for more information related to a particular problem area they found in the report. In this scenario, you cannot write all the queries in advance on behalf of the end users. Without knowing what queries might run, it is impossible to build the perfect index.

Related information

[SQL Create Index statement](#)

Derived key index

You can use the SQL CREATE INDEX statement to create a derived key index using an SQL expression.

Traditionally an index could only specify column names in the key of the index over the table it was based on. With this support, an index can have an expression in place of a column name that can use built-in functions, or some other valid expression. Additionally, you can use the SQL CREATE INDEX statement to create a sparse index using a WHERE condition.

For restrictions and other information about derived indexes, see the Create Index statement and Using derived indexes.

Related reference

[Using derived indexes](#)

SQL indexes can be created where the key is specified as an expression. This type of key is also referred to as a derived key.

Related information

[SQL Create Index statement](#)

Sparse indexes

You can use the SQL CREATE INDEX statement to create a sparse index using SQL selection predicates.

Last release users were given the ability to use the SQL CREATE INDEX statement to create a sparse index using a WHERE condition. With this support, the query optimizer recognizes and considers sparse indexes during its optimization. If the query WHERE selection is a subset of the sparse index WHERE selection, then the sparse index is used to implement the query. Use of the sparse index usually results in improved performance.

Examples

In this example, the query selection is a subset of the sparse index selection and an index scan over the sparse index is used. The remaining query selection (COL3=30) is executed following the index scan.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In this example, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20
```

Related reference

[Using sparse indexes](#)

SQL indexes can be created using WHERE selection predicates. These indexes can also be referred to as sparse indexes. The advantage of a sparse index is that fewer entries are maintained in the index. Only those entries matching the WHERE selection criteria are maintained in the index.

Related information

[SQL Create Index statement](#)

Sparse index optimization

An SQL sparse index is like a select/omit access path. Both the sparse index and the select/omit logical file contain only keys that meet the selection specified. For a sparse index, the selection is specified with a WHERE clause. For a select/omit logical file, the selection is specified in the DDS using the COMP operation.

The reason for creating a sparse index is to provide performance enhancements for your queries. The performance enhancement is done by precomputing and storing results of the WHERE selection in the sparse index. The database engine can use these results instead of recomputing them for a user specified query. The query optimizer looks for any applicable sparse index and can choose to implement the query using a sparse index. The decision is based on whether using a sparse index is a faster implementation choice.

For a sparse index to be used, the WHERE selection in the query must be a subset of the WHERE selection in the sparse index. That is, the set of records in the sparse index must contain all the records to be selected by the query. It might contain extra records, but it must contain all the records to be selected by the query. This comparison of WHERE selection is performed by the query optimizer during optimization. It is like the comparison that is performed for Materialized Query Tables (MQT).

Besides the comparison of the WHERE selection, the optimization of a sparse index is identical to the optimization that is performed for any Binary Radix index.

Refer to section 'Indexes and the Optimizer' for more details on how Binary Radix indexes are optimized.

Related concepts

Indexes & the optimizer

Since the optimizer uses cost based optimization, more information about the database rows and columns makes for a more efficient access plan created for the query. With the information from the indexes, the optimizer can make better choices about how to process the request (local selection, joins, grouping, and ordering).

Related reference

Using sparse indexes

SQL indexes can be created using WHERE selection predicates. These indexes can also be referred to as sparse indexes. The advantage of a sparse index is that fewer entries are maintained in the index. Only those entries matching the WHERE selection criteria are maintained in the index.

Sparse index matching algorithm

This topic is a generalized discussion of how the sparse index matching algorithm works.

The selection in the query must be a subset of the selection in the sparse index in order for the sparse index to be used. This statement is true whether the selection is ANDed together, ORed together, or a combination of the two. For selection where all predicates are ANDed together, all WHERE selection predicates specified in the sparse index must also be specified in the query. The query can contain additional ANDed predicates. The selection for the additional predicates will be performed after the entries are retrieved from the sparse index. See examples A1, A2, and A3 following.

Example A1

In this example, the query selection exactly matches the sparse index selection and an index scan over the sparse index can be used.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

Example A2

In this example, the query selection is a subset of the sparse index selection and an index scan over the sparse index can be used. The remaining query selection (COL3=30) is executed following the index scan.


```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

Example A3

In this example, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20
```

For selection where all predicates are ORed together, all WHERE selection predicates specified in the query, must also be specified in the sparse index. The sparse index can contain additional ORed predicates. All the ORed selection in the query will be executed after the entries are retrieved from the sparse index. See examples O1, O2, and O3 following.

Example O1

In this example, the query selection exactly matches the sparse index selection and an index scan over the sparse index can be used. The query selection is executed following the index scan.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20 or COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20 or COL3=30
```

Example O2

In this example, the query selection is a subset of the sparse index selection and an index scan over the sparse index can be used. The query selection is executed following the index scan.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20 or COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20
```

Example O3

In this example, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20 or COL3=30
```

The previous examples used simple selection, all ANDed, or all ORed together. These examples are not typical, but they demonstrate how the selection of the sparse index is compared to the selection of the query. Obviously, the more complex the selection the more difficult it becomes to determine compatibility.

In the next example T1, the constant 'MN' was replaced by a parameter marker for the query selection. The sparse index had the local selection of COL1='MN' applied to it when it was created. The sparse index matching algorithm matches the parameter marker to the constant 'MN' in the query predicate COL1=?. It verifies that the value of the parameter marker is the same as the constant in the sparse index; therefore the sparse index can be used.

The sparse index matching algorithm attempts to match where the predicates between the sparse index and the query are not the same. An example is a sparse index with a predicate SALARY > 50000, and a query with the predicate SALARY > 70000. The sparse index contains the rows necessary to run the query. The sparse index is used in the query, but the predicate SALARY > 70000 remains as selection in the query (it is not removed).

Example T1

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=? or COL2='TWINS' or COL3='WIN'
```

In the next example T2, the keys of the sparse index match the ORDER BY fields in the query. For the sparse index to satisfy the specified ordering, the optimizer must verify that the query selection is a subset of the sparse index selection. In this example, the sparse index can be used.

Example T2

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL1, COL3)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL2='TWINS'
ORDER BY COL1, COL3
```

Related reference

Using sparse indexes

SQL indexes can be created using WHERE selection predicates. These indexes can also be referred to as sparse indexes. The advantage of a sparse index is that fewer entries are maintained in the index. Only those entries matching the WHERE selection criteria are maintained in the index.

[Details on the MQT matching algorithm](#)

What follows is a generalized discussion of how the MQT matching algorithm works.

Sparse index examples

This topic shows examples of how the sparse index matching algorithm works.

In example S1, the query selection is a subset of the sparse index selection and consequently an index scan over the sparse index is used. The remaining query selection (COL3=30) is executed following the index scan.

Example S1

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In example S2, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

Example S2

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20
```

In example S3, the query selection exactly matches the sparse index selection and an index scan over the sparse index can be used.

Example S3

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In example S4, the query selection is a subset of the sparse index selection and an index scan over the sparse index can be used. The remaining query selection (COL3=30) is executed following the index scan.

Example S4

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In example S5, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

Example S5

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20
```

In example S6, the query selection exactly matches the sparse index selection and an index scan over the sparse index can be used. The query selection is executed following the index scan to eliminate excess records from the sparse index.

Example S6

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20 or COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20 or COL3=30
```

In example S7, the query selection is a subset of the sparse index selection and an index scan over the sparse index can be used. The query selection is executed following the index scan to eliminate excess records from the sparse index.

Example S7

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20 or COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20
```

In example S8, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

Example S8

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20 or COL3=30
```

In the next example S9, the constant 'MN' was replaced by a parameter marker for the query selection. The sparse index had the local selection of COL1='MN' applied to it when it was created. The sparse index matching algorithm matches the parameter marker to the constant 'MN' in the query predicate COL1 =?. It verifies that the value of the parameter marker is the same as the constant in the sparse index; therefore the sparse index can be used.

Example S9

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
Where Col3='WIN' and (Col1=? or Col2='TWINS')
```

In the next example S10, the keys of the sparse index match the order by fields in the query. For the sparse index to satisfy the specified ordering, the optimizer must verify that the query selection is a subset of the sparse index selection. In this example, the sparse index can be used.

Example S10

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL1, COL3)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
Where Col3='WIN' and (Col1='MN' or Col2='TWINS')
ORDER BY COL1, COL3
```

In the next example S11, the keys of the sparse index do not match the order by fields in the query. But the selection in sparse index T2 is a superset of the query selection. Depending on size, the optimizer might choose an index scan over sparse index T2 and then use a sort to satisfy the specified ordering.

Example S11

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL2, COL4)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
Where Col3='WIN' and (Col1='MN' or Col2='TWINS')
ORDER BY COL1, COL3
```

The next example S12 represents the classic optimizer decision: is it better to do an index probe using index IX1 or is it better to do an index scan using sparse index SPR1? Both indexes retrieve the same number of index entries and have the same cost from that point forward. For example, both indexes have the same cost to retrieve the selected records from the dataspace, based on the retrieved entries/keys.

The cost to retrieve the index entries is the deciding criteria. In general, if index IX1 is large then an index scan over sparse index SPR1 has a lower cost to retrieve the index entries. If index IX1 is rather small

then an index probe over index IX1 has a lower cost to retrieve the index entries. Another cost decision is reusability. The plan using sparse index SPR1 is not as reusable as the plan using index IX1 because of the static selection built into the sparse selection.

Example S12

```
CREATE INDEX MYLIB/IX1 on MYLIB/T1 (COL1, COL2, COL3)
```

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
CSELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

Specify PAGESIZE on index creates

You can use the PAGESIZE parameter to specify the access path logical page size used by the system when the access path is created. Use the PAGESIZE parameter when creating keyed files or indexes using the **Create Physical File (CRTPF)** or **Create Logical File (CRTLf)** commands, or the SQL CREATE INDEX statement.

The logical page size is the access path number of bytes that can be moved from auxiliary storage to the job storage pool for a page fault.

Consider using the default of *KEYLEN for this parameter, except in rare circumstances. Then the page size can be determined by the system based on the total length of the keys. When the access path is used by selective queries (for example, individual key lookup), a smaller page size is typically more efficient. When the query-selected keys are grouped in the access path with many records selected, or the access path is scanned, a larger page size is more efficient.

Related information

[Create Logical File \(CRTLf\) command](#)

[Create Physical File \(CRTPF\) command](#)

[SQL Create Index statement](#)

General index maintenance

Whenever indexes are created and used, there is a potential for a decrease in I/O velocity due to maintenance. Therefore, consider the maintenance cost of creating and using additional indexes. For radix indexes with MAINT(*IMMED), maintenance occurs when inserting, updating, or deleting rows.

To reduce the maintenance of your indexes consider:

- Minimizing the number of table indexes by creating composite (multiple column) key indexes. Composite indexes can be used for multiple different situations.
- Dropping indexes during batch inserts, updates, and deletes
- Creating in parallel. Either create indexes, one at a time, in parallel using SMP or create multiple indexes simultaneously with multiple batch jobs
- Maintaining indexes in parallel using SMP

The goal of creating indexes is to improve query performance by providing statistics and implementation choices. Maintain a reasonable balance on the number of indexes to limit maintenance overhead.

Encoded vector indexes

An encoded vector index (EVI) is used to provide fast data access in decision support and query reporting environments.

EVI is a complementary alternative to existing index objects (binary radix tree structure - logical file or SQL index) and are a variation on bitmap indexing. Because of their compact size and relative simplicity, EVIs provide for faster scans of a table that can also be processed in parallel.

An EVI is a data structure that is stored as two components:

- The symbol table contains statistical and descriptive information about each distinct key value represented in the table. Each distinct key is assigned a unique code, either 1 byte, 2 bytes or 4 bytes in size.

By specifying INCLUDE on the create, additional aggregate values can be maintained in real time as an extension of the key portion of the symbol table entry. These aggregated values are over non-key data in the table grouped by the specified EVI key.

- The vector is an array of codes listed in the same ordinal position as the rows in the table. The vector does not contain any pointers to the actual rows in the table.

Advantages of EVIs:

- Require less storage
- May have better build times than radix, especially if the number of unique values in the columns defined for the key is relatively small.
- Provide more accurate statistics to the query optimizer
- Considerably better performance for certain grouping types of queries
- Good performance characteristics for decision support environments.
- Can be further extended for certain types of grouping queries with the addition of INCLUDE values. Provides ready-made numeric aggregate values maintained in real time as part of index maintenance. INCLUDE values become an extension of the EVI symbol table. Multiple include values can be specified over different aggregating columns and maintained in the same EVI provided the group by values are the same. This technique can reduce overall maintenance.

Disadvantages of EVIs:

- Cannot be used in ordering.
- Use for grouping is specialized. Supports:
 - COUNT, DISTINCT requests over key columns
 - aggregate requests over key columns where all other selection can be applied to the EVI symbol table keys
 - INCLUDE aggregates
 - MIN or MAX, if aggregating value is part of the symbol table key.
- Use with joins always done in cooperation with hash table processing.
- Some additional maintenance idiosyncrasies.

Related reference

[Encoded vector index](#)

An encoded vector index is a permanent object that provides access to a table. This access is done by assigning codes to distinct key values and then representing those values in a vector.

Related information

[SQL Create Index statement](#)

[SQL INCLUDE statement](#)

How the EVI works

EVI works in different ways for costing and implementation.

For costing, the optimizer uses the symbol table to collect metadata information about the query.

For implementation, the optimizer can use the EVI in one of the following ways:

- **Selection (WHERE clause)**

The database engine uses the vector to build a dynamic bitmap or list of selected row ids. The bitmap or list contains 1 bit for each row in the table. The bit is turned on for each selected row. Like a bitmap index, these intermediate dynamic bitmaps (or lists) can be ANDed and ORed together to satisfy a query.

For example, a user wants to see sales data for a specific region and time period. You can define an EVI over the `region` and `quarter` columns of the table. When the query runs, the database engine builds dynamic bitmaps using the two EVIs. The bitmaps are ANDed together to produce a single bitmap containing only the relevant rows for both selection criteria.

This ANDing capability drastically reduces the number of rows that the system must read and test. The dynamic bitmaps exist only as long as the query is executing. Once the query is completed, the dynamic bitmaps are eliminated.

- **Grouping or Distinct**

The symbol table within the EVI contains distinct values for the specified columns in the key definition. The symbol table also contains a count of the number of records in the base table that have each distinct value. Queries involving grouping or distinct, based solely on columns in the key, are candidates for a technique that uses the symbol table directly to determine the query result.

The symbol table contains only the key values and their associated counts, unless `INCLUDE` is specified. Therefore, queries involving column function `COUNT` are eligible for this technique. But queries with column functions `MIN` or `MAX` on other non-key columns are not eligible. `MIN` and `MAX` values are not stored in the symbol table.

- **EVI INCLUDE aggregates**

Including additional aggregate values further extends the ability of the symbol table to provide ready-made results. Aggregate data is grouped by the specified columns in the key definition. Therefore, aggregate data must be over columns in the table other than those columns specified as EVI key values.

For performance, these included aggregates are limited to numeric results (`SUM`, `COUNT`, `AVG`, `VARIANCE`) as they can be maintained directly from the inserted or removed row.

`MIN` or `MAX` values would occasionally require other row comparisons during maintenance and therefore are not supported with the `INCLUDE` keyword.

EVI symbol table only access is used to satisfy distinct or grouping requests when the query is run with commitment control `*NONE` or `*CHG`.

`INCLUDE` for additional aggregate values can be used in join queries. When possible, the existence of EVIs with `INCLUDE` aggregates causes the group by process to be pushed down to each table as necessary. See the following EVI `INCLUDE` grouping push down example: [“EVI INCLUDE aggregate example” on page 80](#)

Related reference

[Encoded vector index index-symbol table only access](#)

The encoded vector index can also be used for index-symbol table only access.

[Encoded vector index symbol table scan](#)

An encoded vector index symbol table scan operation is used to retrieve the entries from the symbol table portion of the index.

[Encoded vector index symbol table probe](#)

An encoded vector index symbol table probe operation is used to retrieve entries from the symbol table portion of the index. Scanning the entire symbol table is not necessary.

Index grouping implementation

There are two primary ways to implement grouping using an index: Ordered grouping and pre-summarized processing.

Related information

SQL INCLUDE statement

When to create EVIs

There are several instances to consider creating EVIs.

Consider creating encoded vector indexes when any one of the following is true:

- You want to gather 'live' statistics
- Full table scan is currently being selected for the query
- Selectivity of the query is 20%-70% and using skip sequential access with dynamic bitmaps speed up the scan
- When a star schema join is expected to be used for star schema join queries.
- When grouping or distinct queries are specified against a column, the columns have few distinct values and only the COUNT column function, if any, is used.
- When ready-made aggregate results grouped by the specified key columns would benefit query performance.

Create encoded vector indexes with:

- Single key columns with a low number of distinct values expected
- Keys columns with a low volatility (do not change often)
- Maximum number of distinct values expected using the WITH n DISTINCT VALUES clause
- Single key over foreign key columns for a star schema model

EVI with INCLUDE vs Materialized Query Tables

Although EVIs with INCLUDE are not a substitute for Materialized Query Tables (MQTs), INCLUDE EVIs have an advantage over single table aggregate MQTs (materialized query tables). The advantage is that the ready-made aggregate results are maintained in real time, not requiring explicit REFRESH TABLE requests. For performance and read access to aggregate results, consider turning your single table, aggregate MQTs into INCLUDE EVIs. Keep in mind that the other characteristics of a good EVI are applicable, such as a relatively low number of distinct key values.

As indexes, these EVIs are found during optimization just as any other indexes are found. Unlike MQTs, there is no INI setting to enable and no second pass through the optimizer to cost the application of this form of ready-made aggregate. In addition, EVIs with INCLUDE can be used to populate MQT summary tables if the EVI is a match for a portion of the MQT definition.

Related reference

Encoded vector index symbol table scan

An encoded vector index symbol table scan operation is used to retrieve the entries from the symbol table portion of the index.

Index grouping implementation

There are two primary ways to implement grouping using an index: Ordered grouping and pre-summarized processing.

Related information


SQL INCLUDE statement

EVI maintenance

There are unique challenges to maintaining EVIs. The following table shows a progression of how EVIs are maintained, the conditions under which EVIs are most effective, and where EVIs are least effective, based on the EVI maintenance characteristics.

Table 55. EVI Maintenance Considerations

	Condition	Characteristics
Most Effective	When inserting an existing distinct key value	<ul style="list-style-type: none"> • Minimum overhead • Symbol table key value looked up and statistics updated • Vector element added for new row, with existing byte code • Minimal additional pathlength to maintain any INCLUDED aggregate values (the increment of a COUNT or adding to an accumulating SUM)
	When inserting a <i>new</i> distinct key value - in order, within byte code range	<ul style="list-style-type: none"> • Minimum overhead • Symbol table key value added, byte code assigned, statistics assigned • Vector element added for new row, with new byte code • Minimal additional pathlength to maintain any INCLUDED aggregate values (the increment of a COUNT or adding to an accumulating SUM)
	When inserting a new distinct key value - out of order, within byte code range	<ul style="list-style-type: none"> • Minimum overhead if contained within overflow area threshold • Symbol table key value added to overflow area, byte code assigned, statistics assigned • Vector element added for new row, with new byte code • Considerable overhead if overflow area threshold reached • Access path validated - not available • EVI refreshed, overflow area keys incorporated, new byte codes assigned (symbol table and vector elements updated) • Minimal additional path-length to maintain any INCLUDED aggregate values (the increment of a COUNT or adding to an accumulating SUM)
	When inserting a new distinct key value - out of byte code range	<ul style="list-style-type: none"> • Considerable overhead • Access plan invalidated - not available • EVI refreshed, next byte code size used, new byte codes assigned (symbol table and vector elements updated) • Not applicable to EVIs with INCLUDE, as by definition the max allowed byte code is used
Least Effective		



Related reference

[Encoded vector index](#)

An encoded vector index is a permanent object that provides access to a table. This access is done by assigning codes to distinct key values and then representing those values in a vector.

Related information

[SQL INCLUDE statement](#)

Recommendations for EVI use

Encoded vector indexes are a powerful tool for providing fast data access in decision support and query reporting environments. To ensure the effective use of EVIs, use the following guidelines.

Create EVIs on

- Read-only tables or tables with a minimum of INSERT, UPDATE, DELETE activity.
- Key columns that are used in the WHERE clause - local selection predicates of SQL requests.
- Single key columns that have a relatively small set of distinct values.
- Multiple key columns that result in a relatively small set of distinct values.
- Key columns that have a static or relatively static set of distinct values.
- Non-unique key columns, with many duplicates.

Create EVIs with the maximum byte code size expected

- Use the "WITH n DISTINCT VALUES" clause on the CREATE ENCODED VECTOR INDEX statement.
- If unsure, use a number greater than 65,535 to create a 4 byte code. This method avoids the EVI maintenance involved in switching byte code sizes.
- EVIs with INCLUDE always create with a 4 byte code.

When loading data

- Drop EVIs, load data, create EVIs.
- EVI byte code size is assigned automatically based on the number of actual distinct key values found in the table.
- Symbol table contains all key values, in order, no keys in overflow area.
- EVIs with INCLUDE always use 4 byte code

Consider adding INCLUDE values to existing EVIs

An EVI index with INCLUDE values can be used to supply ready-made aggregate results. The existing symbol table and vector are still used for table selection, when appropriate, for skip sequential plans over large tables, or for index ANDing and ORing plans. If you already have EVIs, consider creating new ones with additional INCLUDE values, and then drop the pre-existing index.

Consider specifying multiple INCLUDE values on the same EVI create

If you need different aggregates over different table values for the same GROUP BY columns specified as EVI keys, define those aggregates in the same EVI. This definition cuts down on maintenance costs and allows for a single symbol table and vector.

For example:

```
Select SUM(revenue) from sales group by Country
```

```
Select SUM(costOfGoods) from sales group by Country, Region
```

Both queries could benefit from the following EVI:

```
CREATE ENCODED VECTOR INDEX eviCountryRegion on Sales(country,region)
INCLUDE(SUM(revenue), SUM(costOfGoods))
```

The optimizer does additional grouping (regrouping) if the EVI key values are wider than the corresponding GROUP BY request of the query. This additional grouping would be the case in the first example query.

If an aggregate request is specified over null capable results, an implicit COUNT over that same result is included as part of the symbol table entry. The COUNT is used to facilitate index maintenance when a requested aggregate needs to reflect. It can also assist with pushing aggregation through a join if the optimizer determines this push is possible. The COUNT is then used to help compensate for fewer join activity due to the pushed down grouping.

Consider EVI INCLUDE and Grouping Sets

EVI INCLUDE support has been expanded to match GROUPING SETs, ROLLUP and CUBE queries.

When EVI INCLUDES are available over a table being aggregated over in a grouping sets query, the query is rewritten to facilitate and match any EVI INCLUDE indexes that might be available. This can result in exceeding good query performance because the table is never accessed. All the aggregate variations necessary to perform the rollup, cube or grouping set query result can be performed over the EVI symbol table with INCLUDE values.


For example on the ROLLUP query below, the grouping is the sum of quantity rolled up at various levels (month, quarter, year) and ONLY the symbol table of the encoded vector index is accessed in the access plan.

```
SELECT year(shipdate) year_ship, quarter(shipdate) quarter, month(shipdate) month_ship,
sum(quantity)
as totquantity
FROM item_fact
GROUP BY ROLLUP (Year(shipdate), Quarter(shipdate), month(shipdate));
```

Here is the EVI INCLUDE create that will facilitate the rollup query.

```
CREATE ENCODED VECTOR INDEX GS_EVI
ON ITEM_FACT
( YEAR ( SHIPDATE ) ASC , QUARTER ( SHIPDATE ) ASC , MONTH ( SHIPDATE ) ASC )
INCLUDE ( SUM ( QUANTITY ) , COUNT ( * ) );
```

The following graphic shows a Visual Explain that illustrates the access and the performance. Instead of accessing potentially millions of rows, the access is over a rather modest size symbol table.



Consider SMP and parallel index creation and maintenance

Symmetrical Multiprocessing (SMP) is a valuable tool for building and maintaining indexes in parallel. The results of using the optional SMP feature of IBM i are faster index build times, and faster I/O velocities while maintaining indexes in parallel. Using an SMP degree value of either *OPTIMIZE or *MAX, additional multiple tasks and additional system resources are used to build or maintain the indexes. With a degree value of *MAX, expect linear scalability on index creation. For example, creating indexes on a 4-processor system can be four times as fast as a 1-processor system.

Checking values in the overflow area

You can also use the **Display File Description (DSPFD)** command (or System i Navigator - Database) to check how many values are in the overflow area. Once the **DSPFD** command is issued, check the overflow area parameter for details on the initial and actual number of distinct key values in the overflow area.

Using CHGLF to rebuild the access path of an index

Use the **Change Logical File (CHGLF)** command with the attribute Force Rebuild Access Path set to YES (FRCRBDAP(*YES)). This command accomplishes the same thing as dropping and recreating the index, but it does not require that you know about how the index was built. This command is especially effective for applications where the original index definitions are not available, or for refreshing the access path.

Related information

[SQL Create Index statement](#)

[SQL INCLUDE statement](#)

[Change Logical File \(CHGLF\) command](#)

Comparing binary radix indexes and encoded vector indexes

Db2 for IBM i makes indexes a powerful tool.

The following table summarizes some of the differences between binary radix indexes and encoded vector indexes:

Comparison value	Binary Radix Indexes	Encoded Vector Indexes
Basic data structure	A wide, flat tree	A Symbol Table and a vector
Interface for creating	Command, SQL, System i Navigator	SQL, System i Navigator
Can be created in parallel	Yes	Yes
Can be maintained in parallel	Yes	Yes
Used for statistics	Yes	Yes
Used for selection	Yes	Yes, with dynamic bitmaps or RRN list
Used for joining	Yes	Yes (with a hash table)
Used for grouping	Yes	Yes
Used for ordering	Yes	No
Used to enforce unique Referential Integrity constraints	Yes	No
Source for predetermined or ready-made numeric aggregate results	No	Yes, with INCLUDE keyword option on create

Indexes & the optimizer

Since the optimizer uses cost based optimization, more information about the database rows and columns makes for a more efficient access plan created for the query. With the information from the indexes, the optimizer can make better choices about how to process the request (local selection, joins, grouping, and ordering).

The CQE optimizer attempts to examine most, if not all, indexes built over a table unless or until it times out. However, the SQE optimizer only considers those indexes that are returned by the Statistics Manager. These include only indexes that the Statistics Manager decides are useful in performing local selection based on the "where" clause predicates. Consequently, the SQE optimizer does not time out.

The primary goal of the optimizer is to choose an implementation that efficiently eliminates the rows that are not interesting or required to satisfy the request. Normally, query optimization is thought of as trying to find the rows of interest. A proper indexing strategy assists the optimizer and database engine with this task.

Instances where an index is not used

Db2 for i does not use indexes in the certain instances.

- For a column that is expected to be updated; for example, when using SQL, your program might include the following:

```
EXEC SQL
  DECLARE DEPTEMP CURSOR FOR
```

```

SELECT EMPNO, LASTNAME, WORKDEPT
FROM CORPDATA.EMPLOYEE
WHERE (WORKDEPT = 'D11' OR
      WORKDEPT = 'D21') AND
      EMPNO = '000190'
FOR UPDATE OF EMPNO, WORKDEPT
END-EXEC.

```

When using the **OPNQRYF** command, for example:

```

OPNQRYF FILE((CORPDATA/EMPLOYEE)) OPTION(*ALL)
QRYSLT('WORKDEPT *EQ 'D11' *OR WORKDEPT *EQ 'D21')
*AND EMPNO *EQ '000190')

```

Even if you do not intend to update the employee department, the system cannot use an index with a key of *WORKDEPT*.

An index can be used if all the index updatable columns are also used within the query as an isolatable selection predicate with an equal operator. In the previous example, the system uses an index with a key of *EMPNO*.

The system can operate more efficiently if the **FOR UPDATE OF** column list only names the column you intend to update: *WORKDEPT*. Therefore, do not specify a column in the **FOR UPDATE OF** column list unless you intend to update the column.

If you have an updatable cursor because of dynamic SQL, or **FOR UPDATE** was not specified and the program contains an **UPDATE** statement, then all columns can be updated.

- For a column being compared with another column from the same row. For example, when using SQL, your program might include the following:

```

EXEC SQL
DECLARE DEPTDATA CURSOR FOR
SELECT WORKDEPT, DEPTNAME
FROM CORPDATA.EMPLOYEE
WHERE WORKDEPT = ADMRDEPT
END-EXEC.

```

When using the **OPNQRYF** command, for example:

```

OPNQRYF FILE (EMPLOYEE) FORMAT(FORMAT1)
QRYSLT('WORKDEPT *EQ ADMRDEPT')

```

Even though there is an index for *WORKDEPT* and another index for *ADMRDEPT*, Db2 for i does not use either index. The index has no added benefit because every row of the table needs to be looked at.

Display indexes for a table

You can display indexes that are created on a table using System i Navigator.

To display indexes for a table, follow these steps:

1. In the System i Navigator window, expand the system that you want to use.
2. Expand **Databases** and the database that you want to work with.
3. Expand **Schemas** and the schema that you want to work with.
4. Right-click a table and select **Show Indexes**.

The Show index window includes the following columns:

Table 57. Columns used in Show index window	
Column name	Description
Name	The SQL name for the index

Table 57. Columns used in Show index window (continued)

Column name	Description
Type	The type of index displayed. Possible values are: <ul style="list-style-type: none"> • Keyed Physical File • Keyed Logical File • Primary Key Constraint • Unique Key Constraint • Foreign Key Constraint • Index
Schema	Schema or library containing the index or access path
Owner	User ID of the owner of this index or access path
System Name	System table name for the index or access path.
Text	The text description of the index or access path
Index partition	Partition detail for the index. Possible values: <ul style="list-style-type: none"> • <blank>, For all partitions • For Each Partition • specific name of the partition
Valid	Whether the access path or index is valid. The possible values are Yes or No.
Creation Date	The timestamp of when the index was created.
Last Build	The last time that the access path or index was rebuilt.
Last Query Use	Timestamp when the access path was last used by the optimizer.
Last Query Statistics Use	Timestamp when the access path was last used for statistics
Query Use Count	Number of times the access path has been used for a query
Query Statistics Use Count	Number of times the access path has been used for statistics
Last Used Date	Timestamp when the access path or index was last used.
Days Used Count	The number of days the index has been used.
Date Reset Days Used Count	The year and date when the days-used count was last set to 0.
Number of Key Columns	The number of key columns defined for the access path or index.
Key Columns	The key columns defined for the access path or index.
Current Key Values	The number of current key values.
Current Size	The size of the access path or index.
Current Allocated Pages	The current number of pages allocated for the access path or index.

Table 57. Columns used in Show index window (continued)

Column name	Description
Logical Page Size	The number of bytes used for the access path or the logical page size of the index. Indexes with larger logical page sizes are typically more efficient when scanned during query processing. Indexes with smaller logical page sizes are typically more efficient for simple index probes and individual key look ups. If the access path or index is an encoded vector, the value 0 is returned.
Duplicate Key Order	How the access path or index handles duplicate key values. Possible values are: <ul style="list-style-type: none"> • Unique - all values are unique. • Unique where not null - all values are unique unless null is specified.
Maximum Key Length	The maximum key length for the access path or index.
Unique Partial Key Values	The number of unique partial keys for the key fields 1 through 4. If the access path is an encoded vector, this number represents the number of full key distinct values.
Overflow Values	The number of overflow values for this encoded vector index.
Key Code Size	The length of the code assigned to each distinct key value of the encoded vector index.
Sparse	Is the index considered sparse. Sparse indexes only contain keys for rows that satisfy the query. Possible values are: <ul style="list-style-type: none"> • Yes • No
Derived Key	Is the index considered derived. A derived key is a key that is the result of an operation on the base column. Possible values are: <ul style="list-style-type: none"> • Yes • No
Partitioned	Is the index partition created for each data partition defined for the table using the specified columns. Possible values are: <ul style="list-style-type: none"> • Yes • No
Maximum Size	The maximum size of the access path or index.
Sort Sequence	The alternate character sorting sequence for National Language Support (NLS).
Language Identifier	The language code for the object.
Estimated Rebuild Time	The estimated time in seconds required to rebuild the access path or index.

Table 57. Columns used in Show index window (continued)

Column name	Description
Held	Is a rebuild of an access path or index held. Possible values are: <ul style="list-style-type: none"> • Yes • No
Maintenance	For objects with key fields or join logical files, the type of access path maintenance used. The possible values are: <ul style="list-style-type: none"> • Do not wait • Delayed • Rebuild
Delayed Maintenance Keys	The number of delayed maintenance keys for the access path or index.
Recovery	When the access path is rebuilt after damage to the access path is recognized. The possible values are: <ul style="list-style-type: none"> • After IPL • During IPL • Next Open
Index Logical Reads	The number of access path or index logical read operations since the last IPL.
WHERE Clause	Specifies the condition to apply for a row to be included in the index.
WHERE Clause Has UDF	Does the WHERE clause have a UDF. Possible values are: <ul style="list-style-type: none"> • Yes • No
Table	Table name of the table that the index is based on.
Table Partition	Partition name of the table that the index is based on.
Table System Name	System name of the table that the index is based on.
Last Rebuild Number Keys	Number of keys in the index when the index was last rebuilt.
Last Rebuild Parallel Degree	Parallel degree used when the index was last rebuilt.
Last Rebuild Time	Amount of time in seconds it took to rebuild the index the last time the index was rebuilt.
Keep in Memory	Is the index kept in memory. Possible values are: <ul style="list-style-type: none"> • Yes • No
Sort Sequence Schema	Schema of the sort sequence table if one is used.
Sort Sequence Name	Name of the sort sequence table if one is used.
Random Reads	The number of reads that have occurred in a random fashion. Random means that the location of the row or key could not be predicted ahead of time.

Table 57. Columns used in Show index window (continued)

Column name	Description
Media Preference	Indicates preference whether the storage for the table, partition, or index is allocated on Solid State Disk (SSD), if available.

Determine unnecessary indexes

You can easily determine which indexes are being used for query optimization.

Before V5R3, it was difficult to determine unnecessary indexes. Using the Last Used Date was not dependable, as it was only updated when the logical file was opened using a native database application (for example, an RPG application). Furthermore, it was difficult to find all the indexes over a physical file. Indexes are created as part of a keyed physical file, keyed logical file, join logical file, SQL index, primary key or unique constraint, or referential constraint. However, you can now easily find all indexes and retrieve statistics on index usage as a result of System i Navigator and IBM i functionality. To assist you in tuning your performance, this function now produces statistics on index usage as well as index usage in a query.

To access index information through the System i Navigator, navigate to: **Database > Schemas > Tables**. Right-click your table and select **Show Indexes**.

You can show all indexes for a schema by right-clicking on **Tables** or **Indexes** and selecting Show indexes.

Note: You can also view the statistics through the Retrieve Member Description (QUSRMBRD) API.

Certain fields available in the **Show Indexes** window can help you to determine any unnecessary indexes. Those fields are:

Last Query Use	States the timestamp when the index was last used to retrieve data for a query.
Last Query Statistic Use	States the timestamp when the index was last used to provide statistical information.
Query Use Count	Lists the number of instances the index was used in a query.
Query Statistics Use	Lists the number of instances the index was used for statistical information.
Last Used Date	The century and date this index was last used.
Days Used Count	The number of days the index was used. If the index does not have a last used date, the count is 0.
Date Reset Days Used Count	The date that the days used count was last reset. You can reset the days used by Change Object Description (CHGOBJD) command.

The fields start and stop counting based on your situation, or the actions you are currently performing on your system. The following list describes what might affect one or both of your counters:

- The SQE and CQE query engines increment both counters. As a result, the statistics field is updated regardless of which query interface is used.
- A save and restore procedure does not reset the statistics counter if the index is restored over an existing index. If an index is restored that does not exist on the system, the statistics are reset.

Related information

[Retrieve Member Description \(QUSRMBRD\) API](#)

[Change Object Description \(CHGOBJD\) command](#)

Reset usage counts

Resetting the usage counts for a table allows you to determine how the changes you made to your indexing strategy affected the indexes and constraints on that table. For example, if your new strategy

causes an index to never be used, you could then delete that index. Resetting usage counts on a table affect all indexes and constraints that are created on that object.

Note: Resetting usage counts for a keyed physical file or a constraint in the Show Indexes window resets the counts of all constraints and keyed access for that file or table.

You can reset index usage counts by right-clicking a specific index in the Indexes folder or in the Show Indexes dialog and selecting **Reset Usage Counts**.

View index build status

You can view a list of indexes that are being built by the database. This view might be helpful in determining when the index becomes usable to your applications.

To display indexes that are being built, follow these steps:

1. In the System i Navigator window, expand the system that you want to use.
2. Expand **Databases**.
3. Expand the database that you want to work with and then expand the Database Maintenance folder. Select **Index Builds**.

Manage index rebuilds

You can manage the rebuild of your indexes using System i Navigator. You can view a list of access paths that are rebuilding and either hold the access path rebuild or change the priority of a rebuild.

To display access paths to rebuild, follow these steps:

1. In the System i Navigator window, expand the system that you want to use.
2. Expand **Databases**.
3. Expand the database that you want to work with and then expand the **Database Maintenance** folder. Select **Index Rebuilds**.

The access paths to rebuild dialog includes the following columns:

<i>Table 58. Columns used in Index rebuilds window</i>	
Column name	Description
Name	Name of access path being rebuilt.
Schema	Schema name where the index is located.
System Name	The system name of the file that owns the index to be rebuilt.
System Schema	System schema name of access path being rebuilt.
Type	The type of index displayed. Possible values are: Keyed Physical File Keyed Logical File Primary Key Unique Key Foreign Key Index

Table 58. Columns used in Index rebuilds window (continued)

Column name	Description
Status	Displays the status of the rebuild. Possible values are: 1-99 – <i>Rebuild Priority</i> Running – <i>Rebuilding</i> Held – <i>Held from be rebuilt</i>
Rebuild Priority	Displays the priority in which the rebuild for this access path is run. Also referred to as sequence number. Possible values are: 1-99: Order to rebuild Held Open
Rebuild Reason	Displays the reason why this access path needs to be rebuilt. Possible values are: Create or build index IPL Runtime error Change file or index sharing Other Not needed Change End of Data Restore Alter table Change table Change file Reorganize Enable a constraint Alter table recovery Change file recovery Index shared Runtime error Verify constraint Convert member Restore recovery

Table 58. Columns used in Index rebuilds window (continued)

Column name	Description
Rebuild Reason Subtype	<p>Displays the subtype reason why this access path needs to be rebuilt. Possible values are:</p> <ul style="list-style-type: none"> Unexpected error Index in use during failure Unexpected error during update, delete, or insert Delayed maintenance overflow or catch-up error Other No event Change End of Data Delayed maintenance mismatch Logical page size mismatch Partial index restore Index conversion Index not saved and restored Partitioning mismatch Partitioning change Index or key attributes change Original index invalid Index attributes change Force rebuild of index Index not restored Asynchronous rebuilds requested Job ended abnormally Alter table Change constraint Index invalid or attributes change Invalid unique index found Invalid constraint index found Index conversion required <p>If there is no subtype, this field displays 0.</p>

Table 58. Columns used in Index rebuilds window (continued)

Column name	Description
Invalidation Reason	<p>Displays the reason why this access path was invalidated. Possible values are:</p> <ul style="list-style-type: none">User requested (See Invalidation Reason type for more information)Create or build IndexLoad (See Invalidation Reason type for more information)Initial Program Load (IPL)Runtime errorModifyJournal failed to build the indexMarked index as fixable during runtimeMarked index as fixable during IPLChange end of data

Table 58. Columns used in Index rebuilds window (continued)

Column name	Description
Invalidation Reason Type	<p>Displays the reason type for why this access path was invalidation.</p> <p>Possible reason types for User requested:</p> <ul style="list-style-type: none"> Invalid because of REORG It is a copy Alter file Converting new member Change to *FRCRBDAP Change to *UNIQUE Change to *REBLD <p>Possible reason type for LOAD</p> <ul style="list-style-type: none"> The index was marked for invalidation but the system crashed before the invalidation could actually occur The index was associated with the overlaid data space header during a load, therefore it was invalidated Index was in IMPI format. The header was converted and now it is invalidated to be rebuilt in RISC format The RISC index was converted to V5R1 format Index invalidated due to partial load Index invalidated due to a delayed maintenance mismatch Index invalidated due to a pad key mismatch Index invalidated due to a significant fields bitmap fix Index invalidated due to a logical page size mismatch Index was not restored. File might have been saved with ACCPTH(*NO) or index did not exist when file was saved. Index was not restored. File might have been saved with ACCPTH(*NO) or index did not exist when file was saved. Index was rebuilt because file was saved in an inconsistent state with SAVACT(*SYSDFN). <p>For other invalidation codes, this field displays 0.</p>
Estimated Rebuild Time	Estimated amount of time in seconds that it takes to rebuild the index access path.
Rebuild Start Time	Time when the rebuild was started.

<i>Table 58. Columns used in Index rebuilds window (continued)</i>	
Column name	Description
Elapsed Rebuild Time	Amount of time that has elapsed in seconds since the start of the rebuild of the access path.
Unique	Indicates whether the rows in the access path are unique. Possible values are: Yes No
Last Query Use	Timestamp when the access path was last used
Last Query Statistics Use	Timestamp when the access path was last used for statistics
Query Use Count	Number of times the access path has been used for a query
Query Statistics Use Count	Number of times the access path has been used for statistics
Partition	Partition detail for the index. Possible values: <ul style="list-style-type: none"> • <blank>, which means For all partitions • For Each Partition • specific name of the partition
Owner	User ID of the owner of this access path.
Parallel Degree	Number of processors to be used to rebuild the index.
Text	Text description of the file owning the index.

You can also use the **Edit Rebuild of Access Paths (EDTRBDAP)** command to manage rebuilding of access paths.

Related information

[Rebuild access paths](#)

[Edit Rebuild of Access Paths \(EDTRBDAP\) command](#)

Indexing strategy

There are two approaches to index creation: proactive and reactive. Proactive index creation involves anticipating which columns are most often used for selection, joining, grouping, and ordering. Then building indexes over those columns. In the reactive approach, indexes are created based on optimizer feedback, query implementation plan, and system performance measurements.

It is useful to initially build indexes based on the database model and applications and not any particular query. As a starting point, consider designing basic indexes based on the following criteria:

- Primary and foreign key columns based on the database model
- Commonly used local selection columns, including columns that are dependent, such as an automobile's make and model
- Commonly used join columns not considered primary or foreign key columns
- Commonly used grouping columns

Related information

[Indexing and statistics strategies for DB2 for i5/OS](#)

Reactive approach to tuning

To perform reactive tuning, build a prototype of the proposed application without any indexes and start running some queries. Or you could build an initial set of indexes and start running the application to see which ones get used and which do not. Even with a smaller database, the slow running queries become obvious quickly.

The reactive tuning method is also used when trying to understand and tune an existing application that is not performing up to expectations. Use the appropriate debugging and monitoring tools, described in the next section, to view the database feedback messages:

- the indexes the optimizer recommends for local selection
- the temporary indexes used for a query
- the query implementation methods the optimizer chose

If the database engine is building temporary indexes to process joins or perform grouping and selection over permanent tables, build permanent indexes over the same columns. This technique is used to eliminate the temporary index creation. In some cases, a temporary index is built over a temporary table, so a permanent index is not able to be built for those tables. You can use the optimization tools listed in the previous section to note the temporary index creation, the reason it was created, and the key columns.

Proactive approach to tuning

Typically you will create an index for the most selective columns and create statistics for the least selective columns in a query. By creating an index, the optimizer knows that the column is selective and it also gives the optimizer the ability to use that index to implement the query.

In a perfect radix index, the order of the columns is important. In fact, it can make a difference as to whether the optimizer uses it for data retrieval at all. As a general rule, order the columns in an index in the following way:

- Equal predicates first. That is, any predicate that uses the "=" operator may narrow down the range of rows the fastest and should therefore be first in the index.
- If all predicates have an equal operator, then order the columns as follows:
 - Selection predicates + join predicates
 - Join predicates + selection predicates
 - Selection predicates + group by columns
 - Selection predicates + order by columns

In addition to the guidelines above, in general, the most selective key columns should be placed first in the index.

Consider the following SQL statement:

```
SELECT b.col1, b.col2, a.col1
FROM table1 a, table2 b
WHERE b.col1='some_value' AND
      b.col2=some_number AND
      a.join_col=b.join_col
GROUP BY b.col1, b.col2, a.col1
ORDER BY b.col1
```

With a query like this, the proactive index creation process can begin. The basic rules are:

- Custom-build a radix index for the largest or most commonly used queries. Example using the query above:

```
radix index over join column(s) - a.join_col and b.join_col
radix index over most commonly used local selection column(s) - b.col2
```

- For ad hoc online analytical processing (OLAP) environments or less frequently used queries, build single-key EVIs over the local selection column(s) used in the queries. Example using the query above:

```
EVI over non-unique local selection columns - b.col1 and b.col2
```

Coding for effective indexes

The following topics provide suggestions to help you design code which allows Db2 for i to take advantage of available indexes:

Avoid numeric conversions

When a column value and a host variable (or constant value) are being compared, try to specify the same data types and attributes. Db2 for i might not use an index for the named column if the host variable or constant value has a greater precision than the precision of the column. If the two items being compared have different data types, Db2 for i needs to convert one or the other of the values, which can result in inaccuracies (because of limited machine precision).

To avoid problems for columns and constants being compared, use the following:

- same data type
- same scale, if applicable
- same precision, if applicable

For example, EDUCLVL is a halfword integer value (SMALLINT). When using SQL, specify:

```
... WHERE EDUCLVL < 11 AND
      EDUCLVL >= 2
```

instead of

```
... WHERE EDUCLVL < 1.1E1 AND
      EDUCLVL > 1.3
```

When using the OPNQRYF command, specify:

```
... QRYSLT('EDUCLVL *LT 11 *AND EDUCLVL *GE 2')
```

instead of

```
... QRYSLT('EDUCLVL *LT 1.1E1 *AND EDUCLVL *GT 1.3')
```

If an index was created over the EDUCLVL column, then the optimizer might not use the index in the second example. The constant precision is greater than the column precision. It attempts to convert the constant to the precision of the column. In the first example, the optimizer considers using the index, because the precisions are equal.

Avoid arithmetic expressions

Do not use an arithmetic expression as an operand to compare to a column in a row selection predicate. The optimizer does not use an index on a column compared to an arithmetic expression. While this technique might not cause the column index to become unusable, it prevents any estimates and possibly the use of index scan-key positioning. The primary thing that is lost is the ability to use and extract any statistics that might be useful in the optimization of the query.

For example, when using SQL, specify the following:

```
... WHERE SALARY > 16500
```

instead of

```
... WHERE SALARY > 15000*1.1
```

Avoid character string padding

Try to use the same data length when comparing a fixed-length character string column value to a host variable or constant value. Db2 for i might not use an index if the constant value or host variable is longer than the column length.

For example, EMPNO is CHAR(6) and DEPTNO is CHAR(3). For example, when using SQL, specify the following:

```
... WHERE EMPNO > '000300' AND  
      DEPTNO < 'E20'
```

instead of

```
... WHERE EMPNO > '000300 ' AND  
      DEPTNO < 'E20 '
```

When using the OPNQRYF command, specify:

```
... QRYSLT('EMPNO *GT "000300" *AND DEPTNO *LT "E20"')
```

instead of

```
... QRYSLT('EMPNO *GT "000300" *AND DEPTNO *LT "E20"')
```

Avoid the use of LIKE patterns beginning with % or _

The percent (%), and underline (_), used in the pattern of a LIKE (OPNQRYF %WLDCRD) predicate, specify a character string like the row column values to select. They can take advantage of indexes when used to denote characters in the middle or at the end of a character string.

For example, when using SQL, specify the following:

```
... WHERE LASTNAME LIKE 'J%SON%'
```

When using the OPNQRYF command, specify the following:

```
... QRYSLT('LASTNAME *EQ %WLDCRD(''J*SON*'')')
```

However, when used at the beginning of a character string, they can prevent Db2 for i from using any indexes that might be defined on the LASTNAME column to limit the number of rows scanned using index scan-key positioning. Index scan-key selection, however, is allowed. For example, in the following queries index scan-key selection can be used, but index scan-key positioning cannot.

In SQL:

```
... WHERE LASTNAME LIKE '%SON'
```

In OPNQRYF:

```
... QRYSLT('LASTNAME *EQ %WLDCRD(''*SON'')')
```

Avoid patterns with a % so that you can get the best performance with key processing on the predicate. If possible, try to get a partial string to search so that index scan-key positioning can be used.

For example, if you were looking for the name "Smithers", but you only type "S%," this query returns all names starting with "S." Adjust the query to return all names with "Smi%". By forcing the use of partial strings, you might get better performance in the long term.

Using derived indexes

SQL indexes can be created where the key is specified as an expression. This type of key is also referred to as a derived key.

For example, look at the following:

```
CREATE INDEX TOTALIX ON EMPLOYEE(SALARY+BONUS+COMM AS TOTAL)
```

In this example, return all the employees whose total compensation is greater than 50000.

```
SELECT * FROM EMPLOYEE  
WHERE SALARY+BONUS+COMM > 50000  
ORDER BY SALARY+BONUS+COMM
```

Since the optimizer uses the index TOTALIX with index probe to satisfy the WHERE selection and the ordering criteria.

Some special considerations to with derived key index usage and matching include:

- There is no matching for index key constants to query host variables. This non-match includes implicit parameter marker conversion performed by the database manager.

```
CREATE INDEX D_IDX1 ON EMPLOYEE (SALARY/12 AS MONTHLY)
```

In this example, return all employees whose monthly salary is greater than 3000.

long months = 12;

```
EXEC SQL SELECT * FROM EMPLOYEE WHERE SALARY/:months > 3000
```

However, in this case the optimizer does not use the index since there is no support for matching the host variable value months in the query to the constant 12 in the index.

Usage of the QAQQINI option PARAMETER_MARKER_CONVERSION with value *NO can be used to prevent conversion of constants to parameter markers. This technique allows for improved derived index key matching. However, because of the performance implications of using this QAQQINI setting, take care with its usage.

- In general, expressions in the index must match the expression in the query:

```
... WHERE SALARY+COMM+BONUS > 50000
```

In this case, the WHERE SALARY+COMM+BONUS is different from the index key SALARY+BONUS+COMM and would not match.

- It is recommended that the derived index keys be kept as simple as possible. The more complex the query expression to match and the index key expression is, the less likely it is that the index is used.
- The CQE optimizer has limited support for matching derived key indexes.

Related reference

[Derived key index](#)

You can use the SQL CREATE INDEX statement to create a derived key index using an SQL expression.

Related information

[SQL Create Index statement](#)

Using sparse indexes

SQL indexes can be created using WHERE selection predicates. These indexes can also be referred to as sparse indexes. The advantage of a sparse index is that fewer entries are maintained in the index. Only those entries matching the WHERE selection criteria are maintained in the index.

In general, the query WHERE selection must be a subset of the sparse index WHERE selection in order for the sparse index to be used.

Here is a simple example of when a sparse index can be used:

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

It is recommended that the WHERE selection in the sparse index is kept as simple as possible. The more complex the WHERE selection, the more difficult it becomes to match the sparse index WHERE selection to the query WHERE selection. Then it is less likely that the sparse index is used. The CQE optimizer does not support sparse indexes. It does support select/omit logical files however. The SQE optimizer matches the CQE optimizer in its support for select/omit logical files and has nearly full support for sparse indexes.

Related reference

[Sparse indexes](#)

You can use the SQL CREATE INDEX statement to create a sparse index using SQL selection predicates.

Related information

[SQL Create Index statement](#)

Using indexes with sort sequence

The following sections provide useful information about how indexes work with sort sequence tables.

Using indexes and sort sequence with selection, joins, or grouping

Before using an existing index, Db2 for i ensures the attributes of the columns (selection, join, or grouping columns) match the attributes of the key columns in the existing index. The sort sequence table is an additional attribute that must be compared.

The query sort sequence table (specified by the SRTSEQ and LANGID) must match the index sort sequence table. Db2 for i compares the sort sequence tables. If they do not match, the existing index cannot be used.

There is an exception to this rule, however. If the sort sequence table associated with the query is a unique-weight sequence table (including *HEX), Db2 for i acts as though no sort sequence table is specified for selection, join, or grouping columns that use the following operators and predicates:

- equal (=) operator
- not equal (^= or <>) operator
- LIKE predicate (OPNQRYF %WLDCRD and *CT)
- IN predicate (OPNQRYF %VALUES)

When these conditions are true, Db2 for i is free to use any existing index where the key columns match the columns and either:

- The index does not contain a sort sequence table or
- The index contains a unique-weight sort sequence table

Note:

1. The table does not need to match the unique-weight sort sequence table associated with the query.
2. Bitmap processing has a special consideration when multiple indexes are used for a table. If two or more indexes have a common key column referenced in the query selection, then those indexes must either use the same sort sequence table or no sort sequence table.

Using indexes and sort sequence with ordering

Unless the optimizer chooses a sort to satisfy the ordering request, the index sort sequence table must match the query sort sequence table.

When a sort is used, the translation is done during the sort. Since the sort is handling the sort sequence requirement, this technique allows Db2 for i to use any existing index that meets the selection criteria.

Index examples

The following index examples are provided to help you create effective indexes.

For the purposes of the examples, assume that three indexes are created.

Assume that an index HEXIX was created with *HEX as the sort sequence.

```
CREATE INDEX HEXIX ON STAFF (JOB)
```

Assume that an index UNQIX was created with a unique-weight sort sequence.

```
CREATE INDEX UNQIX ON STAFF (JOB)
```

Assume that an index SHRIX was created with a shared-weight sort sequence.

```
CREATE INDEX SHRIX ON STAFF (JOB)
```

Index example: Equal selection with no sort sequence table

Equal selection with no sort sequence table (SRTSEQ(*HEX)).

```
SELECT * FROM STAFF  
WHERE JOB = 'MGR'
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))  
QRYSLT('JOB *EQ ''MGR''')  
SRTSEQ(*HEX)
```

The system can use either index HEXIX or index UNQIX.

Index example: Equal selection with a unique-weight sort sequence table

Equal selection with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT * FROM STAFF  
WHERE JOB = 'MGR'
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))  
QRYSLT('JOB *EQ ''MGR''')  
SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The system can use either index HEXIX or index UNQIX.

Index example: Equal selection with a shared-weight sort sequence table

Equal selection with a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT * FROM STAFF  
WHERE JOB = 'MGR'
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))
  QRYSLT('JOB *EQ ''MGR''')
  SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The system can only use index SHRIX.

Index example: Greater than selection with a unique-weight sort sequence table

Greater than selection with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT * FROM STAFF
WHERE JOB > 'MGR'
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))
  QRYSLT('JOB *GT ''MGR''')
  SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The system can only use index UNQIX.

Index example: Join selection with a unique-weight sort sequence table

Join selection with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT * FROM STAFF S1, STAFF S2
WHERE S1.JOB = S2.JOB
```

or the same query using the JOIN syntax.

```
SELECT *
FROM STAFF S1 INNER JOIN STAFF S2
ON S1.JOB = S2.JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE(STAFF STAFF)
  FORMAT(FORMAT1)
  JFLD((1/JOB 2/JOB *EQ))
  SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The system can use either index HEXIX or index UNQIX for either query.

Index example: Join selection with a shared-weight sort sequence table

Join selection with a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT * FROM STAFF S1, STAFF S2
WHERE S1.JOB = S2.JOB
```

or the same query using the JOIN syntax.

```
SELECT *
FROM STAFF S1 INNER JOIN STAFF S2
ON S1.JOB = S2.JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE(STAFF STAFF) FORMAT(FORMAT1)
  JFLD((1/JOB 2/JOB *EQ))
  SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The system can only use index SHRIX for either query.

Index example: Ordering with no sort sequence table

Ordering with no sort sequence table (SRTSEQ(*HEX)).

```
SELECT * FROM STAFF
WHERE JOB = 'MGR'
ORDER BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))
QRYSLT('JOB *EQ ''MGR''')
KEYFLD(JOB)
SRTSEQ(*HEX)
```

The system can only use index HEXIX.

Index example: Ordering with a unique-weight sort sequence table

Ordering with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT * FROM STAFF
WHERE JOB = 'MGR'
ORDER BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))
QRYSLT('JOB *EQ ''MGR''')
KEYFLD(JOB) SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The system can only use index UNQIX.

Index example: Ordering with a shared-weight sort sequence table

Ordering with a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT * FROM STAFF
WHERE JOB = 'MGR'
ORDER BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))
QRYSLT('JOB *EQ ''MGR''')
KEYFLD(JOB) SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The system can only use index SHRIX.

Index example: Ordering with ALWCPYDTA(*OPTIMIZE) and a unique-weight sort sequence table

Ordering with ALWCPYDTA(*OPTIMIZE) and a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT * FROM STAFF
WHERE JOB = 'MGR'
ORDER BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF))
QRYSLT('JOB *EQ ''MGR''')
KEYFLD(JOB)
```



```
SRTSEQ(*LANGIDUNQ) LANGID(ENU)
ALWCPYDTA(*OPTIMIZE)
```

The system can use either index HEXIX or index UNQIX for selection. Ordering is done during the sort using the *LANGIDUNQ sort sequence table.

Index example: Grouping with no sort sequence table

Grouping with no sort sequence table (SRTSEQ(*HEX)).

```
SELECT JOB FROM STAFF
GROUP BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT2)
GRPFLD((JOB))
SRTSEQ(*HEX)
```

The system can use either index HEXIX or index UNQIX.

Index example: Grouping with a unique-weight sort sequence table

Grouping with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT JOB FROM STAFF
GROUP BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT2)
GRPFLD((JOB))
SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The system can use either index HEXIX or index UNQIX.

Index example: Grouping with a shared-weight sort sequence table

Grouping with a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT JOB FROM STAFF
GROUP BY JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT2)
GRPFLD((JOB))
SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The system can only use index SHRIX.

The following examples assume that three more indexes are created over columns JOB and SALARY. The CREATE INDEX statements precede the examples.

Assume an index HEXIX2 was created with *HEX as the sort sequence.

```
CREATE INDEX HEXIX2 ON STAFF (JOB, SALARY)
```

Assume that an index UNQIX2 was created and the sort sequence is a unique-weight sort sequence.

```
CREATE INDEX UNQIX2 ON STAFF (JOB, SALARY)
```

Assume an index SHRIX2 was created with a shared-weight sort sequence.

```
CREATE INDEX SHRIX2 ON STAFF (JOB, SALARY)
```

Index example: Ordering and grouping on the same columns with a unique-weight sort sequence table

Ordering and grouping on the same columns with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY JOB, SALARY
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPFLD(JOB SALARY)
KEYFLD(JOB SALARY)
SRTSEQ(*LANGIDUNQ) LANGID(ENU)
```

The system can use UNQIX2 to satisfy both the grouping and ordering requirements. If index UNQIX2 did not exist, the system creates an index using a sort sequence table of *LANGIDUNQ.

Index example: Ordering and grouping on the same columns with ALWCPYDTA(*OPTIMIZE) and a unique-weight sort sequence table

Ordering and grouping on the same columns with ALWCPYDTA(*OPTIMIZE) and a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY JOB, SALARY
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPFLD(JOB SALARY)
KEYFLD(JOB SALARY)
SRTSEQ(*LANGIDUNQ) LANGID(ENU)
ALWCPYDTA(*OPTIMIZE)
```

The system can use UNQIX2 to satisfy both the grouping and ordering requirements. If index UNQIX2 did not exist, the system does one of the following actions:

- Create an index using a sort sequence table of *LANGIDUNQ or
- Use index HEXIX2 to satisfy the grouping and to perform a sort to satisfy the ordering

Index example: Ordering and grouping on the same columns with a shared-weight sort sequence table

Ordering and grouping on the same columns with a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY JOB, SALARY
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPFLD(JOB SALARY)
KEYFLD(JOB SALARY)
SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The system can use SHRIX2 to satisfy both the grouping and ordering requirements. If index SHRIX2 did not exist, the system creates an index using a sort sequence table of *LANGIDSHR.

Index example: Ordering and grouping on the same columns with ALWCPYDTA(*OPTIMIZE) and a shared-weight sort sequence table

Ordering and grouping on the same columns with ALWCPYDTA(*OPTIMIZE) and a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY JOB, SALARY
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPFLD(JOB SALARY)
KEYFLD(JOB SALARY)
SRTSEQ(*LANGIDSHR) LANGID(ENU)
ALWCPYDTA(*OPTIMIZE)
```

The system can use SHRIX2 to satisfy both the grouping and ordering requirements. If index SHRIX2 did not exist, the system creates an index using a sort sequence table of *LANGIDSHR.

Index example: Ordering and grouping on different columns with a unique-weight sort sequence table

Ordering and grouping on different columns with a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY SALARY, JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPFLD(JOB SALARY)
KEYFLD(SALARY JOB)
SRTSEQ(*LANGIDSHR) LANGID(ENU)
```

The system can use index HEXIX2 or index UNQIX2 to satisfy the grouping requirements. A temporary result is created containing the grouping results. A temporary index is then built over the temporary result using a *LANGIDUNQ sort sequence table to satisfy the ordering requirements.

Index example: Ordering and grouping on different columns with ALWCPYDTA(*OPTIMIZE) and a unique-weight sort sequence table

Ordering and grouping on different columns with ALWCPYDTA(*OPTIMIZE) and a unique-weight sort sequence table (SRTSEQ(*LANGIDUNQ) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY SALARY, JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPFLD(JOB SALARY)
KEYFLD(SALARY JOB)
SRTSEQ(*LANGIDUNQ) LANGID(ENU)
ALWCPYDTA(*OPTIMIZE)
```

The system can use index HEXIX2 or index UNQIX2 to satisfy the grouping requirements. A sort is performed to satisfy the ordering requirements.

Index example: Ordering and grouping on different columns with ALWCPYDTA(*OPTIMIZE) and a shared-weight sort sequence table

Ordering and grouping on different columns with ALWCPYDTA(*OPTIMIZE) and a shared-weight sort sequence table (SRTSEQ(*LANGIDSHR) LANGID(ENU)).

```
SELECT JOB, SALARY FROM STAFF
GROUP BY JOB, SALARY
ORDER BY SALARY, JOB
```

When using the **OPNQRYF** command, specify:

```
OPNQRYF FILE((STAFF)) FORMAT(FORMAT3)
GRPF LD(JOB SALARY)
KEYFLD(SALARY JOB)
SRTSEQ(*LANGIDSHR) LANGID(ENU)
ALWCPYDTA(*OPTIMIZE)
```

The system can use index SHRIX2 to satisfy the grouping requirements. A sort is performed to satisfy the ordering requirements.

Sparse index examples

This topic shows examples of how the sparse index matching algorithm works.

In example S1, the query selection is a subset of the sparse index selection and consequently an index scan over the sparse index is used. The remaining query selection (COL3=30) is executed following the index scan.

Example S1

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In example S2, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

Example S2

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20
```

In example S3, the query selection exactly matches the sparse index selection and an index scan over the sparse index can be used.

Example S3

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In example S4, the query selection is a subset of the sparse index selection and an index scan over the sparse index can be used. The remaining query selection (COL3=30) is executed following the index scan.

Example S4

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

In example S5, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

Example S5

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20
```

In example S6, the query selection exactly matches the sparse index selection and an index scan over the sparse index can be used. The query selection is executed following the index scan to eliminate excess records from the sparse index.

Example S6

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20 or COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20 or COL3=30
```

In example S7, the query selection is a subset of the sparse index selection and an index scan over the sparse index can be used. The query selection is executed following the index scan to eliminate excess records from the sparse index.

Example S7

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20 or COL3=30
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20
```

In example S8, the query selection is not a subset of the sparse index selection and the sparse index cannot be used.

Example S8

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 or COL2=20
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 or COL2=20 or COL3=30
```

In the next example S9, the constant 'MN' was replaced by a parameter marker for the query selection. The sparse index had the local selection of COL1='MN' applied to it when it was created. The sparse index matching algorithm matches the parameter marker to the constant 'MN' in the query predicate COL1=?. It verifies that the value of the parameter marker is the same as the constant in the sparse index; therefore the sparse index can be used.

Example S9

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
Where Col3='WIN' and (Col1=? or Col2='TWINS')
```

In the next example S10, the keys of the sparse index match the order by fields in the query. For the sparse index to satisfy the specified ordering, the optimizer must verify that the query selection is a subset of the sparse index selection. In this example, the sparse index can be used.

Example S10

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL1, COL3)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
Where Col3='WIN' and (Col1='MN' or Col2='TWINS')
ORDER BY COL1, COL3
```

In the next example S11, the keys of the sparse index do not match the order by fields in the query. But the selection in sparse index T2 is a superset of the query selection. Depending on size, the optimizer might choose an index scan over sparse index T2 and then use a sort to satisfy the specified ordering.

Example S11

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL2, COL4)
WHERE COL1='MN' or COL2='TWINS'
```

```
SELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
Where Col3='WIN' and (Col1='MN' or Col2='TWINS')
ORDER BY COL1, COL3
```

The next example S12 represents the classic optimizer decision: is it better to do an index probe using index IX1 or is it better to do an index scan using sparse index SPR1? Both indexes retrieve the same number of index entries and have the same cost from that point forward. For example, both indexes have the same cost to retrieve the selected records from the dataspace, based on the retrieved entries/keys.

The cost to retrieve the index entries is the deciding criteria. In general, if index IX1 is large then an index scan over sparse index SPR1 has a lower cost to retrieve the index entries. If index IX1 is rather small then an index probe over index IX1 has a lower cost to retrieve the index entries. Another cost decision is reusability. The plan using sparse index SPR1 is not as reusable as the plan using index IX1 because of the static selection built into the sparse selection.

Example S12

```
CREATE INDEX MYLIB/IX1 on MYLIB/T1 (COL1, COL2, COL3)
```

```
CREATE INDEX MYLIB/SPR1 on MYLIB/T1 (COL3)
WHERE COL1=10 and COL2=20 and COL3=30
```

```
CSELECT COL1, COL2, COL3, COL4
FROM MYLIB/T1
WHERE COL1=10 and COL2=20 and COL3=30
```

Application design tips for database performance

There are some design tips that you can apply when designing SQL applications to maximize your database performance.

Using live data

The term *live data* refers to the type of access that the database manager uses when it retrieves data without making a copy of the data. Using this type of access, the data, which is returned to the program, always reflects the current values of the data in the database. The programmer can control whether the database manager uses a copy of the data or retrieves the data directly. This control is done by specifying the allow copy data (ALWCOPYDATA) parameter on the precompiler commands or the **Start SQL (STRSQL)** command.

Specifying ALWCOPYDATA(*NO) instructs the database manager to always use live data. In most cases, forcing live data access is a detriment to performance. It severely limits the possible plan choices that the optimizer could use to implement the query. Avoid it in most cases. However, in specialized cases involving a simple query, live data access can be used as a performance advantage. The cursor does not need to be closed and opened again to refresh the data being retrieved.

An example application demonstrating this advantage is one that produces a list on a display. If the display can show only 20 list elements at a time, then, after the initial 20 elements are displayed, the programmer can request that the next 20 rows be displayed. A typical SQL application designed for an operating system other than the IBM i operating system, might be structured as follows:

```
EXEC SQL
  DECLARE C1 CURSOR FOR
  SELECT EMPNO, LASTNAME, WORKDEPT
  FROM CORPDATA.EMPLOYEE
  ORDER BY EMPNO
END-EXEC.

EXEC SQL
  OPEN C1
END-EXEC.

*   PERFORM FETCH-C1-PARA  20 TIMES.

    MOVE EMPNO to LAST-EMPNO.

EXEC SQL
  CLOSE C1
END-EXEC.

*   Show the display and wait for the user to indicate that
*   the next 20 rows should be displayed.

EXEC SQL
  DECLARE C2 CURSOR FOR
  SELECT EMPNO, LASTNAME, WORKDEPT
  FROM CORPDATA.EMPLOYEE
  WHERE EMPNO > :LAST-EMPNO
  ORDER BY EMPNO
END-EXEC.

EXEC SQL
  OPEN C2
END-EXEC.

*   PERFORM FETCH-C21-PARA  20 TIMES.

*   Show the display with these 20 rows of data.

EXEC SQL
  CLOSE C2
END-EXEC.
```

In the preceding example, notice that an additional cursor had to be opened to continue the list and to get current data. This technique can result in creating an additional ODP that increases the processing time

on the system. In place of the preceding example, the programmer can design the application specifying ALWCPYDTA(*NO) with the following SQL statements:

```
EXEC SQL
  DECLARE C1 CURSOR FOR
  SELECT EMPNO, LASTNAME, WORKDEPT
  FROM CORPDATA.EMPLOYEE
  ORDER BY EMPNO
END-EXEC.

EXEC SQL
  OPEN C1
END-EXEC.

*   Display the screen with these 20 rows of data.
*   PERFORM FETCH-C1-PARA 20 TIMES.

*   Show the display and wait for the user to indicate that
*   the next 20 rows should be displayed.

*   PERFORM FETCH-C1-PARA 20 TIMES.

EXEC SQL
  CLOSE C1
END-EXEC.
```

In the preceding example, the query might perform better if the FOR 20 ROWS clause was used on the multiple-row FETCH statement. Then, the 20 rows are retrieved in one operation.

Related information

[Start SQL Interactive Session \(STRSQL\) command](#)

Reducing the number of open operations

The SQL data manipulation language statements must do database open operations in order to create an open data path (ODP) to the data. An open data path is the path through which all input/output operations for the table are performed. In a sense, it connects the SQL application to a table. The number of open operations in a program can significantly affect performance.

A database open operation occurs on:

- An OPEN statement
- SELECT INTO statement
- An INSERT statement with a VALUES clause
- An UPDATE statement with a WHERE condition
- An UPDATE statement with a WHERE CURRENT OF cursor and SET clauses that refer to operators or functions
- SET statement that contains an expression
- VALUES INTO statement that contains an expression
- A DELETE statement with a WHERE condition

An INSERT statement with a select-statement requires two open operations. Certain forms of subqueries could also require one open per subselect.

To minimize the number of opens, Db2 for i leaves the open data path (ODP) open and reuses the ODP if the statement is run again, unless:

- The ODP used a host variable to build a subset temporary index. The optimizer could choose to build a temporary index with entries for only the rows that match the row selection specified in the SQL statement. If a host variable was used in the row selection, the temporary index does not have the entries required for a different host variable value.
- Ordering was specified on a host variable value.

- An **Override Database File (OVRDBF)** or **Delete Override (DLTOVR)** CL command has been issued since the ODP was opened, which affects the SQL statement execution.

Note: Only overrides that affect the name of the table being referred to causes the ODP to be closed within a given program invocation.

- The join is a complex join that requires temporary objects to contain the intermediate steps of the join.
- Some cases involve a complex sort, where a temporary file is required, might not be reusable.
- A change to the library list since the last open has occurred, which changes the table selected by an unqualified referral in system naming mode.
- The join was implemented by the CQE optimizer using hash join.

For embedded static SQL, Db2 for i only reuses ODPs opened by the same statement. An identical statement coded later in the program does not reuse an ODP from any other statement. If the identical statement must be run in the program many times, code it once in a subroutine and call the subroutine to run the statement.

The ODPs opened by Db2 for i are closed when any of the following occurs:

- a CLOSE, INSERT, UPDATE, DELETE, or SELECT INTO statement completes and the ODP required a temporary result that was not reusable or a subset temporary index.
- the **Reclaim Resources (RCLRSC)** command is issued. A **Reclaim Resources (RCLRSC)** is issued when the first COBOL program on the call stack ends or when a COBOL program issues the STOP RUN COBOL statement. **Reclaim Resources (RCLRSC)** does not close the ODPs created for programs precompiled using CLOSQCSR(*ENDJOB). For interaction of **Reclaim Resources (RCLRSC)** with non-default activation groups, see the following books:
 - WebSphere® Development Studio: ILE C/C++ Programmer's Guide
 - WebSphere Development Studio: ILE COBOL Programmer's Guide
 - WebSphere Development Studio: ILE RPG Programmer's Guide
- the last program containing SQL statements on the call stack exits. Exception is for ODPs created for programs precompiled using CLOSQCSR(*ENDJOB) or modules precompiled using CLOSQCSR(*ENDACTGRP).
- a CONNECT (Type 1) statement changes the application server for an activation group, all ODPs created for the activation group are closed.
- a DISCONNECT statement ends a connection to the application server, all ODPs for that application server are closed.
- a released connection is ended by a successful COMMIT, all ODPs for that application server are closed.
- the threshold for open cursors specified by the query options file (QAQQINI) parameter OPEN_CURSOR_THRESHOLD is reached.
- the SQL LOCK TABLE or CL ALCOBJ OBJ((filename *FILE *EXCL)) CONFLICT(*RQSRLS) command closes any pseudo-closed cursors associated with the specified table.
- an application has requested a close, but the data path was left open. The ODP can be forced closed for a specific file by using the ALCOBJ CL command. This close does not force the ODP to close if the application has not requested that the cursor be closed. The syntax for the command is: ALCOBJ OBJ((library/file *FILE *EXCL)) CONFLICT(*RQSRLS).
- an MQT plan expired based on the timestamp.
- an incompatible commitment control change occurred.
- the table size changed beyond tolerance. The optimizer needs to reoptimize based on the new table size.
- a new index or indexes were created. The optimizer can cost a plan created with the new indexes and compare its cost to the previous plan.
- new statistics were created. The optimizer can take advantage of these new statistics to create a more efficient plan.

- host variables are incompatible with a non-reusable MTI, an MQT, or a sparse index used to implement the query.
- data is warm (in memory).
- the OPTIMIZATION_GOAL *All IO or *First IO specified in query options file QAQQINI was changed.
- a hard close was forced.

The optimizer does not recognize that query selectivity has changed due to host variable changes. It continues to use the existing open and access plan. Change of selectivity due to host variables is only evaluated at full open time unless the PSEUDO_OPEN_CHECK_HOST_VARS qaqqini option is altered.

You can control whether the system keeps the ODPs open in the following ways:

- Design the application so a program that issues an SQL statement is always on the call stack
- Use the CLOSQLCSR(*ENDJOB) or CLOSQLCSR(*ENDACTGRP) parameter
- By specifying the OPEN_CURSOR_THRESHOLD and OPEN_CURSOR_CLOSE_COUNT parameters of the query options file (QAQQINI)

You can control whether the optimizer factors in host variable selectivity once in pseudo mode for queries with host variable that have considerable selectivity variability.

- By specifying the PSEUDO_OPEN_CHECK_HOST_VARS parameter of the query options file (QAQQINI)

An open operation occurs for the first execution of each UPDATE WHERE CURRENT OF, when any SET clause expression contains an operator or function. The open can be avoided by coding the function or operation in the host language code.

For example, the following UPDATE causes the system to do an open operation:

```
EXEC SQL
  FETCH EMP1 INTO :SALARY
END-EXEC.

EXEC SQL
  UPDATE CORPDATA.EMPLOYEE
  SET SALARY = :SALARY + 1000
  WHERE CURRENT OF EMP1
END-EXEC.
```

Instead, use the following coding technique to avoid opens:

```
EXEC SQL
  FETCH EMP1 INTO :SALARY
END EXEC.

ADD 1000 TO SALARY.

EXEC SQL
  UPDATE CORPDATA.EMPLOYEE
  SET SALARY = :SALARY
  WHERE CURRENT OF EMP1
END-EXEC.
```

You can determine whether SQL statements result in full opens in several ways. The preferred methods are to use the Database Monitor or by looking at the messages issued while debug is active. You can also use the CL commands **Trace Job (TRCJOB)** or **Display Journal (DSPJRN)**.

Related information

[Reclaim Resources \(RCLRSC\) command](#)

[Trace Job \(TRCJOB\) command](#)

[Display Journal \(DSPJRN\) command](#)

[RPG](#)

[COBOL](#)

[C and C++](#)

Retaining cursor positions

You can improve performance by retaining cursor positions.

Retaining cursor positions for non-ILE program calls

For non-ILE program calls, the close SQL cursor (CLOSQLCSR) parameter allows you to specify the scope of the following:

- The cursors
- The prepared statements
- The locks

When used properly, the CLOSQLCSR parameter can reduce the number of SQL OPEN, PREPARE, and LOCK statements needed. It can also simplify applications by allowing you to retain cursor positions across program calls.

- *ENDPGM** The default for all non-ILE precompilers. With this option, a cursor remains open and accessible only while the program that opened it is on the call stack. When the program ends, the SQL cursor can no longer be used. Prepared statements are also lost when the program ends. Locks, however, remain until the last SQL program on the call stack has completed.
- *ENDSQL** SQL cursors and prepared statements that are created by a program remain open until the last SQL program on the call stack has completed. They cannot be used by other programs, only by a different call to the same program. Locks remain until the last SQL program in the call stack completes.
- *ENDJOB** This option allows you to keep SQL cursors, prepared statements, and locks active for the duration of the job. When the last SQL program on the stack has completed, any SQL resources created by *ENDJOB programs are still active. The locks remain in effect. The SQL cursors that were not explicitly closed by the CLOSE, COMMIT, or ROLLBACK statements remain open. The prepared statements are still usable on subsequent calls to the same program.

Related reference

[Effects of precompile options on database performance](#)

Several precompile options are available for creating SQL programs with improved performance. They are only options because using them could impact the function of the application. For this reason, the default value for these parameters is the value that ensures successful migration of applications from prior releases. However, you can improve performance by specifying other options.

Retaining cursor positions across ILE program calls

For ILE program calls, the close SQL cursor (CLOSQLCSR) parameter allows you to specify the scope of the following:

- The cursors
- The prepared statements
- The locks

When used properly, the CLOSQLCSR parameter can reduce the number of SQL OPEN, PREPARE, and LOCK statements needed. It can also simplify applications by allowing you to retain cursor positions across program calls.

- *ENDACTGRP** The default for the ILE precompilers. With this option, SQL cursors and prepared statements remain open until the activation group that the program is running under ends. They cannot be used by other programs, only by a different call to the same program. Locks remain until the activation group ends.

***ENDMOD** With this option, a cursor remains open and accessible only while the module that opened it is active. When the module ends, the SQL cursor can no longer be used. Prepared statements are also lost when the module ends. Locks, however, remain until the last SQL program in the call stack completes.

General rules for retaining cursor positions for all program calls

Programs compiled with either CLOSQLCSR(*ENDPGM) or CLOSQLCSR(*ENDMOD) must open a cursor every time the program or module is called, in order to access the data. If the SQL program or module is called several times, and you want to take advantage of a reusable ODP, then the cursor must be explicitly closed before the program or module exits.

Using the CLOSQLCSR parameter and specifying *ENDSQL, *ENDJOB, or *ENDACTGRP, you might not need to run an OPEN and a CLOSE statement on every call. In addition to having fewer statements to run, you can maintain the cursor position between calls to the program or module.

The following examples of SQL statements help demonstrate the advantage of using the CLOSQLCSR parameter:

```
EXEC SQL
  DECLARE DEPTDATA CURSOR FOR
  SELECT EMPNO, LASTNAME
  FROM CORPDATA.EMPLOYEE
  WHERE WORKDEPT = :DEPTNUM
END-EXEC.

EXEC SQL
  OPEN DEPTDATA
END-EXEC.

EXEC SQL
  FETCH DEPTDATA INTO :EMPNUM, :LNAME
END-EXEC.

EXEC SQL
  CLOSE DEPTDATA
END-EXEC.
```

If this program is called several times from another SQL program, it is able to use a reusable ODP. This technique means that, as long as SQL remains active between the calls to this program, the OPEN statement does not require a database open operation. However, the cursor is still positioned to the first result row after each OPEN statement, and the FETCH statement will always return the first row.

In the following example, the CLOSE statement has been removed:

```
EXEC SQL
  DECLARE DEPTDATA CURSOR FOR
  SELECT EMPNO, LASTNAME
  FROM CORPDATA.EMPLOYEE
  WHERE WORKDEPT = :DEPTNUM
END-EXEC.

  IF CURSOR-CLOSED IS = TRUE THEN
EXEC SQL
  OPEN DEPTDATA
END-EXEC.

EXEC SQL
  FETCH DEPTDATA INTO :EMPNUM, :LNAME
END-EXEC.
```

If this program is precompiled with the *ENDJOB or *ENDACTGRP option and the activation group remains active, the cursor position is maintained. The cursor position is also maintained when the following occurs:

- The program is precompiled with the *ENDSQL option.
- SQL remains active between program calls.

The result of this strategy is that each call to the program retrieves the next row in the cursor. On subsequent data requests, the OPEN statement is unnecessary and, in fact, fails with a -502 SQLCODE. You can ignore the error, or add code to skip the OPEN. Use a FETCH statement first, and then run the OPEN statement only if the FETCH operation failed.

This technique also applies to prepared statements. A program can first try the EXECUTE, and if it fails, perform the PREPARE. The result is that the PREPARE is only needed on the first call to the program, assuming that the correct CLOSQLCSR option was chosen. If the statement can change between calls to the program, perform the PREPARE in all cases.

The main program might also control cursors by sending a special parameter on the first call only. This special parameter value indicates that because it is the first call, the subprogram performs the OPENS, PREPAREs, and LOCKs.

Note: If you are using COBOL programs, do not use the STOP RUN statement. When the first COBOL program on the call stack ends or a STOP RUN statement runs, a reclaim resource (RCLRSC) operation is done. This operation closes the SQL cursor. The *ENDSQL option does not work as you wanted.

Programming techniques for database performance

By changing the coding of your queries, you can improve their performance.

Use the OPTIMIZE clause

If an application is not going to retrieve the entire result table for a cursor, using the OPTIMIZE clause can improve performance. The query optimizer modifies the cost estimates to retrieve the subset of rows using the value specified on the OPTIMIZE clause.

Assume that the following query returns 1000 rows:

```
EXEC SQL
  DECLARE C1 CURSOR FOR
  SELECT EMPNO, LASTNAME, WORKDEPT
  FROM CORPDATA.EMPLOYEE
  WHERE WORKDEPT = 'A00'
  ORDER BY LASTNAME
  OPTIMIZE FOR 100 ROWS
END EXEC.
```

Note: The values that can be used for the preceding OPTIMIZE clause are 1–9999999 or ALL.

The optimizer calculates the following costs.

The optimize ratio = optimize for n rows value / estimated number of rows in answer set.

Cost using a temporarily created index:

```
Cost to retrieve answer set rows
+ Cost to create the index
+ Cost to retrieve the rows again
  with a temporary index          * optimize ratio
```

Cost using a SORT:

```
Cost to retrieve answer set rows
+ Cost for SORT input processing
+ Cost for SORT output processing * optimize ratio
```

Cost using an existing index:

```
Cost to retrieve answer set rows
using an existing index          * optimize ratio
```

In the previous examples, the estimated cost to sort or to create an index is not adjusted by the optimize ratio. This method allows the optimizer to balance the optimization and preprocessing requirements.

If the optimize number is larger than the number of rows in the result table, no adjustments are made to the cost estimates.

If the OPTIMIZE clause is not specified for a query, a default value is used based on the statement type, value of ALWCPYDTA, or output device.

<i>Table 59. OPTIMIZE FOR n ROWS default value</i>		
Statement Type	ALWCPYDTA(*OPTIMIZE)	ALWCPYDTA(*YES or *NO)
DECLARE CURSOR	The number or rows in the result table.	30 rows or the number of rows in the result table.
Embedded Select	2	2
INTERACTIVE Select output to display	30 rows or the number of rows in the result table.	30 rows or the number of rows in the result table.
INTERACTIVE Select output to printer or database table	The number of rows in the result table.	The number of rows in the result table.

The OPTIMIZE clause influences the optimization of a query:

- To use an existing index (by specifying a small number).
- To enable the creation of an index, or run a sort or hash by specifying many possible rows in the answer set.

Related information

[select-statement](#)

Use FETCH FOR n ROWS

Applications that perform many FETCH statements in succession could be improved by using FETCH FOR n ROWS. With this clause, you can retrieve multiple rows of table data with a single FETCH, putting them into a host structure array or row storage area.

An SQL application that uses a FETCH statement without the FOR n ROWS clause can be improved by using the multiple-row FETCH statement to retrieve multiple rows. After the host structure array or row storage area is filled by the FETCH, the application loops through the data, processing each of the individual rows. The statement runs faster because the SQL run-time was called only once and all the data was simultaneously returned to the application program.

You can change the application program to allow the database manager to block the rows that the SQL run-time retrieves from the tables.

In the following table, the program attempted to FETCH 100 rows into the application. Note the differences in the table for the number of calls to SQL runtime and the database manager when blocking can be performed.

<i>Table 60. Number of Calls Using a FETCH Statement</i>		
	Database Manager Not Using Blocking	Database Manager Using Blocking
Single-Row FETCH Statement	100 SQL calls 100 database calls	100 SQL calls one database call
Multiple-Row FETCH Statement	one SQL runtime call 100 database calls	one SQL runtime call one database call

Related information

[FETCH statement](#)

Improve SQL blocking performance when using FETCH FOR n ROWS

Use these performance techniques to improve SQL blocking performance when using FETCH FOR n ROWS.

You can improve SQL blocking performance with the following:

- Match the attribute information in the host structure array or the descriptor associated with the row storage area with the attributes of the columns retrieved.
- Retrieve as many rows as possible with a single multiple-row FETCH call. The blocking factor for a multiple-row FETCH request is not controlled by the system page sizes or the SEQONLY parameter on the OVRDBF command. It is controlled by the number of rows that are requested on the multiple-row FETCH request.
- Do not mix single- and multiple-row FETCH requests against the same cursor within a program. If one FETCH against a cursor is treated as a multiple-row FETCH, all fetches against that cursor are treated as multiple-row fetches. In that case, each of the single-row FETCH requests is treated as a multiple-row FETCH of one row.
- Do not use the PRIOR, CURRENT, and RELATIVE scroll options with multiple-row FETCH statements. To allow random movement of the cursor by the application, the database manager must maintain the same cursor position as the application. Therefore, the SQL run-time treats all FETCH requests against a scrollable cursor with these options specified as multiple-row FETCH requests.

Use INSERT n ROWS

Applications that perform many INSERT statements in succession could be improved by using INSERT n ROWS. With this clause, you can insert one or more rows of data from a host structure array into a target table. This array must be an array of structures where the elements of the structure correspond to columns in the target table.

An SQL application that loops over an INSERT...VALUES statement (without the n ROWS clause) can be improved by using the INSERT n ROWS statement to insert multiple rows into the table. After the application has looped to fill the host array with rows, a single INSERT n ROWS statement inserts the entire array into the table. The statement runs faster because the SQL runtime was only called once and all the data was simultaneously inserted into the target table.

In the following table, the program attempted to INSERT 100 rows into a table. Note the differences in the number of calls to SQL runtime and to the database manager when blocking can be performed.

Table 61. Number of Calls Using an INSERT Statement

	Database Manager Not Using Blocking	Database Manager Using Blocking
Single-Row INSERT Statement	100 SQL runtime calls 100 database calls	100 SQL runtime calls one database call
Multiple-Row INSERT Statement	1 SQL runtime call 100 database calls	1 SQL runtime call 1 database call

Related information

[INSERT statement](#)

Control database manager blocking

To improve performance, the SQL runtime attempts to retrieve and insert rows from the database manager a block at a time whenever possible.

You can control blocking, if you want. Use the SEQONLY parameter on the CL command **Override Database File (OVRDBF)** before calling the application program that contains the SQL statements. You can also specify the ALWBLK parameter on the CRTSQLxxx commands or use the QSY2.OVERRIDE_TABLE application service.

The database manager does not allow blocking in the following situations:

- The cursor is update or delete capable.
- The length of the row plus the feedback information is greater than 32767. The minimum size for the feedback information is 11 bytes. The feedback size is increased by the number of bytes in the index key columns used by the cursor, and the number of key columns, if any, that are null capable.
- COMMIT(*CS) is specified, and ALWBLK(*ALLREAD) is not specified.
- COMMIT(*ALL) is specified, and the following are true:
 - A SELECT INTO statement or a blocked FETCH statement is not used
 - The query does not use column functions or specify group by columns.
 - A temporary result table does not need to be created.
- COMMIT(*CHG) is specified, and ALWBLK(*ALLREAD) is not specified.
- The cursor contains at least one subquery and the outermost subselect provided a correlated reference for a subquery, or the outermost subselect processed a subquery with an IN, = ANY, or < > ALL subquery predicate operator, which is treated as a correlated reference, and that subquery is not isolatable.

The SQL runtime automatically blocks rows with the database manager in the following cases:

- INSERT

If an INSERT statement contains a select-statement, inserted rows are blocked and not inserted into the target table until the block is full. The SQL runtime automatically does blocking for blocked inserts.

Note: If an INSERT with VALUES is specified, the SQL runtime might not close the internal cursor used to perform the inserts until the program ends. If the same INSERT statement is run again, a full open is not necessary and the application runs much faster.

- OPEN

Blocking is done under the OPEN statement when the rows are retrieved if all the following conditions are true:

- The cursor is only used for FETCH statements.
- No EXECUTE or EXECUTE IMMEDIATE statements are in the program, or ALWBLK(*ALLREAD) was specified, or the cursor is declared with the FOR FETCH ONLY clause.
- COMMIT(*CHG) and ALWBLK(*ALLREAD) are specified, COMMIT(*CS) and ALWBLK(*ALLREAD) are specified, or COMMIT(*NONE) is specified.

Related reference

[OVERRIDE_TABLE procedure](#)

The OVERRIDE_TABLE procedure sets the blocking size for a table.

[Effects of precompile options on database performance](#)

Several precompile options are available for creating SQL programs with improved performance. They are only options because using them could impact the function of the application. For this reason, the default value for these parameters is the value that ensures successful migration of applications from prior releases. However, you can improve performance by specifying other options.

Related information

[Override Database File \(OVRDBF\) command](#)

Optimize the number of columns that are selected with SELECT statements

For each column in the SELECT statement, the database manager retrieves the data from the underlying table and maps it to a host variable in the application program. By minimizing the number of columns that are specified, processing unit resource usage can be conserved.

Even though it is convenient to code SELECT *, it is far better to explicitly code the columns that are required for the application. This technique is especially important for index-only access, or if all the columns participate in a sort operation (as in SELECT DISTINCT and SELECT UNION).

This technique is also important when considering index only access. You minimize the number of columns in a query and increase the odds that an index can be used to completely satisfy the data request.

Related information

[select-statement](#)

Eliminate redundant validation with SQL PREPARE statements

The processing which occurs when an SQL PREPARE statement is run is like the processing which occurs during precompile processing.

The following processing occurs for the statement that is being prepared:

- The syntax is checked.
- The statement is validated to ensure that the usage of objects is valid.
- An access plan is built.

Again when the statement is executed or opened, the database manager revalidates that the access plan is still valid. Much of this open processing validation is redundant with the validation which occurred during the PREPARE processing. The DLYPRP(*YES) parameter specifies whether PREPARE statements in this program completely validates the dynamic statement. The validation is completed when the dynamic statement is opened or executed. This parameter can provide a significant performance enhancement for programs which use the PREPARE SQL statement because it eliminates redundant validation. Programs that specify this precompile option must check the SQLCODE and SQLSTATE after running the OPEN or EXECUTE statement to ensure that the statement is valid. DLYPRP(*YES) does not provide any performance improvement if the INTO clause is used on the PREPARE statement, or if a DESCRIBE statement uses the dynamic statement before an OPEN is issued for the statement.

Related reference

[Effects of precompile options on database performance](#)

Several precompile options are available for creating SQL programs with improved performance. They are only options because using them could impact the function of the application. For this reason, the default value for these parameters is the value that ensures successful migration of applications from prior releases. However, you can improve performance by specifying other options.

Related information

[Prepare statement](#)

Page interactively displayed data with REFRESH(*FORWARD)

In large tables, paging performance is typically degraded because of the REFRESH(*ALWAYS) parameter on the **Start SQL (STRSQL)** command. STRSQL dynamically retrieves the latest data directly from the table. Paging performance can be improved by specifying REFRESH(*FORWARD).

When interactively displaying data using REFRESH(*FORWARD), the results of a select-statement are copied to a temporary table as you page forward through the display. Other users sharing the table can change the rows while you are displaying the select-statement results. If you page backward or forward to rows that have already been displayed, the rows shown are in the temporary table instead of the updated table.

The refresh option can be changed on the Session Services display.

Related information

[Start SQL \(STRSQL\) command](#)

Improve concurrency by avoiding lock waits

The concurrent access resolution option directs the database manager on how to handle cases of record lock conflicts under certain isolation levels.

The concurrent access resolution, when applicable, can have one of the following values:

- **Wait for outcome** (default). This value directs the database manager to wait for the commit or rollback when encountering locked data in the process of being updated or deleted. Locked rows that are in the process of being inserted are not skipped. This option does not apply for read-only queries running under isolation level None or Uncommitted Read.
- **Use currently committed**. This value allows the database manager to use the currently committed version of the data for read-only queries when encountering locked data in the process of being updated or deleted. Locked rows in the process of being inserted can be skipped. This option applies if possible when the isolation level in effect is Cursor Stability and is ignored otherwise.
- **Skip locked data**. This value directs the database manager to skip rows in the case of record lock conflicts. This option is applicable only when the query is running under an isolation level of Cursor Stability or Read Stability and additionally for UPDATE and DELETE queries when the isolation level is None or Uncommitted Read.

The concurrent access resolution values of USE CURRENTLY COMMITTED and SKIP LOCKED DATA can be used to improve concurrency by avoiding lock waits. However, care must be used when using these options because they might affect application functionality. For more information on the USE CURRENTLY COMMITTED option, see [Concurrency](#).

WAIT FOR OUTCOME, USE CURRENTLY COMMITTED, and SKIP LOCKED DATA can be specified as the concurrent-access-resolution-clause in the attribute-string of a PREPARE statement.

Additionally, they can be specified as the concurrent-access-resolution-clause at the statement level on a select-statement, SELECT INTO, searched UPDATE, or searched DELETE statement.

Concurrent access resolution is also specifiable as a precompiler option by using the CONACC parameter on the CRTSQLxxx. The CONACC parameter accepts one of the following values:

- *DFT - specifies that the concurrent access option is not explicitly set for this program. The value that is in effect when the program is invoked is used. The value can be set using the SQL_CONCURRENT_ACCESS_RESOLUTION option in the query options file QAQQINI.
- *CURCMT - use currently committed.
- *WAIT - wait for outcome.

These same options can be set on the RUNSQLSTM and RUNSQL CL commands and by using the SET OPTION SQL statement. Concurrent access resolution can be specified for SQL triggers, functions, and procedures by using the SET OPTION statement.

When the concurrent access resolution option is not directly set by the application, it is set to the value of the SQL_CONCURRENT_ACCESS_RESOLUTION option in the query options file QAQQINI. This option accepts one of the following values:

- *DEFAULT - the default value is set to *WAIT.
- *CURCMT - use currently committed.
- *WAIT - wait for outcome.

Related reference

[QAQQINI query options](#)

There are different options available for parameters in the QAQQINI file.

Related information

[concurrent-access-resolution-clause](#)

Use SELECTIVITY to supply missing information

Some query predicates are inherently difficult for the optimizer to analyze, yet effective optimization depends on accurate information about the data being selected. When other statistics are unavailable, careful use of the SELECTIVITY clause can give the optimizer the information it needs to produce the best access plan for a query.

The optimizer uses a variety of techniques to evaluate and estimate the number of rows a query predicate will select. This information may come from key range estimates, from column statistics, or from inferences based on the cardinality of the data. Sometimes, however, the optimizer has no way to determine how a predicate will apply to the underlying data. In these cases, the SELECTIVITY clause can be added to difficult predicates to provide the extra information that the optimizer needs.

Using SELECTIVITY to correct inaccurate estimates

One common type of incomplete information involves data that is transformed before it is used in a predicate. The transformation could be performed by a user defined function or it could be performed by a built in function. In either case, something about the function prevents the optimizer from making inferences about how the transformation will affect the data.

Consider the following query:

```
SELECT EXCHANGE, S.SYMBOL, NAME, CURRENT_POSITION
FROM SECURITIES S INNER JOIN POSITIONS P ON S.SYMBOL = P.SYMBOL
WHERE RISK_SCORE(EXCHANGE, S.SYMBOL) > .8
      AND EXCHANGE = 'NYSE'
      AND CURRENT_POSITION > 10000
ORDER BY EXCHANGE, S.SYMBOL
```

Because the optimizer does not know the internal logic of the RISK_SCORE function, it uses a default selectivity value. The optimizer will generally assume 33% of the rows satisfy the predicate `RISK_SCORE(EXCHANGE, S.SYMBOL) > .8`. If this assumption is not accurate, it could cause the optimizer to select an inappropriate query plan. To mitigate this, a user who understands the behavior of the RISK_SCORE function and knows that only 1% of the rows in the table will satisfy the predicate can rewrite the WHERE clause as:

```
WHERE RISK_SCORE(EXCHANGE, S.SYMBOL) > .8 SELECTIVITY .01
      AND EXCHANGE = 'NYSE'
      AND CURRENT_POSITION > 10000
```

Another common cause of missing information is when the query selects a small number of rows that have been added to the table very recently. These rows may not yet be reflected in the statistical information automatically maintained by the statistics manager, leading the optimizer to believe that the query selects no rows. As in the above example, the solution is to add a SELECTIVITY clause that accurately represents the percentage of rows selected by the predicate.

To determine when the use of SELECTIVITY could be helpful for a query, use Visual Explain to identify operations with significantly inaccurate *Percent Selectivity* attributes under the *Estimated rows selected and query join info* heading. In many cases, you may have a reasonable intuition for the correct selectivity for the operation. In cases where you need additional information, either **Run and Explain** or use Visual Explain on a plan that is cached in the plan cache, since these options provide you with information about the actual number of rows processed by each operation. These values are given as *Actual Rows Selected Per Plan Step Iteration* and can be compared to the estimated *Rows Selected Per Plan Step Iteration*. Keep in mind that a given operation could process multiple predicates that have been logically combined. Since SELECTIVITY applies to an individual predicate, the effect of SELECTIVITY on the overall *Percent Selectivity* of the operation can be difficult to predict. In general, the effect of adding SELECTIVITY is easiest to predict and understand when the operation has only a single simple predicate (for example, a UDF.)

Using SELECTIVITY to reduce plan volatility

The optimizer uses a variety of factors when re-validating whether a cached access plan should be re-used. One key factor is the selectivity of the query's predicates. The same predicate with different host variable values may select a very different number of rows and thus run best with a different access plan. For this reason, the optimizer will not re-use a cached plan that it estimates will select a significantly different number of rows. Instead, it will re-optimize the query with the new values.

Under some circumstances, you may prefer to use a single plan for all host variable values rather than relying on the dynamic behavior of the optimizer. If the plan is changing due to the selectivity of a host variable value changing, the SELECTIVITY clause can help to lock in a preferred plan.

For example, consider the following query where the plan changes depending on the given value (for example, 1000 rows vs 1,000,000 rows are returned.)

```
SELECT ... FROM EMPLOYEE WHERE SALARY > ?
```

If the preferred plan is known to occur when SALARY > 40000 and it is known that 20% of the rows are selected by this predicate, then this query can be rewritten as:

```
SELECT ... FROM EMPLOYEE WHERE SALARY > ? SELECTIVITY .2
```

When running this query, changes to the host variable value will no longer influence the optimized plan.

Considerations for using SELECTIVITY

While SELECTIVITY can be useful when applied carefully, it can also produce some surprising results. Used without care, SELECTIVITY can cause more harm than benefit.

The optimizer integrates multiple pieces of information - including I/O costs, predicate selectivity, and correlation between columns - to form a coherent model of the data and the environment. By overriding the internal model, a user-provided SELECTIVITY value may introduce inconsistencies into and reduce some of the benefits of the model. As a result, SELECTIVITY can lead to less accurate estimates of I/O costs. It may also prevent the optimizer from detecting correlation between related columns and thereby reduce the accuracy of the estimates for the overall set of predicates.

Furthermore, by locking in a single value, SELECTIVITY prevents the optimizer from automatically responding to changes in the underlying data. As the data grows and changes, the SELECTIVITY value may become inaccurate, requiring you to occasionally re-evaluate and adjust it.

For these reasons, SELECTIVITY is best reserved for use in the limited situations described in this section. Before using SELECTIVITY to solve a query performance problem, follow the other recommendations in this document. For example, creating a new index or column statistic might help the optimizer not simply with one but with many queries. Keep in mind that a hard-to-optimize query may be a sign that re-writing the query or improving the data model is a better long-term strategy. SELECTIVITY should be a last resort when no other options exist for providing the optimizer with the needed information.

General Db2 for i performance considerations

As you code your applications, there are some general tips that can help you optimize performance.

Effects on database performance when using long object names

Long object names are converted internally to system object names when used in SQL statements. This conversion can have some performance impacts. Names of tables, views, indexes, and aliases that are 30 characters or less will generally perform much better than names longer than 30 characters.

Qualify the long object name with a library name and the conversion to the short name happens at precompile time. In this case, there is minimal performance impact when the statement is executed. Otherwise, the conversion is done at execution time and has a small performance impact.

Effects of precompile options on database performance

Several precompile options are available for creating SQL programs with improved performance. They are only options because using them could impact the function of the application. For this reason, the default value for these parameters is the value that ensures successful migration of applications from prior releases. However, you can improve performance by specifying other options.

The following table shows these precompile options and their performance impacts.

Some of these options might be suitable for most of your applications. Use the command **CRTDUPOBJ** to create a copy of the SQL **CRTSQLxxx** command, and the **CHGCMDDF** command to customize the optimal values for the precompile parameters. The **DSPPGM**, **DSPSRVPGM**, **DSPMOD**, or **PRTSQLINF** commands can be used to show the precompile options that are used for an existing program object.

Table 62. Precompile options and their performance impacts

Precompile Option	Optimal Value	Improvements	Considerations
ALWCPYDTA	*OPTIMIZE (the default)	Queries where the ordering or grouping criteria conflicts with the selection criteria.	A copy of the data could be made when the query is opened.
ALWBLK	*ALLREAD (the default)	Additional read-only cursors use blocking.	ROLLBACK HOLD might not change the position of a read-only cursor. Dynamic processing of positioned updates or deletes might fail.
CLOSQLCSR	*ENDJOB, *ENDSQL, or *ENDACTGRP	Cursor position can be retained across program invocations.	Implicit closing of SQL cursor is not done when the program invocation ends.
DLYPRP	*YES	Programs using SQL PREPARE statements could run faster.	Complete validation of the prepared statement is delayed until the statement is run or opened.
TGTRLS	*CURRENT (the default)	The precompiler can generate code that takes advantage of performance enhancements available in the current release.	The program object cannot be used on a system from a previous release.

Related reference

[Effects of the ALWCPYDTA parameter on database performance](#)

Some complex queries can perform better by using a sort or hashing method to evaluate the query instead of using or creating an index.

[Control database manager blocking](#)

To improve performance, the SQL runtime attempts to retrieve and insert rows from the database manager a block at a time whenever possible.

[Retaining cursor positions for non-ILE program calls](#)

For non-ILE program calls, the close SQL cursor (CLOSQLCSR) parameter allows you to specify the scope of the following:

[Eliminate redundant validation with SQL PREPARE statements](#)

The processing which occurs when an SQL PREPARE statement is run is like the processing which occurs during precompile processing.

Effects of the ALWCPYDTA parameter on database performance

Some complex queries can perform better by using a sort or hashing method to evaluate the query instead of using or creating an index.

By using the sort or hash, the database manager is able to separate the row selection from the ordering and grouping process. Bitmap processing can also be partially controlled through this parameter. This separation allows the use of the most efficient index for the selection. For example, consider the following SQL statement:

```
EXEC SQL
  DECLARE C1 CURSOR FOR
  SELECT EMPNO, LASTNAME, WORKDEPT
  FROM CORPDATA.EMPLOYEE
  WHERE WORKDEPT = 'A00'
  ORDER BY LASTNAME
END-EXEC.
```

The above SQL statement can be written in the following way by using the OPNQRYF command:

```
OPNQRYF FILE(CORPDATA/EMPLOYEE)
  FORMAT(FORMAT1)
  QRYSLT(WORKDEPT *EQ 'A00')
  KEYFLD(LASTNAME)
```

In the preceding example, when ALWCPYDTA(*NO) or ALWCPYDTA(*YES) is specified, the database manager could try to create an index from the first index with a column named LASTNAME, if such an index exists. The rows in the table are scanned, using the index, to select only the rows matching the WHERE condition.

If ALWCPYDTA(*OPTIMIZE) is specified, the database manager uses an index with the first index column of WORKDEPT. It then makes a copy of all the rows that match the WHERE condition. Finally, it could sort the copied rows by the values in LASTNAME. This row selection processing is more efficient, because the index used immediately locates the rows to be selected.

ALWCPYDTA(*OPTIMIZE) optimizes the total time that is required to process the query. However, the time required to receive the first row could be increased because a copy of the data must be made before returning the first row of the result table. This initial change in response time could be important for applications that are presenting interactive displays or that retrieve only the first few rows of the query. The Db2 for i query optimizer can be influenced to avoid sorting by using the OPTIMIZE clause.

Queries that involve a join operation might also benefit from ALWCPYDTA(*OPTIMIZE) because the join order can be optimized regardless of the ORDER BY specification.

Related concepts

[Plan cache](#)

The plan cache is a repository that contains the access plans for queries that were optimized by SQE.

Related reference

[Effects of precompile options on database performance](#)

Several precompile options are available for creating SQL programs with improved performance. They are only options because using them could impact the function of the application. For this reason, the default value for these parameters is the value that ensures successful migration of applications from prior releases. However, you can improve performance by specifying other options.

[Radix index scan](#)

A radix index scan operation is used to retrieve the rows from a table in a keyed sequence. Like a table scan, all the rows in the index are sequentially processed, but the resulting row numbers are sequenced based upon the key columns.

[Radix index probe](#)

A radix index probe operation is used to retrieve the rows from a table in a keyed sequence. The main difference between the radix index probe and the scan is that the rows returned are first identified by a probe operation to subset them.

Tips for using VARCHAR and VARGRAPHIC data types in databases

Variable-length column (VARCHAR or VARGRAPHIC) support allows you to define any number of columns in a table as variable length. If you use VARCHAR or VARGRAPHIC support, the size of a table can typically be reduced.

Data in a variable-length column is stored internally in two areas: a fixed-length or ALLOCATE area and an overflow area. If a default value is specified, the allocated length is at least as large as the value. The following points help you determine the best way to use your storage area.

When you define a table with variable-length data, you must decide the width of the ALLOCATE area. If the primary goal is:

- **Space saving:** use ALLOCATE(0).
- **Performance:** the ALLOCATE area must be wide enough to incorporate at least 90% to 95% of the values for the column.

It is possible to balance space savings and performance. In the following example of an electronic telephone book, the following data is used:

- 8600 names that are identified by: last, first, and middle name
- The Last, First, and Middle columns are variable length.
- The shortest last name is two characters; the longest is 22 characters.

This example shows how space can be saved by using variable-length columns. The fixed-length column table uses the most space. The table with the carefully calculated allocate sizes uses less disk space. The table that was defined with no allocate size (with all the data stored in the overflow area) uses the least disk space.

Table 63. Disk space used with variable-length columns

Variety of Support	Last Name Max/Alloc	First Name Max/Alloc	Middle Name Max/Alloc	Total Physical File Size	Number of Rows in Overflow Space
Fixed Length	22	22	22	567 K	0
Variable Length	40/10	40/10	40/7	408 K	73
Variable-Length Default	40/0	40/0	40/0	373 K	8600

In many applications, performance must be considered. If you use the default ALLOCATE(0), it doubles the disk unit traffic. ALLOCATE(0) requires two reads; one to read the fixed-length portion of the row and one to read the overflow space. The variable-length implementation, with the carefully chosen ALLOCATE, minimizes overflow and space and maximizes performance. The size of the table is 28% smaller than the fixed-length implementation. Because 1% of rows are in the overflow area, the access requiring two reads is minimized. The variable-length implementation performs about the same as the fixed-length implementation.

To create the table using the ALLOCATE keyword:

```
CREATE TABLE PHONEDIR
(LAST   VARCHAR(40) ALLOCATE(10),
 FIRST  VARCHAR(40) ALLOCATE(10),
 MIDDLE VARCHAR(40) ALLOCATE(7))
```


If you are using host variables to insert or update variable-length columns, use variable length host variables. Because blanks are not truncated from fixed-length host variables, using fixed-length host variables can cause more rows to spill into the overflow space. This increases the size of the table.

In this example, fixed-length host variables are used to insert a row into a table:

```
01 LAST-NAME PIC X(40).
...
MOVE "SMITH" TO LAST-NAME.
EXEC SQL
  INSERT INTO PHONEDIR
  VALUES(:LAST-NAME, :FIRST-NAME, :MIDDLE-NAME, :PHONE)
END-EXEC.
```

The host-variable LAST-NAME is not variable length. The string "SMITH", followed by 35 blanks, is inserted into the VARCHAR column LAST. The value is longer than the allocated size of 10. 30 of 35 trailing blanks are in the overflow area.

In this example, variable-length host variables are used to insert a row into a table:

```
01 VLAST-NAME.
49 LAST-NAME-LEN PIC S9(4) BINARY.
49 LAST-NAME-DATA PIC X(40).
...
MOVE "SMITH" TO LAST-NAME-DATA.
MOVE 5 TO LAST-NAME-LEN.
EXEC SQL
  INSERT INTO PHONEDIR
  VALUES(:VLAST-NAME, :VFIRST-NAME, :VMIDDLE-NAME, :PHONE)
END-EXEC.
```

The host variable VLAST-NAME is variable length. The actual length of the data is set to 5. The value is shorter than the allocated length. It can be placed in the fixed portion of the column.

Running the **Reorganize Physical File Member (RGZPFM)** command against tables that contain variable-length columns can improve performance. The fragments in the overflow area that are not in use are compacted by the **Reorganize Physical File Member (RGZPFM)** command. This technique reduces the read time for rows that overflow, increases the locality of reference, and produces optimal order for serial batch processing.

Choose the appropriate maximum length for variable-length columns. Selecting lengths that are too long increases the process access group (PAG). A large PAG slows performance. A large maximum length makes SEQONLY(*YES) less effective. Variable-length columns longer than 2000 bytes are not eligible as key columns.

Using LOBs and VARCHAR in the same table

Storage for LOB columns allocated in the same manner as VARCHAR columns. When a column stored in the overflow storage area is referenced, currently all the columns in that area are paged into memory. A reference to a "smaller" VARCHAR column that is in the overflow area can potentially force extra paging of LOB columns. For example, A VARCHAR(256) column retrieved by application has side-effect of paging in two 5 MB BLOB columns that are in the same row. In order to prevent this side-effect, you might want to use ALLOCATE keyword to ensure that only LOB columns are stored in the overflow area.

Related information

[Reorganize Physical File Member \(RGZPFM\) command](#)

[Reorganizing a physical file](#)

[Embedded SQL programming](#)

Using field procedures to provide column level encryption

Field procedures can provide column level encryption in Db2 for i.

A field procedure is a user-written exit routine to transform values in a single column. When values in the column are changed, or new values inserted, the field procedure is invoked for each value. The field

procedure can transform that value (encode it) in any way. The encoded value is then stored. When values are retrieved from the column, the field procedure is invoked for each encoded value. The field procedure decodes each value back to the original value. Any indexes defined on a column that uses a field procedure are built with encoded values.

Field procedures are assigned to a table by the FIELDPROC clause of CREATE TABLE and ALTER TABLE.

A field procedure that is specified for a column is invoked in three general situations:

- For field-definition, when the CREATE TABLE or ALTER TABLE statement that names the procedure is executed. During this invocation, the procedure is expected to:
 - Determine whether the data type and attributes of the column are valid.
 - Verify the literal list, and change it if wanted.
 - Provide the field description of the column.
- For field-encoding, when a column value is field-encoded. That occurs for any value that:
 - is inserted in the column by an SQL INSERT statement, SQL MERGE statement, or native write.
 - is changed by an SQL UPDATE statement, SQL MERGE statement, or native update.
 - is the target column for a copy from a column with an associated field procedure. The field procedure might be invoked to encode the copied data. Examples include SQL Statements ALTER TABLE or CREATE TABLE LIKE/AS and CL commands CPYF and RGZPFM.
 - is compared to a column with a field procedure. The QAQQINI option FIELDPROC_ENCODED_COMPARISON is used to determine if the column value is decoded, or the host variable, constant, or join column is encoded.
 - is the DEFAULT value for a column with an associated field procedure in a CREATE or ALTER TABLE statement.

If there are any **after** or **read** triggers, the field procedure is invoked *before* any of these triggers. If there are any **before** triggers, the field procedure is invoked *after* the before trigger.

- For field-decoding, when a stored value is field-decoded back into its original value. Field-decoding occurs for any value that:
 - is retrieved by an SQL SELECT or FETCH statement, or by a native read.
 - is a column with an associated field procedure that is copied. The field procedure might be invoked to decode the data before making the copy. Examples include SQL Statements ALTER TABLE, CREATE TABLE LIKE/AS, and CL commands CPYF and RGZPFM.
 - is compared to a column with a field procedure. The QAQQINI option FIELDPROC_ENCODED_COMPARISON is used by the optimizer to decide if the column value is decoded, or if the host variable or constant is encoded.

A field procedure is never invoked to process a null value. It is also not invoked for a DELETE operation without a WHERE clause when the table has no DELETE triggers. The field procedure is invoked for empty strings.

Improving performance

For queries that use field procedures, the path length is longer due to the additional processing of calling the field procedure. In order to improve performance of queries, the SQE optimizer:

- attempts to remove decoding operations, based on the QAQQINI FIELDPROC_ENCODED_COMPARISON setting.
- matches existing indexes over columns that have an associated field procedure.
- creates and uses MTIs over columns with field procedures. The MTI will be created as non-resuable which means it will not be shared between queries.
- creates statistics over the encoded values through statistics collection.

The SQE optimizer attempts to do the following optimizations:

- optimization of predicates that compare a field procedure column to a constant or host variable. For example, predicate `FP1(4, C1) = 'abc'` is optimized as `C1 = FP1(0,'abc')`. With this specific example, the optimization is done as long as the QAQQINI option is not `*NONE`.
- remove field procedure decoding operations from join predicates when the same field procedure is applied to both sides of the join predicate, and no compatibility conversion is required. For example, join predicate `FP1(4,T1.C1) > FP1(4,T2.C1)` is rewritten as `T1.C1 > T2.C1`. With this specific example, the optimization is done as long as the QAQQINI option is either `*ALLOW_RANGE` or `*ALL`. This technique is also applied to `=` predicates when the QAQQINI option is `*ALLOW_EQUAL`.
- remove field procedures from GROUP BY and ORDER BY clauses. For example, `ORDER BY FP1(4,C1)` is rewritten as `ORDER BY C1` if the QAQQINI setting is either `*ALLOW_RANGE` or `*ALL`.

The CQE optimizer does not look at the QAQQINI option, which means it always runs in `*NONE` mode. `*NONE` mode requires that all references to the column are decoded before any other operation is performed. A CQE query does not match any existing indexes when the column has an associated field procedure. If an index is required, a temporary index is built with the index keys decoded.

Related reference

[QAQQINI query options](#)

There are different options available for parameters in the QAQQINI file.

Related information

[Defining field procedures](#)

[CREATE TABLE](#)

Field procedure examples

The following examples show various field procedure-related optimizations done by the SQE optimizer.

The examples show the FieldProc name along with the encoding (field procedure function code 0) or decoding (field procedure function code 4) in the pseudo-SQL. These codes indicate how the optimizer is optimizing the field procedure calls.

Given the following table and index:

```
CREATE TABLE T1 (col1 CHAR(10), col2 CHAR(10) FIELDPROC 'FP1')
CREATE INDEX IX1 on T1(col2)
```

Example 1

A user query written as:

```
SELECT col1, col2 FROM T1 WHERE col2 = 'abc'
```

Is represented by the optimizer as:

```
SELECT col1, FP1(4, col2) FROM T1 WHERE FP1(4,col2) = 'abc'
```

Note the FP1 with the decode operation around the COL2 references in the SELECT list and the WHERE clause.

Assuming the QAQQINI FIELDPROC_ENCODED COMPARISON is set to `*ALLOW_EQUAL`, `*ALLOW_RANGE` or `*ALL`:

The query optimizer rewrites the query as:

```
SELECT col1, 'abc' FROM T1 WHERE col2 = FP1(0, 'abc')
```

This rewrite allows the query optimizer to use the encoded index IX1 to implement the WHERE clause and only cause one invocation of the field procedure for the query.

Example 2

```
SELECT col2 FROM T1 ORDER BY col2
```

Is represented by the query optimizer as:

```
SELECT FP1(4, col2) FROM T1 ORDER BY FP1(4, col2)
```

The optimized version removes the FieldProc from the ORDER BY clause assuming that the field procedure QAQQINI option is set to *ALLOW_RANGE or *ALL:

```
SELECT FP1(4, col2) FROM T1 ORDER BY col2
```

Example 3

```
Select col2, COUNT(*) FROM T1 GROUP BY col2
```

Is represented by the query optimizer as:

```
Select FP1(4, col2), COUNT(*) FROM T1 GROUP BY FP1(4, col2)
```

The optimized version removes the field procedure invocation from the GROUP BY clause column col2, allowing it to group the encoded data and only run the field procedure once per group. The decoded grouped data is returned to the user. This optimization is done if the field procedure QAQQINI option is set to *ALLOW_RANGE or *ALL:

```
SELECT FP1(4, col2), COUNT(*) FROM T1 GROUP BY col2
```

IS NULL/IS NOT NULL predicate does not require calling the field procedure field-decode option 4. The field procedure cannot change the nullability of the field.

Db2 for i Services

There are many system-provided views, procedures, and functions.

These are grouped in the following categories.

Application Services

These procedures provide interfaces that are useful for application development.

DELIMIT_NAME scalar function

The DELIMIT_NAME function returns a name with delimiters if the delimiters are needed for use in an SQL statement.

►► DELIMIT_NAME — (— *name* —) ◄◄

The schema is QSYS2.

name A character or graphic string expression that identifies a name. The string must contain only characters allowed in an SQL identifier. If the string is longer than 128 characters, it will be truncated to 128 characters.

The result of the function is a varying length character string that contains *name* correctly delimited. This includes delimiting reserved words. If *name* is the null value or an empty string, null is returned.

Example

- Delimit these names:

```
VALUES DELIMIT_NAME('ABC'),
       DELIMIT_NAME('abc'),
       DELIMIT_NAME('test"name'),
       DELIMIT_NAME('test'name2'),
       DELIMIT_NAME('NEW')
```

Returns the values:

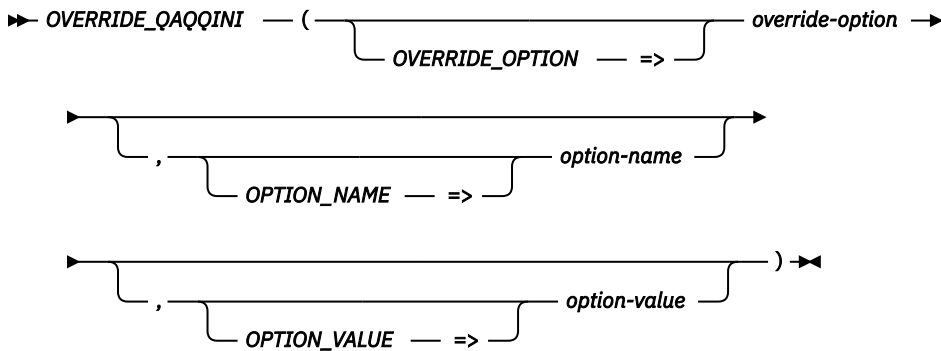
```
ABC
"abc"
"test"name"
"test'name2"
"NEW"
```

OVERRIDE_QAQQINI procedure

The `OVERRIDE_QAQQINI` procedure creates and modifies a temporary version of the QAQQINI file.

The temporary QAQQINI file will be created in QTEMP. It inherits all query options that are already in place for the job. The `OVERRIDE_QAQQINI` procedure can be called multiple times to establish job specific QAQQINI settings.

The procedure can also be called to discard the temporary customization settings.



The schema is QSYS2.

Authorization: For most QAQQINI options, none is required. For the following three options, *JOBCTL or QIBM_DB_SQLADM function usage is required. These options are more restrictive because they can affect the performance of other jobs.

- QUERY_TIME_LIMIT when the *option-value* is not 0.
- STORAGE_LIMIT
- PARALLEL_DEGREE

override-option

An integer value that indicates the function to perform.

- 1 Create the QAQQINI override file. A procedure call with this *override-option* value must be run before option 2 can be used to change QAQQINI options.
- 2 Set a QAQQINI option to the specified value. See [“QAQQINI query options”](#) on page 190 for the list of options and values.
- 3 Discard the temporary QAQQINI file.

option-name

A character or graphic string expression that identifies the name of the QAQQINI option to be changed.

This parameter is required when *override-option* is 2. For options 1 and 3, it can be omitted or passed as the empty string.

option-value

A character or graphic string expression that identifies the value to assign to the QAQQINI option identified by *option-name*.

This parameter is required when *override-option* is 2. For options 1 and 3, it can be omitted or passed as the empty string.

Example

- Establish the temporary override for QAQQINI. The job's current QAQQINI values will be used as the initial values.

```
CALL QSYS2.OVERRIDE_QAQQINI(1);
```

- Improve the chances that you can acquire an exclusive lock for a database file. Multiple calls to the procedure can be used to change more than one QAQQINI value.

```
CALL QSYS2.OVERRIDE_QAQQINI(2, 'PREVENT_ADDITIONAL_CONFLICTING_LOCKS', '*YES');
```

- Discard the temporary QAQQINI file and revert to using the job's version of the QAQQINI file.

```
CALL QSYS2.OVERRIDE_QAQQINI(3);
```

OVERRIDE_TABLE procedure

The `OVERRIDE_TABLE` procedure sets the blocking size for a table.

► `OVERRIDE_TABLE` (— *schema-name* — , — *table-name* — , — *blocking-size* —) ►

The schema is `QSYS2`.

schema-name A character string expression containing the name of the schema.

table-name A character string expression containing the name of the table.

blocking-size A character string expression containing the blocking size. It can be a specific byte count or a special value of `*BUF32KB`, `*BUF64KB`, `*BUF128KB`, or `*BUF256KB`.

Example

- Override the `EMPLOYEE` table to use 256K blocking for sequential processing.

```
CALL QSYS2.OVERRIDE_TABLE('CORPDATA', 'EMP', '*BUF256KB');
```

- Remove the override.

```
CALL QSYS2.OVERRIDE_TABLE('CORPDATA', 'EMP', 0);
```

Related reference

[Control database manager blocking](#)

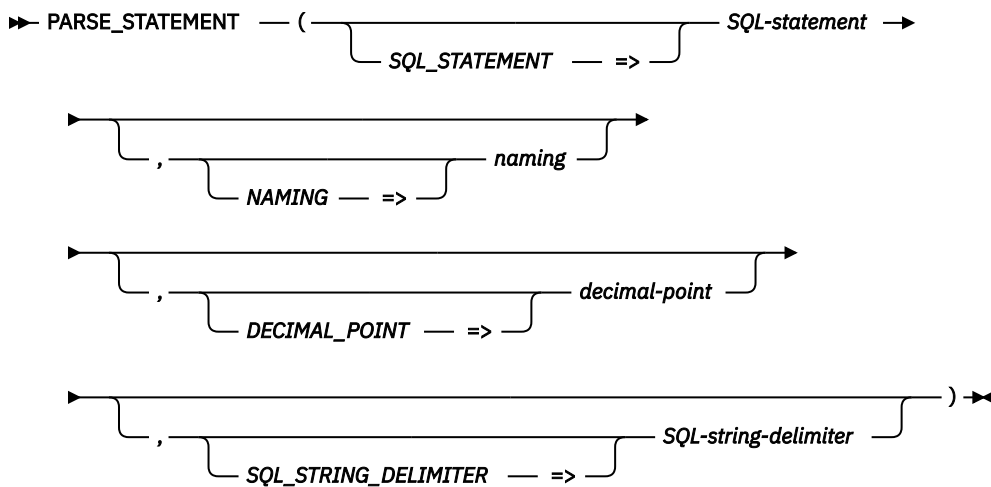
To improve performance, the SQL runtime attempts to retrieve and insert rows from the database manager a block at a time whenever possible.

PARSE_STATEMENT table function

The `PARSE_STATEMENT` table function returns a list of object and column names that are used in an SQL query, data change statement, or other statement where a query or expression is specified.

Authorization:

- None required.



The schema is QSYS2.

SQL-statement A character or graphic string expression that contains a valid SQL statement. The maximum string length is 2 megabytes.

naming A character or graphic string expression that defines the naming rule for the statement.

***SYS** System naming rules apply. This is the default.

***SQL** SQL naming rules apply.

decimal-point A character or graphic string expression that defines the decimal point for numeric constants in *SQL-statement*.

***PERIOD or .** The decimal point is the period. This is the default.

***COMMA or ,** The decimal point is the comma.

SQL-string-delimiter A character or graphic string expression that defines the string delimiter for strings in *SQL-statement*. Delimited identifiers in the SQL statement will use the opposite character.

***APOSTSQL or '** The apostrophe character (') is used to delimit strings. This is the default.

***QUOTESQL or "** The quote character (") is used to delimit strings.

When the SQL statement is parsed, object names are identified and a result row is returned for every name. This is done at the SQL parser level where names are identified strictly by where they appear in the syntax. The following restrictions apply:

- Names used in data change statements and in any query construct are returned.
- For CALL, the procedure being called is returned.
- For DDL statements, the following additional items are returned:
 - For CREATE INDEX, the table on which the index is being created.
 - For CREATE TABLE, any table referenced using the LIKE clause.
 - For CREATE TABLE and ALTER TABLE, any table or column in the REFERENCES clause.
 - For CREATE TRIGGER, the table or view on which the trigger is being defined.
 - For DROP, the object being dropped.

Any statement that does not contain these constructs, a query, or an expression returns no rows.

- Names in a *routine-body*, *triggered-action*, and *trigger-body* are not returned. To see these references, use QSYS2.SYSPROGRAMSTMTSTAT to find all the statements for the generated program or service program and pass each of them as an argument to this table function.
- If the SQL statement is the null value, an empty string, a string of all blanks, or contains a syntax error, no row is returned.

The result of the function is a table containing a row for each name reference with the format shown in the following table. All the columns are null capable.

Table 64. PARSE_STATEMENT table function

Column Name	Data Type	Description
NAME_TYPE	VARCHAR(8)	Type of object name. ALIAS This is an alias name. COLUMN This is a column name, or a global variable name not returned as a DDL TARGET OBJECT. FUNCTION This is a function name. INDEX This is an index name. MASK This is a mask name. PACKAGE This is a package name. PERM This is a permission name. PROC This is a procedure name. ROUTINE This is a routine name for a DDL operation. SCHEMA This is a schema name. SEQUENCE This is a sequence name. SPECIFIC This is a routine specific name for a DDL operation. TABLE This is a table, view, or alias name. TRIGGER This is a trigger name. TYPE This is a user-defined type name. VARIABLE This is a variable name for a DDL operation. VIEW This is a view name for a DDL operation. XSROBJ This is an XSROBJECT name.
NAME	VARCHAR(128)	The object name. Contains null if NAME_TYPE is COLUMN without a table qualifier.
SCHEMA	VARCHAR(128)	The schema name. Contains null if NAME is not qualified with a schema name.
RDB	VARCHAR(128)	The relational database name. Contains null if NAME is not qualified with a relational database name.
COLUMN_NAME	VARCHAR(128)	The column name. Contains null if NAME_TYPE is not COLUMN.

Table 64. PARSE_STATEMENT table function (continued)

Column Name	Data Type	Description	
USAGE_TYPE	VARCHAR(17)	How this name is used in the statement.	
		DDL SOURCE OBJECT	Name identifies the table an index or trigger is being created on, or the table referenced by CREATE TABLE LIKE.
		DDL TARGET OBJECT	Name is the primary object of a DDL statement.
		EXPRESSION	Name is referenced in an index key expression.
		PARAMETER DEFAULT	Name is referenced in a parameter default expression.
		QUERY	Name is referenced as part of a query construct.
		REFERENCES	Name is part of a foreign key constraint referencing another table.
		TARGET PROCEDURE	This is the procedure that is the target of a CALL statement.
NAME_START_POSITION	INTEGER	TARGET TABLE	This is the table that will be affected for an insert, update, delete, merge, or truncate statement. Also set for any explicitly specified columns from the target table for insert, update, and merge.
		Position within the <i>SQL-statement</i> string that this name begins. For qualified TABLE names, this is the position where the RDB or schema name begins. For all other name types, this is the position of the name.	

Table 64. PARSE_STATEMENT table function (continued)

Column Name	Data Type	Description
SQL_STATEMENT_TYPE	VARCHAR(32)	Type of SQL statement. <ul style="list-style-type: none"> • ALTER FUNCTION • ALTER PROCEDURE • ALTER TABLE • CALL • CREATE FUNCTION • CREATE INDEX • CREATE MASK • CREATE PERMISSION • CREATE PROCEDURE • CREATE TABLE • CREATE TRIGGER • CREATE VARIABLE • CREATE VIEW • DECLARE CURSOR • DECLARE GLOBAL TEMPORARY TABLE • DELETE • DROP ALIAS • DROP FUNCTION • DROP INDEX • DROP MASK • DROP PACKAGE • DROP PERMISSION • DROP PROCEDURE • DROP ROUTINE • DROP SCHEMA • DROP SEQUENCE • DROP TABLE • DROP TRIGGER • DROP TYPE • DROP VARIABLE • DROP VIEW • DROP XSROBJECT • EXECUTE IMMEDIATE • INSERT • MERGE • PREPARE • QUERY • SET • SET CURRENT TEMPORAL SYSTEM_TIME • TRUNCATE • UPDATE • VALUES INTO

Example

For every program and service program in library APPLIB, find all the references to table names in static SQL statements.

```
WITH program_statements(naming_mode, dec_point, string_delim, stmt_text,
  system_program_name, program_type)
AS (SELECT a.naming, a.decimal_point, a.sql_string_delimiter, b.statement_text,
  a.system_program_name, a.program_type
  FROM qsys2.sysprogramstat a INNER JOIN
  qsys2.sysprogramstmtstat b ON a.program_schema = b.program_schema AND
```

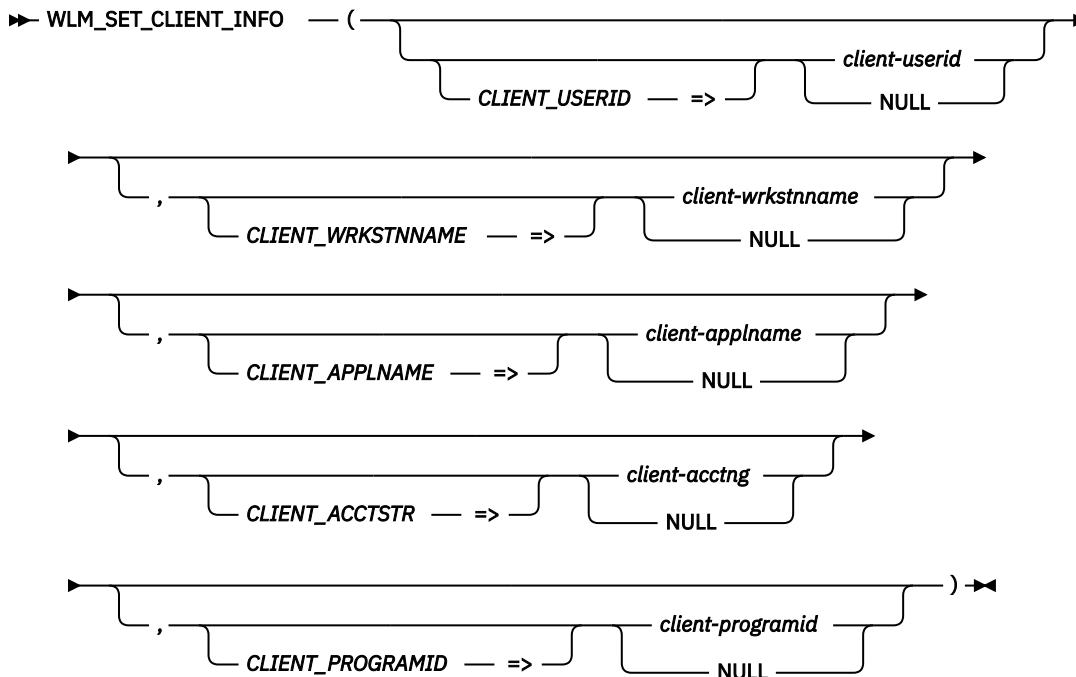
```

a.program_name = b.program_name AND
a.module_name = b.module_name
WHERE a.number_statements > 0 AND
      a.program_schema = 'APPLIB' AND b.program_schema = 'APPLIB')
SELECT system_program_name, program_type, c.schema, c.name, stmt_text
FROM program_statements,
TABLE(qsys2.parse_statement(stmt_text, naming_mode, dec_point, string_delim)) c
WHERE c.name_type = 'TABLE'
ORDER BY c.schema, c.name;

```

WLM_SET_CLIENT_INFO procedure

The WLM_SET_CLIENT_INFO procedure sets values for the SQL client special registers.



The schema is SYSPROC.

If a parameter is omitted or the NULL value is specified for a parameter, the value of the corresponding special register is not changed.

- client-userid** A character or graphic string expression containing the value to set for the CLIENT USERID special register for the current connection.
- client-wrkstnname** A character or graphic string expression containing the value to set for the CLIENT WRKSTNNAME special register for the current connection.
- client-applname** A character or graphic string expression containing the value to set for the CLIENT APPLNAME special register for the current connection.
- client-acctng** A character or graphic string expression containing the value to set for the CLIENT ACCTNG special register for the current connection.
- client-programid** A character or graphic string expression containing the value to set for the CLIENT PROGRAMID special register for the current connection.

Example

- Change the values of the CURRENT CLIENT_USERID and CURRENT CLIENT_PROGRAMID special registers for the current connection.

```
CALL SYSPROC.WLM_SET_CLIENT_INFO(CLIENT_USERID => 'PGM1USER',
                                CLIENT_PROGRAMID => 'PGM1');
```

Performance Services

These services include procedures that provide interfaces to work with indexes and a view to see information about database monitors.

Related reference

[QAQQINI file override support](#)

If you find working with the QAQQINI query options file cumbersome, consider using the QSYS2.OVERRIDE_QAQQINI procedure. Instead of creating, managing, and using a QAQQINI *FILE object directly, this procedure can be called to work with a temporary version of the INI file. It uses user-specified options and values. The support relies upon the QTEMP library, so any changes affect only the job which calls the procedure.

ACT_ON_INDEX_ADVICE procedure

The ACT_ON_INDEX_ADVICE procedure creates new indexes for a table based on indexes that have been advised for the table.

```
►► ACT_ON_INDEX_ADVICE ( ( — schema-name — , — table-name — , ►►
                        ► — times_advised — , — mti_used — , — average_estimate — ) ►►
```

The schema is SYSTOOLS.

<i>schema-name</i>	A character string containing the system name of the schema containing the table.
<i>table-name</i>	A character string containing the system name of the table. If NULL is passed, this parameter is not used to select the target index advice.
<i>times-advised</i>	The number of times an index should have been advised before creating a permanent index. If NULL is passed, this parameter is not used to select the target index advice.
<i>mti-used</i>	The number of times a maintained temporary index (MTI) has been used because a matching permanent index did not exist. If NULL is passed, this parameter is not used to select the target index advice.
<i>average-estimate</i>	The average estimated number of seconds needed to execute the query that drove the index advice. If NULL is passed, this parameter is not used to select the target index advice.

For each potential index meeting the specified criteria, a CREATE INDEX statement will be run to generate the permanent index. A radix index will be named name_RADIX_INDEX_n. An EVI index will be named name_EVI_INDEX_n. The *name* represents the table name and *n* is a unique number. The row containing this advised index is removed from the QSYS2.SYSIXADV table.

Examples

For schema PRODLIB, find all instances of index advice where a maintained temporary index was used more than 1000 times and create permanent SQL indexes.

```
CALL SYSTOOLS.ACT_ON_INDEX_ADVICE('PRODLIB', NULL, NULL, 1000, NULL)
```

Related reference

[SYSTOOLS](#)

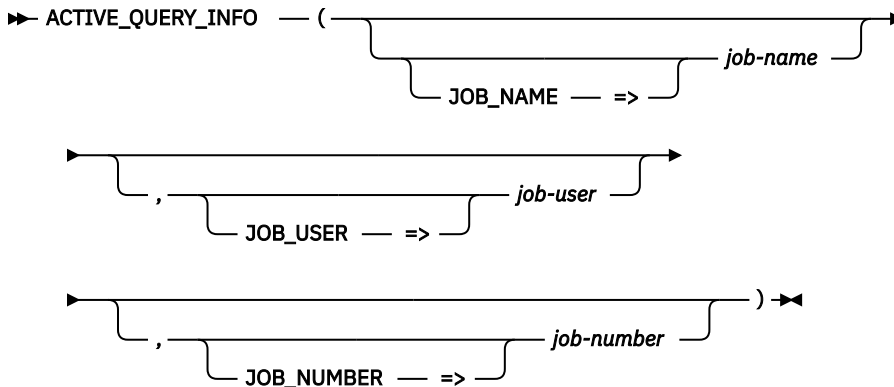
SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

ACTIVE_QUERY_INFO table function

The ACTIVE_QUERY_INFO table function returns information about active SQL Query Engine (SQE) queries. An active query is either open or pseudo-closed.

Authorization: None required to see information for the current job.

Otherwise, the caller must have *JOBCTL special authority or QIBM_DB_SQLADM function usage.



The schema is QSYS2.

The values specified on the input parameters are ANDed together.

job-name A character or graphic string expression that contains an unqualified job name that determines the job information to be returned.

The job name can end with a wildcard character. For example, 'QPADEV*' indicates that any job name starting with the characters 'QPADEV' is a match.

The string can be the following special values:

- * Information for the current job is returned. The *job-user* and *job-number* parameters cannot be specified.
- *ALL** Information for all job names is returned.

If this parameter is not specified, is an empty string, or is the null value, information for all jobs is returned.

job-user A character or graphic string expression that contains a user name that determines the job information to be returned.

The job user name can end with a wildcard character. For example, 'Q*' indicates that any job user name starting with the character 'Q' is a match.

The string can be the following special value:

- *ALL** Information for all job user names is returned.

If this parameter is not specified, is an empty string, or is the null value, information for all users is returned.

job-number A character or graphic string expression that contains a job number that determines the job information to be returned.

The string can be the following special value:

- *ALL** Information for all job numbers is returned.

If this parameter is not specified, is an empty string, or is the null value, information for all job numbers is returned.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 65. ACTIVE_QUERY_INFO table function

Column Name	Data Type	Description
QUALIFIED_JOB_NAME	VARCHAR(28)	The qualified job name.
JOB_NAME	VARCHAR(10)	The name of the job.
JOB_USER	VARCHAR(10)	The user profile that started the job.
JOB_NUMBER	VARCHAR(6)	The job number of the job.
QUERY_TYPE	VARCHAR(6)	The type of query. HLL High Level Language (HLL) open done using SQE. HLL opens include opens done via RPG, C, and COBOL NATIVE Native query SQL SQL query
PSEUDO_CLOSED	VARCHAR(3)	The current state of the SQL cursor associated with the query. NO Query is open YES Query is pseudo-closed Contains the null value if the query is not an SQL query.
QRO_HASH	VARCHAR(8)	The QRO_HASH which uniquely identifies an SQE query.
PLAN_IDENTIFIER	DECIMAL(20,0)	The Plan Identifier of the SQE plan.
FULL_OPEN_TIMESTAMP	TIMESTAMP	The full open timestamp.
LAST_PSEUDO_OPEN_TIMESTAMP	TIMESTAMP	The timestamp of the last pseudo-open. Contains the null value if the query has not been pseudo opened.
LIBRARY_NAME	VARCHAR(10)	The name of the library that contains FILE_NAME.
FILE_NAME	VARCHAR(10)	The database file name of the first table referenced by the query.
NUMBER_OF_PSEUDO_CLOSSES	BIGINT	The number of complete runs for this full open of the query. Returns 0 if this query has not been pseudo-closed.
CURRENT_ROW_COUNT	BIGINT	If this cursor is not pseudo-closed, the current number of rows fetched for this run of the query. Contains the null value if this cursor is currently pseudo-closed.
CURRENT_RUNTIME	BIGINT	If this cursor is not pseudo-closed, the current runtime in microseconds for this run of the query. Contains the null value if this cursor is currently pseudo-closed.
CURRENT_TEMPORARY_STORAGE	BIGINT	The current amount of temporary storage, in MB, used by the query. This size does not include storage used by MTIs.
CURRENT_DATABASE_READS	BIGINT	If this cursor is not pseudo-closed, the current number of asynchronous and synchronous database reads for this run of the query. Contains the null value if this cursor is currently pseudo-closed.
CURRENT_PAGE_FAULTS	BIGINT	If this cursor is not pseudo-closed, the current number of page faults for this run of the query. Contains the null value if this cursor is currently pseudo-closed.
MTI_COUNT	BIGINT	The number of Maintained Temporary Indexes (MTIs). Returns 0 if no MTIs are used by the query.
MTI_SIZE	BIGINT	The current size of MTIs, in MB. This size includes shared MTIs. Contains the null value if no MTIs are used by the query.

Table 65. ACTIVE_QUERY_INFO table function (continued)

Column Name	Data Type	Description
AVERAGE_ROW_COUNT	BIGINT	The average number of rows fetched for pseudo-closed runs of this query. Contains the null value if there are no prior runs.
AVERAGE_RUNTIME	BIGINT	The average runtime in microseconds for pseudo-closed runs of this query. Contains the null value if there are no prior runs.
AVERAGE_TEMPORARY_STORAGE	BIGINT	The average amount of temporary storage in MB for pseudo-closed runs of this query. This size does not include storage used by MTIs. Contains the null value if there are no prior runs.
AVERAGE_DATABASE_READS	BIGINT	The average number of asynchronous and synchronous database reads for pseudo-closed runs of this query. Contains the null value if there are no prior runs.
AVERAGE_PAGE_FAULTS	BIGINT	The average number of page faults for pseudo-closed runs of this query. Contains the null value if there are no prior runs.

Examples

- Retrieve information about all active SQE queries on the system.

```
SELECT *
FROM TABLE(QSYS2.ACTIVE_QUERY_INFO( ));
```

- Find which QZDASOINIT jobs are using MTIs.

```
SELECT JOB_NAME, MTI_COUNT, MTI_SIZE
FROM TABLE(QSYS2.ACTIVE_QUERY_INFO(
  JOB_NAME => 'QZDASOINIT'))
WHERE MTI_COUNT > 0;
```

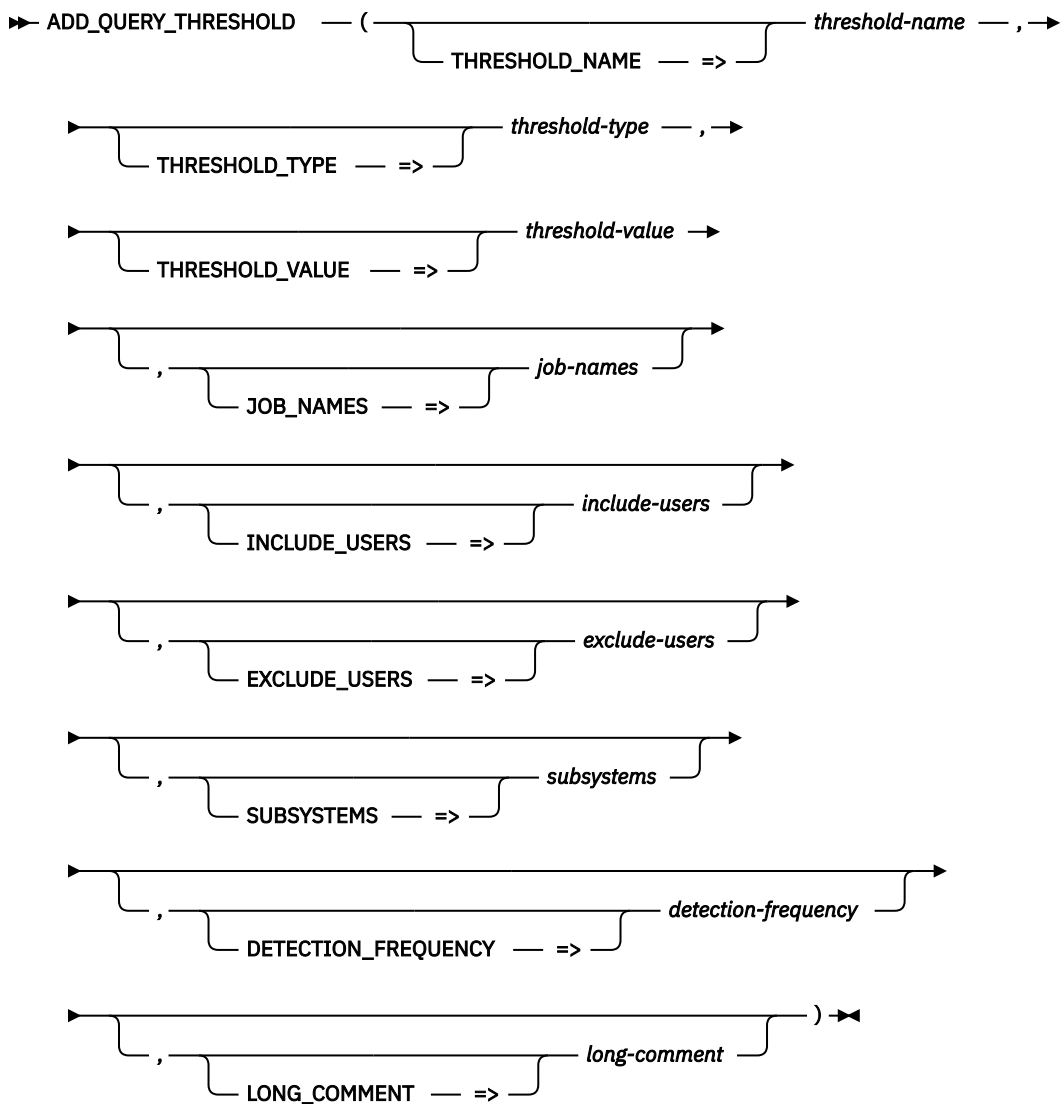
ADD_QUERY_THRESHOLD procedure

The ADD_QUERY_THRESHOLD procedure defines a new threshold to be used by the Query Supervisor. Adding a threshold only affects subsequently executed queries. Queries that are currently running continue to use the thresholds that were in effect when the query execution was initiated. The Query Supervisor only supervises user jobs. It does not affect system jobs or queries run by the operating system. The thresholds only apply to resources used by the SQL Query Engine (SQE).

Query Supervisor thresholds can include using job name, user name, and subsystem filters. The filters are combined to restrict the set of jobs for the rule. For jobs that meet the filtering criteria, when the threshold is met or exceeded, any exit programs registered with the QIBM_QQQ_QRY_SUPER exit point are called. An exit program can be used to implement actions such as query logging or real-time notification. It can also direct SQE to terminate the query. The exit program interface is described here: [Query Supervisor Exit Program](#).

When the first threshold is defined for a system, any active jobs, including the job that runs the initial ADD_QUERY_THRESHOLD, will not be affected by Query Supervisor thresholds.

Authorization: The caller must have *JOBCTL special authority or be authorized to the QIBM_DB_SQLADM function usage ID.



The schema is QSYS2.

threshold-name A character or graphic string that provides a name for the threshold. The name can be up to 30 characters long and can contain any characters including blanks. The name cannot be the same as an existing threshold name.

threshold-type A character or graphic string that specifies the type of the threshold to be enforced.

CPU TIME The total processing unit time used by the query, in seconds.

ELAPSED TIME The total clock time, in seconds.

TEMPORARY STORAGE The amount of storage, in megabytes (MB), that the query uses.

TOTAL IO COUNT The total number of I/O operations.

threshold-value A big integer value that contains the threshold value, in units defined by the specified *threshold-type*. The value must be greater than 0.

job-names A character or graphic string that specifies up to 100 job names separated by either a blank or a comma. Each job name can end with a wildcard character. For example, 'QPADEV*' indicates that any job name starting with the characters 'QPADEV' is a match.

Only jobs running with a matching job name are eligible to be supervised. Can contain the following special value:

***ALL** All jobs are supervised. This is the default.

include-users A character or graphic string that specifies up to 100 user names separated by either a blank or a comma. A group profile does not expand to include all profiles that are members of the group. Each name can end with a wildcard character. For example, 'ADM*' indicates that any user name starting with the characters 'ADM' is a match. Queries where the effective user of the thread matches one of these names are eligible to be supervised. If one or more names are specified for *include-users*, the *exclude-users* parameter must have a value of *NONE. Can contain the following special value:

***ALL** Jobs for all users are supervised. This is the default.

exclude-users A character or graphic string that specifies up to 100 user names separated by either a blank or a comma. A group profile does not expand to include all profiles that are members of the group. Each name can end with a wildcard character. For example, 'ADM*' indicates that any user name starting with the characters 'ADM' is a match. Queries where the effective user of the thread matches one of these names are not supervised. If one or more names are specified for *exclude-users*, the *include-users* parameter must have a value of *ALL. Can contain the following special value:

***NONE** No jobs are explicitly eliminated from being supervised based on the user name. This is the default.

subsystems A character or graphic string that specifies up to 100 subsystem names separated by either a blank or a comma. Each name can end with a wildcard character. For example, 'RSBS*' indicates that any subsystem name starting with the characters 'RSBS' is a match.

Only jobs running in a subsystem that match one of these names are eligible to be supervised. Can contain the following special value:

***ALL** All jobs are supervised. This is the default.

detection-frequency An integer value that specifies the minimum time, in seconds, between calls to exit programs registered with the QIBM_QQQ_QRY_SUPER exit point for this threshold for a specific thread. During the *detection-frequency* time interval following a threshold detection, the Query Supervisor does not process additional instances for this threshold in the same thread. Each time a query is run in a thread, a specific threshold can be processed at most once during the execution of that query.

The *detection-frequency* provides a protection from having a specific threshold recognized and processed more frequently than desired. If the same threshold is encountered within the same thread and the *detection-frequency* time period has not yet completed, the associated exit programs are not called and there is no external evidence that the Query Supervisor did not perform threshold detection.

The default is 600 (10 minutes).

long-comment A character or graphic string that describes this Query Supervisor threshold. It can be up to 2000 characters long.

Examples

- Set a threshold value for total CPU time to 30 seconds for all jobs. When this value is reached, any registered exit programs will be called.

```
CALL QSYS2.ADD_QUERY_THRESHOLD(THRESHOLD_NAME => 'MAXTIME',  
                               THRESHOLD_TYPE => 'CPU TIME',
```



```

THRESHOLD_VALUE => 30,
LONG_COMMENT    => 'Maximum runtime for all jobs');

```

- Set a threshold value for temporary storage to 10 megabytes for all jobs except for those running with the BATCHRUN user profile. Set a detection frequency of 2 minutes. When the storage threshold value is reached, any registered exit programs will be called. If the same thread reaches this threshold for a different query within 2 minutes, the registered exit programs will not be called.

```

CALL QSYS2.ADD_QUERY_THRESHOLD(THRESHOLD_NAME    => 'Temp storage',
                               THRESHOLD_TYPE    => 'TEMPORARY STORAGE',
                               THRESHOLD_VALUE    => 10,
                               EXCLUDE_USERS     => 'BATCHRUN',
                               DETECTION_FREQUENCY => 120);

```

DATABASE_MONITOR_INFO view

The DATABASE_MONITOR_INFO view returns information about database monitors and plan cache event monitors on the server. Database monitors are started using the Start Database Monitor (STRDBMON) command. The QSYS2.START_PLAN_CACHE_EVENT_MONITOR procedure is used to start a plan cache event monitor. SQL Performance Monitors within IBM i Navigator are synonymous with database monitors and are included in this view.

The following table describes the columns in the view. The system name is DBMON_INFO. The schema is QSYS2.

Table 66. DATABASE_MONITOR_INFO view

Column Name	System Column Name	Data Type	Description
MONITOR_ID	MONITOR_ID	CHAR(10)	The system-assigned monitor ID for this monitor.
MONITOR_TYPE	MONTYPE	VARCHAR(7)	Type of monitor. PUBLIC A monitor is considered public when the STRDBMON JOB parameter indicates that jobs other than the current job should be monitored. Public monitors remain active until they are explicitly ended using the End Database Monitor (ENDDBMON) command. PRIVATE A private monitor occurs when the STRDBMON JOB parameter indicates to monitor only the current job. The monitor is ended as part of job termination processing, if needed. Only a private monitor that is active in the current connection will be returned. EVENT An SQL plan cache event monitor intercepts plans as they are moved from the plan cache into a database monitor file.
MONITOR_STATUS	STATUS	VARCHAR(8)	Status of this monitor. ACTIVE Monitor is active. INACTIVE For a PUBLIC or PRIVATE monitor, it is inactive and can become ACTIVE. For an EVENT monitor, entries are no longer being collected. CLOSING The PUBLIC or PRIVATE monitor is not active or is in the processing of ending. It is not known if the entry can be reused for monitoring.
MONITOR_RECORD_TYPE	RCDTYPE	VARCHAR(6)	Type of database records in this monitor. DETAIL Both basic and detail database monitor records. An EVENT monitor always has a value of DETAIL. BASIC Only basic database monitor records
MONITOR_LIBRARY	MONLIB	VARCHAR(10)	Library for this monitor.
MONITOR_FILE	MONFILE	VARCHAR(10)	The file to which the database activity detail is written for this monitor.
MONITOR_MEMBER	MONMBR	VARCHAR(10)	Member for this monitor.

Table 66. DATABASE_MONITOR_INFO view (continued)

Column Name	System Column Name	Data Type	Description
IASP_NUMBER	IASPNUMBER	SMALLINT	The independent auxiliary storage pool (IASP) number for the monitor file.
MONITOR_MEMBER_OPTION	MBROPT	VARCHAR(7) Nullable	Value used for the member replace option the last time this monitor was started. <ul style="list-style-type: none"> REPLACE ADD Contains the null value for an EVENT monitor.
NUMBER_ROWS	CARD	BIGINT Nullable	The number of rows in the database monitor file. Contains the null value if information is not available.
DATA_SIZE	SIZE	BIGINT Nullable	The total size, in bytes, of the database monitor file. Contains the null value if information is not available.
MONITOR_JOB_FILTER	JOB	VARCHAR(32)	Qualified job name for this monitor. For an EVENT monitor, this is the job that started the monitor. Following the qualified job name is the filter operator that applies to the job name. This is either *EQ or *NE. The special value of *ALL indicates all jobs on the system are monitored. A generic name is allowed for both the job name and the user name.
HOST_VARIABLE	HOSTVAR	VARCHAR(9) Nullable	How host variables are handled in this database monitor. <p>BASIC Host variables are written in the QQQ3010 database monitor record.</p> <p>SECURE No host variables are captured and no QQQ3010 record is written.</p> <p>CONDENSED Host variable values are captured in the QQQ1000 database monitor record in column QQDBCLOB1. No QQQ3010 record is written.</p> Contains the null value for an EVENT monitor.
FORCE_RECORDS	FRCRCD	SMALLINT Nullable	The number of records to be held in the buffer before forcing the records to be written to the file when running with a private monitor. Contains the null value if the system calculates the value or for an EVENT monitor.
RUN_THRESHOLD_FILTER	RUNTHLD	INTEGER Nullable	The filtering threshold, in seconds, based on the estimated run time of SQL statements in this monitor. Contains the null value if a run time threshold is not used for filtering or for an EVENT monitor.
STORAGE_THRESHOLD_FILTER	STGTHLD	INTEGER Nullable	The filtering threshold, in megabytes, based on the estimated temporary storage usage of SQL statements in this monitor. Contains the null value if a temporary threshold is not used for filtering or for an EVENT monitor.
INCLUDE_SYSTEM_SQL	INCSYSSQL	VARCHAR(3)	Monitor includes records for system-generated SQL statements. <p>YES Monitor records are generated for both user-specified and system-generated SQL statements.</p> <p>NO Monitor records are generated for only user-specified SQL statements.</p> <p>INI For a PUBLIC or PRIVATE monitor, records are generated based on the value of the SQL_DBMON_OUTPUT option in the QAQQINI query options.</p>

Table 66. DATABASE_MONITOR_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FILE_FILTER	FTRFILE	VARCHAR(2728) Nullable	<p>A list of up to 10 qualified file references that are used for filtering. Following each file name is the filter operator that applies to the file name. This is either *EQ or *NE. When more than one file is listed, a comma and a single blank separate the entries. Either the file name or the library name can be a generic name.</p> <p>A special value of *ALL for the file name indicates all files in the library.</p> <p>Contains the null value if no database files are used for filtering.</p>
USER_FILTER	FTRUSER	VARCHAR(158) Nullable	<p>A list of up to 10 user profiles that are used for filtering. Following each user profile name is the filter operator that applies to the user profile. This is either *EQ or *NE. When more than one profile is listed, a comma and a single blank separate the entries. A profile name can be a generic name.</p> <p>Contains the null value if the user profile is not used for filtering.</p>
TCPIP_FILTER	FTRINTNETA	VARCHAR(254) Nullable	<p>The TCP/IP address or host name is used for filtering.</p> <p>This is an IPv4, IPv6, or IP host domain name, or the special value of *LOCAL.</p> <p>Contains the null value if the TCP/IP address or host name is not used for filtering or for an EVENT monitor.</p>
LOCAL_PORT_FILTER	FTRLCLPORT	INTEGER Nullable	<p>Filtering is based on the local TCP/IP port number. Monitor records will be created for TCP/IP database server jobs running on behalf of the specified local TCP/IP port. Jobs named QRWTSRVR and QZDASOINIT are examples of these server jobs.</p> <p>The IBM i well defined port numbers are documented here: Port numbers for host servers and server mapper.</p> <p>Contains the null value if the port number is not used for filtering or for an EVENT monitor.</p>
QUERY_GOVERNOR_FILTER	FTRQRYGOVR	VARCHAR(11) Nullable	<p>The query governor is used for filtering.</p> <p>ALL Monitor records will be collected when a query governor limit is exceeded.</p> <p>CONDITIONAL Monitor records will be conditionally collected when a query governor limit is exceeded.</p> <p>Contains the null value if the query governor is not used for filtering or for an EVENT monitor.</p>
CLIENT_ACCTNG_FILTER	FTRCLTACG	VARCHAR(128) Nullable	<p>The CURRENT CLIENT_ACCTNG special register is used for filtering.</p> <p>Contains the null value if the CURRENT CLIENT_ACCTNG special register is not used for filtering or for an EVENT monitor.</p>
CLIENT_APPLNAME_FILTER	FTRCLTAPP	VARCHAR(128) Nullable	<p>The CURRENT CLIENT_APPLNAME special register is used for filtering.</p> <p>Contains the null value if the CURRENT CLIENT_APPLNAME special register is not used for filtering or for an EVENT monitor.</p>
CLIENT_PROGRAMID_FILTER	FTRCLTPGM	VARCHAR(128) Nullable	<p>The CURRENT CLIENT_PROGRAMID special register is used for filtering.</p> <p>Contains the null value if the CURRENT CLIENT_PROGRAMID special register is not used for filtering or for an EVENT monitor.</p>
CLIENT_USERID_FILTER	FTRCLTUSR	VARCHAR(128) Nullable	<p>The CURRENT CLIENT_USERID special register is used for filtering.</p> <p>Contains the null value if the CURRENT CLIENT_USERID special register is not used for filtering or for an EVENT monitor.</p>
CLIENT_WRKSTNNAME_FILTER	FTRCLTWS	VARCHAR(128) Nullable	<p>The CURRENT CLIENT_WRKSTNNAME special register is used for filtering.</p> <p>Contains the null value if the CURRENT CLIENT_WRKSTNNAME special register is not used for filtering or for an EVENT monitor.</p>

Table 66. DATABASE_MONITOR_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_CODE_FILTER	FTRSQLCODE	VARCHAR(7) Nullable	<p>How the SQLCODE result from a statement execution is used for filtering.</p> <p>NONZERO Any SQL statement with an SQLCODE value that is non-zero is included in the monitor.</p> <p>ERROR Any SQL statement with an SQLCODE that is less than zero is collected in the monitor.</p> <p>WARNING Any SQL statement with an SQLCODE that is greater than zero is collected in the monitor.</p> <p>SQLCODE Any SQL statement with an SQLCODE that exactly matches the value in the SQLCODE_VALUE column is collected in the monitor.</p> <p>Contains the null value if the SQLCODE for a statement is not used for filtering or for an EVENT monitor.</p>
SQLCODE_VALUE	SQLCODEVAL	INTEGER Nullable	<p>The positive or negative SQLCODE value to use for filtering.</p> <p>Contains the null value if the SQL_CODE_FILTER column contains a value other than SQLCODE.</p>

Examples

Example 1: Get the MONITOR_ID for all the active PUBLIC monitors and the file names associated with the MONITOR_IDs.

```
SELECT MONITOR_ID, MONITOR_LIBRARY, MONITOR_FILE
FROM QSYS2.DATABASE_MONITOR_INFO
WHERE MONITOR_STATUS = 'ACTIVE' AND
MONITOR_TYPE = 'PUBLIC'
```

Example 2: Find the active monitors that have outfiles larger than 1Gig.

```
SELECT MONITOR_LIBRARY, MONITOR_FILE, NUMBER_ROWS, DATA_SIZE
FROM QSYS2.DATABASE_MONITOR_INFO
WHERE MONITOR_STATUS = 'ACTIVE' AND
DATA_SIZE > 1073741824
```

Example 3: Find any active monitors that are filtering based upon a specific SQLCODE (FTRSQLCODE).

```
SELECT MONITOR_ID, MONITOR_LIBRARY, MONITOR_FILE, SQLCODE_VALUE
FROM QSYS2.DATABASE_MONITOR_INFO
WHERE MONITOR_STATUS = 'ACTIVE' AND
SQL_CODE_FILTER = 'SQLCODE'
```

Example 4: Get the MONITOR_ID for a user's SQL plan cache event monitor and use it to end the active event monitor.

```
CALL QSYS2.END_PLAN_CACHE_EVENT_MONITOR (SELECT MONITOR_ID
FROM QSYS2.DATABASE_MONITOR_INFO
WHERE MONITOR_TYPE = 'EVENT' AND
MONITOR_LIBRARY = 'USERLIB')
```

HARVEST_INDEX_ADVICE procedure

The HARVEST_INDEX_ADVICE procedure generates one or more CREATE INDEX statements in source file members for a specified table based on indexes that have been advised for the table.

```
►► HARVEST_INDEX_ADVICE — ( — schema-name — , — table-name — , ►  
  
    ► — times_advised — , — mti_used — , — average_estimate — , — output-library — , — output-file ►  
  
    ► — ) ►►
```

The schema is SYSTOOLS.

<i>schema-name</i>	A character string containing the system name of the schema containing the table.
<i>table-name</i>	A character string containing the system name of the table.
<i>times-advised</i>	The number of times an index should have been advised before creating a permanent index. Pass a value of 1 to not limit index creation by the number of times advised.
<i>mti-used</i>	The number of times a maintained temporary index (MTI) has been used because a matching permanent index did not exist. Pass a value of 0 to not limit index creation by MTI use.
<i>average-estimate</i>	The average estimated number of seconds needed to execute the query that drove the index advice. Pass a value of 0 to not limit index creation by the average query estimate.
<i>output-library</i>	A character string value containing the name of the library for the output source file.
<i>output-file</i>	A character string value containing the name of the output source file. The file must exist and must be a source physical file.

For each potential index meeting the specified criteria, a CREATE INDEX statement to create the permanent index will be generated in a member in the source file provided to this procedure. A radix index will be named name_RADIX_INDEX_n. An EVI index will be named name_EVI_INDEX_n. The *name* represents the table name and *n* is a unique number. The row containing this advised index is removed from the QSYS2.SYSIXADV table.

Example

Harvest create index statements for file TOYSTORE/SALES into source physical file QGPL/INDEXSRC. Then use RUNSQLSTM to create the indexes.

```
BEGIN  
  DECLARE NOT_FOUND CONDITION FOR '02000';  
  DECLARE ERROR_COUNT INTEGER DEFAULT 0;  
  DECLARE AT_END INT DEFAULT 0;  
  DECLARE V_INDEX_COUNT INTEGER DEFAULT 0;  
  DECLARE V_PARTITION_NAME VARCHAR(10);  
  DECLARE INDEX_SOURCE_CURSOR CURSOR FOR  
    SELECT TABLE_PARTITION FROM QSYS2.SYSPARTITIONSTAT  
      WHERE TABLE_SCHEMA = 'QGPL' AND TABLE_NAME = 'INDEXSRC';  
  DECLARE CONTINUE_HANDLER FOR SQLEXCEPTION SET ERROR_COUNT = ERROR_COUNT + 1;  
  
  CALL QSYS2.QCMDXEC('DLTF FILE(QGPL/INDEXSRC)');  
  CALL QSYS2.QCMDXEC('CRTSRCPF FILE(QGPL/INDEXSRC) RCDLEN(10000)');  
  -- Create the source file with a big record length,  
  -- to allow the create index statement to fit on one line  
  CALL SYSTOOLS.HARVEST_INDEX_ADVICE('TOYSTORE', 'SALES', 100, 50, 5, 'QGPL', 'INDEXSRC');  
  BEGIN  
    DECLARE V_RUNSQLSTM_TEXT VARCHAR(500);  
    DECLARE CONTINUE_HANDLER FOR SQLEXCEPTION SET AT_END = 1;  
    DECLARE CONTINUE_HANDLER FOR NOT_FOUND SET AT_END = 1;  
    OPEN INDEX_SOURCE_CURSOR;  
    FETCH FROM INDEX_SOURCE_CURSOR INTO V_PARTITION_NAME;  
    WHILE ( AT_END = 0 ) DO  
      -- Now that QGPL/INDEXSRC has been populated with members named HARVnnnn,
```

```

-- the RUNSQLSTM command can be used to execute the CREATE INDEX statement.
-- By using ERRLLVL = 30, we'll ignore any failures.
SET V_RUNSQLSTM_TEXT = 'RUNSQLSTM SRCFILE(QGPL/INDEXSRC) SRCMBR('
    CONCAT RTRIM(V_PARTITION_NAME)
    CONCAT ') COMMIT(*NONE) NAMING(*SQL) ERRLLVL(30) MARGINS(10000)';
CALL QSYS2.QCMDEXC(V_RUNSQLSTM_TEXT);
FETCH FROM INDEX_SOURCE_CURSOR INTO V_PARTITION_NAME;
END WHILE;
CLOSE INDEX_SOURCE_CURSOR;
END;
END;

```

Related reference


SYSTOOLS

SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

MTI_INFO table function

The MTI_INFO table function returns information about Maintained Temporary Indexes (MTIs).

Authorization: The caller must have *JOBCTL special authority or QIBM_DB_SQLADM function usage.



The schema is QSYS2.

table-schema A character or graphic string expression that contains the schema name of the base table for the MTIs to be returned. The schema name may be either the SQL or the system schema name.

Can contain the following special value:

***ALL** Information for MTIs in all schemas with a base table specified by *table-schema* is returned. This is the default.

If this parameter contains an empty string or the null value, *ALL is used.

table-name A character or graphic string expression that contains the table name of the base table for the MTIs to be returned. The table name may be either the SQL or the system table name.

Can contain the following special value:

***ALL** Information for all MTIs with a base table in *table-schema* is returned. This is the default.

If this parameter contains an empty string or the null value, *ALL is used.

The result of the function is a table containing multiple rows with the format shown in the following table. Each row contains information for one MTI on the system. All columns are nullable.

Table 67. MTI_INFO table function

Column Name	Data Type	Description
TABLE_SCHEMA	VARCHAR(128)	Schema name of the base table that the MTI is created over.
TABLE_NAME	VARCHAR(128)	Name of the base table that the MTI is created over.
REFERENCE_COUNT	BIGINT	The current number of references to this MTI. This includes plans in the plan cache and open queries that are using this MTI. When this count goes to 0, the MTI will be deleted.

Table 67. MTI_INFO table function (continued)

Column Name	Data Type	Description
KEYS	INTEGER	Number of keys.
KEY_DEFINITION	DBCLOB(10000) CCSID 1200	Key definition.
STATE	VARCHAR(20)	The current state of the MTI. CREATED MTI is in a valid state but has not been populated. This can occur if the query subtree that uses the MTI has not been run. CREATING MTI is currently being created. DELETING MTI is currently being deleted. MAPPING ERROR A selection or mapping error occurred during index build. A mapping or selection error can occur if the key definition or the sparse definition contains an expression and that expression resulted in an error when being run. An example of that is a divide by zero. POPULATING The MTI is currently populating. REBUILD REQUIRED The index portion of the MTI will need to be re-populated on the next use of the MTI. VALID MTI is in a valid state and populated. nnnn A 4 digit number indicating an error status to be used by IBM support.
MTI_SIZE	BIGINT	The current size of the MTI in bytes.
CREATE_TIME	TIMESTAMP	The timestamp of when the MTI was created.
LAST_BUILD_START_TIME	TIMESTAMP	Start of the last index build. Contains the null value if the index portion of the MTI has not been populated.
LAST_BUILD_END_TIME	TIMESTAMP	End of the last index build. Contains the null value if the index portion of the MTI has not been populated or is currently being populated.
REUSABLE	VARCHAR(3)	MTI is reusable across queries. NO MTI can only be used for this query. YES MTI is reusable.
SPARSE	VARCHAR(3)	MTI is sparse. NO MTI is not sparse. YES MTI is sparse.
SPARSE_DEFINITION	DBCLOB(10000) CCSID 1200	The predicate used to create the sparse MTI. Contains the null value if SPARSE is NO.
QRO_HASH	VARCHAR(8)	An internally generated identifier which identifies the SQE query which originally created the MTI.
PLAN_IDENTIFIER	DECIMAL(20,0)	Identifies the plan within the plan cache which originally created the MTI.
USER_NAME	VARCHAR(10)	The effective user of the thread that created the MTI.
QUALIFIED_JOB_NAME	VARCHAR(28)	The fully qualified job name of the job that created the MTI.
JOB_NAME	VARCHAR(10)	The name of the job.
JOB_USER	VARCHAR(10)	The user profile that started the job.
JOB_NUMBER	VARCHAR(6)	The job number of the job.
MTI_NAME	VARCHAR(128)	A generated name that identifies the MTI.
LIBRARY_NAME	VARCHAR(10)	The name of the library that contains FILE_NAME.
FILE_NAME	VARCHAR(10)	The database file name of the base table that the MTI is created over.

Example

- Return information about all MTIs that exist over the APPLIB/EMPLOYEE table.

```
SELECT * FROM TABLE(QSYS2.MTI_INFO('APPLIB', 'EMPLOYEE'));
```

QUERY_SUPERVISOR view

The QUERY_SUPERVISOR view contains the threshold rules defined for the Query Supervisor.

Authorization: The caller must have *JOBCTL special authority or be authorized to the QIBM_DB_SQLADM function usage ID.

The following table describes the columns in the view. The system name is QRY_SUPER. The schema is QSYS2.

Table 68. QUERY_SUPERVISOR view

Column Name	System Column Name	Data Type	Description
THRESHOLD_NAME	NAME	VARGRAPHIC(30) CCSID 1200	The name of the threshold.
THRESHOLD_TYPE	TYPE	VARCHAR(30)	The type of threshold CPU TIME The total processing unit time used by the query, in seconds. ELAPSED TIME The total clock time, in seconds. TEMPORARY STORAGE The amount of storage, in megabytes (MB), that the query uses. TOTAL IO COUNT The total number of I/O operations.
THRESHOLD_VALUE	VALUE	BIGINT	The value associated with THRESHOLD_TYPE, in the appropriate units.
JOB_NAMES	JOB_NAMES	VARCHAR(1099) Nullable	A list of job names supervised by this threshold. Each name in the list is padded with blanks to fill ten characters, with a single comma separating the name entries. Contains the null value if all job names are supervised.
INCLUDE_USERS	INCL_USERS	VARCHAR(1099) Nullable	A list of user names supervised by this threshold. Each name in the list is padded with blanks to fill ten characters, with a single comma separating the name entries. Contains the null value if all user names are supervised.
EXCLUDE_USERS	EXCL_USERS	VARCHAR(1099) Nullable	A list of user names that are not supervised by this threshold. Each name in the list is padded with blanks to fill ten characters, with a single comma separating the name entries. Contains the null value if no user names are excluded.
SUBSYSTEMS	SUBSYSTEMS	VARCHAR(1099) Nullable	A list of subsystems supervised by this threshold. Each name in the list is padded with blanks to fill ten characters, with a single comma separating the name entries. Contains the null value if all subsystems are supervised.

Table 68. QUERY_SUPERVISOR view (continued)

Column Name	System Column Name	Data Type	Description
DETECTION_FREQUENCY	FREQUENCY	INTEGER	<p>The minimum frequency, in seconds, that a notification is sent for this threshold within a specific thread. During the time interval following a threshold detection, the Query Supervisor does not process additional instances for this threshold, in the same thread. Each time a query is run in a thread, a specific threshold can be processed at most once during the execution of that query.</p> <p>The <i>detection-frequency</i> value provides a protection from having a specific threshold recognized and processed more frequently than desired. If the same threshold is encountered within the same thread and the <i>detection-frequency</i> time period has not yet completed, the associated exit programs are not called and there is no external evidence that the Query Supervisor did not perform threshold detection.</p>
LONG_COMMENT	REMARKS	VARGRAPHIC(2000) CCSID 1200 Nullable	<p>Description of this threshold rule.</p> <p>Contains the null value if the rule has no descriptive text.</p>

Example

- List all the thresholds defined for the Query Supervisor.

```
SELECT *
FROM QSYS2.QUERY_SUPERVISOR ORDER BY THRESHOLD_TYPE, THRESHOLD_VALUE DESC;
```

REMOVE_INDEXES procedure

The REMOVE_INDEXES procedure drops any indexes meeting the specified criteria.

►► REMOVE_INDEXES (— *schema-name* — , — *times_used* — , — *index-age* —) ◀◀

The schema is SYSTOOLS.

schema-name A character string containing the system name of the schema containing the indexes to be evaluated. If the NULL value is passed, the entire database is processed.

times_used A big integer value indicating the number of times an index has been used.

index-age A character string containing an SQL labeled duration, such as '2 MONTHS'.

The procedure will evaluate all indexes for the specified *schema-name* value and drop any index that does not meet the *times-used* and *index-age* threshold. If the number of times the index has been used by a query and used for statistics is less than the *times-used* value, the index is considered under utilized and is a candidate to be dropped. An index that has existed at least the length of time indicated by *index-age* is also a candidate to be dropped. Any index that meets both criteria is dropped.

Only index names that have names like name_RADIX_INDEX_x or name_EVI_INDEX_x will be considered by this procedure.

Examples

- Remove any index in MYLIB that is older than a month that has never been used.

```
CALL SYSTOOLS.REMOVE_INDEXES('MYLIB', 1, '1 MONTH')
```

- Remove all indexes from all schemas on the system that have existed for at least two weeks and haven't been used at least 100 times.

```
CALL SYSTOOLS.REMOVE_INDEXES(NULL, 100, '14 DAYS')
```

Related reference

SYSTOOLS

SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

REMOVE_QUERY_THRESHOLD procedure

The REMOVE_QUERY_THRESHOLD procedure removes a threshold that is monitored by the Query Supervisor. Removing a threshold only affects subsequently executed queries. Queries that are currently running will continue to use the thresholds that were in effect when the query execution was initiated.

Authorization: The caller must have *JOBCTL special authority or be authorized to the QIBM_DB_SQLADM function usage ID.

➤ REMOVE_QUERY_THRESHOLD — (————— *threshold-name* —) ➤
 └─── THRESHOLD_NAME — => ───┘

The schema is QSYS2.

threshold-name A character or graphic string that specifies the name of an existing threshold to be removed.

Example

Remove the MAXTIME threshold rule.

```
CALL QSYS2.REMOVE_QUERY_THRESHOLD(THRESHOLD_NAME => 'MAXTIME');
```

RESET_TABLE_INDEX_STATISTICS procedure

The RESET_TABLE_INDEX_STATISTICS procedure clears usage statistics for indexes defined over a table or tables and optionally deletes rows from the index advice tracking table.

Authorization: The counts will only be reset when the caller has *OBJMGT and *OBJOPR authority on the table. For each index found over the table, *OBJOPR is required. If the user does not have the required authority to the table, the object is skipped and no warning is returned. If the user does not have the required authority to the index, the object is skipped and an SQL warning is returned. To delete index advice, the DELETE privilege is required on QSYS2/SYSIXADV. Index advice is only deleted when the caller has the required authority to the table and index.

➤ RESET_TABLE_INDEX_STATISTICS — (————— *schema-name* — , —→
 └─── SCHEMA_NAME — => ───┘

 └─── *table-name* —→
 └─── TABLE_NAME — => ───┘

) ➤
 └─── *delete-advice* —
 └─── DELETE_ADVICE — => ───┘

The schema is QSYS2.

This procedure will zero the QUERY_USE_COUNT and QUERY_STATISTICS_COUNT usage statistics for all indexes over the specified tables. These counts can also be reset using the Change Object Description (CHGOBJD) CL command, but the command requires an exclusive lock.

schema-name A character string expression for the name of the schema or schemas to use. The name is case-sensitive and must not be delimited. Wildcard characters (_ and %) are allowed in the string following the rules for the SQL LIKE predicate.

table-name A character string expression for the name of the table or tables to use. The name is case-sensitive and must not be delimited. Wildcard characters (_ and %) are allowed in the string following the rules for the SQL LIKE predicate.

delete-advice A character string expression that indicates whether this procedure should remove rows from the index advice tracking table.

NO Index advice for the table is not affected. This is the default.

YES This procedure will delete rows from the index advice tracking table (QSYS2/SYSIXADV) that correspond to *schema-name* and *table-name*. To be removed, an index must exist and the user must be authorized to the index.

The procedure writes information related to every index processed into an SQL global temporary table. The fields LAST_QUERY_USE, LAST_STATISTICS_USE, LAST_USE_DATE, and NUMBER_DAYS_USED are not affected. The following query will display the results of the last call to the procedure:

```
SELECT * FROM SESSION.SQL_INDEXES_RESET
```

The table that is created contains the following columns:

Table 69. SQL_INDEXES_RESET result table

Column Name	System Column Name	Data Type	Description
TABLE_SCHEMA	DBNAME	VARCHAR(128)	Schema name of table.
TABLE_NAME	NAME	VARCHAR(128)	Name of table.
TABLE_PARTITION	TABLE00001	VARCHAR(128)	Name of the table partition or member.
PARTITION_TYPE	PARTI00001	CHAR(1)	The type of table partitioning.
PARITION_NUMBER	PARTI00002	INTEGER	The partition number of this partition.
NUMBER_DISTRIBUTED_PARTITIONS	NUMBE00001	INTEGER	If the table is a distributed table, contains the total number of partitions.
INDEX_SCHEMA	INDEX00001	VARCHAR(128)	Schema name of index.
INDEX_NAME	INDEX_NAME	VARCHAR(128)	Name of index.
INDEX_MEMBER	INDEX00002	VARCHAR(128)	Partition or member name of index.
INDEX_TYPE	INDEX_TYPE	CHAR(11)	Type of index.
LAST_QUERY_USE	LAST_00002	TIMESTAMP	The timestamp of the last time the SQL index was used in a query since the last time the usage statistics were reset.
LAST_STATISTICS_USE	LAST_00003	TIMESTAMP	The timestamp of the last time the SQL index was used by the optimizer for statistics since the last time the usage statistics were reset.
QUERY_USE_COUNT	QUERY00001	BIGINT	The number of times the SQL index was used in a query since the last time the usage statistics were reset.
QUERY_STATISTICS_COUNT	QUERY00002	BIGINT	The number of times the SQL index was used by the optimizer for statistics since the last time the usage statistics were reset.
SYSTEM_TABLE_SCHEMA	SYS_DNAME	CHAR(10)	System table schema name.
SYSTEM_TABLE_NAME	SYS_TNAME	CHAR(10)	System table name.
SYSTEM_TABLE_MEMBER	SYSTE00001	CHAR(10)	System member name.

Examples

- Zero the statistics for all indexes over table TOYSTORE.SALES

```
CALL QSYS2.RESET_TABLE_INDEX_STATISTICS ('TOYSTORE', 'SALES')
```

- Zero the statistics for all indexes over any table in schema TOYSTORE whose name starts with the letter S.

```
CALL QSYS2.RESET_TABLE_INDEX_STATISTICS ('TOYSTORE', 'S%')
```

Plan Cache Services

These services include procedures to assist you in performing database administration (DBA) and database engineering (DBE) tasks.

These procedures allow for programmatic access to the SQL plan cache and can be used for tasks such as scheduling plan cache captures or pre-starting an event monitor.

CHANGE_PLAN_CACHE_SIZE procedure

The CHANGE_PLAN_CACHE_SIZE procedure provides a way to change the maximum amount of storage used by the plan cache.

```
►► CHANGE_PLAN_CACHE_SIZE — ( — integer-expression — ) ►◄
```

The schema is QSYS2.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.

integer-expression A numeric expression that specifies the size, in megabytes, that the plan cache cannot exceed. Once set, the maximum size of the plan cache will be retained across IPLs and IBM i operating system upgrades. If the value is zero, the plan cache is reset to its default value which allows the plan cache to be auto-sized by the database.

Example

- Change the plan cache maximum size to 3072 megabytes:


```
CALL QSYS2.CHANGE_PLAN_CACHE_SIZE(3072);
```

CLEAR_PLAN_CACHE procedure

The CLEAR_PLAN_CACHE procedure either clears all plans or removes the specified plans from the SQL plan cache without requiring a system IPL.

Authorization: The CLEAR_PLAN_CACHE procedure requires that the authorization ID associated with the statement has *JOBCTL special authority or QIBM_DB_SQLADM function usage.

```
►► CLEAR_PLAN_CACHE — ( —>
```



The schema is QSYS2.

qro-hash A character or graphic string expression that identifies a set of plans for a specific SQE query. If *qro-hash* is not specified, the null value is used.

plan-identifier A numeric value which uniquely identifies a plan within the plan cache. If *plan-identifier* is not specified, the null value is used.

If *qro-hash* is specified and *plan-identifier* is not specified, the procedure will clear all the plans from the SQL plan cache with that *qro-hash*. If a *qro-hash* and a *plan-identifier* are specified, that specific plan will be removed. If no matching plan is found, the procedure returns without error.

If no parameters are specified, or a null value is specified for both *qro-hash* and *plan-identifier*, the procedure will clear all plans in the SQL plan cache that exist at the time the procedure is run.

Besides clearing the plan information, any Maintained Temporary Indexes (MTIs) not currently in use by a query will be deleted as part of the clear operation. Queries run while the CLEAR_PLAN_CACHE procedure is running may have their plans removed, but the queries themselves will not incur a failure related to plan removal. After the clear is complete, as queries are re-optimized, they will be inserted into the plan cache.

Removal of all plans is intended for use primarily in performance test and quality assurance environments. It provides database performance analysts with a way to create a consistent environment from which to evaluate potential database performance changes.

Removal of a specific plan is intended for limited use, such as when directed by IBM service or when needing to force the optimizer to re-optimize a specific query.

The time the CLEAR_PLAN_CACHE procedure takes to run will vary depending on the plan cache size. To avoid tying up an interactive job, it is recommended that the procedure should be submitted in a batch job using a combination of the Submit Job (SBMJOB) and Run SQL (RUNSQL) CL commands.

Notes

The QRO hash is an internally generated identifier for an SQE query. In general, this identifier will be unique for each SQE query and uses implicit schema qualification among other data to generate the QRO hash. If the SQE optimizer generates multiple plans for the same query, then multiple plans will have the same QRO hash. However, every plan will have a unique plan identifier. The QRO hash for a statement may change on release boundaries or after loading PTFs. The QRO hash is externalized:

- In a Visual Explain
- From Show Statements exploration of the SQL Plan Cache and SQL Plan Cache Snapshots
- In the QQC83 column of the 3014 record of a database monitor or plan cache snapshot file
- As information passed to a Query Supervisor exit program.
- As returned from the QSYS2.ACTIVE_JOB_INFO table function

The plan identifier is a unique number that is generated when the plan is optimized. The plan identifier is externalized:

- The statement number of a Visual Explain of a plan cache snapshot or a Visual Explain from the Show Statements exploration of the SQL Plan Cache
- From Show Statements exploration of the SQL Plan Cache and SQL Plan Cache Snapshots
- In the QQCNT column of the 1000 record of a plan cache snapshot file.
- As information passed to a Query Supervisor exit program.

Example

- Submit a job to clear the plan cache.

```
SBMJOB CMD(RUNSQL SQL('CALL QSYS2.CLEAR_PLAN_CACHE()') COMMIT(*NONE) NAMING(*SQL))
```

- Clear all plans that have a QRO hash of 'D0C257FD'.

```
CALL QSYS2.CLEAR_PLAN_CACHE(QRO_HASH => 'D0C257FD');
```

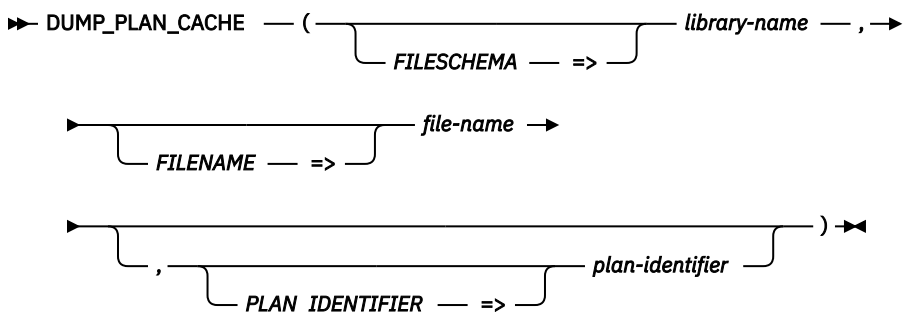
- Clear one specific plan.

```
CALL QSYS2.CLEAR_PLAN_CACHE(QRO_HASH => '1239A9F0',  
                             PLAN_IDENTIFIER => 1305582);
```

DUMP_PLAN_CACHE procedure

The DUMP_PLAN_CACHE procedure creates a database monitor file (snapshot) from the plan cache.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.



The schema is QSYS2.

library-name A character or graphic string expression that identifies the name of the library containing *file-name*. The special value of *CURLIB can be used.

file-name A character or graphic string expression that identifies the name of the resulting database monitor file. If the file does not exist, it is created. If the file exists, the new plan cache snapshot will be appended to it.

plan-identifier A numeric value which uniquely identifies a plan within the plan cache.

The file has the same definition and authorities as the QSYS/QAQQDBMN file. If the file does not exist, it is created with the public authority set to *EXCLUDE. See [Database monitor SQL table format](#) for more information.

If *plan-identifier* is specified, only that specific plan is dumped. If *plan-identifier* is not specified, all user-initiated plans in the plan cache are dumped.

The time the DUMP_PLAN_CACHE procedure takes to run will vary depending on the plan cache size. To avoid tying up an interactive job, it is recommended that the procedure should be submitted in a batch job using a combination of the Submit Job (SBMJOB) and Run SQL (RUNSQL) CL commands.

Notes

The plan identifier is a unique number that is generated when the plan is optimized. The plan identifier is externalized in several ways:

- The statement number of a Visual Explain of a plan cache snapshot or a Visual Explain from the Show Statements exploration of the SQL Plan Cache
- From Show Statements exploration of the SQL Plan Cache and SQL Plan Cache Snapshots
- In the QUCNT column of the 1000 record of a plan cache snapshot file.
- As information passed to a Query Supervisor exit program.

Example

- Dump the plan cache to a database performance monitor file called SNAPSHOT1 in library SNAPSHOTS.

```
CALL QSYS2.DUMP_PLAN_CACHE('SNAPSHOTS', 'SNAPSHOT1');
```

- Dump a specific plan to a database performance monitor file called QUERY1 in library SNAPSHOTS.

```
CALL QSYS2.DUMP_PLAN_CACHE('SNAPSHOTS', 'QUERY1', 126783);
```

DUMP_PLAN_CACHE_PROPERTIES procedure

The DUMP_PLAN_CACHE_PROPERTIES creates a file containing the properties of the plan cache.

► DUMP_PLAN_CACHE_PROPERTIES — (— *library-name* — , — *file-name* —) ◄

The schema is QSYS2.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.

library-name A character or graphic string expression that identifies the name of the library containing *file-name*.

file-name A character or graphic string expression that identifies the name of the resulting properties file. It must be a valid system file name. If the file does not exist, it is created with public authority for the file the same as the create authority specified for the library in which the file is created. Use the QSYS2.LIBRARY_INFO table function or the Display Library Description (DSPLIBD) command to show the library's create authority. If the file exists, the new plan cache properties will be appended to it.

The file definition matches the definition of QSYS2/QDBOPPCGEN, which has the following definition.

Table 70. QDBOPPCGEN table

Column Name	System Column Name	Data Type	Description
HEADING	HEADING	VARCHAR(128)	Description of the value for this row.
VALUE	VALUE	VARCHAR(128)	The current value of the item.
VALUE_UNITS	VALUEUNITS	VARCHAR(128)	The type of unit that applies to the VALUE columns. MB The number is in megabytes. % The number represents a percentage. If there is no unit, the value is a single blank.
ENTRY_NUMBER	ENTRYNBR	INTEGER	A unique identifying number for this row. The number assigned to a value will be the same in every generated plan cache properties file.
ENTRY_TYPE	ENTRY_TYPE	CHAR(2)	The type of entry. DA Detail entry HE Heading entry AD Modifiable detail entry
GRAPHABLE	GRAPHABLE	CHAR(1)	The data value is meaningful to be graphed.
MINVALUE	MINVALUE	VARCHAR(128)	The minimum value allowed for this item
MAXVALUE	MAXVALUE	VARCHAR(128)	The maximum value allowed for this item.
DFTVALUE	DFTVALUE	VARCHAR(128)	The initial value provided by the system for this value.
CHANGE_WARNING	CHGWARN	VARCHAR(512)	Not used.

Example

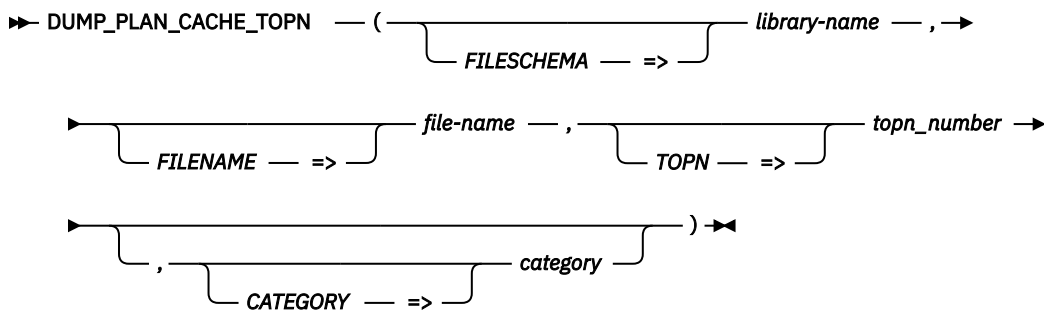
- Dump the plan cache properties to a file called PCPROP1 in library SNAPSHOTS.

```
CALL QSYS2.DUMP_PLAN_CACHE_PROPERTIES('SNAPSHOTS','PCPROP1');
```

DUMP_PLAN_CACHE_TOPN procedure

The DUMP_PLAN_CACHE_TOPN procedure creates a database monitor file (snapshot) from the plan cache containing the user-initiated queries based on the specified *category* option.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.



The schema is QSYS2.

- library-name** A character or graphic string expression that identifies the name of the library containing *file-name*.
- file-name** A character or graphic string expression that identifies the name of the resulting database monitor file. If the file does not exist, it is created. If the file exists, the new plan cache properties will be appended to it.
- topn-number** An integer expression representing the number of queries to dump
- category** A character or graphic string expression that identifies the type of top N queries to be dumped.
 - CPU** Dump the TOPN queries with the most accumulated CPU time.
 - DATABASE READS** Dump the TOPN queries with the most accumulated database read I/Os. This includes both synchronous and asynchronous reads.
 - RUNTIME** Dump the TOPN queries with the longest accumulated runtime. This is the default.
 - STORAGE** Dump the TOPN queries with the largest temporary storage usage.

The file has the same definition as the QSYS/QAQQDBMN file. If the file does not exist, it is created with the public authority set to *EXCLUDE. See [Database monitor SQL table format](#) for more information.

Example

- Capture the 20 queries with the largest elapsed time and dump the details into a snapshot file named SNAPSHOTS/TOPN121413

```
CALL QSYS2.DUMP_PLAN_CACHE_TOPN('SNAPSHOTS', 'TOPN121413', 20);
```

DUMP_SNAP_SHOT_PROPERTIES procedure

The DUMP_SNAP_SHOT_PROPERTIES returns a result set containing the properties for an existing plan cache snapshot.

►► DUMP_SNAP_SHOT_PROPERTIES (— *library-name* — , — *file-name* —) ►►

The schema is QSYS2.

Authorization: None required.

- library-name** A character or graphic string expression that identifies the name of the library containing *file-name*.
- file-name** A character or graphic string expression that identifies the name of an existing database monitor file.

The result set has the following definition.

Column Name	Data Type	Description
HEADING	VARCHAR(128)	Description of the value for this row.
VALUE	VARCHAR(128)	The current value of the item.
VALUE_UNITS	VARCHAR(128)	The type of unit that applies to the VALUE column. MB The number is in megabytes. Returns null if no unit applies.
ENTRY_NUMBER	INTEGER	A unique identifying number for this row. The number assigned to a value will be the same in every generated plan cache properties file.
ENTRY_TYPE	CHAR(2)	The type of entry. DA Detail entry HE Heading entry AD Modifiable detail entry

Example

- Return the plan cache properties of an existing snapshot.

```
CALL QSYS2.DUMP_SNAP_SHOT_PROPERTIES('SNAPSHOTS', 'TOPN121413');
```

END_ALL_PLAN_CACHE_EVENT_MONITORS procedure

The END_ALL_PLAN_CACHE_EVENT_MONITORS procedure ends all active plan cache event monitors.

►► END_ALL_PLAN_CACHE_EVENT_MONITORS — (—) ◄◄

The schema is QSYS2.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.

Example

- End all active plan cache event monitors:

```
CALL QSYS2.END_ALL_PLAN_CACHE_MONITORS();
```

END_PLAN_CACHE_EVENT_MONITOR procedure

The END_PLAN_CACHE_EVENT_MONITOR procedure ends the event monitor identified by the given *monitor-ID*.

►► END_PLAN_CACHE_EVENT_MONITOR — (— *monitor-ID* —) ◄◄

The schema is QSYS2.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.

monitor-ID A character or graphic string expression that identifies the monitor to be ended.

Example

- End the plan cache monitor identified by PLANCACHE1.

```
CALL QSYS2.END_PLAN_CACHE_EVENT_MONITOR('PLANC00001');
```

IMPORT_PC_EVENT_MONITOR procedure

The IMPORT_PC_EVENT_MONITOR procedure imports an event monitor file for usage in the SQL Performance Center and IBM i Navigator.

►► IMPORT_PC_EVENT_MONITOR ((— *library-name* — , — *file-name* — , — *imported-name* —)) ►◄

The schema is QSYS2.

- library-name*** A character or graphic string expression that identifies the name of the library containing *file-name*.
- file-name*** A character or graphic string expression that identifies the name of an existing event monitor file.
- imported-name*** The name of the plan cache event monitor to import. This name is used to identify the plan cache event monitor in IBM i Navigator.

Example

- Import a plan cache event monitor SNAPSHOTS/MON1 and name it NEWMON.

```
CALL QSYS2.IMPORT_PC_EVENT_MONITOR('SNAPSHOTS','MON1','NEWMON');
```

IMPORT_PC_SNAPSHOT procedure

The IMPORT_PC_SNAPSHOT procedure imports a plan cache snapshot file for usage in the SQL Performance Center and IBM i Navigator.

►► IMPORT_PC_SNAPSHOT ((— *library-name* — , — *file-name* — , — *imported-name* —)) ►◄

The schema is QSYS2.

- library-name*** A character or graphic string expression that identifies the name of the library containing *file-name*.
- file-name*** A character or graphic string expression that identifies the name of an existing plan cache snapshot file.
- imported-name*** The name of the plan cache snapshot to import. This name is used to identify the plan cache snapshot in IBM i Navigator.

Example

- Import a plan cache snapshot file SNAPSHOTS/SNAP121413.

```
CALL QSYS2.IMPORT_PC_SNAPSHOT('SNAPSHOTS','SNAP121413','Snapshot captured on 12/14/2013');
```

REMOVE_PC_EVENT_MONITOR procedure

The REMOVE_PC_EVENT_MONITOR procedure deletes a file containing an event monitor.

►► REMOVE_PC_EVENT_MONITOR ((— *library-name* — , — *file-name* —)) ►◄

The schema is QSYS2.

- library-name*** A character or graphic string expression that identifies the name of the library containing *file-name*.

file-name A character or graphic string expression that identifies the name of an existing event monitor file.

Example

- Remove the plan cache event monitor PRUNEDP1 in SNAPSHOTS.

```
CALL QSYS2.REMOVE_PC_EVENT_MONITOR('SNAPSHOTS','PRUNEDP1');
```

REMOVE_PC_SNAPSHOT procedure

The REMOVE_PC_SNAPSHOT procedure deletes a file containing a plan cache snapshot.

► REMOVE_PC_SNAPSHOT — (— *library-name* — , — *file-name* —) ◄

The schema is QSYS2.

library-name A character or graphic string expression that identifies the name of the library containing *file-name*.

file-name A character or graphic string expression that identifies the name of an existing plan cache snapshot file.

Example

- Remove the plan cache snapshot file called PC1 in library SNAPSHOTS.

```
CALL QSYS2.REMOVE_PC_SNAPSHOT('SNAPSHOTS','PC1');
```

REMOVE_PERFORMANCE_MONITOR procedure

The REMOVE_PERFORMANCE_MONITOR procedure deletes a file containing a performance monitor.

► REMOVE_PERFORMANCE_MONITOR — (— *library-name* — , — *file-name* —) ◄

The schema is QSYS2.

library-name A character or graphic string expression that identifies the name of the library containing *file-name*.

file-name A character or graphic string expression that identifies the name of an existing performance monitor file.

Example

- Remove the performance monitor MON1 in SNAPSHOTS.

```
CALL QSYS2.REMOVE_PERFORMANCE_MONITOR('SNAPSHOTS','MON1');
```

START_PLAN_CACHE_EVENT_MONITOR procedure

The START_PLAN_CACHE_EVENT_MONITOR procedure starts an event monitor to capture plans as they are removed from the cache and generates performance information into a database monitor file.

► START_PLAN_CACHE_EVENT_MONITOR — (— *library-name* — , — *file-name* —)

◄ — *monitor-ID* — ◄

The schema is QSYS2.

Authorization: *JOBCTL special authority or QIBM_DB_SQLADM function usage is required.

library-name A character or graphic string expression that identifies the name of the library to contain *file-name*. *library-name* cannot be QTEMP.

file-name A character or graphic string expression that identifies the name of the database monitor file. It must be a valid system file name. If the file does not exist, it is created.
Initially the file is created and populated with the starting record id 3018 (column QQRID = 3018).

monitor-ID An optional character string output variable that will contain the 10 character identifier of the event monitor that was started.

The event monitor stays active until one of the following occurs:

- It is ended by one of the end event monitor procedure calls.
- It is ended using the IBM i Navigator interface.
- An IPL of the system occurs.
- The specified database monitor file is deleted or otherwise becomes unavailable.

Example

- Start an event monitor and place plan information into a database monitor file called PRUNEDP1 in library SNAPSHOTS:

```
CALL QSYS2.START_PLAN_CACHE_EVENT_MONITOR('SNAPSHOTS','PRUNEDP1');
```

- Start an event monitor and place plan information into a database monitor file called PRUNEDPLANS1 in library SNAPSHOTS. Capture the monitor id into host variable HVmonid for use later:

```
CALL QSYS2.START_PLAN_CACHE_EVENT_MONITOR('SNAPSHOTS','PRUNEDPLANS1',:HVmonid);
```

Utility Services

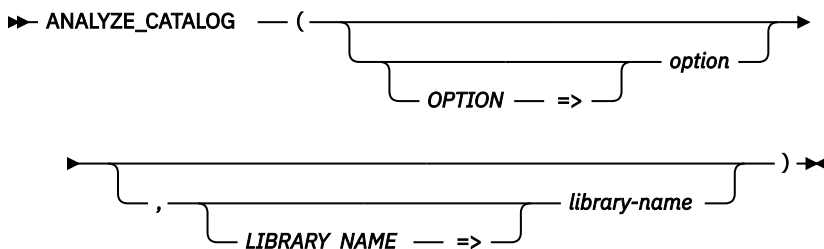
These procedures provide interfaces to monitor and work with SQL in jobs on the current system or to compare constraint and routine information across systems.

ANALYZE_CATALOG table function

The ANALYZE_CATALOG table function returns one row for each inconsistency it finds in the database catalog.

If no rows are returned, no inconsistencies were identified. Multiple rows can be returned for the same object.

Authorization: The caller must have *ALLOBJ special authority.

►► ANALYZE_CATALOG (

The schema is QSYS2.

option A character or graphic string expression that identifies the type of catalog analysis to perform.

DBXREF Analysis will be done for objects in the database cross reference files. This is the default.

Verification is done for constraints defined for *FILE objects.

The analysis done by this function goes beyond what RCLDBXREF OPTION(*CHECK) provides. It verifies that the constraint definitions in every database file object are correctly recorded in the database catalog.

DBXREF SERVER Analysis will be done of the database cross reference servers. The database cross reference servers will be verified to ensure they are functioning properly. One or more result rows with DBXREF SERVER in the CATEGORY column indicate the current database cross reference server status.

library-name A character or graphic string expression that identifies the library to be analyzed. This parameter only applies when *option* is DBXREF.

***ALL** All *FILE objects within all libraries will be analyzed, including any in independent auxiliary storage pools that are accessible. This is the default.

name All *FILE objects within this library will be analyzed.

The result of the function is a table containing one row for each reported item with the format shown in the following table. All the columns are nullable.

Table 71. ANALYZE_CATALOG table function

Column Name	Data Type	Description
LIBRARY_NAME	VARCHAR(10)	Name of the library containing the database object. Can contain the null value when CATEGORY is DBXREF SERVER.
OBJECT_NAME	VARCHAR(128)	Name of the database object. Can contain the null value when CATEGORY is DBXREF SERVER.
OBJECT_TYPE	VARCHAR(30)	Type of database object. Can contain the null value when CATEGORY is DBXREF SERVER.
SEVERITY	VARCHAR(7)	The severity of the reported item. ERROR This row indicates an error that should be corrected. INFO This row provides information about processing that was performed. WARNING This row indicates a problem that should be investigated, but it does not indicate a serious problem.
CATEGORY	VARCHAR(20)	The category of this catalog row. CONSTRAINT An inconsistency was found with a *FILE object constraint. DBXREF SERVER An inconsistency was found with the database cross reference server jobs. RECLAIM The CATALOG_NAME for this row is known to be inconsistent. Either RCLDBXREF *FIX or RCLSTG SELECT(*DBXREF) is required. Use RCLDBXREF *CHECK to determine the action that needs to be taken.
DESCRIPTION	VARCHAR(250)	A text description of the inconsistency in the database catalog.

Table 71. ANALYZE_CATALOG table function (continued)

Column Name	Data Type	Description
DETAIL	VARCHAR(128)	Additional detail about the inconsistency. <ul style="list-style-type: none"> When CATEGORY is CONSTRAINT, this is the name of the constraint. When CATEGORY is DBXREF SERVER, this can contain the number of entries in the cross reference queue Contains the null value if this row does not have additional detail information.
CATALOG_LIBRARY	VARCHAR(10)	The name of the library containing the catalog. Contains the null value if this row is not related to a specific catalog.
CATALOG_NAME	VARCHAR(10)	The catalog table that contains inconsistent information. Contains the null value if this row is not related to a specific catalog.

Example

- Determine if there are any constraint inconsistencies for any database files on the system.

```
SELECT * FROM TABLE(QSYS2.ANALYZE_CATALOG(OPTION => 'DBXREF'));
```

- Determine if there are any constraint inconsistencies for database files in library APPLIB.

```
SELECT * FROM TABLE(QSYS2.ANALYZE_CATALOG(OPTION => 'DBXREF',
LIBRARY_NAME => 'APPLIB'));
```

- Examine how the database cross reference servers are running. If they are behind on processing, the DESCRIPTION and DETAIL columns will indicate the number of entries in the queue that are waiting to be processed.

```
SELECT * FROM TABLE(QSYS2.ANALYZE_CATALOG(OPTION => 'DBXREF SERVER'));
```

CANCEL_SQL procedure

The CANCEL_SQL procedure requests cancellation of an SQL statement for the specified job.

➤ CANCEL_SQL — (— *job-name* —) ➤

The schema is QSYS2.

job-name A character string containing the qualified job name to be cancelled. It must be in upper case.

The CANCEL_SQL() procedure provides an alternative to end job immediate. It supports all application and interactive SQL environments.

When an SQL cancel is requested, an asynchronous request is sent to the job identified by *job-name*. If the job is processing an interruptible, long-running machine operation, analysis is done within the job to determine whether it is safe to cancel the statement. When it is determined to be safe to cancel the statement, an SQL0952 escape message is sent, causing the statement to terminate.

If it isn't safe to end the SQL statement, or if there is no active SQL statement, the request to cancel is ignored. The caller of the cancel procedure will observe a successful return which only indicates that the caller had the necessary authority to request a cancel and that the target job exists. The caller of the CANCEL_SQL() procedure has no programmatic means of determining that the cancel request resulted in a cancelled SQL statement.

If the cancel request occurs during the act of committing or rolling back a commitment-control transaction, the request is ignored.

Authorization: The CANCEL_SQL procedure requires that the authorization ID associated with the statement has *JOBCTL special authority.

Errors: The procedure will fail with a SQL0443 if the target job is not found. The procedure will fail with SQL0443 and SQL0552 if the caller does not have *JOBCTL user special authority.

Commitment control: When the target application is running without commitment control (COMMIT = *NONE), the cancelled SQL statement will terminate without rolling back the partial results of the statement. If the cancelled statement is a query, the query ends. However, if the cancelled statement was a long-running INSERT, UPDATE, or DELETE SQL statement, the changes made prior to cancellation remain intact.

If the target application is using transaction management, the SQL statement is running under a transaction savepoint level. When a long running INSERT, UPDATE, or DELETE SQL statement is cancelled, the changes made prior to cancellation are rolled back.

In both cases, the application receives control back with an indication that the SQL statement failed. It is up to the application to determine the next action.

Example

Safely cancel a job running an SQL statement.

```
CALL QSYS2.CANCEL_SQL('483456/QUSER/QZDASOINIT')
```

CHECK_SYSCST procedure

The CHECK_SYSCST procedure compares entries in the QSYS2.SYSCST table between two systems.

```
►► CHECK_SYSCST ( ( — remote-rdb-name — , — schema-name — ) , — avoid-result-set — )
```

The schema is SYSTOOLS.

remote-rdb-name A character string containing the name of the remote database.

schema-name A character string containing the name of the schema to compare.

avoid-result-set An integer value that indicates whether a result set should be returned. The default is 0.

1 No result set is returned.

0 Result set is returned.

This procedure will return a result set to the caller. If no result set is requested, the differences are logged to the SESSION.SYSCSTDIFF table.

The result set that is returned or the table that is created contains the following columns:

Table 72. SYSCSTDIFF result set

Column Name	System Column Name	Data Type	Description
SERVER_NAME	SRVRNAME	VARCHAR(18)	Name of server where the request was run.
CONSTRAINT_SCHEMA	CDNAME	VARCHAR(128)	Name of the schema containing the constraint.
CONSTRAINT_NAME	RELNAME	VARCHAR(128)	Name of the constraint.
CONSTRAINT_TYPE	TYPE	VARCHAR(11)	Type of constraint.
TABLE_SCHEMA	TDBNAME	VARCHAR(128)	Name of schema containing the table.
TABLE_NAME	TBNAME	VARCHAR(128)	Name of the table which the constraint is created over.
SYSTEM_TABLE_SCHEMA	SYS_DNAME	CHAR(10)	System name of schema containing the table.
SYSTEM_TABLE_NAME	SYS_TNAME	CHAR(10)	System name of the table which the constraint is created over.

Table 72. SYSCSTDIFF result set (continued)

Column Name	System Column Name	Data Type	Description
CONSTRAINT_KEYS	COLCOUNT	SMALLINT	Specifies the number of key columns if this is a UNIQUE, PRIMARY KEY, or FOREIGN KEY constraint.
CONSTRAINT_STATE	CST_STATE	VARCHAR(11)	Indicates whether the constraint is established or defined.
ENABLED	ENABLED	VARCHAR(3)	Indicates whether the constraint is enabled or disabled.
CHECK_PENDING	CHECK00001	VARCHAR(3)	Indicates whether the constraint is in check pending state.

Example

Find all the differences in constraint settings between the current system and LP01UT18 for the CORPDB_EX schema:

```
CALL SYSTOOLS.CHECK_SYSCST('LP01UT18', 'CORPDB_EX')
```

Related reference

SYSTOOLS

SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

CHECK_SYSRoutine procedure

The CHECK_SYSRoutine procedure compares entries in the QSYS2.SYSROUTINES table between two systems.

►► CHECK_SYSRoutine — (►►

►► *remote-rdb-name* , *schema-name* [, *avoid-result-set*] ►►

The schema is SYSTOOLS.

remote-rdb-name A character string containing the name of the remote database.

schema-name A character string containing the name of the schema to compare.

avoid-result-set An integer value that indicates whether a result set should be returned. The default is 0.

1 No result set is returned.

0 Result set is returned.

This procedure will return a result set to the caller. If no result set is requested, the differences are logged to the SESSION.SYSRTNDIFF table.

The result set that is returned or the table that is created contains the following columns:

Table 73. SYSRTNDIFF result set

Column Name	System Column Name	Data Type	Description
SERVER_NAME	SRVRNAME	VARCHAR(18)	Name of server where the request was run.
ROUTINE_CREATED	RTNCREATE	TIMESTAMP	The timestamp when the routine was created.
ROUTINE_DEFINER	DEFINER	VARCHAR(128)	Name of the user that defined the routine.
LAST_ALTERED	ALTEREDTS	TIMESTAMP	Timestamp when routine was last altered.
SPECIFIC_SCHEMA	SPECSHEMA	VARCHAR(128)	Schema name of the routine instance.
SPECIFIC_NAME	SPECNAME	VARCHAR(128)	Specific name of the routine instance.
ROUTINE_SCHEMA	RTNSHEMA	VARCHAR(128)	Name of the schema that contains the routine.
ROUTINE_NAME	RTNNAME	VARCHAR(128)	Name of the routine.

Table 73. SYSRTNDIFF result set (continued)

Column Name	System Column Name	Data Type	Description
ROUTINE_TYPE	RTNTYPE	VARCHAR(9)	Type of the routine.
ROUTINE_BODY	BODY	VARCHAR(8)	Type of the routine body.
EXTERNAL_NAME	EXTNAME	VARCHAR(279)	External program name for routine.
IN_PARMS	IN_PARMS	SMALLINT	Identifies the number of input parameters.
OUT_PARMS	OUT_PARMS	SMALLINT	Identifies the number of output parameters.
INOUT_PARMS	INOUT_PARM	SMALLINT	Identifies the number of input/output parameters.
SQL_DATA_ACCESS	DATAACCESS	VARCHAR(8)	Identifies whether a routine contains SQL and whether it reads or modifies data.
PARM_SIGNATURE	SIGNATURE	VARCHAR(2048)	The signature of the routine.

Example

Compare the current system to a remote system to find which routines differ, when they were created, and who created them.

```
CALL SYSTOOLS.CHECK_SYSRoutine('LP01UT18', 'CORPDB_EX')
```

Related reference

SYSTOOLS

SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

DUMP_SQL_CURSORS procedure

The DUMP_SQL_CURSORS procedure lists the open cursors for a job.

►► DUMP_SQL_CURSORS — (→

► — *job-name* — , — *library-name* — , — *table-name* — , — *output-option* —) ►►

The schema is QSYS2.

job-name A character string containing a qualified job name or a value of '*' to indicate the current job

library-name A character string containing a system library name for the procedure output. An empty string is allowed.

table-name A character string containing a system table name for the procedure output. An empty string is allowed.

output-option An integer value that indicates how to return the output.

- 1 Ignore *library-name* and *table-name* parameters and return a result set.
- 2 Ignore *library-name* and *table-name* parameters and place the results in table QTEMP/SQL_CURSORS.
- 3 Place the results in table *table-name* in library *library-name*. If the table doesn't exist, it will be created. If the table already exists, the results will be appended to the existing table.
- 4 Place the results in table *table-name* in library *library-name*. If the table doesn't exist, it will not be created.

The result set that is returned or the table that is created contains the following columns:

Table 74. DUMP_SQL_CURSORS result table

Column Name	System Column Name	Data Type	Description
SQL_IDENTITY	SQL_I00001	INTEGER	Unique identifier for the row.
DUMPTIME	DUMPTIME	TIMESTAMP	Timestamp when row was inserted.
DUMP_BY_USER	DUMPUSER	VARCHAR(18)	User ID used to insert row.
CURSOR_NAME	CSRNAME	VARCHAR(128)	Name of the cursor.
PSEUDO_CLOSED	PSEUDO	VARCHAR(3)	Pseudo close state of the cursor. Valid values are: YES Cursor is currently pseudo closed. NO Cursor is currently opened.
STATEMENT_NAME	STMTNAME	VARCHAR(128)	Name of the statement corresponding to the cursor
OBJECT_NAME	OBJNAME	CHAR(10)	Object containing the current SQL statement. Blank if current SQL statement is not in a program, service program, or package.
OBJECT_LIBRARY	OBJLIB	CHAR(10)	Library for object containing the current SQL statement. Blank if current SQL statement is not in a program, service program, or package.
OBJECT_TYPE	OBJTYPE	CHAR(10)	Type of object containing the current SQL statement. Blank if current SQL statement is not in a program, service program, or package.
JOBNAME	JOBNAME	CHAR(28)	System job name for the cursor. Contains * if current job was specified for <i>job-name</i> argument.

Example

Populate QGPL/SQLCSR1 table with open SQL cursors for the current job.

```
CALL QSYS2.DUMP_SQL_CURSORS('*' , 'QGPL' , 'SQLCSR1' , 3);
```

END_IDLE_SQE_THREADS procedure

The END_IDLE_SQE_THREADS procedure will end any idle SQE threads for the current job.

Authorization: None required.

►► END_IDLE_SQE_THREADS — (—) ◄◄

The schema is QSYS2.

When the SQE database engine runs a query using Symmetric Multiprocessing (SMP) or a query that contains a fenced UDF or UDTF, system threads are started. These threads may be left active for a period of time so that they are eligible to be re-used. This can cause problems with an application when a non-threadsafe command is executed after the query is run. When this happens the application receives an error message, CPF180B "Function XXXX is not allowed in a job which has multiple threads." Running the END_IDLE_SQE_THREADS procedure will end all SQE threads for the job that are not actively being used, so this error can be avoided.

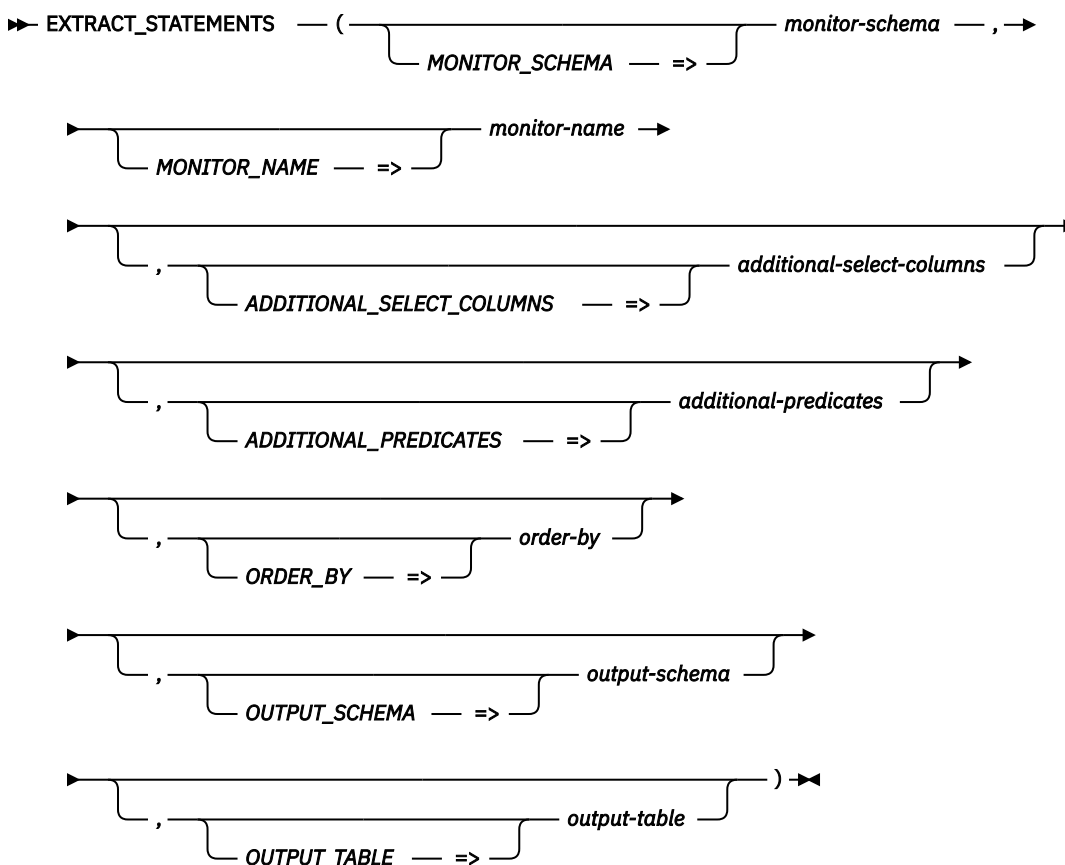
Example

- End the idle SQE threads.

```
CALL QSYS2.END_IDLE_SQE_THREADS();
```

EXTRACT_STATEMENTS procedure

The EXTRACT_STATEMENTS procedure returns details from an SQL plan cache snapshot, an SQL performance monitor, or an SQL plan cache event monitor in the form of an SQL table or a result set.



The schema is QSYS2.

- | | |
|----------------------------------|---|
| monitor-schema | A character or graphic string expression that identifies the name of the library containing <i>monitor-name</i> . |
| monitor-name | A character or graphic string expression that identifies the name of an existing database monitor file. |
| additional-select-columns | A character or graphic string of up to 5000 characters containing additional columns or expressions to be appended to the generated SELECT clause. A value of *AUDIT will cause the procedure to return the merged statement and columns that are normally interesting for auditing.

If <i>additional-select-columns</i> is not specified or has the null value, no additions are made to the generated SELECT clause. |
| additional-predicates | A character or graphic string of up to 5000 characters containing additional predicates to be appended to the generated WHERE clause.

If <i>additional-predicates</i> is not specified or has the null value, no additions are made to the generated WHERE clause. |
| order-by | A character or graphic string of up to 5000 characters containing additional options to be appended to the end of the generated query. This can include the ORDER BY clause or other clauses such as FETCH FIRST n ROWS. |

If *order-by* is not specified or has the null value, no additions are made to the end of the query.

output-schema The schema name for the output table.

If *output-schema* is not specified, the null value is used.

output-name The table name to contain the output. If the table identified by *output-schema* and *output-table* does not exist, it will be created. If the table exists, the result of this procedure call will be appended to the table. For an existing table, the number of selected columns must match the selected columns when the table was generated.

If *output-name* is not specified, the null value is used.

If *output-schema* or *output-table* have the null value, a result set containing the extracted statement information is returned.

Examples

- Extract the 100 most recent statements from monitor APRIL1014:

```
CALL QSYS2.DUMP_PLAN_CACHE('SNAPSHOTS', 'APRIL2014');

CALL QSYS2.EXTRACT_STATEMENTS('SNAPSHOTS', 'APRIL2014', '*AUDIT',
  'AND QQC21 NOT IN
  (''CH'', ''CL'', ''CN'', ''DE'', ''DI'', ''DM'', ''HC'', ''HH'', ''JR'', ''FE'',
  ''PD'', ''PR'', ''PD''),
  ' ORDER BY QQSTIM DESC FETCH FIRST 100 ROWS ONLY ');
```

- Extract all the queries where the query took longer than one second:

```
CALL QSYS2.DUMP_PLAN_CACHE('SNAPSHOTS', 'APRIL2014');

CALL QSYS2.EXTRACT_STATEMENTS('SNAPSHOTS', 'APRIL2014',
  ADDITIONAL_SELECT_COLUMNS => 'DECIMAL(QQI6)/1000000.0 as Total_time,
  QVC102 as Current_User_Profile ',
  ADDITIONAL_PREDICATES => ' AND QQI6 > 1000000 ',
  ORDER_BY => ' ORDER BY QQI6 DESC ');
```

FIND_AND_CANCEL_QSQSRVR_SQL procedure

The FIND_AND_CANCEL_QSQSRVR_SQL procedure finds a set of jobs with SQL activity and safely cancels them.

►► FIND_AND_CANCEL_QSQSRVR_SQL — (— *job-name* —) ◄◄

The schema is QSYS2.

job-name A character string containing a qualified job name.

The FIND_AND_CANCEL_QSQSRVR_SQL procedure uses the FIND_QSQSRVR_JOBS and CANCEL_SQL procedures to determine the set of jobs that have SQL activity for the provided *job-name*. Each of these jobs is made a target of an SQL cancel request.

Example

Cancel all the QSQSRVR jobs used by a specific job.

```
CALL QSYS2.FIND_AND_CANCEL_QSQSRVR_SQL('564321/APPUSER/APPJOBNAME')
```

FIND_QSQRVR_JOBS procedure

The FIND_QSQRVR_JOBS procedure returns information about a QSQRVR job.

►► FIND_QSQRVR_JOBS — (— *job-name* —) ►►

The schema is QSYS2.

job-name A character string containing a qualified job name.

If the specified job is active and is set up to use SQL server mode, the procedure determines which QSQRVR jobs are being used by the application in the form of active SQL server mode connections. The procedure collects and returns work management, performance, and SQL information. It returns two SQL result sets, one containing summary information and one containing detailed SQL server mode job information.

Authorization: To invoke FIND_QSQRVR_JOBS you need *JOBCTL special authority, QIBM_DB_SQLADM function usage, or QIBM_DB_SYSMON function usage.

The results of the procedure call are saved in two temporary tables, QTEMP.QSQRVR_SUMMARY and QTEMP.QSQRVR_DETAIL. When called from within IBM i Navigator Run SQL Scripts, two results sets are displayed. When called from other interfaces, you need to query the temporary tables to see the data.

The result sets that are returned or the tables that are created contain the following columns:

Table 75. FIND_QSQRVR_JOBS result set 1

Column Name	System Column Name	Data Type	Description
SQL_IDENTITY	SQL_I00001	INTEGER	Unique identifier for this row.
NUMBER_OF_ACTIVE_JOBS	NUMJOBS	INTEGER	Number of QSQRVR jobs active for this job.
SERVER_MODE_JOB	SRVRJOB	CHAR(28)	The fully qualified QSQRVR job name for an active SQL Server Mode connection established by <i>job-name</i> .
SERVER_MODE_CONNECTING_JOB	CONNJOB	CHAR(28)	The fully qualified job name of the application job. This value matches what was input for <i>job_name</i> .
TOTAL_PROCESSING_TIME	TOTALCPU	BIGINT	The total amount of CPU time (in milliseconds) that has been used by all server jobs.
TEMP_MEG_STORAGE	TEMPMSTG	INTEGER	The total amount of auxiliary storage (in megabytes) that is currently allocated to all server jobs.
PAGE_FAULTS	FAULTS	BIGINT	The total number of times an active program referenced an address that was not in main storage for all server jobs.
IO_REQUESTS	IOREQS	BIGINT	The total number of auxiliary I/O requests performed by the job across all routing steps for all server jobs. This includes both database and non-database paging.

Table 76. FIND_QSQRVR_JOBS result set 2

Column Name	System Column Name	Data Type	Description
SQL_IDENTITY	SQL_I00001	INTEGER	Unique identifier for this row.
JOB_NAME	JOBNAME	CHAR(10)	Job name.
USER_NAME	USERNAME	CHAR(10)	User ID for the job.
JOB_NUMBER	JOBNUM	CHAR(6)	Job number.
JOB_INTERNAL_IDENTIFIER	JOBID	CHAR(16)	Internal identifier assigned to job.
CURRENT_USERNAME	CURRUSER	CHAR(10)	The user profile that the thread is currently running under.
SUBSYSTEM_DESCRIPTION_NAME	SBSNAME	CHAR(10)	Name of subsystem where job is running.
RUN_PRIORITY	PRIORITY	INTEGER	The highest run priority allowed for any thread within this job.
SYSTEM_POOL_IDENTIFIER	POOLID	INTEGER	The identifier of the system-related pool from which the job's main storage is allocated.

Table 76. FIND_QSQRVR_JOBS result set 2 (continued)

Column Name	System Column Name	Data Type	Description
TOTAL_PROCESSING_TIME	TOTALCPU	BIGINT	The amount of CPU time (in milliseconds) that has been currently used by this job.
PAGE_FAULTS	FAULTS	BIGINT	The number of times an active program referenced an address that was not in main storage during the current routing step of the specified job.
IO_REQUESTS	IOREQS	BIGINT	The number of auxiliary I/O requests performed by the job across all routing steps. This includes both database and non-database paging.
MEMORY_POOL_NAME	POOLNAME	CHAR(10)	The name of the memory pool in which the job started running.
TEMP_MEG_STORAGE	TEMPMSTG	INTEGER	The amount of auxiliary storage (in megabytes) that is currently allocated to this job.
TIME_SLICE	TSlice	INTEGER	The maximum amount of processor time (in milliseconds) given to each thread in this job before other threads in this job and in other jobs are given the opportunity to run.
DEFAULT_WAIT	DFTWAIT	INTEGER	The default maximum time (in seconds) that a thread in the job waits for a system instruction to acquire a resource.
SQL_APPLICATION_LIBRARY	SQLLIB	CHAR(10)	The library name for the SQL statement object.
SQL_APPLICATION_PROGRAM	SQLPGM	CHAR(10)	The program, service program, or package name of the object which contains the last SQL statement executed in the job.
SQL_APPLICATION_TYPE	APPTYPE	CHAR(10)	The object type.
SERVER_MODE_CONNECTING_JOB	CONNJOB	CHAR(28)	The qualified job name of the job which established the SQL Server Mode connection.
SERVER_MODE_CONNECTED_THREAD	CONNTHD	CHAR(10)	The thread identifier of the last thread to use this connection.
STATUS_OF_CURRENT_SQL_STMT	STMTSTAT	CHAR(10)	Status of the SQL statement. Values are ACTIVE or COMPLETED.
SQL_STATEMENT	SQLSTMT	VARCHAR(1000)	First 1000 characters of the SQL statement.

GENERATE_SQL procedure

The GENERATE_SQL procedure generates the SQL data definition language statements required to recreate a database object. The results are returned in the specified database source file member, source stream file, or as a result set.

The database source file member or integrated file system (IFS) stream file will contain the generated SQL statements. If the output source file is QTEMP/Q_GENSQL with a member name of Q_GENSQL, the source file is returned as a result set as well.

Authorization:

When writing to a database source physical file the user must have:

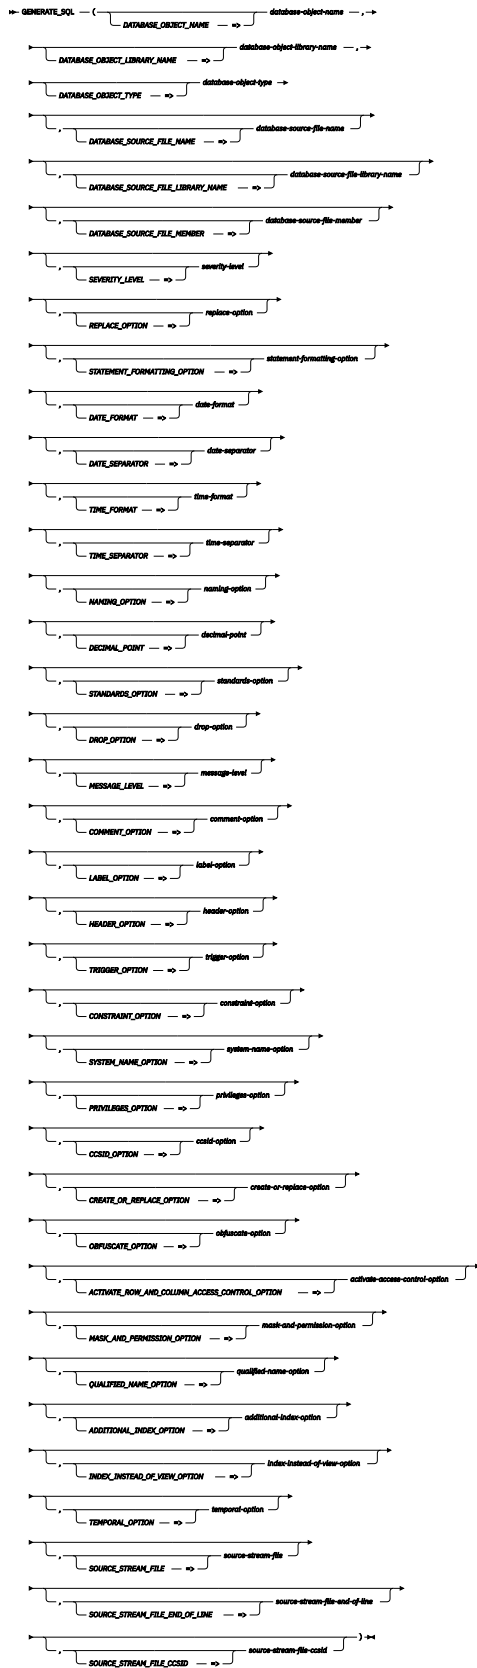
- *EXECUTE to the library containing the source physical file
- To add a member:
 - *OBJOPR and *ADD to the source physical file
- To replace a member:
 - *OBJOPR, *DELETE, *ADD, and either *OBJMGT or *OBJALTER to the source physical file

When writing to an IFS stream file:

- Execute (*X) data authority to each directory preceding the stream file being written and
- When the stream file exists:
 - Write (*W) data authority to the stream file
- When the stream file does not exist:
 - Write and Execute (*WX) authority to the parent directory of the stream file

To generate source for an object type, the following authorities are needed:

- TABLE, VIEW, CONSTRAINT, or TRIGGER
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *FILE object
- INDEX
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *FILE object, and
 - *OBJOPR to the base *FILE object
- MASK or PERMISSION
 - *EXECUTE and *OBJOPR to the library, and at least one of the following:
 - *OBJOPR to the *FILE object
 - QIBM_DB_SECADM function usage
- ALIAS
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *FILE object
- FUNCTION or PROCEDURE
 - *OBJOPR to the library
- TYPE
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *SQLUDT object
- SCHEMA
 - *OBJOPR and either *READ or *EXECUTE to the *LIB object
- SEQUENCE
 - *EXECUTE and *OBJOPR to the library, and
 - *USE to the *DTAARA object
- VARIABLE
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *SRVPGM object
- XSR
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *SQLXSR object



The schema is QSYS2.

database-object-name A character or graphic string expression that identifies the name of the database object for which DDL will be generated. Either the SQL name or the system name may be

specified. The name is case sensitive. Delimiters must not be specified. For example, a file with a name of "abc" must be specified as abc. A file with a name of ABC must be specified in upper case. If the object type is a FUNCTION or PROCEDURE, this name must be the specific name of the function or procedure. If TABLE or VIEW is specified for the object type, the object name may identify an alias. In this case, the object that the alias points to will be generated. A CREATE ALIAS statement will be generated only if ALIAS is specified for the object type.

A '%' wildcard character can be used to select multiple objects of the same type. For example, a name of 'TSTV%' will process all objects of *database-object-type* that start with the characters 'TSTV'

database-object-library-name A character or graphic string expression that identifies the name of the library containing the object for which DDL will be generated. Either the SQL name or the system name may be specified. The name is case sensitive. Delimiters must not be specified. This name is ignored if the specified object type is SCHEMA. A '%' wildcard character can be used to select multiple libraries.

database-object-type A character or graphic string expression that identifies the type of the database object or object attribute for which DDL is generated. You can use these special values for the object type:

ALIAS	The object is an SQL alias.
CONSTRAINT	The object attribute is a constraint.
FUNCTION	The object is an SQL function.
INDEX	The object is an SQL index.
MASK	The object is an SQL column mask.
PERMISSION	The object is an SQL row permission.
PROCEDURE	The object is an SQL procedure.
SCHEMA	The object is an SQL schema.
SEQUENCE	The object is an SQL sequence.
TABLE	The object is an SQL table or physical file.
TRIGGER	The object attribute is a trigger.
TYPE	The object is an SQL type.
VARIABLE	The object is an SQL global variable.
VIEW	The object is an SQL view or logical file.
XSR	The object is an XML schema repository object.

database-source-file-name A character or graphic string expression that identifies the name of the source file that contains the SQL statements generated by the procedure. The name must be a valid system name. The name is case sensitive. If delimiters are required for the name to be valid, they must be specified. For example, a file with a name of "abc" must be specified with the surrounding quotes. A file with a name of ABC must be specified in upper case. The record length of the specified source file must be greater than or equal to 92. Can contain the following special value:

***STMF** The output is to be written to *source-stream-file*.

If *database-source-file-name* is not specified, Q_GENSQL will be used.

database-source-file-library-name A character or graphic string expression that identifies the name of the library containing the source file that contains the SQL statements generated by the procedure. The name must be a valid system name. The name is case sensitive. If delimiters are required for

the name to be valid, they must be specified. You can use these special values for the library name:

***CURLIB** The job's current library
***LIBL** The library list

If *database-source-file-library-name* is not specified, QTEMP will be used.

database-source-file-member

A character or graphic string expression that identifies the name of the source file member that contains the SQL statements generated by the procedure. The name must be a valid system name. The name is case sensitive. If delimiters are required for the name to be valid, they must be specified. You can use these special values for the member name:

***FIRST** The first database physical file member found.
***LAST** The last database physical file member found.

If values are provided for *database-source-file-library-name*, *database-source-file-name*, and *database-source-file-member* the object must exist.

If *database-source-file-member* is not specified, Q_GENSQL will be used.

severity-level

The severity level at which the operation fails. If errors occur that have a severity level greater than this value, the operation ends. The valid values are in the range 0 through 39 inclusive. Any severity 40 error will cause the procedure to fail.

If *severity-level* is not specified, 39 will be used.

replace-option

The replace option for the database source file member or source stream file. The valid values are:

- 0** The resulting SQL statements are appended to the end of the database source file member or source stream file.
- 1** The database source file member or source stream file is cleared prior to adding the resulting SQL statements. If this option is chosen, the clear might happen even if an error is returned from the procedure.

If *replace-option* is not specified, 1 will be used.

statement-formatting-option

The formatting option used in the generated SQL statements. The valid values are:

- 0** No additional formatting characters are added to the generated SQL statements.
- 1** Additional end-of-line characters and tab characters are added to the generated SQL statements.

If *statement-formatting-option* is not specified, 1 will be used.

date-format

The date format used for date constants in a generated SQL CREATE TABLE statement. The date format may not apply to date constants that are in ISO, EUR, USA, or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION.

If *date-format* is not specified, ISO will be used.

date-separator

The date separator used for date constants in a generated SQL CREATE TABLE statement. The date separator may not apply to date constants that are in ISO, EUR, USA, or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION statement.

If *date-separator* is not specified, - will be used.

time-format

The format used for time constants in a generated SQL CREATE TABLE statement. The time format may not apply to time constants that are in ISO, EUR, USA, or JIS format in

a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION statement.

If *time-format* is not specified, ISO will be used.

time-separator

The time separator used for time constants in a generated SQL CREATE TABLE statement. The time separator may not apply to time constants that are in ISO, EUR, USA, or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION statement.

If *time-separator* is not specified, . will be used.

naming-option

The naming convention used for qualified names in the generated SQL statements. The valid values are:

SQL schema.table syntax

SYS library/file syntax

If *naming-option* is not specified, SQL will be used.

decimal-point The decimal point used for numeric constants. The valid values are:

. Period separator

, Comma separator

If *decimal-point* is not specified, . will be used.

standards-option

The standards option specifies whether the generated SQL statements should contain Db2 for i extensions or whether the statements should conform to the Db2 family SQL or to the ANS and ISO SQL standards. The valid values are:

0 Db2 for i extensions may be generated in SQL statements.

1 The generated SQL statements must conform to SQL statements common to the Db2 family.

2 The generated SQL statements must conform to the ANSI and ISO SQL standards.

If *standards-option* is not specified, 0 will be used.

drop-option

The drop option specifies whether DROP (or ALTER) SQL statements should be generated prior to the CREATE statement to drop the specified object. The valid values are:

0 DROP statements should not be generated.

1 DROP statements should be generated.

If *drop-option* is not specified, 0 will be used.

message-level

The severity level at which the messages are generated. If errors occur that have a severity level greater than this value, a message is generated in the output. The valid values are in the range 0 through 39 inclusive. The message level must be less than or equal to the severity level.

If *message-level* is not specified, 0 will be used.

comment-option

The comment option specifies whether COMMENT SQL statements should be generated if a comment exists on the specified database object. If comments are not supported by the specified database object, the comment option is ignored. The valid values are:

0 COMMENT SQL statements should not be generated.

1 COMMENT SQL statements should be generated. If the specified database object type is a table or view, COMMENT SQL statements will also be generated for columns of the table or view.

- 2** COMMENT SQL statements should be generated. If the specified database object has no comment and its type is INDEX, SEQUENCE, TABLE, TYPE, VARIABLE, VIEW, or XSR, the system object text will be used for the comment.
- If the specified database object type is a table or view, COMMENT SQL statements will also be generated for columns of the table or view.

If *comment-option* is not specified, **1** will be used.

label-option The label option specifies whether LABEL SQL statements should be generated if a label exists on the specified database object. If labels are not supported by the specified database object, the label option is ignored. The valid values are:

- 0** LABEL SQL statements should not be generated.
- 1** LABEL SQL statements should be generated. If the specified database object type is a table or view, LABEL SQL statements will also be generated for columns of the table or view.

If *label-option* is not specified, **1** will be used.

header-option The header option specifies whether a header should be generated prior to the CREATE statement. The header consists of comments that describe the version, date and time, the relational database, and some of the options used to generate the SQL statements. The valid values are:

- 0** A header should not be generated.
- 1** A header should be generated.

If *header-option* is not specified, **1** will be used.

trigger-option The trigger option specifies whether triggers should be generated when the object type is a TABLE or VIEW. The valid values are:

- 0** Triggers should not be generated.
- 1** Triggers should be generated.

If *trigger-option* is not specified, **1** will be used.

constraint-option The constraint option specifies whether constraints should be generated when the object type is a TABLE. The valid values are:

- 0** Constraints should not be generated.
- 1** Constraints should be generated.
- 2** Constraints should be generated as part of the CREATE TABLE statement.

If *constraint-option* is not specified, **1** will be used.

system-name-option The system name option specifies whether a FOR SYSTEM NAME clause should be generated for the system name when it is different from the SQL name and the object type is an INDEX, TABLE, VIEW, SEQUENCE, or VARIABLE. The valid values are:

- 0** A FOR SYSTEM NAME clause should not be generated.
- 1** A FOR SYSTEM NAME clause should be generated.

If *system-name-option* is not specified, **1** will be used.

privileges-option The privileges option specifies whether GRANT SQL statements should be generated on the specified database object. If privileges are not supported by the specified database object, the privileges option is ignored. The valid values are:

- 0 GRANT SQL statements should not be generated.
- 1 GRANT SQL statements should be generated.

If *privileges-option* is not specified, 1 will be used.

ccsid-option The CCSID option specifies whether the CCSID attribute should be generated for column definitions when the object type is a TABLE. The valid values are:

- 0 CCSID attribute should not be generated.
- 1 CCSID attribute should be generated.

If *ccsid-option* is not specified, 1 will be used.

create-or-replace-option The create or replace option specifies whether CREATE OR REPLACE should be generated for the specified database object on the CREATE statement. This option is ignored if the specified database object does not support CREATE OR REPLACE. The valid values are:

- 0 CREATE OR REPLACE should not be generated.
- 1 CREATE OR REPLACE should be generated.

If *create-or-replace-option* is not specified, 0 will be used.

obfuscate-option The obfuscate option specifies whether an obfuscated SQL statement should be returned for SQL functions, SQL procedures, or SQL triggers that were not created using obfuscated statements. This option is ignored if the standards option is not '0'. This option is also ignored if the object is not an SQL function, procedure, or trigger. This option is ignored if the object is already obfuscated. Setting Obfuscate option = 0 cannot be used as a means of obtaining the unobfuscated SQL statement for an obfuscated object. The valid values are:

- 0 An obfuscated statement should not be generated.
- 1 An obfuscated statement should be generated for SQL functions, SQL procedures, or SQL triggers.

If *obfuscate-option* is not specified, 0 will be used.

activate-access-control-option The activate row and column access control option specifies whether an ALTER TABLE to activate row and column access control should be generated when the object type is a TABLE. This option is ignored if the standards option is not '0' or '1'. The valid values are:

- 0 Activate row and column access control should not be generated.
- 1 Activate row and column access control should be generated.

If *activate-access-control-option* is not specified, 1 will be used.

mask-and-permission-option The mask and permission option specifies whether row permissions and column masks should be generated when the object type is a TABLE. This option is ignored if the standards option is not '0' or '1'. The valid values are:

- 0 Permissions and masks should not be generated.
- 1 Permissions and masks should be generated.

If *mask-and-permission-option* is not specified, 1 will be used.

qualified-name-option The qualified name option specifies whether qualified or unqualified names should be generated for the specified database object. The valid values are:

- 0 Qualified object names should be generated. Unqualified names within the body of SQL routines will remain unqualified.
- 1 Unqualified object names should be generated when the a library is found which matches the database object library name. Any SQL object or column reference that is RDB qualified will be generated in its fully qualified form. For example, rdb-name.schema-name.table-name and rdb-name.schema-name.table-name.column-name references will retain their full qualification.

If *qualified-name-option* is not specified, 0 will be used.

additional-index-option

The additional index option specifies whether additional CREATE INDEX statements will be generated for DDS-created keyed physical or logical files. The valid values are:

- 0 Additional CREATE INDEX statements will not be generated.
- 1 An additional CREATE INDEX statement will be generated that matches the index for a DDS-created keyed physical file. If the physical file has a PRIMARY KEY constraint, a CREATE INDEX statement is not generated.

An additional CREATE INDEX statement will be generated that matches the index for a DDS-created keyed logical file. If a value of '1' is specified for the index instead of view option, an additional CREATE INDEX statement is not generated. Additional CREATE INDEX statements will also be generated that match the join indexes of a DDS-created join logical file.

If *additional-index-option* is not specified, 0 will be used.

index-instead-of-view-option

The index instead of view option specifies whether a CREATE INDEX or CREATE VIEW statement will be generated for a DDS-created keyed logical file. The valid values are:

- 0 A CREATE VIEW statement will be generated.
- 1 A CREATE INDEX statement will be generated that matches the index for a DDS-created keyed logical file.

If *index-instead-of-view-option* is not specified, 0 will be used.

temporal-option

The temporal option specifies whether a CREATE TABLE and an ALTER TABLE statement will be generated when the object type is a TABLE and the table is defined as a temporal table. This option is ignored if the object is not a temporal table or if the standards option is not '0' or '1'. The valid values are:

- 0 A CREATE TABLE statement will be generated.
- 1 A CREATE TABLE statement will be generated and an ALTER TABLE statement will be generated to add versioning.
- 2 Only an ALTER TABLE statement will be generated to add versioning.

If *temporal-option* is not specified, 0 will be used.

source-stream-file

A character or graphic string expression that identifies the source stream file that contains the SQL statements generated by the procedure. This parameter is ignored unless *database-source-file-name* has a value of *STMF.

If the file does not exist, it will be created using the CCSID specified by *source-stream-file-ccsid*. As lines are written to the file, *source-stream-file-end-of-line* determines the end of line sequence, if any, to be appended to each line.

When *source-stream-file* is used, values provided for *database-source-file-library-name* and *database-source-file-member* are ignored.

Writing to the QSYS.LIB file system is not supported.

source-stream-file-end-of-line

A character or graphic string expression that defines the end of line character(s) which will be appended to the end of each line when writing to *source-stream-file*. The carriage return character is always X'0D'. Based on the CCSID of the source stream file, the line feed character is X'25' for an EBCDIC CCSID and X'0A' for ASCII and UTF-8 CCSIDs. The valid values are:

- CR** A carriage return is appended.
- CRLF** A carriage return and line feed are appended.
- LF** A line feed is appended.
- LFCR** A line feed and carriage return are appended.

If *source-stream-file-end-of-line* is not specified, CRLF will be used.

This parameter is ignored when writing to a source file.

source-stream-file-ccsid

An integer value that defines the CCSID to use if a new source stream file is created. If *source-stream-file* is specified and the source stream file does not exist, it will be created with this CCSID value. A value of 0 indicates that the default job CCSID is to be used. A CCSID of 65535 cannot be specified.

If *source-stream-file-ccsid* is not specified, 0 will be used.

This parameter is ignored when writing to a source file.

Examples

- Generate DDL for all tables in a schema and return the source as a result set.

```
CALL QSYS2.GENERATE_SQL('%', 'SAMPLE_CORPDB', 'TABLE', REPLACE_OPTION => '0');
```

- Generate DDL for all indexes starting with 'X' within the SAMPLE_CORPDB schema, place the output in a file named DDLSOURCE/GENFILE member INDEXSRC.

```
CALL QSYS2.GENERATE_SQL('X%', 'SAMPLE_CORPDB', 'INDEX',  
    'GENFILE', 'DDLSOURCE', 'INDEXSRC',  
    REPLACE_OPTION => '0');
```

- Generate DDL for a single table and include the constraints within a CREATE OR REPLACE TABLE statement.

```
CALL QSYS2.GENERATE_SQL('EMPLOYEE', 'SAMPLE_CORPDB', 'TABLE',  
    'GENFILE', 'DDLSOURCE', 'MASTERSRC',  
    CREATE_OR_REPLACE_OPTION => '1',  
    CONSTRAINT_OPTION => '2');
```

- Generate DDL for MYLIB.MYTABLE, placing the output in source stream file /usr/mytable_ddl.

```
CALL QSYS2.GENERATE_SQL(DATABASE_OBJECT_NAME => 'MYTABLE',  
    DATABASE_OBJECT_LIBRARY_NAME => 'MYLIB',  
    DATABASE_OBJECT_TYPE => 'TABLE',  
    DATABASE_SOURCE_FILE_NAME => '*STMF',  
    SOURCE_STREAM_FILE => '/usr/mytable_ddl');
```

GENERATE_SQL_OBJECTS procedure

The GENERATE_SQL_OBJECTS procedure generates the SQL data definition language (DDL) statements required to recreate a set of database objects. The results are returned in the specified database source file member, source stream file, or as a result set. The procedure will generate the objects so that dependent objects are generated after depended on objects.

The database source file member or integrated file system (IFS) stream file will contain the generated SQL statements. If the output source file is QTEMP/Q_GENSQOBJ, the source file is returned as a result set as well.

Authorization:

When writing to a database source physical file the user must have:

- *EXECUTE to the library containing the source physical file
- To add a member:
 - *OBJOPR and *ADD to the source physical file
- To replace a member:
 - *OBJOPR, *DELETE, *ADD, and either *OBJMGT or *OBJALTER to the source physical file


When writing to an IFS stream file:

- Execute (*X) data authority to each directory preceding the stream file being written and
- When the stream file exists:
 - Write (*W) data authority to the stream file
- When the stream file does not exist:
 - Write and Execute (*WX) authority to the parent directory of the stream file

To generate source for an object type, the following authorities are needed:

- TABLE, VIEW, CONSTRAINT, or TRIGGER
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *FILE object
- INDEX
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *FILE object, and
 - *OBJOPR to the base *FILE object
- MASK or PERMISSION
 - *EXECUTE and *OBJOPR to the library, and at least one of the following:
 - *OBJOPR to the *FILE object
 - QIBM_DB_SECADM function usage
- ALIAS
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *FILE object
- FUNCTION or PROCEDURE
 - *OBJOPR to the library
- TYPE
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *SQLUDT object
- SCHEMA
 - *OBJOPR and either *READ or *EXECUTE to the *LIB object
- SEQUENCE
 - *EXECUTE and *OBJOPR to the library, and
 - *USE to the *DTAARA object

- VARIABLE
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *SRVPGM object
- XSR
 - *EXECUTE and *OBJOPR to the library, and
 - *OBJOPR to the *SQLXSR object



The schema is QSYS2.

system-table-name A character or graphic string expression that identifies the name of a table that contains the names and types of the database objects for which DDL will be generated. The

system name of the table must be specified. The name is case sensitive. Delimiters must be specified if they are required. For example, a file with a name of "abc" must be specified as "abc". A file with a name of ABC must be specified in upper case. Wildcard characters are not supported.

The specified table must contain three columns that contain the names and types of the objects for which DDL will be generated. The column names of the table must be OBJECT_SCHEMA, OBJECT_NAME, and SQL_OBJECT_TYPE, in that order.

For example, create the table like this:

```
CREATE TABLE QTEMP.INORDER (OBJECT_SCHEMA VARCHAR(258),
                             OBJECT_NAME VARCHAR(258),
                             SQL_OBJECT_TYPE CHAR(10));
```

The contents of the columns must be specified according to the following rules for the corresponding parameters in the QSQGNDDL API. Each row in the table must identify an object that is distinct from every other object in the table.

OBJECT_SCHEMA Identifies the schema name of an object for which DDL will be generated. The name must be delimited if delimiters are required in an SQL statement. This name is ignored if the specified object type is SCHEMA. *LIBL and *CURLIB are not allowed.

OBJECT_NAME Identifies the name of an object for which DDL will be generated. The name must be delimited if delimiters are required in an SQL statement. If the object type is a FUNCTION or PROCEDURE, this name must be the specific name of the function or procedure. If TABLE or VIEW is specified for the object type, the object name must not identify an alias.

SQL_OBJECT_TYPE Identifies the SQL object type of an object for which DDL will be generated.

ALIAS	The object is an SQL alias.
CONSTRAINT	The object attribute is a constraint.
FUNCTION	The object is an SQL function.
INDEX	The object is an SQL index.
MASK	The object is an SQL column mask.
PERMISSION	The object is an SQL row permission.
PROCEDURE	The object is an SQL procedure.
SCHEMA	The object is an SQL schema.
SEQUENCE	The object is an SQL sequence.
TABLE	The object is an SQL table or physical file.
TRIGGER	The object attribute is a trigger.
TYPE	The object is an SQL type.
VARIABLE	The object is an SQL global variable.
VIEW	The object is an SQL view or logical file.
XSR	The object is an XML schema repository object.

system-table-schema

A character or graphic string expression that identifies the name of the library of the table that contains the names and types of the database objects for which DDL will be generated. The system name of the schema must be specified. The name is case sensitive. Delimiters must be specified if they are required. For example, a schema with

a name of "lib1" must be specified as "lib1". A schema with a name of LIB1 must be specified in upper case. Wildcard characters are not supported. *LIBL and *CURLIB are not allowed.

The default is QTEMP.

database-source-file-name

A character or graphic string expression that identifies the name of the source file that contains the SQL statements generated by the procedure. The name must be a valid system name. The name is case sensitive. If delimiters are required for the name to be valid, they must be specified. For example, a file with a name of "abc" must be specified with the surrounding quotes. A file with a name of ABC must be specified in upper case. The record length of the specified source file must be greater than or equal to 92. Can contain the following special value:

***STMF** The output is to be written to *source-stream-file*.

If *database-source-file-name* is not specified, Q_GENSQOBJ will be used.

database-source-file-library-name

A character or graphic string expression that identifies the name of the library containing the source file that contains the SQL statements generated by the procedure. The name must be a valid system name. The name is case sensitive. If delimiters are required for the name to be valid, they must be specified.

If *database-source-file-library-name* is not specified, QTEMP will be used.

database-source-file-member

A character or graphic string expression that identifies the name of the source file member that contains the SQL statements generated by the procedure. The name must be a valid system name. The name is case sensitive. If delimiters are required for the name to be valid, they must be specified.

If values are provided for *database-source-file-library-name*, *database-source-file-name*, and *database-source-file-member* the object must exist.

If *database-source-file-member* is not specified, Q_GENSQOBJ will be used.

severity-level

The severity level at which the operation fails. If errors occur that have a severity level greater than this value, the operation ends. The valid values are in the range 0 through 39 inclusive. Any severity 40 error will cause the procedure to fail.

If *severity-level* is not specified, 39 will be used.

replace-option

The replace option for the database source file member or source stream file. The valid values are:

0 The resulting SQL statements are appended to the end of the database source file member or source stream file.

1 The database source file member or source stream file is cleared prior to adding the resulting SQL statements. If this option is chosen, the clear might happen even if an error is returned from the procedure.

If *replace-option* is not specified, 1 will be used.

statement-formatting-option

The formatting option used in the generated SQL statements. The valid values are:

0 No additional formatting characters are added to the generated SQL statements.

1 Additional end-of-line characters and tab characters are added to the generated SQL statements.

If *statement-formatting-option* is not specified, 1 will be used.

date-format

The date format used for date constants in a generated SQL CREATE TABLE statement. The date format may not apply to date constants that are in ISO, EUR, USA,

or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION.

If *date-format* is not specified, ISO will be used.

date-separator

The date separator used for date constants in a generated SQL CREATE TABLE statement. The date separator may not apply to date constants that are in ISO, EUR, USA, or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION statement.

If *date-separator* is not specified, - will be used.

time-format

The format used for time constants in a generated SQL CREATE TABLE statement. The time format may not apply to time constants that are in ISO, EUR, USA, or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION statement.

If *time-format* is not specified, ISO will be used.

time-separator

The time separator used for time constants in a generated SQL CREATE TABLE statement. The time separator may not apply to time constants that are in ISO, EUR, USA, or JIS format in a CREATE VIEW, CREATE TRIGGER, CREATE FUNCTION, CREATE PROCEDURE, CREATE MASK, or CREATE PERMISSION statement.

If *time-separator* is not specified, . will be used.

naming-option

The naming convention used for qualified names in the generated SQL statements. The valid values are:

SQL schema.table syntax

SYS library/file syntax

If *naming-option* is not specified, SQL will be used.

decimal-point

The decimal point used for numeric constants. The valid values are:

. Period separator

, Comma separator

If *decimal-point* is not specified, . will be used.

standards-option

The standards option specifies whether the generated SQL statements should contain Db2 for i extensions or whether the statements should conform to the Db2 family SQL or to the ANS and ISO SQL standards. The valid values are:

0 Db2 for i extensions may be generated in SQL statements.

1 The generated SQL statements must conform to SQL statements common to the Db2 family.

2 The generated SQL statements must conform to the ANSI and ISO SQL standards.

If *standards-option* is not specified, 0 will be used.

drop-option

The drop option specifies whether DROP (or ALTER) SQL statements should be generated prior to the CREATE statement to drop the specified object. The valid values are:

0 DROP statements should not be generated.

1 DROP statements should be generated.

If *drop-option* is not specified, 0 will be used.

message-level

The severity level at which the messages are generated. If errors occur that have a severity level greater than this value, a message is generated in the output. The valid

values are in the range 0 through 39 inclusive. The message level must be less than or equal to the severity level.

If *message-level* is not specified, 0 will be used.

comment-option

The comment option specifies whether COMMENT SQL statements should be generated if a comment exists on the specified database object. If comments are not supported by the specified database object, the comment option is ignored. The valid values are:

- 0** COMMENT SQL statements should not be generated.
- 1** COMMENT SQL statements should be generated. If the specified database object type is a table or view, COMMENT SQL statements will also be generated for columns of the table or view.
- 2** COMMENT SQL statements should be generated. If the specified database object has no comment and its type is INDEX, SEQUENCE, TABLE, TYPE, VARIABLE, VIEW, or XSR, the system object text will be used for the comment.
If the specified database object type is a table or view, COMMENT SQL statements will also be generated for columns of the table or view.

If *comment-option* is not specified, 1 will be used.

label-option

The label option specifies whether LABEL SQL statements should be generated if a label exists on the specified database object. If labels are not supported by the specified database object, the label option is ignored. The valid values are:

- 0** LABEL SQL statements should not be generated.
- 1** LABEL SQL statements should be generated. If the specified database object type is a table or view, LABEL SQL statements will also be generated for columns of the table or view.

If *label-option* is not specified, 1 will be used.

header-option

The header option specifies whether a header should be generated prior to the first generated statement. The header consists of comments that describe the version, date and time, the relational database, and some of the options used to generate the SQL statements. The valid values are:

- 0** A header should not be generated.
- 1** A header should be generated.

If *header-option* is not specified, 1 will be used.

trigger-option

The trigger option specifies whether triggers should be generated when the object type is a TABLE or VIEW. The valid values are:

- 0** Triggers should not be generated.
- 1** Triggers should be generated.

If *trigger-option* is not specified, 1 will be used.

constraint-option

The constraint option specifies whether constraints should be generated when the object type is a TABLE. The valid values are:

- 0** Constraints should not be generated.
- 1** Constraints should be generated.
- 2** Constraints should be generated as part of the CREATE TABLE statement.

If *constraint-option* is not specified, 1 will be used.

system-name-option The system name option specifies whether a FOR SYSTEM NAME clause should be generated for the system name when it is different from the SQL name and the object type is an INDEX, TABLE, VIEW, SEQUENCE, or VARIABLE. The valid values are:

- 0 A FOR SYSTEM NAME clause should not be generated.
- 1 A FOR SYSTEM NAME clause should be generated.

If *system-name-option* is not specified, 1 will be used.

privileges-option The privileges option specifies whether GRANT SQL statements should be generated on the specified database object. If privileges are not supported by the specified database object, the privileges option is ignored. The valid values are:

- 0 GRANT SQL statements should not be generated.
- 1 GRANT SQL statements should be generated.

If *privileges-option* is not specified, 1 will be used.

ccsid-option The CCSID option specifies whether the CCSID attribute should be generated for column definitions when the object type is a TABLE. The valid values are:

- 0 CCSID attribute should not be generated.
- 1 CCSID attribute should be generated.

If *ccsid-option* is not specified, 1 will be used.

create-or-replace-option The create or replace option specifies whether CREATE OR REPLACE should be generated for the specified database object on the CREATE statement. This option is ignored if the specified database object does not support CREATE OR REPLACE. The valid values are:

- 0 CREATE OR REPLACE should not be generated.
- 1 CREATE OR REPLACE should be generated.

If *create-or-replace-option* is not specified, 0 will be used.

obfuscate-option The obfuscate option specifies whether an obfuscated SQL statement should be returned for SQL functions, SQL procedures, or SQL triggers that were not created using obfuscated statements. This option is ignored if the standards option is not '0'. This option is also ignored if the object is not an SQL function, procedure, or trigger. This option is ignored if the object is already obfuscated. Setting Obfuscate option = 0 cannot be used as a means of obtaining the unobfuscated SQL statement for an obfuscated object. The valid values are:

- 0 An obfuscated statement should not be generated.
- 1 An obfuscated statement should be generated for SQL functions, SQL procedures, or SQL triggers.

If *obfuscate-option* is not specified, 0 will be used.

activate-access-control-option The activate row and column access control option specifies whether an ALTER TABLE to activate row and column access control should be generated when the object type is a TABLE. This option is ignored if the standards option is not '0' or '1'. The valid values are:

- 0 Activate row and column access control should not be generated.
- 1 Activate row and column access control should be generated.

If *activate-access-control-option* is not specified, 1 will be used.

mask-and-permission-option

The mask and permission option specifies whether row permissions and column masks should be generated when the object type is a TABLE. This option is ignored if the standards option is not '0' or '1'. The valid values are:

- 0 Permissions and masks should not be generated.
- 1 Permissions and masks should be generated.

If *mask-and-permission-option* is not specified, 1 will be used.

qualified-name-option

The qualified name option specifies whether qualified or unqualified names should be generated for the specified database object. The valid values are:

- 0 Qualified object names should be generated. Unqualified names within the body of SQL routines will remain unqualified.
- 1 Unqualified object names should be generated when the a library is found which matches the database object library name. Any SQL object or column reference that is RDB qualified will be generated in its fully qualified form. For example, rdb-name.schema-name.table-name and rdb-name.schema-name.table-name.column-name references will retain their full qualification.

If *qualified-name-option* is not specified, 0 will be used.

additional-index-option

The additional index option specifies whether additional CREATE INDEX statements will be generated for DDS-created keyed physical or logical files. The valid values are:

- 0 Additional CREATE INDEX statements will not be generated.
- 1 An additional CREATE INDEX statement will be generated that matches the index for a DDS-created keyed physical file. If the physical file has a PRIMARY KEY constraint, a CREATE INDEX statement is not generated.
An additional CREATE INDEX statement will be generated that matches the index for a DDS-created keyed logical file. If a value of '1' is specified for the index instead of view option, an additional CREATE INDEX statement is not generated. Additional CREATE INDEX statements will also be generated that match the join indexes of a DDS-created join logical file.

If *additional-index-option* is not specified, 0 will be used.

index-instead-of-view-option

The index instead of view option specifies whether a CREATE INDEX or CREATE VIEW statement will be generated for a DDS-created keyed logical file. The valid values are:

- 0 A CREATE VIEW statement will be generated.
- 1 A CREATE INDEX statement will be generated that matches the index for a DDS-created keyed logical file.

If *index-instead-of-view-option* is not specified, 0 will be used.

temporal-option

The temporal option specifies whether a CREATE TABLE and an ALTER TABLE statement will be generated when the object type is a TABLE and the table is defined as a temporal table. This option is ignored if the object is not a temporal table or if the standards option is not '0' or '1'. The valid values are:

- 0 A CREATE TABLE statement will be generated.
- 1 A CREATE TABLE statement will be generated and an ALTER TABLE statement will be generated to add versioning.
- 2 Only an ALTER TABLE statement will be generated to add versioning.

If *temporal-option* is not specified, 0 will be used.

source-stream-file

A character or graphic string expression that identifies the source stream file that contains the SQL statements generated by the procedure. This parameter is ignored unless *database-source-file-name* has a value of *STMF.

If the file does not exist, it will be created using the CCSID specified by *source-stream-file-ccsid*. As lines are written to the file, *source-stream-file-end-of-line* determines the end of line sequence, if any, to be appended to each line.

When *source-stream-file* is used, values provided for *database-source-file-library-name* and *database-source-file-member* are ignored.

Writing to the QSYS.LIB file system is not supported.

source-stream-file-end-of-line

A character or graphic string expression that defines the end of line character(s) which will be appended to the end of each line when writing to *source-stream-file*. The carriage return character is always X'0D'. Based on the CCSID of the source stream file, the line feed character is X'25' for an EBCDIC CCSID and X'0A' for ASCII and UTF-8 CCSIDs. The valid values are:

- | | |
|-------------|---|
| CR | A carriage return is appended. |
| CRLF | A carriage return and line feed are appended. |
| LF | A line feed is appended. |
| LFCR | A line feed and carriage return are appended. |

If *source-stream-file-end-of-line* is not specified, CRLF will be used.

This parameter is ignored when writing to a source file.

source-stream-file-ccsid

An integer value that defines the CCSID to use if a new source stream file is created. If *source-stream-file* is specified and the source stream file does not exist, it will be created with this CCSID value. A value of 0 indicates that the default job CCSID is to be used. A CCSID of 65535 cannot be specified.

If *source-stream-file-ccsid* is not specified, 0 will be used.

This parameter is ignored when writing to a source file.

Notes

- If an error occurs while generating the DDL for an object, the source file or source stream file will contain the error and processing will continue to the next object. After processing the last object, a warning SQLSTATE '01H52' will be returned.
- Objects are generated in the following order:
 1. Schemas
 2. Types
 3. Sequences
 4. Aliases
 5. Non-MQT tables and any constraints and indexes on those tables
 6. Functions
 7. Procedures
 8. Variables
 9. Views, DDS-created logical files and MQTs and any constraints and indexes on those tables
 10. Triggers
 11. Masks
 12. Permissions

13. XSR objects

Restrictions

- One use of this procedure is to create a clone of a set of objects in another library by using `QUALIFIED_NAME_OPTION=>1`, setting the current schema and path, and then running the generated script.
 - If a depended on object is not included in the list of objects for which DDL will be generated, errors may occur when attempting to run the generated script. For example, if view V1 is based on table T1, but only V1 is specified, the attempt to run the generated script will fail because T1 was not included.
 - The QSQGNDDL API, on which this procedure is based, generates a qualified name in some cases. Thus, it may be necessary to make minor modifications to the script prior to running it. For more information see the Qualified name option parameter in [Generate Data Definition Language \(QSQGNDDL\) API](#).
- A function or procedure that has a parameter with a DEFAULT clause that references a variable, view, or MQT will not create when running the generated script. This is because variables, views, and MQTs are generated after functions and procedures. Note that references to variables, views, and MQTs within the body of function or procedure are soft dependencies and will not prevent the create.
- A variable that contains a DEFAULT clause that references a view or MQT will not create when running the generated script. This is because views and MQTs are generated after variables.

Examples

- Generate ordered DDL for the objects listed in the QTEMP.INORDER file.

```
CALL QSYS2.GENERATE_SQL_OBJECTS('INORDER', 'QTEMP');
```

RELATED_OBJECTS table function

The RELATED_OBJECTS table function returns a list of all objects that depend on the specified database file, either directly or indirectly.

The list contains all objects that are directly dependent on the input database file. Each identified view and global variable is processed recursively to obtain its dependents until no more dependents are found.

If the input file is an SQL alias, a program described file, or the file is not found, no rows are returned.

The dependency information is collected from the database cross reference tables and the SQL catalog. Some dependencies might not be returned based on how the dependent object was specified in the original statement. For example, if a CREATE FUNCTION statement references an unqualified table, the dependency will not be complete in the SQL catalog view (SYSROUTINEDEP) and subsequently will not be returned by this function.

Authorization: None required.

```
➤ RELATED_OBJECTS ( ( library-name , ➤  
  LIBRARY_NAME => )  
  ( file-name ) ➤  
  FILE_NAME => ) ➤
```

The schema is SYSTOOLS.

library-name A character or graphic string expression that identifies the library that contains *file-name*. It must exist on the current server.

file-name A character or graphic string expression that identifies the database file to list related objects for. This must be the system name of the file.

The result of the function is a table containing rows with the format shown in the following table. All columns are nullable.

Table 77. RELATED_OBJECTS table function

Column Name	Data Type	Description
SOURCE_SCHEMA_NAME	VARCHAR(128)	The name of the schema that contains SOURCE_SQL_NAME.
SOURCE_SQL_NAME	VARCHAR(128)	The name of the object that this row is dependent on.
SQL_OBJECT_TYPE	VARCHAR(24)	The SQL type of the dependent object. <ul style="list-style-type: none"> ALIAS An SQL alias defined for SOURCE_SQL_NAME FOREIGN KEY The child table in a referential constraint with SOURCE_SQL_NAME FUNCTION A function that references SOURCE_SQL_NAME HISTORY TABLE A history table associated with the temporal table defined by SOURCE_SQL_NAME INDEX An SQL index defined on SOURCE_SQL_NAME KEYED LOGICAL FILE A native keyed logical file defined on SOURCE_SQL_NAME LOGICAL FILE A native logical file defined on SOURCE_SQL_NAME MASK A column mask on a column in SOURCE_SQL_NAME or that references SOURCE_SQL_NAME in its definition MATERIALIZED QUERY TABLE A materialized query table that references SOURCE_SQL_NAME PERMISSION A row permission that is defined on SOURCE_SQL_NAME or that references SOURCE_SQL_NAME in its definition PROCEDURE A procedure that references SOURCE_SQL_NAME TEXT INDEX An OmniFind text index built over SOURCE_SQL_NAME TRIGGER An SQL trigger that is defined on SOURCE_SQL_NAME or that references SOURCE_SQL_NAME in its definition VARIABLE A global variable that references SOURCE_SQL_NAME in its default expression VIEW A view that references SOURCE_SQL_NAME in its definition XML SCHEMA An XML schema that uses SOURCE_SQL_NAME
SCHEMA_NAME	VARCHAR(128)	The SQL schema containing SQL_NAME.
SQL_NAME	VARCHAR(128)	The SQL name of the dependent object. <ul style="list-style-type: none"> • When SQL_OBJECT_TYPE is FOREIGN KEY, this is the name of the child table
LIBRARY_NAME	VARCHAR(10)	The system library name of the dependent object. <ul style="list-style-type: none"> • When SQL_OBJECT_TYPE is TRIGGER, the library containing the trigger program <p>Contains the null value when SQL_OBJECT_TYPE is FOREIGN KEY, FUNCTION, PROCEDURE, and TEXT INDEX.</p>

Table 77. RELATED_OBJECTS table function (continued)

Column Name	Data Type	Description
SYSTEM_NAME	VARCHAR(279)	The related system name of the dependent object. <ul style="list-style-type: none"> When SQL_OBJECT_TYPE is FOREIGN KEY, the constraint name When SQL_OBJECT_TYPE is FUNCTION or PROCEDURE, the external name When SQL_OBJECT_TYPE is MASK or PERMISSION, the mask or permission name When SQL_OBJECT_TYPE is TRIGGER, the name of the trigger program Contains the null value when SQL_OBJECT_TYPE is TEXT INDEX.
OBJECT_OWNER	VARCHAR(10)	The owner of the object. When SQL_OBJECT_TYPE is FUNCTION, MASK, PERMISSION, PROCEDURE, TRIGGER, or VARIABLE this is the object's definer. Contains the null value when SQL_OBJECT_TYPE is FOREIGN KEY or TEXT INDEX.
LONG_COMMENT	VARGRAPHIC(2000) CCSID 1200	The SQL long comment for the object. Contains the null value when the comment is not available.
OBJECT_TEXT	VARGRAPHIC(50) CCSID 1200	The system text description for the object. Contains the null value when the text is not available.
LAST_ALTERED	TIMESTAMP	The timestamp when the object was last altered. For objects that cannot be altered, this is the create timestamp. Contains the null value when SQL_OBJECT_TYPE is FOREIGN KEY.

Examples

- List all the objects that are dependent on the ORDERS file in APPLIB.

```
SELECT * FROM TABLE(SYSTOOLS.RELATED_OBJECTS(
                                LIBRARY_NAME=>'APPLIB',
                                FILE_NAME =>'ORDERS'));
```

- List all the objects that are dependent on the ORDERS file in APPLIB. For each object, return how many recursive steps were required to find the object.

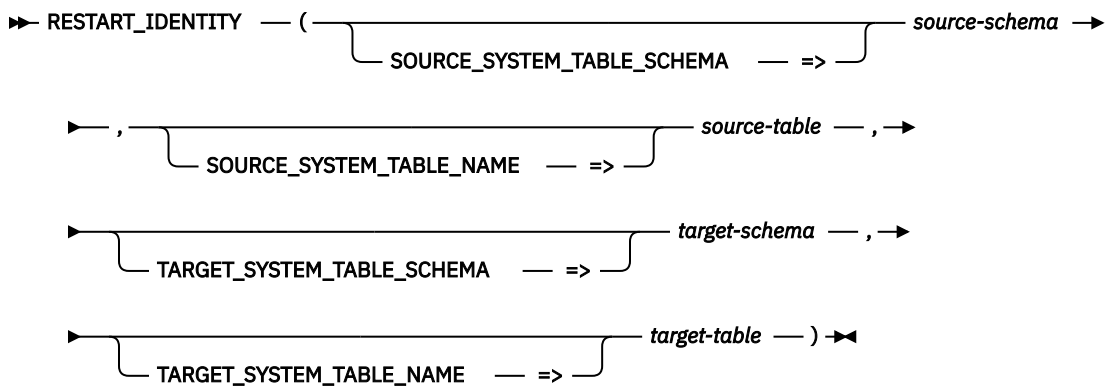
```
SELECT LEVEL, RO.*
FROM (VALUES('APPLIB', 'ORDERS')) I(IN_LIB, IN_FILE),
     TABLE(QSYS2.OBJECT_STATISTICS(IN_LIB, '*FILE', IN_FILE)),
     TABLE(SYSTOOLS.RELATED_OBJECTS(IN_LIB, IN_FILE)) RO
START WITH SOURCE_SCHEMA_NAME = OBJLONGSCHEMA AND
           SOURCE_SQL_NAME = OBJLONGNAME
CONNECT BY SOURCE_SQL_NAME = PRIOR SQL_NAME AND
           SOURCE_SCHEMA_NAME = PRIOR SCHEMA_NAME
ORDER BY 1, 2, 3;
```

RESTART_IDENTITY procedure

The RESTART_IDENTITY procedure examines the *source-table* and determines the identity column and its next value. The next value and column name are used to configure the *target-table* to use the same next value.

Authorization: The caller must have at least one of the following:

- For the target table being modified,
 - The system authority *OBJALTER on the table, and
 - The system authority *EXECUTE on the schema containing the table
- *ALLOBJ authority



The schema is QSYS2.

- source-schema** A character or graphic string for the schema name containing *source-file*. It must be a system schema name.
- source-table** A character or graphic string for the table name that has the identity value to copy. It must be a system table name. The table must contain an identity column.
- target-schema** A character or graphic string for the schema name containing *target-table*. It must be a system schema name.
- target-table** A character or graphic string for the table name that is to have its identity column value reset. It must be a system table name. The table must contain an identity column that has the same name as the identity column in *source-table*.

Example

Set the identity column in NEWTABLE to have the same next value as the identity column in OLDTABLE

```
CALL QSYS2.RESTART_IDENTITY(SOURCE_SYSTEM_TABLE_SCHEMA => 'OLDLIB',
                           SOURCE_SYSTEM_TABLE_NAME => 'OLDTABLE',
                           TARGET_SYSTEM_TABLE_SCHEMA => 'NEWLIB',
                           TARGET_SYSTEM_TABLE_NAME => 'NEWTABLE')
```

SWAP_DYNUSRPRF procedure

The `SWAP_DYNUSRPRF` procedure switches the SQL `DYNUSRPRF` setting for a program or service program. The possible values are `*OWNER` and `*USER`. Calling this procedure will change the value for a program or service program to the opposite setting. The switch will be performed for all modules in an ILE program or service program.

The following query shows the current value of the `DYNUSRPRF` setting for any modules that are part of an ILE program or service program:

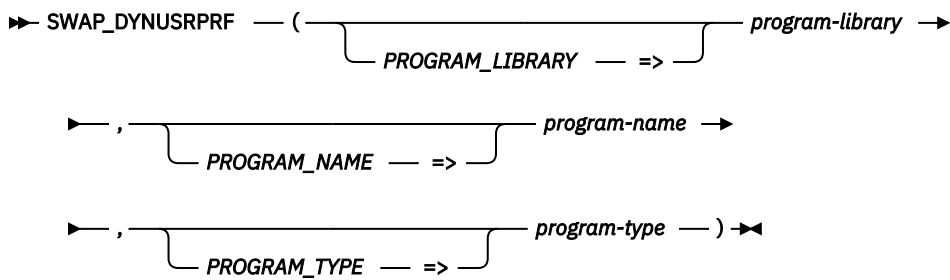
```
SELECT SQL_DYNAMIC_USER_PROFILE FROM QSYS2.BOUND_MODULE_INFO
WHERE PROGRAM_LIBRARY = 'MYLIB' AND PROGRAM_NAME = 'MYPGM';
```

For an OPM program, run the following query:

```
SELECT SQL_DYNAMIC_USER_PROFILE FROM QSYS2.PROGRAM_INFO
WHERE PROGRAM_LIBRARY = 'MYLIB' AND PROGRAM_NAME = 'MYPGM';
```

Authorization: The caller must have at least one of the following:

- Authorization to the Database Security Administrator function of the IBM i. The Change Function Usage (`CHGFCNUSG`) command, with a function ID of `QIBM_DB_SECADM`, can be used to change the list of users allowed to use the function.
- `*ALLOBJ` authority



The schema is QSYS2.

program-library A character or graphic string containing the name of the library containing the program. Can be one of the following special values:

- *CURLIB** The job's current library is used.
- *LIBL** The library list is used.

program-name A character or graphic string containing the program name.

program-type The object type for *program-name*.

- *PGM** The object is a program.
- *SRVPGM** The object is a service program.

Example

Switch the value of the dynamic user profile for program MYLIB/MYPGM.

```
CALL QSYS2.SWAP_DYNUSRPRF('MYLIB', 'MYPGM', '*PGM');
```

SYSFILES view

The `SYSFILES` view contains information about database files. Additional information is available in the `QSYS2.SYSTABLES` view.

The information returned is similar to the detail seen from the Display File Description (DSPFD) command and the Retrieve Database File Description (QDBRTVFD) API.

Authorization: The caller must have:

- *OBJOPR authority to the file, and
- *EXECUTE authority to the library containing the file.

The following table describes the columns in the view. The schema is QSYS2.

Table 78. `SYSFILES` view

Column Name	System Column Name	Data Type	Description
SYSTEM_TABLE_SCHEMA	LIB_NAME	VARCHAR(10)	The library containing the file.
SYSTEM_TABLE_NAME	FILE_NAME	VARCHAR(10)	The name of the file.
TABLE_SCHEMA	TABSHEMA	VARCHAR(128)	The SQL schema name of the library.
TABLE_NAME	TABNAME	VARCHAR(128)	The SQL name of the file.
IASP_NUMBER	IASPNUMBER	INTEGER	The independent auxiliary storage pool (IASP) number.
TEXT_DESCRIPTION	TEXT	VARGRAPHIC(50) CCSID 1200 Nullable	The descriptive text for the file. Contains the null value if there is no descriptive text.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
NATIVE_TYPE	DDS_TYPE	VARCHAR(8)	Type of file. LOGICAL This is a DDS created logical file or an SQL view or index. PHYSICAL This is a physical file or an SQL table.
FILE_TYPE	FILETYPE	VARCHAR(6)	Type of file. DATA This is a data file. SOURCE This is a source physical file.
SQL_OBJECT_TYPE	SQL_TYPE	VARCHAR(5) Nullable	SQL object type. INDEX This is an SQL index. TABLE This is an SQL table. VIEW This is an SQL view. Contains the null value if this is not an SQL object.
LAST_ALTERED_TIMESTAMP	ALTEREDTS	TIMESTAMP(0)	Timestamp when the file was last altered or created.
FILE_LEVEL_ID	FILE_LVLID	CHAR(13)	File level identifier. This is the date of the file in CYYMMDDHHMMSS format.
LEVEL_CHECK	LVLCHK	VARCHAR(3)	Record format level check (LVLCHK). NO The record format level identifiers are not checked when the file is opened. YES The record format level identifiers are checked when the file is opened. See Level checking for some additional information about how level checking works.
FILE_OWNER	OWNER	VARCHAR(10) Nullable	The owner of the object. Contains the null value if the object owner is not available.
CREATE_PUBLIC_AUTHORITY	PUB_AUTH	VARCHAR(10)	The public authority that the file was created with (AUT). This is not the current public authority for the file. *ALL Public all authority *CHANGE Public change authority *EXCLUDE Public exclude authority *USE Public use authority authorization-list-name The name of the authorization list whose authority is used for the file.
NUMBER_MEMBERS	MEMBERS	INTEGER	Number of members.
MAXIMUM_MEMBERS	MAXMBRS	INTEGER Nullable	Maximum members (MAXMBRS). 1 to 32767 The maximum number of members for the file. Contains the null value if no maximum is specified; 32767 is used (*NOMAX).
MAXIMUM_RECORD_LENGTH	RECLNGTH	INTEGER	Maximum record length. This is the length of the record in the file's record format that contains the largest number of bytes.
NUMBER_KEY_FIELDS	KEY_FLDS	INTEGER	Number of key fields for the file.
MAXIMUM_KEY_LENGTH	KEY_LEN	INTEGER	Maximum key length for the file.
TRIGGER_COUNT	TRIG_CNT	INTEGER	Number of triggers.
CONSTRAINT_COUNT	CONSTR_CNT	INTEGER Nullable	Number of constraints for the file. Contains the null value if this is not a physical file.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
NUMBER_BASED_ON_FILES	BASE_FILES	INTEGER Nullable	Number of files directly referenced by a logical file, view, or index. Contains the null value for physical files.
ALLOW_READ	ALWREAD	VARCHAR(3)	Allow read operation. NO Records are not allowed to be read from the file. YES Records are allowed to be read from the file.
ALLOW_WRITE	ALWWRITE	VARCHAR(3)	Allow write operation. NO Records are not allowed to be written to the file. YES Records are allowed to be written to the file.
ALLOW_UPDATE	ALWUPD	VARCHAR(3)	Allow update operation (ALWUPD). NO Records are not allowed to be updated in the file. YES Records are allowed to be updated in the file.
ALLOW_DELETE	ALWDLT	VARCHAR(3)	Allow delete operation (ALWDLT). NO Records are not allowed to be deleted from the file. YES Records are allowed to be deleted from the file.
MAXIMUM_FILE_WAIT_TIME	WAITFILE	INTEGER	Maximum file wait time (WAITFILE). -1 The default wait time specified in the class description is used (*CLS). 0 The program does not wait for the file; an immediate allocation is required (*IMMED). 1 through 32767 The maximum number of seconds a program waits for the file.
MAXIMUM_RECORD_WAIT_TIME	WAITRCD	INTEGER	Maximum record wait time (WAITRCD). -2 The default wait time allowed by the system is used. This is 32767 seconds (*NOMAX). -1 The program does not wait for the record, an immediate allocation is required (*IMMED). 1 through 32767 The maximum number of seconds a program waits for the record.
FORCE_WRITE_RATIO	FRCRATIO	INTEGER Nullable	Records to force a write (FRCRATIO). 1 through 32767 When the cumulative count of inserted, updated, and deleted records has reached this value, the changed records are forced to storage. Contains the null value is there is no force write ratio.
SELECT_OMIT	SELECTOMIT	VARCHAR(3)	Select/omit setting. NO The file is not a select/omit logical file. YES The file is a select/omit logical file.
PROGRAM_DESCRIBED	PGM_DESC	VARCHAR(3)	Program described file setting. NO The file is not program described. YES The file is program described.
DISTRIBUTED	DIST_FILE	VARCHAR(3)	Distributed file setting. NO The file is not distributed. YES The file is distributed.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
FILE_VRM	FILE_VRM	CHAR(6)	File version, release, and modification level. VxRyMz, where x is the version, y the release, and z the modification level. This is the release where the file was created.
EARLIEST_POSSIBLE_RELEASE	MINRLS	CHAR(6) Nullable	Earliest supported version, release, and modification level. New database support used in the file will prevent the file from being saved to a prior version, release, and modification level. The value is formatted as VxRyMz, where x is the version, y the release, and z the modification level. Contains the null value if the value is prior to V2.
ALLOW_NULL_KEYS	ALWNULLK	VARCHAR(3) Nullable	Allow null value key setting (ALWNULL). NO Null value keys are not allowed. YES Null value keys are allowed. Contains the null value if ACCESS_PATH_KEYED is NO.
ALLOW_NULL_DATA	ALWNULLD	VARCHAR(3)	Allow null value data (ALWNULL). NO The file record format(s) do not allow null value fields. YES The file record format(s) allow null value fields.
PRIMARY_KEY	PRIKEY	VARCHAR(3)	Primary key (*PRIKEY). NO The access path for the file is not a primary key. YES The access path for the file is a primary key.
UNIQUE_CONSTRAINT	UNQCST	VARCHAR(3)	Unique constraint. NO The access path for the file is not a unique constraint. YES The access path for the file is a unique constraint.
VOLATILE	VOLATILE	VARCHAR(3)	SQL volatile table setting. NO The file is not an SQL volatile table. YES The file is an SQL volatile table.
KEEP_IN_MEMORY	KEEPINMEM	VARCHAR(3)	The memory preference of the file. NO The file does not have the keep in memory indication set. YES The file has the keep in memory indication set.
MEDIA_PREFERENCE	UNIT	VARCHAR(4)	Preferred storage unit for the file (UNIT). *ANY No storage media is preferred. Storage will be allocated from any available storage media. *SSD Solid state disk storage media is preferred. Storage may be allocated from solid state disk storage media, if available.
SOURCE_FILE_LIBRARY	SRCLIB	VARCHAR(10) Nullable	Library containing SOURCE_FILE. Contains the null value if a source file was not used.
SOURCE_FILE	SRCFILE	VARCHAR(10) Nullable	Name of the source file containing the DDS used to create the file. Contains the null value if a source file was not used.
SOURCE_FILE_MEMBER	SRCMBR	VARCHAR(10) Nullable	Source file member name within SOURCE_FILE used to create the file. Contains the null value if a source file was not used.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
ACCESS_PATH_KEYED	ACCESSPATH	VARCHAR(3)	<p>Whether the file has a keyed sequence access path.</p> <p>NO The file does not have a keyed sequence access path. The file is processed using arrival sequence.</p> <p>YES The file has a keyed sequence access path.</p>
ACCESS_PATH_TYPE	AP_TYPE	VARCHAR(14) Nullable	<p>Access path type.</p> <p>EVI Encoded vector with a 1-, 2-, or 4-byte vector.</p> <p>KEYED FCFO Keyed sequence access path with duplicate keys allowed. Duplicate keys are accessed in first-changed-first-out (FCFO) order.</p> <p>KEYED FIFO Keyed sequence access path with duplicate keys allowed. Duplicate keys are accessed in first-in-first-out (FIFO) order.</p> <p>KEYED LIFO Keyed sequence access path with duplicate keys allowed. Duplicate keys are accessed in last-in-first-out (LIFO) order.</p> <p>KEYED NO ORDER Keyed sequence access path with duplicate keys allowed. No order is guaranteed when accessing duplicate keys.</p> <p>KEYED UNIQUE Keyed sequence access path with no duplicate keys allowed (UNIQUE).</p> <p>Contains the null value if ACCESS_PATH_KEYED is NO.</p>
ACCESS_PATH_MAINTENANCE	MAINT	VARCHAR(6) Nullable	<p>Access path maintenance (MAINT).</p> <p>*DLY Delayed maintenance</p> <p>*IMMED Immediate maintenance</p> <p>*REBLD Rebuild maintenance</p> <p>Contains the null value if ACCESS_PATH_KEYED is NO.</p>
ACCESS_PATH_SIZE	ACCPHYSIZ	VARCHAR(7) Nullable	<p>Access path size (ACCPHYSIZ).</p> <p>*MAX1TB All access paths associated with this file will be allowed to occupy a maximum of 1 terabyte (1 099 511 627 776 bytes) of auxiliary storage</p> <p>*MAX4GB All access paths associated with this file will be allowed to occupy a maximum of 4 gigabytes (4 294 966 272 bytes) of auxiliary storage.</p> <p>Contains the null value if ACCESS_PATH_KEYED is NO.</p>
LOGICAL_PAGE_SIZE	PAGESIZE	INTEGER Nullable	<p>Access path page size.</p> <p>-1 Access path is a 4 gigabyte access path</p> <p>0 System determines page size from the key length of the access path</p> <p>8 Page size is 8 kilobytes</p> <p>16 Page size is 16 kilobytes</p> <p>32 Page size is 32 kilobytes</p> <p>64 Page size is 64 kilobytes</p> <p>128 Page size is 128 kilobytes</p> <p>256 Page size is 256 kilobytes</p> <p>512 Page size is 512 kilobytes</p> <p>Contains the null value if ACCESS_PATH_KEYED is NO.</p>

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
FORCE_KEYED_ACCESS_PATH	FRCACCPH	VARCHAR(3) Nullable	Force keyed access path (FRCACCPH). NO The access path and changed records are not forced to auxiliary storage when the access path is changed. YES The access path and changed records are forced to auxiliary storage when the access path is changed (*YES). Contains the null value if ACCESS_PATH_KEYED is NO.
ACCESS_PATH_JOURNALED	JOURNALED	VARCHAR(3) Nullable	Access path journaled. NO The access path is not journaled. YES The access path is journaled. Contains the null value if ACCESS_PATH_KEYED is NO.
ACCESS_PATH_RECOVERY	RECOVER	VARCHAR(7) Nullable	Access path recovery (RECOVER). *AFTIPL The file access path is built after the IPL is completed. *IPL The file access path is built during the IPL. *NO The file access path is built when the file is next opened. Contains the null value if ACCESS_PATH_KEYED is NO.
SRTSEQ_IND	SRTSEQ_IND	INTEGER	Sort sequence table (SRTSEQ) indicators. 1 No sort sequence table is used for the file, and the hexadecimal value of the characters will be used to determine the sort sequence (*HEX). 2 A sort sequence table was specified for the file. 3 No sort sequence table for the file; however, an alternate collating sequence table was specified.
SORT_SEQUENCE_LIBRARY	SRTSEQ_LIB	VARCHAR(10) Nullable	The library containing the sort sequence table or alternate collating sequence table. Can contain the special value *LIBL. Contains the null value if SRTSEQ_IND is 1 or if there is no library.
SORT_SEQUENCE	SRTSEQ	VARCHAR(10) Nullable	The sort sequence table or alternate collating sequence table associated with the file. Contains the null value if SRTSEQ_IND is 1.
LANGUAGE_IDENTIFIER	LANGID	CHAR(3) Nullable	Language identifier (LANGID). Contains the null value if there is no language identifier.
ROUNDING_MODE	DECFLTRND	VARCHAR(8) Nullable	Rounding mode to be used for decimal floating point (DECFLOAT) calculations. CEILING DOWN FLOOR HALFDOWN HALFEVEN HALFUP UP Contains the null value if the file has no decimal floating point fields or this is not an SQL view, SQL index with a derived key expression, SQL materialized query table, or logical file.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
DECFLOAT_WARNINGS	DECFLTWRN	VARCHAR(3) Nullable	Indicates whether warnings should be returned from decimal floating point calculations. NO Warnings are not returned. YES Warnings are returned. Contains the null value if the file has no decimal floating point fields or this is not an SQL view, SQL index with a derived key expression, SQL materialized query table, or logical file.
NUMBER_RECORD_FORMATS	FORMATS	INTEGER	Total number of record formats.
FORMAT_LEVEL_ID	FMTLVLID	CHAR(13) Nullable	The record format level ID. Contains the null value if file has more than one format or if no value is available.
FORMAT_NAME	FMT_NAME	VARCHAR(10) Nullable	The name of the record format. Contains the null value if file has more than one format.
RECORD_LENGTH	RCD_LEN	INTEGER Nullable	The length of the record format. Contains the null value if file has more than one format.
NUMBER_FIELDS	FIELDS	INTEGER Nullable	The number of fields in the record format. Contains the null value if file has more than one format.
COMMON_CCSID	CCSID	INTEGER Nullable	The CCSID used when all fields with a character, open, and graphic data type use the same CCSID. Contains the null value if file has more than one format or if all character, open, and graphic fields do not use the same CCSID.
CONTAINS_EXPLICIT_CCSID	EXPLICIT	VARCHAR(3) Nullable	Explicit CCSID setting. NO A CCSID was not specified for any character type fields in the format. YES A CCSID was specified for one or more character type fields in the format. Contains the null value if file has more than one format.
CONTAINS_MULTIPLE_CCSDS	MULTIPLE	VARCHAR(3)	Multiple CCSID setting. NO The file has only one CCSID for its character type fields or the file has no character type fields. YES The file has more than one CCSID for its character type fields
CONTAINS_UNICODE	UNICODE	VARCHAR(3) Nullable	Format contains UTF-8, UTF-16, or UCS-2 fields. NO The file record format does not contain UTF-8, UTF-16, or UCS-2 fields. YES The file record format contains UTF-8, UTF-16, or UCS-2 fields. Contains the null value if file has more than one format.
CONTAINS_VARYING_LENGTH	VARLEN	VARCHAR(3)	File contains variable length fields (VARLEN). NO The file record format does not contain variable length fields. YES The file record format contains variable length fields.
CONTAINS_DATETIME	DATETIME	VARCHAR(3)	File contains date/time/timestamp fields. NO The file record format does not contain date, time, or timestamp fields. YES The file record format contains date, time, or timestamp fields.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
CONTAINS_GRAPHIC	GRAPHIC	VARCHAR(3)	File contains graphic fields. NO The file record format does not contain graphic fields. YES The file record format contains graphic fields.
CONTAINS_LOB	LOB	VARCHAR(3)	File contains large object fields. These are the SQL data types character large object (CLOB), double-byte character large object (DBCLOB), and binary large object (BLOB). NO The file record format does not have a large object field. YES The file record format has a large object field.
CONTAINS_ROWID	ROWID	VARCHAR(3)	File contains ROWID fields. NO The file record format does not have a ROWID column. YES The file record format has a ROWID column.
CONTAINS_UDT	UDT	VARCHAR(3)	File contains user-defined type fields. NO The file record format does not have a user-defined type field. YES The file record format has a user-defined type field.
CONTAINS_DATA LINK	DATA LINK	VARCHAR(3)	File contains DataLink fields. NO The file record format does not have a DataLink field. YES The file record format has a DataLink field.
CONTAINS_DATA LINK_ FILE_LINK_CONTROL	DATA LINKFL	VARCHAR(3)	File contains DataLink with FILE LINK CONTROL fields. NO The file record format does not have a DataLink field with FILE LINK CONTROL. YES The file record format has a DataLink field with FILE LINK CONTROL.
CONTAINS_NULL	NULLABLE	VARCHAR(3) Nullable	File contains null capable fields. NO The file record format does not have any null capable fields. YES The file record format has null capable fields. Contains the null value if file has more than one format.
CONTAINS_DEFAULT	DFT	VARCHAR(3) Nullable	The physical file format contains fields with explicit default values. NO The file record format does not have any fields with explicit default values. YES The file record format has fields with explicit default values. Contains the null value if file has more than one format.
CONTAINS_IDENTITY	IDENTITY	VARCHAR(3)	File contains an identity column. NO The file record format does not have an identity column. YES The file record format has an identity column.
CONTAINS_ROW_CHANGE_TIMESTAMP	ROW_CHANGE	VARCHAR(3)	File contains a row change timestamp column. NO The file record format does not have a row change timestamp column. YES The file record format has a row change timestamp column.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
CONTAINS_USER_DEFINED_FUNCTION	UDF	VARCHAR(3)	File contains a user-defined function. NO The file does not use a user-defined function. YES The file uses a user-defined function.
Values for the following columns are returned when NATIVE_TYPE is PHYSICAL. Otherwise, the columns will contain the null value.			
ALLOCATE_STORAGE	ALLOCATE	VARCHAR(3) Nullable	The allocate storage setting (ALLOCATE). NO New members added to the file allow the system to determine storage space that is allocated for the member (ALLOCATE(*NO)). YES New members added to the file use the initial number of records to determine storage space that is allocated for the member (ALLOCATE(*YES)).
CONTIGUOUS_STORAGE	CONTIG	VARCHAR(3) Nullable	The contiguous storage setting (CONTIG). NO Storage should not be attempted to be allocated contiguously (CONTIG(*NO)). YES Storage should be attempted to be allocated contiguously (CONTIG(*YES)). Contains the null value if ALLOCATE_STORAGE is NO.
MAXIMUM_DELETED_PERCENTAGE	DLTPCT	INTEGER Nullable	Maximum percentage of deleted records allowed (DLTPCT). 1 to 100 The threshold percentage of deleted records a member can contain before a message is sent to the history log. See Deleted records for additional information. Contains the null value if the number of deleted records is not checked when the member is closed (*NONE).
INITIAL_RECORDS	SIZE_INIT	INTEGER Nullable	Initial number of records (SIZE). This is the number of records that can be inserted before an automatic extension occurs. Contains the null value if the number of records that can be inserted into each member is not limited by the user. The system determines the maximum member size (*NOMAX).
INCREMENT_RECORDS	SIZE_INCR	INTEGER Nullable	Increment number of records (SIZE). This is the maximum number of records that can be inserted into the member after an automatic extension occurs. Contains the null value when INITIAL_RECORDS is null.
MAXIMUM_INCREMENTS	SIZE_MAX	INTEGER Nullable	Maximum number of increments (SIZE). This is the maximum number of automatic extensions that can be made to the member. Contains the null value when INITIAL_RECORDS is 0.
REUSE_DELETED_RECORDS	REUSEDLT	VARCHAR(3) Nullable	Reuse deleted records (REUSEDLT). NO Deleted records are not reused on subsequent writes or inserts. YES Deleted records can be reused on subsequent writes or inserts.
MATERIALIZED_QUERY_TABLE	MQT	VARCHAR(3) Nullable	SQL materialized query table setting. NO This is not an SQL materialized query table. YES This is an SQL materialized query table.
PARTITIONED_TABLE	PARTITION	VARCHAR(3) Nullable	Partitioned table setting. NO This is not a partitioned table. YES This is a partitioned table.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
ROW_AND_COLUMN_ACCESS_CONTROL	RCAC	VARCHAR(3) Nullable	Row and column access control setting. NO Access controls are not defined for the file. YES Access controls are defined for the file. The QSYS2.SYSCONTROLS view contains detailed information about row and column access controls.
Values for the following columns are returned when NATIVE_TYPE is LOGICAL. Otherwise, the columns will contain the null value.			
TOTAL_SELECT_OMIT	TOTAL_SO	INTEGER Nullable	Total number of select/omit statements for all record formats.
FMTSLR_LIBRARY	FMTSLR_LIB	VARCHAR(10) Nullable	Record format selector program library (FMTSLR) Contains the null value if there is no record format selector program.
FMTSLR_PROGRAM	FMTSLR_PGM	VARCHAR(10) Nullable	Record format selector program (FMTSLR) Contains the null value if there is no record format selector program.
IS_JOIN_LOGICAL	JFILE	VARCHAR(3) Nullable	Join logical file setting (JFILE). NO The file is not a join logical file. YES The file is a join logical file.
DYNAMIC_SELECTION	DYNSLT	VARCHAR(3) Nullable	Dynamic selection setting (DYNSLT). NO The selection and omission tests specified for the file are done when the access path is updated. YES The selection and omission tests specified for the file are done when the file is read.
WITH_CHECK_OPTION	CHECK	VARCHAR(8) Nullable	With check option. CASCADED The cascaded check option was specified. NO No check option was specified or this is not an SQL view. LOCAL The local check option was specified.
PHYSICAL_LOB	PF_LOB	VARCHAR(3) Nullable	Whether this logical file has no large object fields, but the based-on physical file has a large object field. NO The logical file and based-on physical file either both have large object fields or both do not. YES The logical file has no large object fields, but the based-on physical file has a large object field.
PHYSICAL_DATA LINK	PF_DATA LINK	VARCHAR(3) Nullable	Whether this logical file has no DataLink fields, but the based-on physical file has a DataLink field. NO The logical file and based-on physical file either both have DataLink fields or both do not. YES The logical file has no DataLink fields, but the based-on physical file has a DataLink field.
INDEX_COLUMN_IS_EXPRESSION	IXEXP	VARCHAR(3) Nullable	If the SQL index key column is an expression. NO The key column is not an expression or this is not an SQL index. YES The key column is an expression.

Table 78. SYSFILES view (continued)

Column Name	System Column Name	Data Type	Description
INDEX_EXPRESSION_HAS_UDF	IXEXPUDF	VARCHAR(3) Nullable	If the SQL index key column is an expression and the expression contains a user-defined function (UDF). NO The key column is not an expression or the expression does not contain a user-defined function, or this is not an SQL index. YES The key column is an expression and the expression contains a UDF.
INDEX_HAS_SEARCH_CONDITION	SPARSE	VARCHAR(3) Nullable	If the index has a search condition. NO The index does not have a search condition, or this is not an SQL index YES The index has a search condition.
INDEX_SEARCH_CONDITION_HAS_UDF	SPARSE_UDF	VARCHAR(3) Nullable	If the index search condition contains a user-defined function. NO The index is not sparse or does not contain a user-defined function, or this is not an SQL index YES The index is sparse and the search condition contains a UDF.

Example


Examine all the files in APPLIB1.

```
SELECT * FROM QSYS2.SYSFILES WHERE LIB_NAME = 'APPLIB1';
```

VALIDATE_DATA, VALIDATE_DATA_FILE, and VALIDATE_DATA_LIBRARY table functions

The VALIDATE_DATA, VALIDATE_DATA_FILE, and VALIDATE_DATA_LIBRARY table functions examine the data in database physical file members to verify it is properly formed. This can be used to identify invalid data, such as invalid decimal values, that were inserted into a DDS-created physical file.

Authorization: The user must have *EXECUTE authority to the library and *OBJOPR and *READ authority to the file.



library-name An character string expression that specifies the name of the library containing *file-name*.

file-name A character string expression that specifies the name of a database physical file.

member-name A character string expression that specifies the name of a member in *file-name*. The special values *FIRST and *LAST are supported. If this parameter is omitted, *FIRST is used.

The VALIDATE_DATA_FILE table function validates all members in a physical file. It is identical to VALIDATE_DATA except it only has *library-name* and *file-name* parameters.

The VALIDATE_DATA_LIBRARY table function validates all members in all physical files in a library. It is identical to VALIDATE_DATA except it only has a *library-name* parameter.

The result of the function is a table containing a row for each row in a member that has invalid data in at least one column. The columns of the result table are described in the following table. The result columns are nullable.

Table 79. VALIDATE_DATA, VALIDATE_DATA_FILE, and VALIDATE_DATA_LIBRARY table function

Column Name	Data Type	Description
VALIDATE_TIME	TIMESTAMP	The timestamp when this row was generated.
LIBRARY_NAME	VARCHAR(10)	The library containing the file.
FILE_NAME	VARCHAR(10)	The physical file containing the member.
MEMBER_NAME	VARCHAR(10)	The member with the invalid data.
RELATIVE_RECORD_NUMBER	BIGINT	The relative record number (RRN) of the row in MEMBER_NAME with the invalid data.
SQL_WARNING	INTEGER	The SQLCODE of the warning that indicated this error.
REASON_CODE	INTEGER	The reason code from the warning that indicated this error.
COLUMN_NAME	VARCHAR(128)	The name of the column with the invalid data.
WARNING_TEXT	VARCHAR(1000)	The text from the warning that indicated this error.

Note

This function is provided in the SYSTOOLS schema as an example of how to examine all the rows in a table by using an SQL table function. Creating customized versions of this table function to better suit a specific need is encouraged. Use the Insert Generated SQL feature in IBM i Access Client Solutions (ACS) to extract the source for this function. Then modify it and create a new procedure in a user-specified schema.

Example

- Validate the content of all members in FILE1. If no rows are returned, all the data appears to be good.

```
SELECT * FROM TABLE(SYSTOOLS.VALIDATE_DATA_FILE('APPLIB', 'FILE1'));
```

IBM i Services

There are many system services that can be accessed through system-provided SQL views, procedures, and functions. These provide an SQL interface to access, transform, order, and subset the information without needing to code to a system API.

Information about the supported IBM i releases, including the database PTF groups where each service was introduced or enhanced, is available in the IBM i Technology Updates wiki. See this page for all the details: <https://www.ibm.com/support/pages/node/1119123>

Application Services

These procedures, functions, and views provide interface information that can be used by applications.

ACTIVATION_GROUP_INFO table function

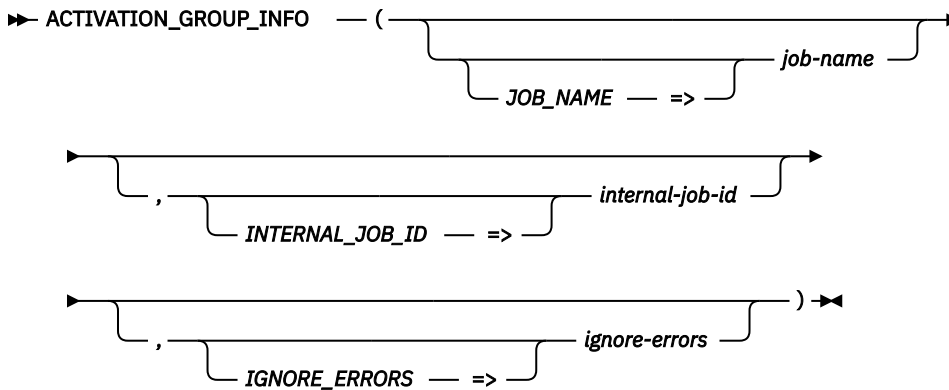
The ACTIVATION_GROUP_INFO table function returns all the activation groups that are associated with a job and their attributes.

This information is similar to what can be accessed through the Display Job (DSPJOB) CL command and the Open List of Activation Group Attributes (QWVOLAGP) API.

Authorization:

None required for a job where the caller's user profile is the same as the job user identity of the job for which the information is being returned.

Otherwise, the caller must have *JOBCTL special authority.



The schema is QSYS2.

job-name A character string that identifies the qualified name of a job in the form *job-number/job-user/job-name*.

The special value of '*' indicates the current job.

If *job-name* is provided, *internal-job-id* must be omitted.

internal-job-id A binary string that contains an internal job identifier. This can be passed for quicker performance. This value is returned in the INTERNAL_JOB_ID column by the ACTIVE_JOB_INFO table function.

If *internal-job-id* is provided, *job-name* must be omitted.

ignore-errors A character string that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing one row for every activation group for the job with the format shown in the following table. All columns are nullable.

Table 80. ACTIVATION_GROUP_INFO table function

Column Name	Data Type	Description
ACTIVATION_GROUP_NAME	VARCHAR(10)	The name of the activation group. Can contain one of the following special values: *DFACTGRP The activation group is one of the default activation groups. *NEW The activation group is *NEW.
ACTIVATION_GROUP_NUMBER	DECIMAL(20,0)	The activation group number.
STORAGE_MODEL	VARCHAR(10)	The storage model of the activation group. *SINGLVL The activation group is single level store. *TERASPACE The activation group is teraspace.
STATE	VARCHAR(6)	The state of the activation group. SYSTEM The activation group is in system state. USER The activation group is in user state.
NUMBER_OF_ACTIVATIONS	INTEGER	The total number of program activations in this activation group.

Table 80. ACTIVATION_GROUP_INFO table function (continued)

Column Name	Data Type	Description
PROGRAM_LIBRARY	VARCHAR(10)	The name of the library that contains the program that caused this activation group to be created. Contains the null value when PROGRAM is null.
PROGRAM	VARCHAR(10)	The name of the program that caused this activation group to be created. Contains the null value when the activation group is one of the default activation groups or if the program no longer exists in the system.
PROGRAM_TYPE	VARCHAR(6)	The type of program that caused this activation group to be created. JAVA A Java program. PGM A program. SRVPGM A service program. Contains the null value when PROGRAM is null.
SHARED_ACTIVATION_GROUP	VARCHAR(3)	Whether the activation group is shared. A shared activation group is an activation group that belongs to more than one job at the same time. NO The activation group is not shared with other jobs. YES The activation group is shared with other jobs.
IN_USE	VARCHAR(3)	Whether the activation group is eligible to be reclaimed. An activation group can be reclaimed using the Reclaim Activation Group (RCLACTGRP) command. NO The activation group is not in use and is eligible to be reclaimed. YES The activation group is in use and cannot be reclaimed.
STATIC_STORAGE_SIZE	BIGINT	The total amount of static storage allocated to the activation group, in bytes. If the size exceeds 4,294,967,295 bytes, 4,294,967,295 will be returned.
NUMBER_OF_HEAPS	INTEGER	The total number of heaps that are allocated by this activation group.
HEAP_STORAGE_SIZE	BIGINT	The total amount of heap storage that is allocated to the activation group, in bytes. If the size exceeds 4,294,967,295 bytes, 4,294,967,295 will be returned.

Example

- List the activation group information for the current job.

```
SELECT * FROM TABLE(QSYS2.ACTIVATION_GROUP_INFO('*'));
```

- List the activation group information for a job based on its internal job identifier.

```
SELECT * FROM TABLE(QSYS2.ACTIVATION_GROUP_INFO(
INTERNAL_JOB_ID => BX'00D10003007E3F00A784A7CBD6E89001'));
```

BINDING_DIRECTORY_INFO view

The BINDING_DIRECTORY_INFO view returns information about the object entries in binding directories.

The values returned for the columns in the view are similar to the values returned by the Display Binding Directory (DSPBNDDIR) CL command

Authorization: The caller must have:

- *USE authority to the library containing the binding directory, and
- *OBJOPR and *READ authority to the binding directory.

To return the create timestamp for a binding directory entry, the caller must have:

- *EXECUTE authority to the library containing the binding directory entry, and
- Some authority to the binding directory entry.

The following table describes the columns in the view. The system name is BNDDIR_INF. The schema is QSYS2.

Table 81. BINDING_DIRECTORY_INFO view

Column Name	System Column Name	Data Type	Description
BINDING_DIRECTORY_LIBRARY	BNDDIR_LIB	VARCHAR(10)	The library containing the binding directory.
BINDING_DIRECTORY	BNDDIR	VARCHAR(10)	The binding directory name.
ENTRY_LIBRARY	ENTRY_LIB	VARCHAR(10)	The library containing ENTRY. Can contain the special value *LIBL.
ENTRY	ENTRY	VARCHAR(10)	The name of the binding directory entry.
ENTRY_TYPE	ENTRY_TYPE	VARCHAR(7)	The object type of the binding directory entry. *MODULE The object is a module. *SRVPGM The object is a service program.
ENTRY_ACTIVATION	ENTRY_ACT	VARCHAR(6) Nullable	The activation control of the bound service program. *DEFER Activation of the bound service program may be deferred until a function it exports is called. *IMMED Activation of the bound service program takes place immediately when the program or service program it is bound to is activated. Contains the null value if ENTRY_TYPE is *MODULE.
ENTRY_CREATE_TIMESTAMP	ENTRY_TS	TIMESTAMP(0) Nullable	The timestamp when the object was created. Contains the null value if the entry timestamp is not available.

Examples

- Return a list of the entries for all binding directories in library TESTLIB that start with 'TEST'.

```
SELECT * FROM QSYS2.BINDING_DIRECTORY_INFO
  WHERE BINDING_DIRECTORY_LIBRARY = 'TESTLIB'
        AND BINDING_DIRECTORY LIKE 'TEST%'
  ORDER BY BINDING_DIRECTORY_LIBRARY,
           BINDING_DIRECTORY,
           ENTRY_LIBRARY,
           ENTRY;
```

- List all of the binding directories that have the entry MYLIB/HELLO, type *MODULE.

```
SELECT BINDING_DIRECTORY,
       BINDING_DIRECTORY_LIBRARY
  FROM QSYS2.BINDING_DIRECTORY_INFO
  WHERE ENTRY = 'HELLO'
        AND ENTRY_LIBRARY = 'MYLIB'
        AND ENTRY_TYPE = '*MODULE'
  ORDER BY 1, 2;
```

BOUND_MODULE_INFO view

The BOUND_MODULE_INFO view returns information about modules bound into an ILE program or service program.

The values returned for the columns in the view are closely related to the values returned for *MODULE detail on the DSPPGM (Display Program) and DSPSRVPGM (Display Service Program) CL commands and the List ILE Program Information (QBNLPGMI) and the List Service Program Information (QBNLSPGM) APIs.

Authorization: The caller must have:

- *USE authority to the program or service program, and
- *EXECUTE authority to the library containing the program or service program.

The following table describes the columns in the view. The system name is MODULE_INF. The schema is QSYS2.

Table 82. BOUND_MODULE_INFO view

Column Name	System Column Name	Data Type	Description
PROGRAM_LIBRARY	PGM_LIB	VARCHAR(10)	The library containing the program or service program.
PROGRAM_NAME	PGM_NAME	VARCHAR(10)	The program or service program containing the module.
OBJECT_TYPE	OBJ_TYPE	VARCHAR(7)	Object type for PROGRAM_NAME. *PGM A program. *SRVPGM A service program.
BOUND_MODULE_LIBRARY	BDMODLIB	VARCHAR(10)	The library containing the module bound into PROGRAM_NAME at bind time.
BOUND_MODULE	BDMOD	VARCHAR(10)	The module bound into PROGRAM_NAME. This is a copy of the module that was bound into PROGRAM_NAME. It is not the *MODULE object on the system.
MODULE_ATTRIBUTE	ATTRIBUTE	VARCHAR(10) Nullable	The language in which the module is written. Can contain the null value, for example, for a module created by a compilation process internal to IBM.
MODULE_CREATE_TIMESTAMP	CREATE_TS	TIMESTAMP(0) Nullable	The timestamp when the module was created. Contains the null value if no timestamp is available.
SOURCE_FILE_LIBRARY	SRCLIB	VARCHAR(10) Nullable	The name of the library that contains the source file used to create the module. Contains the null value if no source file was used to create the module.
SOURCE_FILE	SRCFILE	VARCHAR(10) Nullable	The name of the source file used to create the module. Contains the null value if no source file was used to create the module.
SOURCE_FILE_MEMBER	SRCMBR	VARCHAR(10) Nullable	The name of the member in the source file. Contains the null value if no source file was used to create the module.
SOURCE_CHANGE_TIMESTAMP	SRC_CHGTS	TIMESTAMP(0) Nullable	The timestamp when the source that was used to create this module was last changed. Contains the null value if no source file was used to create the module.
MODULE_CCSID	TGTCCSID	INTEGER	The coded character set identifier (CCSID) for this module.

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SORT_SEQUENCE_LIBRARY	SRTSEQ_LIB	VARCHAR(10) Nullable	<p>The name of the library that contained the sort sequence table used when the module was compiled. This does not apply to SQL statements in the module. Can contain the following special values:</p> <p>*LIBL The sort sequence table is found in the library list when PROGRAM_NAME runs this module.</p> <p>*CURLIB The sort sequence table is found in the current library when PROGRAM_NAME runs this module.</p> <p>Contains the null value if SORT_SEQUENCE contains a special value or is null.</p>
SORT_SEQUENCE	SRTSEQ	VARCHAR(10) Nullable	<p>The name of the sort sequence table used when the module was compiled. This does not apply to SQL statements in the module. Can contain the following special values:</p> <p>*HEX No sort sequence is used.</p> <p>*JOB RUN The sort sequence value associated with the job at the time PROGRAM_NAME runs this module is used.</p> <p>*LANGIDSHR The shared sort sequence for the language identifier is used.</p> <p>*LANGIDUNQ The unique sort sequence for the language identifier is used.</p> <p>Contains the null value if the module does not contain any sort sequence information.</p>
LANGUAGE_ID	LANGID	VARCHAR(7) Nullable	<p>The language identifier used when the module was compiled. This does not apply to SQL statements in the module. Can contain the following special value:</p> <p>*JOB RUN The language identifier associated with the job at the time PROGRAM_NAME runs this module.</p> <p>Contains the null value if the module does not contain any language identification information.</p>
DEBUG_DATA	DEBUG_DATA	VARCHAR(4)	<p>Whether debug data was generated when this module was created. If debug data exists, the module may be debugged using the source debugger.</p> <p>*NO Debug data was not generated.</p> <p>*YES Debug data was generated.</p>

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
OPTIMIZATION_LEVEL	OPTIMIZE	INTEGER	<p>Optimization level for the module. Optimization improves the performance of a module, but reduces the ability to immediately change and accurately display the value of variables during debugging.</p> <p>10 No additional optimization is performed on the generated code. Variables can be displayed and changed when the program is being debugged. With no optimization of the code, this value provides the lowest level of module performance. 10 is also known as *NONE.</p> <p>20 Some optimization is performed on the generated code. When the module optimized at this level is being debugged, the variables can be displayed but not changed. 20 is also known as *BASIC.</p> <p>30 More optimization is performed in addition to those performed at optimization level 20. Variables cannot be changed but can be displayed while the program is being debugged. However, the displayed value of the variable during debugging may not be its actual value. 30 is also known as *FULL.</p> <p>40 Maximum level of optimization. This level includes all the optimizations performed at optimization level 30. In addition, it includes optimization that disables call and instruction tracing. Thus, tracing of modules created at this optimization level cannot be done.</p>
MAX_OPTIMIZATION_LEVEL	MAX_OPTLVL	INTEGER Nullable	<p>The highest level of optimization this module could have at bind time. If observability has been removed from the module, this maximum optimization level value might not be the same as the one specified at module creation.</p> <p>The values are identical to the values for the OPTIMIZATION_LEVEL column.</p> <p>Contains the null value if the module is not restricted to a maximum optimization level. The module can be retranslated to any of the supported optimization levels.</p>
OBJECT_CONTROL_LEVEL	OBJSSI	CHAR(8) Nullable	<p>The object control level for the module at the time it was bound into PROGRAM_NAME.</p> <p>Contains the null value if there is no object control level for the module.</p>
RELEASE_CREATED_ON	CREATE_ON	CHAR(6)	<p>The version, release, and modification level of the operating system on which the module was created, in VxRxMx format.</p>
TARGET_RELEASE	TGTRLS	CHAR(6)	<p>The version, release, and modification level of the operating system for which the module was created, in VxRxMx format.</p>
CREATION_DATA	CREATEDATA	VARCHAR(6)	<p>Whether the bound module has all the creation data and if that data is observable or unobservable.</p> <p>*NO Not all the creation data is present in the bound module.</p> <p>*UNOBS The creation data is present in the bound module but not all of that data is observable.</p> <p>*YES The creation data is present in the bound module and all of that data is observable.</p>

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
TERASPACE_STORAGE_ENABLED	TERASPACE	VARCHAR(4)	<p>The teraspace storage capability for this bound module.</p> <p>*NO The module is not teraspace storage enabled.</p> <p>*YES The module is teraspace storage enabled.</p>
STORAGE_MODEL	STGMDL	VARCHAR(10)	<p>Where the automatic and static storage for this bound module is allocated at run time.</p> <p>*INHERIT Automatic and static storage are allocated from either single-level storage or teraspace, depending on the activation.</p> <p>*SINGLVL Automatic and static storage are allocated from single-level storage.</p> <p>*TERASPACE Automatic and static storage are allocated from teraspace.</p>
NUMBER_PROCEDURES	PROCS	INTEGER	<p>The number of procedures defined in the module. This number includes the program entry procedure (PEP), if one was generated by the compiler for this module.</p>
NUMBER_PROCEDURES_BLOCK_REORDERED	PROCS_BRO	INTEGER	<p>The number of procedures defined in the module that are block reordered.</p> <p>If the module does not have block-order profiling data applied, this value will be zero.</p>
NUMBER_PROCEDURES_BLOCK_ORDER_MEASURED	PROCS_BOM	INTEGER	<p>The number of procedures defined in the module that had block-order profiling data collected at the time block-order profiling data was applied.</p> <p>If the module does not have block-order profiling data applied, this value will be zero.</p>
PROFILING_DATA	PRFDTA	VARCHAR(10)	<p>The profiling data attribute for this bound module.</p> <p>*APYBLKORD Block-order profiling data is applied to the module. See NUMBER_PROCEDURES_BLOCK_REORDERED for the current number of procedures in this module that are block reordered.</p> <p>*COL The collection of profiling data is enabled. Any block-order profiling data has been removed for the module.</p> <p>*NOCOL The collection of profiling data is not enabled and block-order profiling data is not applied to the module.</p>
ALLOW_RTVCLSRC	RTVCLSRC	VARCHAR(4)	<p>Whether the module has CL source data that can be retrieved from the program.</p> <p>*NO There is no CL source for this module.</p> <p>*YES There is CL source for this module.</p>
USER_MODIFIED	USER_MOD	VARCHAR(4) Nullable	<p>Indicates whether the module was changed by the user at bind time.</p> <p>*NO The user did not change the module.</p> <p>*YES The user changed the module.</p> <p>Contains the null value if the information is not available.</p>

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LIC_OPTIONS	LICOPT	VARGRAPHIC(500) CCSID 1200 Nullable	The Licensed Internal Code options that are in use by the module. Contains the null value if no LIC options were used for the module.
LICENSED_PROGRAM	LICPGM	VARCHAR(13) Nullable	If the module was part of a licensed program at bind time, contains the product number and the level of the licensed program. Contains the null value if the module was not part of a licensed program at bind time.
PTF_NUMBER	PTF	CHAR(5) Nullable	The program temporary fix (PTF) that resulted in the creation of the module. Contains the null value for user-created modules.
APAR_ID	APAR_ID	CHAR(6) Nullable	The module was changed as the result of the authorized program analysis report (APAR) with this identification number. Contains the null value if the module was not changed at bind time.
SQL_STATEMENT_COUNT	NBRSTMTS	INTEGER	The number of SQL statements contained in the module. Contains 0 if there are no SQL statements in the module.
SQL_RELATIONAL_DATABASE	RDB	VARCHAR(18) Nullable	The default relational database that was specified on the SQL precompile. Can contain the following special value: *LOCAL The module can only access data on the local system. Contains the null value if no package was created for the module by the SQL precompiler or if the module does not contain SQL statements.
SQL_COMMITMENT_CONTROL	ISOLATION	VARCHAR(5) Nullable	The level of commitment control that was specified on the SQL precompile. *ALL Read stability. *CHG Uncommitted read. *CS Cursor stability. *NONE No commit. *RR Repeatable read. Contains the null value if the module does not contain SQL statements.
SQL_NAMING	NAMING	VARCHAR(4) Nullable	The convention used for naming objects in SQL statements. *SQL The SQL naming convention is used. *SYS The system naming convention is used. Contains the null value if the module does not contain SQL statements.

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_DATE_FORMAT	DATFMT	VARCHAR(4) Nullable	The date format attribute. *DMY Day/month/year format (dd/mm/yy). *EUR European format (dd.mm.yyyy). *ISO International Standards Organization format (yyyy-mm-dd). *JIS Japanese Industrial Standard format (yyyy-mm-dd). *JUL Julian format (yy/dds). *MDY Month/day/year format (mm/dd/yy). *USA USA format (mm/dd/yyyy). *YMD Year/month/day format (yy/mm/dd). Contains the null value if the module does not contain SQL statements.
SQL_DATE_SEPARATOR	DATSEP	CHAR(1) Nullable	The date separator attribute. Contains the null value if the module does not contain SQL statements.
SQL_TIME_FORMAT	TIMFMT	VARCHAR(4) Nullable	The time format attribute. *EUR European format (hh.mm.ss). *HMS Hours/minutes/seconds format (hh:mm:ss). *ISO International Standards Organization format (hh.mm.ss). *JIS Japanese Industrial Standard format (hh.mm.ss). *USA USA format (hh:mm a.m. or p.m.). Contains the null value if the module does not contain SQL statements.
SQL_TIME_SEPARATOR	TIMSEP	CHAR(1) Nullable	The time separator attribute. Contains the null value if the module does not contain SQL statements.
SQL_SORT_SEQUENCE_LIBRARY	SQL_SSEQLB	VARCHAR(10) Nullable	The name of the library that is used to locate the SQL sort sequence table. The following special values can be returned: *CURLIB The SQL sort sequence table is found by looking in the current library. *LIBL The SQL sort sequence table is found by looking in the library list. Contains the null value if SQL_SORT_SEQUENCE contains a special value or if the module does not contain SQL statements.

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_SORT_SEQUENCE	SQL_SRTSEQ	VARCHAR(10) Nullable	<p>The sort sequence table name specified when the module was compiled. The following special values can be returned:</p> <p>*HEX No SQL sort sequence is used for the SQL statements.</p> <p>*JOBRUN The SQL sort sequence is the SRTSEQ value associated with the job at the time the SQL statements within the module are run.</p> <p>*LANGIDSHR The shared SQL sort sequence for the language identifier (LANGID) is used for the SQL statements.</p> <p>*LANGIDUNQ The unique SQL sort sequence for the language identifier (LANGID) is used for the SQL statements.</p> <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_LANGUAGE_ID	SQL_LANGID	VARCHAR(7) Nullable	<p>The language identifier specified when the module was compiled. The following special value can be returned:</p> <p>*JOBRUN The language identifier is the LANGID associated with the job at the time the module is run.</p> <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_DEFAULT_SCHEMA	DFTRBCOL	VARCHAR(10) Nullable	<p>The schema name used for unqualified names of tables, views, indexes, and SQL packages in static statements.</p> <p>Contains the null value if there is no default schema name or if the module does not contain SQL statements.</p>
SQL_PATH	SQLPATH	VARCHAR(3483) Nullable	<p>The list of libraries used during resolution of functions, procedures, and data types within SQL statements. The list is in the form of repeating library names, each surrounded by double quotes and separated by commas.</p> <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_DYNAMIC_USER_PROFILE	DYNUSRPRF	VARCHAR(10) Nullable	<p>The user profile used for dynamic SQL statements. The following special values can be returned:</p> <p>*OWNER Local dynamic SQL statements are run under the profile of the program or service program's owner. Distributed dynamic SQL statements are run under the profile of the SQL package's owner.</p> <p>*USER Local dynamic SQL statements are run under the profile of the job or thread. Distributed dynamic SQL statements are run under the profile of the application server job.</p> <p>Contains the null value if the module does not contain SQL statements.</p>

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_ALLOW_COPY_DATA	ALWCPYDTA	VARCHAR(9) Nullable	<p>Whether a copy of the data can be used in the implementation of an SQL query.</p> <p>*NO A copy of the data is not allowed.</p> <p>*OPTIMIZE A copy of the data is allowed whenever it might result is better performance.</p> <p>*YES A copy of the data is allowed, but only when necessary.</p> <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_CLOSE_SQL_CURSOR	CLOSQLCSR	VARCHAR(10) Nullable	<p>Specifies the CLOSQLCSR attribute.</p> <p>*ENDACTGRP SQL cursors are closed, SQL prepared statements are implicitly discarded, and LOCK TABLE locks are released when the activation group ends.</p> <p>*ENDMOD SQL cursors are closed and SQL prepared statements are implicitly discarded when the module is exited. LOCK TABLE locks are released when the first SQL program on the call stack ends.</p> <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_DELAY_PREPARE	DLYPRP	VARCHAR(4) Nullable	<p>Indicates the delay prepare attribute.</p> <p>*NO Dynamic statement validation is performed when the dynamic statements are prepared.</p> <p>*YES Dynamic statement validation is delayed until the dynamic statements are used.</p> <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_ALLOW_BLOCK	ALWBLK	VARCHAR(8) Nullable	<p>Whether blocking is used to improve the performance of certain SQL statements.</p> <p>*ALLREAD Rows are blocked for read-only cursors.</p> <p>*NONE Rows are not blocked for retrieval of data for cursors.</p> <p>*READ Rows are blocked for read-only retrieval of data for cursors when:</p> <ul style="list-style-type: none"> The commitment control value is *NONE. The cursor is declared with a FOR READ ONLY clause or there are no dynamic statements that could run a positioned UPDATE or DELETE statement for the cursor. <p>Contains the null value if the module does not contain SQL statements.</p>
SQL_PACKAGE_LIBRARY	SQLPKGLIB	VARCHAR(10) Nullable	<p>The name of the library the SQL package is in.</p> <p>Contains the null value if the module is not distributed or if the module does not contain SQL statements.</p>

Table 82. BOUND_MODULE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_PACKAGE	SQLPKG	VARCHAR(10) Nullable	The name of the SQL package created on the relational database specified on the RDB parameter of the command that created this module. Contains the null value if the module is not distributed or if it does not contain SQL statements.
SQL_RDB_CONNECTION_METHOD	RDBCNNMTH	VARCHAR(4) Nullable	Specifies the semantics used for CONNECT statements: *DUW CONNECT (Type 2) semantics are used to support distributed unit of work. *RUW CONNECT (Type 1) semantics are used to support remote unit of work. Contains the null value if the module is not distributed or does not contain SQL statements.

Examples

- Find any bound modules that include source changes from the last 7 days.

```
SELECT *
FROM QSYS2.BOUND_MODULE_INFO
WHERE PROGRAM_LIBRARY = 'QGPL'
AND SOURCE_CHANGE_TIMESTAMP > CURRENT_TIMESTAMP - 7 DAYS
ORDER BY SOURCE_CHANGE_TIMESTAMP DESC;
```

BOUND_SRVPGM_INFO view

The BOUND_SRVPGM_INFO view returns information about service programs bound into an ILE program or service program.

The values returned for the columns in the view are closely related to the values returned for *SRVPGM detail on the DSPPGM (Display Program) and DSPSRVPGM (Display Service Program) CL commands and the List ILE Program Information (QBNLPGMI) and the List Service Program Information (QBNLSPGM) APIs.

Authorization: The caller must have:

- *READ authority to the program or service program, and
- *EXECUTE authority to the library containing the program or service program.

The following table describes the columns in the view. The system name is SRVPGM_INF. The schema is QSYS2.

Table 83. BOUND_SRVPGM_INFO view

Column Name	System Column Name	Data Type	Description
PROGRAM_LIBRARY	PGM_LIB	VARCHAR(10)	Library containing the program or service program.
PROGRAM_NAME	PGM_NAME	VARCHAR(10)	Program or service program name.
OBJECT_TYPE	OBJ_TYPE	VARCHAR(7)	Object type for PROGRAM_NAME. *PGM A program. *SRVPGM A service program.

Table 83. BOUND_SRVPGM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
BOUND_SERVICE_PROGRAM_LIBRARY	BDSRVPGMLB	VARCHAR(10)	The name of the library containing the service program bound to PROGRAM_NAME at bind time. This is the library name in which the activation expects to find the service program at run time. Can contain the following special value: *LIBL Indicates the library list is used at the time the service program is needed.
BOUND_SERVICE_PROGRAM	BDSRVPGM	VARCHAR(10)	The name of the service program bound to PROGRAM_NAME.
BOUND_SERVICE_PROGRAM_SIGNATURE	SIGNATURE	BINARY(16)	The current signature of the service program at the time the service program was bound to PROGRAM_NAME.
BOUND_SERVICE_PROGRAM_ACTIVATION	SRVPGM_ACT	VARCHAR(6)	Specifies when the bound service program is activated. *DEFER The bound service program activation is deferred until a procedure it exports is called. *IMMED The bound service program activation happens when PROGRAM_NAME is activated.

Example

- Examine whether service programs in APPLIB are taking advantage of deferred service program activation.

```
SELECT BOUND_SERVICE_PROGRAM_ACTIVATION, COUNT(*) AS BOUND_SERVICE_PROGRAM_ACTIVATION_COUNT
FROM QSYS2.BOUND_SRVPGM_INFO
WHERE PROGRAM_LIBRARY = 'APPLIB'
GROUP BY BOUND_SERVICE_PROGRAM_ACTIVATION
ORDER BY 2 DESC;
```


CLEAR_DATA_QUEUE procedure

The CLEAR_DATA_QUEUE procedure clears all messages from the specified data queue or clears messages that match the key provided.

This procedure provides function similar to the Clear Data Queue (QCLRDTAQ) API.

Authorization: The caller must have:

- *EXECUTE authority for the library, and
- *OBJOPR and *READ authority for the *DTAQ



The schema is QSYS2.

data-queue A character or graphic string containing the name of the data queue.

data-queue-library A character or graphic string containing the name of the library containing the data queue. Can be one of the following special values:

- *CURLIB** The job's current library is used.
- *LIBL** The library list is used. This is the default.

key-data A character string containing the data to use as the key for selecting messages to be removed from the data queue. This parameter can only be specified for a keyed data queue. If this parameter is not specified, all messages will be cleared from the data queue.

The length of *key-data* must be the length specified on the `KEYLEN` parameter on the Create Data Queue (`CRTDTAQ`) command. The `KEY_LENGTH` column of the `QSYS2.DATA_QUEUE_INFO` view contains this value.

When this parameter is specified, *key-order* must also be specified.

key-order The comparison criteria between the keys of messages on the data queue and the *key-data* parameter. Valid values are:

- EQ** Equal
- GE** Greater than or equal
- GT** Greater than
- LE** Less than or equal
- LT** Less than
- NE** Not equal

This parameter is ignored if *key-data* is not specified.

Example

Clear all entries from data queue DQ1 in library TESTLIB.

```
CALL QSYS2.CLEAR_DATA_QUEUE('DQ1', 'TESTLIB');
```

DATA_AREA_INFO table function

The DATA_AREA_INFO table function returns a row for the specified data area, including a data area in QTEMP or the *LDA, *PDA, or *GDA data areas.

The values returned are closely related to the values returned by the Retrieve Data Area (RTVDTAARA) CL command and the Retrieve Data Area (QWCRDTAA) API.

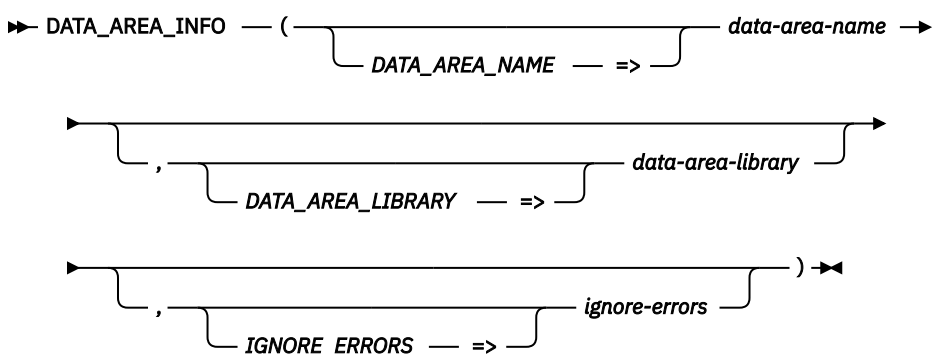
Authorization: The caller must have:

- *USE authority to the data area, and
- *EXECUTE authority to the library containing the data area.

For a DDM data area, the caller must be able to connect to the remote system.

To return information for the data area on the remote system, the caller must have:

- *USE authority to the data area, and
- *EXECUTE authority to the library containing the data area.



The schema is QSYS2.

data-area-name A character or graphic string expression that identifies the data area.
Can contain the following special values:

- *GDA** Group data area.
- *LDA** Local data area.
- *PDA** Program initialization parameter data area.

data-area-library A character or graphic string expression that identifies the library containing the data area. If *data-area-library* is not specified, *LIBL is used. If *data-area-name* is one of the special values, *data-area-library* is ignored.

Can contain the following special values:

- *CURLIB** The current library is used.
- *LIBL** The library list is used.

ignore-errors A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned. This is the default.

YES A warning is returned.

A partial row is returned when an error is encountered. Columns other than the data area name (and the library name if explicitly specified as an input parameter) will be null.

The result of the function is a table containing one row with the format shown in the following table. All the columns are nullable.

Table 84. DATA_AREA_INFO table function

Column Name	Data Type	Description
DATA_AREA_LIBRARY	VARCHAR(10)	Library containing the data area. Contains the null value if this request is for *GDA, *LDA, or *PDA.
DATA_AREA_NAME	VARCHAR(10)	Name of the data area.
DATA_AREA_TYPE	VARCHAR(5)	The type of data area. *CHAR A character data area. *DEC A decimal data area. *LGL A logical data area.
LENGTH	INTEGER	Specifies the length of the data area. <ul style="list-style-type: none"> For a character data area, it is the maximum number of characters. For a decimal data area it is the total number of digits, including decimal positions. For a logical data area, it is always 1.
DECIMAL_POSITIONS	INTEGER	Specifies the number of digits to the right of the decimal point for a decimal data area. Contains the null value if not a decimal data area.
DATA_AREA_VALUE	VARCHAR(2000)	The value currently assigned to the data area. For a decimal data area, the SQL default decimal point is used. See Decimal point .
DATA_AREA_BINARY_VALUE	VARBINARY(2000)	The value currently assigned to the data area as binary data.

Example

- Return the current value of the TESTDATA data area. Use the library list to locate the data area.

```
SELECT DATA_AREA_VALUE FROM TABLE(QSYS2.DATA_AREA_INFO(
    DATA_AREA_NAME => 'TESTDATA',
    DATA_AREA_LIBRARY => '*LIBL'));
```

DATA_AREA_INFO view

The DATA_AREA_INFO view returns the values of data areas.

The values returned for the columns in the view are closely related to the values returned by the Retrieve Data Area (RTVDTAARA) CL command and the Retrieve Data Area (QWCRDTAA) API.

Authorization: The caller must have:

- *USE authority to the data area, and
- *EXECUTE authority to the library containing the data area.

For a DDM data area, the caller must be able to connect to the remote system.

To return information for the data area on the remote system, the caller must have:

- *USE authority to the data area, and
- *EXECUTE authority to the library containing the data area.

The following table describes the columns in the view. The system name is DTAARA_INF. The schema is QSYS2.

Table 85. DATA_AREA_INFO view

Column Name	System Column Name	Data Type	Description
DATA_AREA_LIBRARY	DTAARA_LIB	VARCHAR(10)	Library containing the data area.

Table 85. DATA_AREA_INFO view (continued)

Column Name	System Column Name	Data Type	Description
DATA_AREA_NAME	DTAARA	VARCHAR(10)	Name of the data area.
DATA_AREA_TYPE	TYPE	VARCHAR(5) Nullable	The type of data area. *CHAR A character data area. *DEC A decimal data area. *LGL A logical data area. Contains the null value if this is a DDM data area that was unable to access the data.
LENGTH	LENGTH	INTEGER Nullable	Specifies the length of the data area. <ul style="list-style-type: none"> For a character data area, it is the maximum number of characters. For a numeric data area it is the total number of digits, including decimal positions. For a logical data area, it is always 1. Contains the null value if this is a DDM data area that was unable to access the data.
DECIMAL_POSITIONS	SCALE	INTEGER Nullable	Specifies the number of digits to the right of the decimal point for a decimal data area. Contains the null value if not a decimal data area or if this is a DDM data area that was unable to access the data.
DATA_AREA_VALUE	VALUE	VARCHAR(2000) Nullable	The value currently assigned to the data area as character data. For a decimal data area, the SQL default decimal point is used. See Decimal point . Contains the null value if this is a DDM data area that was unable to access the data.
DATA_AREA_BINARY_VALUE	BIN_VALUE	VARBINARY(2000) Nullable	The value currently assigned to the data area as binary data. Contains the null value if this is a DDM data area that was unable to access the data.
SQL_SEQUENCE	SEQUENCE	VARCHAR(3)	This data area is defined as an SQL sequence. NO This is not an SQL sequence. YES This is an SQL sequence.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The text description of the data area. Contains the null value if the data area has no description or if this is a DDM data area.

Example

- Return a list of values for all data areas in MYLIB.

```
SELECT DATA_AREA_NAME, DATA_AREA_VALUE FROM QSYS2.DATA_AREA_INFO
WHERE DATA_AREA_LIBRARY = 'MYLIB' ;
```

DATA_QUEUE_ENTRIES table function

The DATA_QUEUE_ENTRIES table function returns one or more messages from the specified data queue. The messages are not removed from the data queue. The message data is returned as character, UTF-8, and binary data.

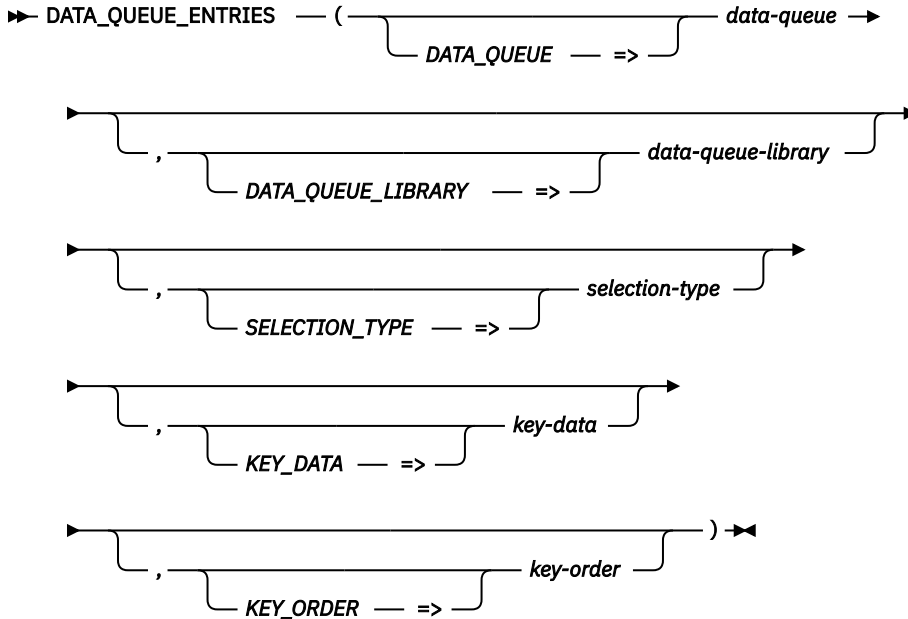
Distributed data management (DDM) data queues are not supported.

The MESSAGE_DATA, MESSAGE_DATA_UTF8, and MESSAGE_DATA_BINARY columns contain identical values. It is up to the user to determine which data type is most appropriate for working with the result data.

The values returned for the result columns of the table function are closely related to the values returned by the Retrieve Data Queue Message (QMHRDQM) API.

Authorization: The caller must have:

- *EXECUTE authority for the library, and
- *OBJPR and *READ authority for the *DTAQ



The schema is QSYS2.

data-queue A character or graphic string containing the name of the data queue.

data-queue-library A character or graphic string containing the name of the library containing the data queue. Can be one of the following special values:

- *CURLIB** The job's current library is used.
- *LIBL** The library list is used. This is the default.

selection-type A character string indicating how the messages are to be returned. If this parameter is omitted, ALL is used.

- ALL** All messages are to be returned in the order based on the type of data queue. FIFO queues are returned in FIFO order, LIFO queues are returned in LIFO order, and keyed queues are returned in ascending key order.
- FIRST** The first message is to be returned. This option is not allowed for a keyed data queue.
- KEY** Messages meeting the key criteria are to be returned. This option is only valid for a keyed data queue. When this value is specified, *key-data* and *key-order* must also be specified.
- LAST** The last message is to be returned. This option is not allowed for a keyed data queue.
- REVERSE** All messages are to be returned in reverse order of the type of data queue. For example, LIFO queues are returned in FIFO order. This option is not allowed for a keyed data queue.

key-data A character string containing the data to use as the key for receiving a message from the data queue. This parameter is required for a keyed data queue. It must not be specified for a non-keyed data queue.

The length of *key-data* must be the length specified on the KEYLEN parameter on the Create Data Queue (CRTDTAQ) command. The KEY_LENGTH column of the QSYS2.DATA_QUEUE_INFO view contains this value.

When this parameter is specified, *key-order* must also be specified.

key-order The comparison criteria between the keys of messages on the data queue and the *key-data* parameter. Valid values are:

- EQ** All messages with a key equal to *key-data* are to be returned.
- GE** All messages with a key greater than or equal to *key-data* are to be returned.
- GT** All messages with a key greater than *key-data* are to be returned.
- LE** All messages with a key less than or equal to *key-data* are to be returned.
- LT** All messages with a key less than *key-data* are to be returned.
- NE** All messages with a key not equal to *key-data* are to be returned.

This parameter is ignored if *key-data* is not specified.

The result of the function is a table containing one or more rows with the format shown in the following table. If no messages are selected, no rows are returned. All the columns are nullable.

Table 86. DATA_QUEUE_ENTRIES table function

Column Name	Data Type	Description
ORDINAL_POSITION	INTEGER	The relative position of this row in the result data set.
DATA_QUEUE_LIBRARY	VARCHAR(10)	The library in which the data queue was found.
DATA_QUEUE	VARCHAR(10)	The name of the data queue.
MESSAGE_DATA	CLOB(64512)	The message received from the data queue as character data.
MESSAGE_DATA_UTF8	CLOB(64512) CCSID 1208	The message received from the data queue represented as character data in CCSID 1208.
MESSAGE_DATA_BINARY	BLOB(64512)	The message received from the data queue in binary form. This is the raw form of the data.
KEY_DATA	VARCHAR(256)	For a keyed data queue, the key value of the returned message. This is the actual key value, which could be different than the <i>key-data</i> parameter value. Contains the null value if this is not a keyed data queue.
MESSAGE_ENQUEUE_TIMESTAMP	TIMESTAMP(0)	The date and time that the message was placed on the data queue.
SENDER_JOB_NAME	VARCHAR(28)	The qualified job name of the sender. Contains the null value if no sender information is available for the message.
SENDER_CURRENT_USER	VARCHAR(10)	The current user profile of the sender. Contains the null value if no sender information is available for the message.

Example

Look at all the messages in data queue DQ1 in TESTLIB.

```
SELECT * FROM TABLE(QSYS2.DATA_QUEUE_ENTRIES(
                                DATA_QUEUE => 'DQ1',
                                DATA_QUEUE_LIBRARY => 'TESTLIB'))
ORDER BY ORDINAL_POSITION;
```

DATA_QUEUE_INFO view

The DATA_QUEUE_INFO view returns a row for every data queue.

The information returned is similar to the information available through the Retrieve Data Queue Description (QMHQRDQD) API.

Authorization: The caller must have:

- *OBJOPR and *READ authority to the data queue, and
- *EXECUTE authority to the library containing the data queue.

The following table describes the columns in the view. The system name is DTAQ_INFO. The schema is QSYS2.

Table 87. DATA_QUEUE_INFO view

Column name	System column name	Data type	Description
DATA_QUEUE_LIBRARY	DTAQ_LIB	VARCHAR(10)	Library containing the data queue.
DATA_QUEUE_NAME	DTAQ	VARCHAR(10)	Name of the data queue.
DATA_QUEUE_TYPE	DTAQ_TYPE	VARCHAR(8)	This will be set to one of the following values: DDM The data queue is a DDM data queue. STANDARD The data queue is a standard data queue.
MAXIMUM_MESSAGE_LENGTH	MSG_LENGTH	INTEGER Nullable	The maximum length allowed for messages. This is the value that was specified with the MAXLEN keyword on the CRTDTAQ command. Contains the null value for a DDM data queue.
SEQUENCE	SEQUENCE	VARCHAR(5) Nullable	The sequence in which messages can be removed from the queue. FIFO First-in first-out KEYED Keyed LIFO Last-in first-out Contains the null value for a DDM data queue.
KEY_LENGTH	KEY_LENGTH	INTEGER Nullable	The length of the message reference key for a keyed data queue, in bytes. Values are from 1 to 256. Contains the null value if this is not a keyed queue or is a DDM data queue.
INCLUDE_SENDER_ID	SENDER_ID	VARCHAR(3) Nullable	Indicates if the queue was created to include the sender ID with sent messages. NO The sender ID is not included when data is sent to the data queue. YES The sender ID is included when data is sent to the data queue. Contains the null value for a DDM data queue.
CURRENT_MESSAGES	CUR_MSGS	INTEGER Nullable	The number of messages currently on the data queue. Contains the null value for a DDM data queue.
MAXIMUM_MESSAGES	MAX_MSGS	INTEGER Nullable	The maximum number of messages that will fit into the data queue. Contains the null value for a DDM data queue.

Table 87. DATA_QUEUE_INFO view (continued)

Column name	System column name	Data type	Description
SPECIFIED_MAXIMUM_MESSAGES	SPEC_MAX	INTEGER Nullable	The maximum number of messages that was specified on the SIZE keyword of the CRTDTAQ command. Can contain the following special values: -1 *MAX16MB was specified for the data queue size. -2 *MAX2GB was specified for the data queue size. Contains the null value for a DDM data queue.
INITIAL_MESSAGE_ALLOCATION	INIT_ALOC	INTEGER Nullable	The number of messages that will fit into the storage allocated for the data queue when it is created or when it is automatically reclaimed. Contains the null value for a DDM data queue.
CURRENT_MESSAGE_ALLOCATION	CUR_ALOC	INTEGER Nullable	The number of entries that will fit into the data queue before it is extended. When the queue is extended, additional storage is allocated for the queue. The data queue can be extended until it reaches the value for the maximum number of entries allowed. Contains the null value for a DDM data queue.
FORCE	FORCE	VARCHAR(3) Nullable	Whether the data queue is forced to auxiliary storage when entries are sent or received. NO The data queue is not forced to auxiliary storage after entries are sent or received. YES The data queue is forced to auxiliary storage after entries are sent or received. Contains the null value for a DDM data queue.
AUTOMATIC_RECLAIM	RECLAIM	VARCHAR(3) Nullable	Whether or not the data queue has the amount of storage allocated for the queue reclaimed when the queue is empty. NO Storage is not reclaimed. YES Storage is reclaimed when the queue is empty. The amount of storage allocated will be set to the initial number of entries. Contains the null value for a DDM data queue.
LAST_RECLAIM_TIMESTAMP	RECLAIM_TS	TIMESTAMP Nullable	The date and time that the last automatic reclaim was done. Contains the null value for a DDM data queue or when no reclaim has occurred for a standard data queue.
ENFORCE_DATA_QUEUE_LOCKS	DTAQ_LOCKS	VARCHAR(3) Nullable	Identifies whether IBM-supplied data queue operations will enforce a lock on the data queue. This attribute cannot be specified on the Create Data Queue (CRTDTAQ) CL Command. The default when a data queue is created is for locks to be ignored. A data queue can be locked with the Allocate Object (ALCOBJ) CL Command. When locks are enforced, performance can be degraded due to the additional locking performed by all data queue operations. NO Locks on the data queue are ignored by IBM-supplied data queue operations. YES Locks on the data queue are enforced by IBM-supplied data queue operations. Contains the null value for a DDM data queue.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The text description of the data queue. Contains the null value if there is no text description.

Table 87. DATA_QUEUE_INFO view (continued)

Column name	System column name	Data type	Description
REMOTE_DATA_QUEUE_LIBRARY	RMT_DTAQL	VARCHAR(10) Nullable	<p>The name of the library for the remote data queue on the target system. This is the data queue name that was specified on the RMTDTAQ parameter of the CRTDTAQ command. Can contain the following special values:</p> <p>*CURLIB The current library will be used to find the data queue.</p> <p>*LIBL The library list will be used to find the data queue.</p> <p>Contains the null value for a standard data queue.</p>
REMOTE_DATA_QUEUE	RMT_DTAQ	VARCHAR(10) Nullable	<p>The name of the remote data queue on the target system. This is the data queue name that was specified on the RMTDTAQ parameter of the CRTDTAQ command.</p> <p>Contains the null value for a standard data queue.</p>
REMOTE_LOCATION	REMOTE_LOC	VARCHAR(8) Nullable	<p>The name of the remote location that is used with this object. This is the name that was specified on the RMTLOCNAME parameter of the CRTDTAQ command. Can contain the following special value:</p> <p>*RDB The remote location information from the relational database entry returned in the relational database entry name field is used to determine the remote system.</p> <p>Contains the null value for a standard data queue.</p>
RELATIONAL_DATABASE_NAME	RDBNAME	VARCHAR(18) Nullable	<p>The name of the relational database entry that identifies the target system or target ASP group. This is the name that was specified on the RDB parameter of the CRTDTAQ command.</p> <p>Contains the null value for a standard data queue or a data queue that is not an RDB type DDM data queue.</p>
APPC_DEVICE_DESCRIPTION	APPC_DEVD	VARCHAR(10) Nullable	<p>The name of the APPC device description on the source system that is used with this DDM data queue. This is the name that was specified on the DEV parameter of the CRTDTAQ command. Can contain the following special value:</p> <p>*LOC The device associated with the remote location is used.</p> <p>Contains the null value for a standard data queue and for RDB type DDM data queues.</p>
LOCAL_LOCATION	LOCAL_LOC	VARCHAR(8) Nullable	<p>The name of the local location. This is the name that was specified on the LCLLOCNAME parameter of the CRTDTAQ command. Can contain the following special values:</p> <p>*LOC The device associated with the remote location is used.</p> <p>*NETATR The LCLLOCNAME value specified in the system network attributes is used.</p> <p>Contains the null value for a standard data queue and for RDB type DDM data queues.</p>
MODE	MODE	VARCHAR(8) Nullable	<p>The mode name used with the remote location name to communicate with the target system. This is the name that was specified on the MODE parameter of the CRTDTAQ command. Can contain the following special value:</p> <p>*NETATR The mode name specified in the system network attributes is used.</p> <p>Contains the null value for a standard data queue and for RDB type DDM data queues.</p>

Table 87. DATA_QUEUE_INFO view (continued)

Column name	System column name	Data type	Description
REMOTE_NETWORK_ID	REMOTE_NET	VARCHAR(8) Nullable	<p>The remote network identifier in which the remote location used to communicate with the target system. This is the name that was specified on the RMTNETID parameter of the CRTDTAQ command. Can contain the following special values:</p> <p>*LOC The remote network ID associated with the remote location is used.</p> <p>*NETATR The remote network ID value specified in the system network attributes is used.</p> <p>*NONE No remote network ID is used.</p> <p>Contains the null value for a standard data queue and for RDB type DDM data queues.</p>

Example

Find the number of messages currently on data queue DQ1 in TESTLIB.

```
SELECT CURRENT_MESSAGES FROM QSYS2.DATA_QUEUE_INFO
WHERE DATA_QUEUE_LIBRARY = 'TESTLIB' AND DATA_QUEUE_NAME = 'DQ1';
```

DB_TRANSACTION_INFO view

The DB_TRANSACTION_INFO view returns one row for each commitment definition.

The values returned for the columns in the view are similar to the values returned by the Work with Commitment Definitions (WRKCMTDFN) CL command and by the Database Transactions and Global Transactions lists in ACS.

Authorization: The caller must have *JOBCTL user special authority, or QIBM_DB_SQLADM or QIBM_DB_SYSMON function usage authority.

The following table describes the columns in the view. The system name is TRANS_INFO. The schema is QSYS2.

Table 88. DB_TRANSACTION_INFO view

Column Name	System Column Name	Data Type	Description
COMMITMENT_DEFINITION	COMMIT_DFN	VARCHAR(10)	The name of the commitment definition.
COMMITMENT_DEFINITION_DESCRIPTION	TEXT	VARCHAR(50) Nullable	<p>The text description of the commitment definition.</p> <p>Contains the null value if there is no text description.</p>
COMMITMENT_DEFINITION_ID	CMTDEF_ID	VARCHAR(10) FOR BIT DATA	<p>The ID associated with the commitment definition. Can contain the following special values:</p> <p>*DFACTGRP This is the commitment definition for the default activation group.</p> <p>*JOB The commitment definition is shared by activation groups within a job.</p>
JOB_NAME	JOB_NAME	VARCHAR(28)	The qualified job name that is using commitment control.

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
THREAD_ASSOCIATION_STATUS	THD_STATUS	VARCHAR(11) Nullable	<p>The thread association status of the *TNSOBJ scoped commitment definition.</p> <p>*ATTACHED There are threads attached to the *TNSOBJ scoped commitment definition.</p> <p>*UNATTACHED There are no threads attached to the *TNSOBJ scoped commitment definition.</p> <p>Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.</p>
LOGICAL_UNIT_OF_WORK_ID	LUWID	VARCHAR(39)	<p>The identifier (ID) of the current logical unit of work (LUW). This value can be used to find associated commitment definitions on this system and on other systems. The logical unit of work ID (LUWID) is a character string containing the network-qualified logical unit (LU) name, the instance number, and the sequence number.</p> <p>The network-qualified LU name consists of a network ID with a maximum of 8 characters, a period delimiter, followed by a LU name with a maximum of 8 characters. The instance number is a 12 character value, each character representing a single hexadecimal digit. The sequence number is a decimal value with values ranging from 1 through 65535.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOGICAL_UNIT_OF_WORK_STATE	LUW_STATE	VARCHAR(20)	<p>The current state of the logical unit of work (LUW).</p> <p>COMMIT IN PROGRESS A commit operation is in progress. The system is attempting to commit the resources for this commitment definition and all downstream resources associated with this LUW.</p> <p>COMMITTED A commit operation is in progress. The resources for this commitment definition and all downstream resources associated with this LUW have been committed.</p> <p>LAST AGENT PENDING A commit operation is in progress. The resources for this commitment definition have been prepared and the system is waiting for the last agent selected by this commitment definition to return the commit or rollback decision.</p> <p>PREPARE IN PROGRESS A commit or prepare operation is in progress. The system is attempting to prepare the resources for this commitment definition and to prepare all downstream resources associated with this LUW.</p> <p>PREPARED A commit operation is in progress. The resources for this commitment definition and all downstream resources associated with this LUW have been prepared.</p> <p>RESET A commit, rollback, or prepare operation is not in progress for this commitment definition.</p> <p>ROLLBACK IN PROGRESS A rollback operation is in progress. The system is attempting to roll back the committable changes for this commitment definition and for all the committable changes associated with this commitment definition.</p> <p>ROLLBACK REQUIRED A rollback operation is required. The resources registered to this commitment definition cannot be used until the LUW is rolled back.</p> <p>VOTE READ ONLY A commit operation is in progress. The commitment definition had no changes to commit, the Vote Read Only commitment option is set to 'Yes', and all downstream locations voted read only. This commitment definition and any downstream locations will not participate in the committed wave during this commit operation.</p>
STATE_TIMESTAMP	STATE_TIME	TIMESTAMP(0) Nullable	<p>The timestamp when the logical unit of work started or became undecided.</p> <p>Contains the null value if the commitment definition has performed no work and is not undecided.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCK_SPACE_ID	LOCKID	CHAR(20)	The lock space identifier for this commitment definition.
ACTIVATION_GROUP	ACTGRP	BIGINT Nullable	The activation group for this commitment definition. Contains the null value if the commitment definition is the job commitment definition (*JOB), an explicitly named commitment definition, or a commitment definition that represents an XA transaction branch with lock space scoped locks.
ASPGRP	ASPGRP	VARCHAR(10)	The ASP group for this commitment definition. If *SYSBAS, the commitment definition resides on the system auxiliary storage pool.
RESOURCE_LOCATION	RSC_LOC	VARCHAR(6)	Indicates whether local or remote resources are currently under commitment control. BOTH Both local and remote resources are under commitment control. LOCAL Local resources are under commitment control. NONE No resources are under commitment control. REMOTE Remote resources are under commitment control. API resources, and relational database resources with remote location name *ARDPGM, are considered to be local, even though they may represent objects on a remote system. Because these resources are accessed and updated by user programs, the system cannot determine whether these programs access objects on a remote system.
DEFAULT_LOCK_LEVEL	DFT_LCKL	VARCHAR(4)	The level of record locking on the records in each database file under commitment control for the commitment definition when files are opened using traditional system interfaces. For files opened using SQL interfaces, the effective isolation level for each SQL operation determines the level of record locking that is used. *ALL Every record that is accessed in a file under commitment control is locked until the logical unit of work is committed or rolled back. *CHG Every record that is changed in a file under commitment control is locked until the logical unit of work is committed or rolled back. *CS Every record that is changed in a file under commitment control is locked until the logical unit of work is committed or rolled back. Records that are accessed but not changed are locked only until a different record is accessed.
USER_NAME	USER_NAME	VARCHAR(10)	The user profile that started the transaction.

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCAL_CHANGES_PENDING	LOCAL_PEND	VARCHAR(3)	<p>Indicates whether this commitment definition has local changes pending. A local pending change means:</p> <ul style="list-style-type: none"> • One or more record level changes are pending. SQL statements run against the relational database might or might not have committable changes to the local database. • One or more object level changes are pending. • An API resource is registered that does not allow save-while-active requests to perform normally or does not allow independent ASPs to be quiesced. <p>NO This commitment definition has no local changes pending.</p> <p>YES This commitment definition has local changes pending.</p>
LOCAL_JOURNAL_CHANGES_PENDING	LOCAL_JRN	VARCHAR(3)	<p>Indicates this commitment definition has journals with changes pending.</p> <p>This means that one or more local journals have changes pending.</p> <p>NO This commitment definition has no local journal changes pending.</p> <p>YES This commitment definition has local journal changes pending.</p>
LOCAL_RECORD_CHANGES_PENDING	LOCAL_RCD	VARCHAR(3)	<p>Indicates whether this commitment definition has records with changes pending.</p> <p>This means that one or more local records have insert, update, or delete changes pending.</p> <p>NO This commitment definition has no local record level changes pending.</p> <p>YES This commitment definition has local record level changes pending.</p>
LOCAL_OBJECT_CHANGES_PENDING	LOCAL_OBJ	VARCHAR(3)	<p>Indicates whether this commitment definition has objects with changes pending.</p> <p>This means that one or more local objects have changes pending.</p> <p>NO This commitment definition has no local object level changes pending.</p> <p>YES This commitment definition has local object level changes pending.</p>
LOCAL_API_CHANGES_PENDING	LOCAL_API	VARCHAR(3)	<p>Indicates whether this commitment definition has local API changes pending.</p> <p>This means that one or more API resources are registered that do not allow save-while-active requests to perform normally or do not allow independent ASPs to be quiesced.</p> <p>NO This commitment definition has no local API resource changes pending.</p> <p>YES This commitment definition has local API resource changes pending.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCAL_RDB_CHANGES_PENDING	LOCAL_RDB	VARCHAR(3)	<p>Indicates whether this commitment definition has local RDB changes pending.</p> <p>This means that one or more record or object level changes are pending for a local RDB resource. SQL statements run against the relational database may or may not have committable changes to the local database.</p> <p>NO This commitment definition has no local RDB changes pending.</p> <p>YES This commitment definition has local RDB changes pending.</p>
REMOTE_CHANGES_PENDING	RMT_PEND	VARCHAR(3)	<p>Indicates whether this commitment definition has remote changes pending. A remote pending change means:</p> <ul style="list-style-type: none"> • One or more record or object level changes are pending for an RDB resource. SQL statements run against the relational database may or may not have committable changes to the remote database. • One or more record or object level changes are pending for a TCP/IP connection resource. • One or more record or object level changes are pending for a remote file resource. • One or more record or object level changes are pending for a conversation resource. <p>NO This commitment definition has no remote changes pending.</p> <p>YES This commitment definition has remote changes pending.</p>
REMOTE_RDB_CHANGES_PENDING	RMT_RDB	VARCHAR(3)	<p>Indicates whether this commitment definition has remote RDB changes pending. This means:</p> <ul style="list-style-type: none"> • One or more record or object level changes are pending for a remote RDB resource. SQL statements run against the relational database may or may not have committable changes to the remote database. • One or more record or object level changes are pending for a TCP/IP connection resource. <p>NO This commitment definition has no remote RDB changes pending.</p> <p>YES This commitment definition has remote RDB changes pending.</p>
REMOTE_FILE_CHANGES_PENDING	RMT_FILE	VARCHAR(3)	<p>Indicates whether this commitment definition has remote file changes pending. This means that one or more record or object level changes are pending for a remote file resource.</p> <p>NO This commitment definition has no remote file changes pending.</p> <p>YES This commitment definition has remote file changes pending.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
REMOTE_CONVERSATION_CHANGES_PENDING	RMT_CONV	VARCHAR(3)	<p>Indicates whether this commitment definition has remote conversation changes pending. This means that one or more record or object level changes are pending for a conversation resource. SQL statements run against the relational database may or may not have committable changes to the remote database.</p> <p>NO This commitment definition has no remote conversation changes pending.</p> <p>YES This commitment definition has remote conversation changes pending.</p>
ROLE	ROLE	VARCHAR(19) Nullable	<p>Indicates the role that this location is playing in the transaction program network.</p> <p>AGENT This location is a leaf in the transaction program network and it has no subordinate locations.</p> <p>CASCADER This location is an intermediate node in the transaction program network and it has subordinate locations. The commit or rollback operation did not start at this location.</p> <p>INITIATOR This location started the commit or rollback operation and is at the root of the transaction program network.</p> <p>LAST AGENT This location is a leaf in the transaction program network and it is the last agent, which decides whether to commit or to roll back the logical unit of work.</p> <p>LAST AGENT CASCADER This location is a middle node in the transaction program network and it is the last agent, which decides whether the logical unit of work is committed or rolled back. This location also can delegate another location of its choice to be the last agent.</p> <p>LOCAL This location is the only location with resources registered to the commitment definition.</p> <p>X/OPEN AGENT This location is a leaf in the transaction program network. The commit operation was initiated by an X/Open transaction manager.</p> <p>X/OPEN CASCADER This location is a middle node in the transaction program network. The commit operation was initiated by an X/Open transaction manager and this location has subordinate locations.</p> <p>Contains the null value if the role is not recognized.</p>
RESYNC_IN_PROGRESS	RESYNCING	VARCHAR(3)	<p>Indicates whether this commitment definition is resynchronizing resources that are associated with the logical unit of work.</p> <p>NO This commitment definition is not resynchronizing resources.</p> <p>YES This commitment definition is resynchronizing resources.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
HEURISTIC_OPERATION	HEUR_OPER	VARCHAR(20) Nullable	<p>The user-specified operation (resulting from a heuristic decision) performed against this commitment definition's resources and against all of the downstream resources associated with the logical unit of work.</p> <p>COMMIT IN PROGRESS The user forced a commit operation that has not yet completed.</p> <p>FORCED COMMIT The user forced a commit operation.</p> <p>FORCED ROLLBACK The user forced a rollback operation.</p> <p>ROLLBACK IN PROGRESS The user forced a rollback operation that has not yet completed.</p> <p>Contains the null value if the user has not specified a heuristic operation.</p>
JOB_ACTIVE	JOB_ACTIVE	VARCHAR(6)	<p>Indicates whether this commitment definition is part of an active job.</p> <p>NO This commitment definition is not part of an active job. The job was ended before all logical units of work were complete.</p> <p>SERVER The job that the commitment definition was part of has ended, and the commitment definition has been made part of a database server job. The SERVER_JOB_NAME column contains the name of the database server job.</p> <p>YES This commitment definition is part of an active job.</p>
SERVER_JOB_NAME	SERVER_JOB	VARCHAR(28) Nullable	<p>Qualified job name of the database server job of which this commitment definition is a part.</p> <p>Contains the null value if the commitment definition is not part of a database server job.</p>
LOCK_SCOPE	LOCK_SCOPE	VARCHAR(7) Nullable	<p>Indicates where the locks acquired on behalf of the commitment definition are scoped.</p> <p>*ACTGRP The commitment definition has an activation group level scope.</p> <p>*EXPL The commitment definition is explicitly named scope.</p> <p>*JOB The commitment definition has a job level scope.</p> <p>*TNSOBJ The commitment definition has a transaction object level scope (XA).</p> <p>Contains the null value if the lock scope is not recognized.</p>
LOCK_WAIT_TIME	LOCK_WAIT	INTEGER	<p>The maximum number of seconds that the system waits to acquire a lock requested on behalf of this commitment definition. Lock wait time values that are specified or defaulted using system interfaces other than the xa_open API are used if they are smaller than this value, or if this value is zero.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCK_LIMIT	LOCK_LIMIT	BIGINT	<p>Indicates the maximum number of records which can be locked within a transaction. If multiple journals are involved in the transaction, this limit applies to each journal, not the transaction as a whole.</p> <p>The COMMITMENT_CONTROL_LOCK_LIMIT QAQQINI option can be used to configure this value before commitment control is started.</p>
NESTING_LEVEL_DEPTH	NEST_LEVEL	INTEGER	The current nesting depth of all nested transactions for this commitment definition.
NUMBER_SAVEPOINTS	SAVEPOINTS	INTEGER	The number of active savepoints.
COMMIT_ROLLBACK_END	CR_END	VARCHAR(3)	<p>Indicates if COMMIT, ROLLBACK, and ENDCMTCTL are allowed to be executed.</p> <p>NO COMMIT, ROLLBACK, and ENDCMTCTL are not allowed.</p> <p>YES COMMIT, ROLLBACK, and ENDCMTCTL are allowed.</p>
EXPLICIT_COMMIT_ROLLBACK	EXPL_CR	VARCHAR(3)	<p>Indicates if an explicit COMMIT or ROLLBACK was executed.</p> <p>NO An explicit COMMIT or ROLLBACK has not been executed.</p> <p>YES An explicit COMMIT or ROLLBACK has been executed.</p>
SQL_DYNAMIC_COMMIT	SQL_DYNC	VARCHAR(3)	<p>Indicates if an SQL Dynamic COMMIT was executed.</p> <p>NO An SQL dynamic COMMIT has not been executed.</p> <p>YES An SQL dynamic COMMIT has been executed.</p>
SQL_DYNAMIC_ROLLBACK	SQL_DYNR	VARCHAR(3)	<p>Indicates if an SQL Dynamic ROLLBACK was executed.</p> <p>NO An SQL dynamic ROLLBACK has not been executed.</p> <p>YES An SQL dynamic ROLLBACK has been executed.</p>
SQL_HOLD_COMMIT	SQL_HOLD	VARCHAR(3)	<p>Indicates the SQL HOLD value used on the call to COMMIT.</p> <p>NO SQLHOLD(NO) specified on COMMIT call.</p> <p>YES SQLHOLD(YES) specified on COMMIT call.</p>
COMMIT_DURABLE	DURABLE	VARCHAR(3)	<p>Indicates whether commit operations are guaranteed to be durable.</p> <p>NO The commit operation is not durable. Atomicity of the transaction is guaranteed, but one or more transactions may be lost in the event of a system failure.</p> <p>YES The commit operation is durable. When the transaction is committed, the transaction is guaranteed to persist on the system.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
NUMBER_COMMITS	COMMIT_OPS	BIGINT Nullable	The total number of commit operations performed since this commitment definition was created. This number includes commit operations initiated by both the user and the operating system. Returns 2,147,483,648 if the number of commits has exceeded 2,147,483,647. Contains the null value if an IPL has been performed on the system since the commitment definition was created.
NUMBER_ROLLBACKS	ROLLB_OPS	BIGINT Nullable	The total number of rollback operations performed since this commitment definition was created. This number includes rollback operations initiated by both the user and the operating system. Returns 2,147,483,648 if the number of rollbacks has exceeded 2,147,483,647. Contains the null value if an IPL has been performed on the system since the commitment definition was created.
DEFAULT_JOURNAL_LIBRARY	DFT_JRNLIB	VARCHAR(10) Nullable	The name of the library in which the default journal is located. Contains the null value if there is no default journal.
DEFAULT_JOURNAL_NAME	DFT_JRN	VARCHAR(10) Nullable	The name of the default journal that was specified when the commitment definition was created. All logical unit of work (LUW) entries are placed in this journal. Contains the null value if there is no default journal.
NOTIFY_OBJECT_TYPE	NOTIFY_TYP	VARCHAR(7) Nullable	Type of notify object for the commitment definition. <p>*DTAARA The notify object is a data area.</p> <p>*FILE The notify object is a database file member</p> <p>*MSGQ The notify object is a message queue.</p> <p>Contains the null value if there is no notification object.</p>
NOTIFY_OBJECT_LIBRARY	NOTIFY_LIB	VARCHAR(10) Nullable	The name of the library in which the notify object for the commitment definition can be located. Contains the null value if there is no notification object.
NOTIFY_OBJECT	NOTIFY_OBJ	VARCHAR(10) Nullable	The name of the object to which notification is sent regarding the status of the commitment definition. Contains the null value if there is no notification object.
NOTIFY_OBJECT_MEMBER	NOTIFY_MBR	VARCHAR(10) Nullable	The name of the database file member that is the notify object for the commitment definition. Contains the null value if there is no notification object or NOTIFY_OBJECT_TYPE is not *FILE.

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
WAIT_FOR_OUTCOME	WAIT_OUTC	VARCHAR(17) Nullable	<p>This commitment option indicates how this system handles resynchronization if a communications or remote system failure occurs during a commit or rollback operation. This value is set to WAIT when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API. Note: If the commitment definition has a Wait for Outcome value of WAIT or inherits a value of WAIT, the value of the ACCEPT_VOTE_RELIABLE commitment option is ignored, and the system behaves as though the ACCEPT_VOTE_RELIABLE commitment option is NO.</p> <p>INHERIT OR NOWAIT When this system is the initiator of the commit or rollback operation, the INHERIT OR NOWAIT value has the same effect as the NOWAIT value. When this system is not the initiator and the initiator supports the presumed abort protocol, the Wait for Outcome value is inherited from the initiator. When this system is not the initiator and the initiator does not support the presumed abort protocol, the INHERIT OR NOWAIT value has the same effect as the NOWAIT value.</p> <p>INHERIT OR WAIT When this system is the initiator of the commit or rollback operation, the INHERIT OR WAIT value has the same effect as the WAIT value. When this system is not the initiator and the initiator supports the presumed abort protocol, the Wait for Outcome value is inherited from the initiator. When this system is not the initiator and the initiator does not support the presumed abort protocol, the INHERIT OR WAIT value has the same effect as the WAIT value.</p> <p>NOWAIT This system attempts to resynchronize one time before allowing the commit or rollback operation to complete. If the initial attempt fails, the resynchronization is completed in a database server job, and the application is not notified of the result of the resynchronization. Resynchronizations required with the last agent location always wait regardless of the current status of the wait for outcome option.</p> <p>WAIT This system completes resynchronization before allowing the commit or rollback operation to complete.</p> <p>Contains the null value if the Wait for Outcome option was not specified for the commitment definition.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PROBLEM_ACTION	ACTION	VARCHAR(8)	<p>Indicates what this system does if another system, which controls whether the logical unit of work should commit or rollback, sends this system an unrecognized message or detects damage in the logical unit of work. This value is set to ROLLBACK when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API.</p> <p>COMMIT The changes associated with this logical unit of work are committed.</p> <p>ROLLBACK The changes associated with this logical unit of work are rolled back.</p>
VOTE_READ_ONLY_PERMITTED	VOTE_RO	VARCHAR(3)	<p>Indicates whether this system can vote read-only in response to a commit operation started on another system. Note: If this system votes read-only, control is not returned to the application until this system receives a message from the starting system containing an indicator that the next logical unit of work has started. The length of the delay in regaining control depends on the application running on the starting system. This value is set to NO when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API.</p> <p>NO This system cannot vote read-only.</p> <p>YES This system can vote read-only.</p>
END_JOB_ACTION	ENDJOB_ACT	VARCHAR(8) Nullable	<p>Indicates what this system does if the logical unit of work is in an undecided state and the job is ending abnormally. This value is set to WAIT when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API.</p> <p>COMMIT This system commits the changes made to the logical of work.</p> <p>ROLLBACK This system rolls back the changes made to the logical unit of work.</p> <p>WAIT The system gets the commit or rollback decision from the system that started the operation. Based on that decision, this system commits or rolls back the changes made to the logical unit of work before ending the job. Heuristic operations can be performed if the time spent waiting is unacceptable.</p> <p>Contains the null value if the end job action is not known.</p>
LAST_AGENT_PERMITTED	LAST_AGENT	VARCHAR(6) Nullable	<p>Indicates whether a last agent can be selected when one is eligible to be selected during a commit operation. A last agent is eligible to be selected at the location that initiates a commit operation, and at locations that are selected as a last agent by the location that propagates the commit operation to that location. This value is set to SYSTEM when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API.</p> <p>NO The system is not allowed to select a last agent.</p> <p>SYSTEM The system is allowed to select a last agent.</p> <p>Contains the null value if the last agent permitted value is not known.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LEAVE_OUT_OK	LEAVE_OUT	VARCHAR(3)	<p>The commitment option that indicates whether it is OK to leave this system out of a commit operation initiated at another system. If this system indicates it can be left out, no communications flows are sent to this system during subsequent commit or rollback operations until a data flow is received from the initiator. Also, control is not returned to the application until the data flow is received. This value is set to NO when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API.</p> <p>NO This system may not be left out of subsequent logical units of work.</p> <p>YES This system may be left out of subsequent logical units of work.</p>
ACCEPT_VOTE_RELIABLE	ACCEPT_VR	VARCHAR(3)	<p>Indicates whether this system accepts the vote reliable indicator if it is received from its agents during the prepare wave of a commit operation. The vote reliable indicator indicates that it is unlikely that a heuristic decision will be made at the agent if a failure occurs before the committed wave completes. If an agent sends the vote reliable indicator, and this location accepts it, performance is improved because one communications flow is eliminated and control is returned to the application before the committed wave is completed for that agent. However, if a heuristic decision is made at that agent which causes heuristic damage, the application at this location will not receive an error message if the Accept vote reliable commitment option is set to YES. This value is set to YES when commitment control is started and can be changed using the Change Commitment Options (QTNCHGCO) API.</p> <p>If the commitment definition has a Wait for outcome value of Wait or inherits a value of Wait, the value of the Accept vote reliable commitment option is ignored, and the system behaves as though the Accept vote reliable commitment option is No.</p> <p>NO This system does not accept the vote reliable indicator.</p> <p>YES This system accepts the vote reliable indicator.</p>
RECYCLE_COUNT	RECYCLECNT	INTEGER	<p>The number of times that this commitment definition has been recycled as a spare.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
RECEIVED_DRDA_SYNCTYPE	RCV_SYNC	VARCHAR(17) Nullable	<p>The two-phase commitment control synchronization type received by commitment control on the target-side of a DRDA connection.</p> <p>COMMITTED The application requester committed its unit or work.</p> <p>FORGET The application requester issued a forget request.</p> <p>MIGRATE The application requester issued a migrate request.</p> <p>NEW UNIT OF WORK The application requester is requesting a new unit of work.</p> <p>PREPARE TO COMMIT The application requester is requesting the unit of work be prepared to commit.</p> <p>REQUEST TO COMMIT The application requester is requesting the unit of work be committed.</p> <p>ROLLBACK The application requester is requesting the unit of work be rolled back.</p> <p>Contains the null value if no two-phase commitment control requests have been received.</p>
RETURNED_DRDA_SYNCTYPE	RET_SYNC	VARCHAR(17) Nullable	<p>The two-phase commitment control synchronization type returned by commitment control on the target-side of a DRDA connection.</p> <p>COMMITTED The application server committed its unit or work.</p> <p>FORGET The application server returned a forget request.</p> <p>MIGRATE The application server returned a migrate request.</p> <p>REQUEST TO COMMIT The application server is returning the unit of work be committed.</p> <p>ROLLBACK The application server is returning the unit of work be rolled back.</p> <p>Contains the null value if no two-phase commitment control requests have been returned.</p>
XA_CONNECTION_TYPE	XA_TYPE	VARCHAR(8) Nullable	<p>Indicates whether this commitment definition was created in an SQL Server Job for the purposes of running XA transactions.</p> <p>CLI The XA transactions are being performed in an SQL Server Job over a CLI connection.</p> <p>EMBEDDED The XA transactions are being performed in an SQL Server Job over an embedded SQL connection.</p> <p>NONE The XA transactions are not performed in an SQL Server Job.</p> <p>Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
XA_RMID	XA_RMID	INTEGER Nullable	The resource manager ID (RMID). The unique number is assigned by the transaction manager to identify this instance of the XA resource manager within the thread of control. This RMID is passed on subsequent calls to XA routines to identify the resource manager. The identifier remains constant until the transaction manager in this thread closes the resource manager. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_TRANSACTION_MANAGER	XA_MANAGER	VARCHAR(10) Nullable	The name of the transaction manager that started the XA transaction branch represented by this commitment definition. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_RDB_NAME	XA_RDB	VARCHAR(18) Nullable	The relational database name (RDB) associated with the XA transaction represented by this commitment definition. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_TRANSACTION_BRANCH_STATE	XA_STATE	VARCHAR(23) Nullable	Indicates the state of the XA transaction branch represented by this commitment definition. ACTIVE One or more threads of control are actively associated with the transaction branch. IDLE No threads of control are actively associated with the transaction branch. PREPARED The transaction branch has been prepared. ROLLBACK ONLY The transaction branch is required to roll back. HEURISTICALLY COMPLETED The transaction branch has been heuristically committed or rolled back. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_XID_FORMAT_ID	XA_FMTID	VARBINARY(4) Nullable	The hex value of the format identifier portion of the XA transaction identifier (XID). Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_XID_GLOBAL_TRANSACTION_ID	XA_GTRID	VARBINARY(64) Nullable	The hex value of the global transaction identifier (GTRID) portion of the XA transaction identifier (XID). Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_XID_BRANCH_QUALIFIER	XA_BQUAL	VARBINARY(64) Nullable	The hex value of the branch qualifier (BQUAL) portion of the XA transaction identifier (XID). Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
XA_OPEN_SQL_HOLD	XA_SQLHOLD	CHAR(1) Nullable	<p>Indicates how SQL cursors are affected by some XA operations. Specified on the db2xa_open() API xainfo string SQLHOLD keyword.</p> <p>A</p> <ul style="list-style-type: none"> • db2xa_commit() and db2xa_rollback(): Cursors are not affected since db2xa_end() already closed them. • db2xa_end() : All cursors are closed. <p>C</p> <ul style="list-style-type: none"> • db2xa_commit() and db2xa_rollback(): Cursors declared WITH HOLD are held open. Cursors not declared WITH HOLD are closed. • db2xa_end(): All cursors are held open. <p>E</p> <ul style="list-style-type: none"> • db2xa_commit(): Cursors declared WITH HOLD are held open. Cursors not declared WITH HOLD are closed. • db2xa_rollback(): All cursors are closed. • db2xa_end(): Cursors declared WITH HOLD are held open. Cursors not declared WITH HOLD are closed. <p>L</p> <ul style="list-style-type: none"> • db2xa_commit(): If the relational database resides on an IBM i, all cursors are left open. Otherwise, cursors declared WITH HOLD are left open and cursors not declared WITH HOLD are closed. • db2xa_rollback(): If the relational database resides on an IBM i, all cursors are left open. Otherwise, all cursors are closed. • db2xa_end(): All cursors are held open. <p>N</p> <ul style="list-style-type: none"> • db2xa_commit(): Cursors declared WITH HOLD are held open. Cursors not declared WITH HOLD are closed. • db2xa_rollback(): All cursors are closed. • db2xa_end(): All cursors are held open. <p>Y</p> <ul style="list-style-type: none"> • db2xa_commit() and db2xa_rollback(): If the relational database resides on an IBM i, all cursors are left open. Otherwise, the operation will fail. • db2xa_end(): All cursors are held open. <p>Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition or it is a *TNSOBJ scoped commitment definition but not one where SQLHOLD was specified on the db2xa_open() API.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
XA_OPEN_TBLCS	XA_TBLCS	CHAR(1) Nullable	<p>Indicates whether loosely coupled threads of control (those working on transaction branches with the same global transaction identifier (GTRID), but different branch qualifiers (BQUALS)), share locks. This is what was specified on xa_open() API.</p> <p>N Locks are not shared. Resource deadlock may occur between the transaction branches.</p> <p>S Locks are shared. Resource deadlock will not occur between the transaction branches.</p> <p>Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition or it is a *TNSOBJ scoped commitment definition but not one where TBLCS was specified on the xa_open() API.</p>
XA_THREAD_OF_CONTROL	XA_THDCTL	VARCHAR(10) Nullable	<p>The thread of control for the XA transaction branch represented by this commitment definition.</p> <p>CONNECTION The SQL connection is the XA thread of control. The XA transaction branch represented by this commitment definition was started by the SQLSetConnectAttr API or by a remote system that established an SQL connection to this system. All SQL work requested using the connection that started the transaction branch becomes part of the transaction branch regardless of requesting thread.</p> <p>THREAD The thread is the XA thread of control. The XA transaction branch represented by this commitment definition was started by the xa_start or db2xa_start API. All SQL work requested by the thread that started the transaction branch becomes part of the transaction branch regardless of which SQL connection is used to perform that work.</p> <p>Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.</p>
XA_THREAD_ASSOCIATION_COUNT	XA_THDCNT	BIGINT Nullable	<p>The number of threads or SQL connections currently associated with the XA transaction branch represented by this commitment definition. The associations may be active or suspended.</p> <p>Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.</p>

Table 88. DB_TRANSACTION_INFO view (continued)

Column Name	System Column Name	Data Type	Description
XA_LOCK_SHARING	XA_LOCKSHR	VARCHAR(3) Nullable	Specifies whether locks are shared with loosely coupled transaction branches. A loosely coupled transaction branch is one with a transaction identifier (XID) that has the same global transaction identifier (GTRID) but a different branch qualifier (BQUAL). NO This commitment definition does not share locks with loosely coupled transaction branches. YES This commitment definition shares locks with loosely coupled transaction branches. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_LOCK_SPACE_HANDLE	XA_LOCKSP	INTEGER Nullable	The lock space associated space handle. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition.
XA_TRANSACTION_TIMEOUT	XA_TIMEOUT	INTEGER Nullable	The number of seconds that an XA transaction is allowed to exist before being automatically rolled back by the system. Contains the null value if the commitment definition is not a *TNSOBJ scoped commitment definition or if the number of seconds is 0.

Example

- Show all jobs that have local work that has not been committed. The rows that are returned are not ordered.

```
SELECT *
FROM QSYS2.DB_TRANSACTION_INFO
WHERE LOCAL_CHANGES_PENDING = 'YES';
```

- Show all jobs that have a commitment definition. Order the returned rows by job name, then job user, then by job number.

```
SELECT *
FROM QSYS2.DB_TRANSACTION_INFO
ORDER BY SUBSTR(SUBSTR(JOB_NAME,8),POSSTR(SUBSTR(JOB_NAME,8),'/')+1), --job name
SUBSTR(JOB_NAME,8,POSSTR(SUBSTR(JOB_NAME,8),'/')-1), -- job user
SUBSTR(JOB_NAME,1,6) -- job number
;
```

ENVIRONMENT_VARIABLE_INFO view

The ENVIRONMENT_VARIABLE_INFO view contains information about environment variables.

The values returned for the columns in the view are similar to the values returned by the WRKENVVAR CL command or [Get All System-Level Environment Variables API](#). Refer to the API for more detailed information.

Authorization: None required.

The following table describes the columns in the view. The system name is ENV_VARS. The schema is QSYS2.

Table 89. ENVIRONMENT_VARIABLE_INFO view

Column Name	System Column Name	Data Type	Description
ENVIRONMENT_VARIABLE_TYPE	VAR_TYPE	VARCHAR(6)	The type of environment variable. SYSTEM Defined as a system level environment variable. JOB Defined as a job level environment variable. This variable and value only apply to the current connection. PASE Defined as an IBM® Portable Application Solutions Environment for i (PASE for i) environment variable. This variable and value only apply to the current job. PASE variables are not returned unless the PASE environment has been started.
ENVIRONMENT_VARIABLE_NAME	VAR_NAME	VARGRAPHIC(128) CCSID 1200	The name of the environment variable. If the name is longer than 128 characters, it will be truncated with no warning. If ENVIRONMENT_VARIABLE_CCSID is 65535, the content of this column is set using the job default CCSID.
ENVIRONMENT_VARIABLE_VALUE	VAR_VALUE	VARGRAPHIC(1024) CCSID 1200 Nullable	The current value of the environment variable. If the value is longer than 1024 characters, it will be truncated with no warning. If ENVIRONMENT_VARIABLE_CCSID is 65535, the content of this column is set using the job default CCSID. Contains null if there is no value.
ENVIRONMENT_VARIABLE_BINARY_NAME	VAR_BNAME	VARBINARY(128)	The name of the environment variable in binary form. This is the raw value for the name. If the name is longer than 128 characters, it will be truncated with no warning.
ENVIRONMENT_VARIABLE_BINARY_VALUE	VAR_BVALUE	VARBINARY(1024) Nullable	The current value of the environment variable. This is the raw value for the value. If the value is longer than 1024 characters, it will be truncated with no warning. Contains null if there is no value.
ENVIRONMENT_VARIABLE_CCSID	VAR_CCSID	INTEGER	The CCSID value associated with the environment variable.

Example

Look at all system level environment variables and their values for this connection:

```
SELECT ENVIRONMENT_VARIABLE_NAME, ENVIRONMENT_VARIABLE_VALUE
FROM QSYS2.ENVIRONMENT_VARIABLE_INFO
WHERE ENVIRONMENT_VARIABLE_TYPE = 'SYSTEM'
```

EXIT_POINT_INFO view

The EXIT_POINT_INFO view returns information about exit points.

The values returned for the columns in the view are closely related to the values returned by the Work with Registration Information (WRKREGINF) CL command and by the Retrieve Exit Information (QUSRTEI, QusRetrieveExitInformation) API.

Authorization: None required.

The following table describes the columns in the view. The system name is EXIT_POINT. The schema is QSYS2.

Table 90. EXIT_POINT_INFO view

Column Name	System Column Name	Data Type	Description
EXIT_POINT_NAME	EXIT_NAME	VARCHAR(20)	The name of the exit point.

Table 90. EXIT_POINT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
EXIT_POINT_FORMAT	EXIT_FMT	CHAR(8)	The format name associated with the exit point.
REGISTERED	REGISTERED	VARCHAR(3)	Whether the exit point is registered with the registration facility. NO The exit point is unregistered. YES The exit point is registered.
ALLOW_DEREGISTRATION	DEREGISTER	VARCHAR(3)	Whether the exit point can be deregistered. NO The exit point cannot be deregistered. YES The exit point can be deregistered.
ALLOW_CHANGE	CHANGE	VARCHAR(3)	Whether the exit point controls can be changed. NO The exit point controls cannot be changed. YES The exit point controls can be changed.
EXIT_PROGRAMS	EXIT_PGMS	INTEGER	The current number of exit programs associated with the exit point.
MAXIMUM_EXIT_PROGRAMS	MAX_PGMS	INTEGER Nullable	The maximum number of exit programs that the exit point allows. Contains the null value if the maximum number of exit programs is *NOMAX.
TEXT_DESCRIPTION	TEXT	VARCHAR(132) Nullable	The descriptive text for the exit point. Contains the null value if no descriptive text is available.
ADD_EXIT_PROGRAM_LIBRARY	ADD_LIB	VARCHAR(10) Nullable	The library in which ADD_EXIT_PROGRAM resides. Contains the null value if no add exit program is defined.
ADD_EXIT_PROGRAM	ADD_PGM	VARCHAR(10) Nullable	The preprocessing exit program name that is called by the registration facility when the Add Exit Program API is called for the exit point. Contains the null value if no add exit program is defined.
ADD_EXIT_PROGRAM_FORMAT	ADD_FMT	CHAR(8) Nullable	The format name for ADD_EXIT_PROGRAM. Contains the null value if no add exit program is defined.
REMOVE_EXIT_PROGRAM_LIBRARY	RMV_LIB	VARCHAR(10) Nullable	The library in which REMOVE_EXIT_PROGRAM resides. Contains the null value if no remove exit program is defined.
REMOVE_EXIT_PROGRAM	RMV_PGM	VARCHAR(10) Nullable	The preprocessing exit program name that is called by the registration facility when the Remove Exit Program API is called for the exit point. Contains the null value if no remove exit program is defined.

Table 90. EXIT_POINT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
REMOVE_EXIT_PROGRAM_FORMAT	RMV_FMT	CHAR(8) Nullable	The format name for REMOVE_EXIT_PROGRAM. Contains the null value if no remove exit program is defined.
RETRIEVE_EXIT_PROGRAM_LIBRARY	RTV_LIB	VARCHAR(10) Nullable	The library in which RETRIEVE_EXIT_PROGRAM resides. Contains the null value if no retrieve exit program is defined.
RETRIEVE_EXIT_PROGRAM	RTV_PGM	VARCHAR(10) Nullable	The preprocessing exit program name that is called by the registration facility when the Retrieve Exit Information API is called for the exit point. Contains the null value if no retrieve exit program is defined.
RETRIEVE_EXIT_PROGRAM_FORMAT	RTV_FMT	CHAR(8) Nullable	The format name for RETRIEVE_EXIT_PROGRAM. Contains the null value if no retrieve exit program is defined.

Example

- List all the security exit points.

```
SELECT *
FROM QSYS2.EXIT_POINT_INFO
WHERE EXIT_POINT_NAME LIKE 'QIBM_QSY%';
```

EXIT_PROGRAM_INFO view

The EXIT_PROGRAM_INFO view returns information about exit programs.

The values returned for the columns in the view are closely related to the values returned by the Work with Registration Information (WRKREGINF) CL command and by the Retrieve Exit Information (QUSRTVEI, QusRetrieveExitInformation) API.

Authorization: None required.

The following table describes the columns in the view. The system name is EXIT_PGM. The schema is QSYS2.

Table 91. EXIT_PROGRAM_INFO view

Column Name	System Column Name	Data Type	Description
EXIT_POINT_NAME	EXIT_NAME	VARCHAR(20)	The exit point name.
EXIT_POINT_FORMAT	EXIT_FMT	CHAR(8)	The exit point format name associated with the exit point.
REGISTERED	REGISTERED	VARCHAR(3)	Whether the exit point is registered with the registration facility. NO The exit point is unregistered. YES The exit point is registered.

Table 91. EXIT_PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
COMPLETE	COMPLETE	VARCHAR(3)	<p>Whether the information returned for the exit point is complete and accurate. Incomplete information may occur when the exit point's retrieve preprocessing program indicates that the information it returned is incomplete or inaccurate.</p> <p>NO The exit point entry information is not complete or accurate. The remaining columns should be ignored.</p> <p>YES The exit point entry information is complete and accurate.</p>
EXIT_PROGRAM_NUMBER	NUMBER	INTEGER	The exit program number associated with the exit program. This number determines the processing sequence of exit programs associated with an exit point, where the lowest number is processed first.
EXIT_PROGRAM_LIBRARY	LIBRARY	VARCHAR(10)	The library in which EXIT_PROGRAM resides.
EXIT_PROGRAM	PROGRAM	VARCHAR(10)	The name of the exit program.
TEXT_DESCRIPTION	TEXT	VARCHAR(132) Nullable	<p>The descriptive text for the exit program .</p> <p>Contains the null value if no descriptive text is available.</p>
EXIT_PROGRAM_DATA	DATA	VARCHAR(2048) Nullable	<p>The data that is associated with the exit program.</p> <p>Contains the null value when there is no exit program data.</p>
EXIT_PROGRAM_DATA_CCSD	DATA_CCSD	INTEGER Nullable	<p>The coded character set identifier (CCSID) to use for working with the exit program data.</p> <p>Contains the null value when EXIT_PROGRAM_DATA is null.</p>
THREADSAFE	THREADSAFE	VARCHAR(3) Nullable	<p>The thread safety status of the exit program entry.</p> <p>NO The exit program entry is not threadsafe.</p> <p>YES The exit program entry is threadsafe.</p> <p>Contains the null value when the threadsafe status of the exit program entry is not known.</p>
MULTITHREADED_JOB_ACTION	JOB_ACTION	VARCHAR(6)	<p>The action to take when calling an exit program in a multithreaded job.</p> <p>*MSG Run the exit program in the current multithreaded job, but send an informational message. CPI3C80 can be used as the informational message.</p> <p>*NORUN Do not run the exit program in the current multithreaded job. Depending on the exit point, do one of the following:</p> <ol style="list-style-type: none"> 1. Send an escape message and do not call the exit program. CPF3C80 can be used as the escape message. 2. Send an informational message and do not call the exit program. CPF3C80 can be used as the informational message. 3. Call the exit program in a non-multithreaded job. <p>*RUN Run the exit program in the current multithreaded job.</p>

Table 91. EXIT_PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
QMLTTHDACN_SYSTEM_VALUE	QMLTTHDACN	VARCHAR(3)	<p>Indicates whether the QMLTTHDACN system value was used in determining the multithreaded job action.</p> <p>NO The QMLTTHDACN system value was not used to determine the multithreaded job action.</p> <p>YES The QMLTTHDACN system value was used to determine the multithreaded job action.</p>

Example

- List all the exit programs associated with security exit points.

```
SELECT *
FROM QSYS2.EXIT_PROGRAM_INFO
WHERE EXIT_POINT_NAME LIKE 'QIBM_QSY%'
ORDER BY EXIT_POINT_NAME, EXIT_PROGRAM_NUMBER;
```

LPRINTF procedure

The LPRINTF procedure writes an informational message to the joblog. It is severity 0 and has no message identifier.

Authorization: None required.

► LPRINTF — (————— *print-string* —) ◄

└── PRINT_STRING — => ─┘

The schema is SYSTOOLS.

print-string A character or graphic string containing the information to be written to the joblog. It can be up to 1000 characters long.

Note

This procedure is provided in the SYSTOOLS schema as an example of how to write to the joblog using an SQL procedure. Creating customized versions of this procedure to better suit a specific need is encouraged. Use the Insert Generated SQL feature in IBM i Access Client Solutions (ACS) to extract the source for this procedure. Then modify it and create a new procedure in a user-specified schema.

Example

Write a message to the joblog.

```
CALL SYSTOOLS.LPRINTF('This message sent on '
                      CONCAT DAYOFWEEK(CURRENT DATE) CONCAT ' at '
                      CONCAT CURRENT TIME);
```

Results in this string in the joblog:

```
This message sent on Saturday at 04:21 PM
```

PROGRAM_EXPORT_IMPORT_INFO view

The PROGRAM_EXPORT_IMPORT_INFO view returns the data and procedure that are exported or imported for an ILE program or service program.

The values returned for the columns in the view are closely related to the values returned for *PROCEXP, *DTAEXP, *ACTGRPEXP, and *ACTGRPIMP detail on the DSPSRVPGM (Display Service Program) CL command and *ACTGRPEXP and *ACTGRPIMP detail on the DSPPGM (Display Program) CL command. It is similar to the information returned by the List Service Program Information (QBNLSPGM) and List Program Information (QBNLPGMI) APIs.

Authorization: The caller must have:

- *READ authority to the program or service program, and
- *EXECUTE authority to the library containing the program or service program.

The following table describes the columns in the view. The system name is EXPIMP_INF. The schema is QSYS2.

Table 92. PROGRAM_EXPORT_IMPORT_INFO view

Column Name	System Column Name	Data Type	Description
PROGRAM_LIBRARY	PGM_LIB	VARCHAR(10)	The library containing the program or service program.
PROGRAM_NAME	PGM_NAME	VARCHAR(10)	The program or service program to which this export or import applies.
OBJECT_TYPE	OBJ_TYPE	VARCHAR(7)	Object type for PROGRAM_NAME. *PGM A program. *SRVPGM A service program.
SYMBOL_NAME	NAME	VARGRAPHIC(8192) CCSID 1200	The name of a procedure or data export or import.
SYMBOL_USAGE	USAGE	VARCHAR(10)	The type of export or import. *ACTGRPEXP The data export to the activation group. *ACTGRPIMP A data import that is resolved by a weak export that had been exported to the activation group. *DTAEXP A data item exported from this service program. *PROCEXP A procedure exported from this service program.
ARGUMENT_OPTIMIZATION	ARGOPT	VARCHAR(4) Nullable	Whether the service program procedure export uses argument optimization. *NO The procedure export does not use argument optimization. *YES The procedure export uses argument optimization. Contains the null value if SYMBOL_USAGE is not *PROCEXP.
DATA_ITEM_SIZE	DATA_SIZE	INTEGER Nullable	The size of the data item export, in bytes. Contains the null value if SYMBOL_USAGE is not *ACTGRPEXP.

Example

- Show all the procedure exports for service program APP_PGM1 in APPLIB.

```

SELECT *
FROM QSYS2.PROGRAM_EXPORT_IMPORT_INFO
WHERE PROGRAM_LIBRARY = 'APPLIB' AND
      PROGRAM_NAME = 'APP_PGM1' AND
      OBJECT_TYPE = '*SRVPGM' AND
      SYMBOL_USAGE = '*PROCEXP';

```

- Find the service program in QSYS that exports printf to joblog (Qp0zLprintf).

```

SELECT *
FROM QSYS2.PROGRAM_EXPORT_IMPORT_INFO
WHERE PROGRAM_LIBRARY = 'QSYS' AND SYMBOL_NAME = 'Qp0zLprintf';

```

PROGRAM_INFO view

The PROGRAM_INFO view returns information about programs.

The values returned for the columns in the view are closely related to the values returned by the DSPPGM (Display Program) and DSPSRVPGM (Display Service Program) CL commands (the *BASIC, *SIZE, and *SIGNATURE details for ILE programs and service programs) and the Retrieve Program Information (QCLRPGMI) and Retrieve Service Program Information (QBNRSPGM) APIs.

Authorization: The caller must have:

- *READ authority to the program or service program, and
- *EXECUTE authority to the library containing the program or service program.

The following table describes the columns in the view. The first set of columns, through CONVERSION_DETAIL, apply to ILE and OPM programs and service programs. The next group of columns apply to ILE programs and service programs. The final group of columns apply to OPM programs. The system name is PGM_INFO. The schema is QSYS2.

Table 93. PROGRAM_INFO view

Column Name	System Column Name	Data Type	Description
PROGRAM_LIBRARY	PGM_LIB	VARCHAR(10)	The library containing the program or service program.
PROGRAM_NAME	PGM_NAME	VARCHAR(10)	The program or service program.
PROGRAM_TYPE	PGM_TYPE	VARCHAR(3)	The type of program or service program. ILE An Integrated Language Environment program or service program. OPM An original program model program.
OBJECT_TYPE	OBJ_TYPE	VARCHAR(7)	Object type for PROGRAM_NAME. *PGM A program. *SRVPGM A service program.
CREATE_TIMESTAMP	CREATE_TS	TIMESTAMP(0)	The timestamp when the program or service program was created.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The user text, if any, used to briefly describe the program or service program. Contains the null value if there is no text description.
PROGRAM_OWNER	OWNER	VARCHAR(10) Nullable	The name of the program or service program owner's user profile. Contains the null value if no program owner is available.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PROGRAM_ATTRIBUTE	ATTRIBUTE	VARCHAR(10) Nullable	<p>For an ILE program, the ILE language used for the module containing the program entry procedure (PEP).</p> <p>For a service program, the language in which the modules of the service program are written.</p> <p>For an OPM program, the language the program is written in.</p> <p>Contains the null value if there is no attribute value or if a service program contains modules generated by different compilers.</p>
USER_PROFILE	USRPRF	VARCHAR(6)	<p>The value specified for the USRPRF option on the command used to create the program or service program.</p> <p>*OWNER The program or service program runs under both the current user's and the owner's user profiles.</p> <p>*USER The program or service program runs under the current user's user profile.</p>
USE_ADOPTED_AUTHORITY	USEADPAUT	VARCHAR(4)	<p>The value specified for the USEADPAUT option on the command used to change the program or service program.</p> <p>*NO Does not use program adopted authority from previous call levels when this program or service program is running</p> <p>*YES Uses program adopted authority from previous call levels when this program or service program is running.</p>
PROGRAM_STATE	STATE	VARCHAR(8)	<p>The state of the program or service program.</p> <p>*INHERIT The program or service program runs under (inherits) the same state as its caller.</p> <p>*SYSTEM The program or service program can call user-domain or system-domain programs.</p> <p>*USER The program or service program can call only user-domain programs.</p>
PROGRAM_DOMAIN	DOMAIN	VARCHAR(7)	<p>The domain of the program or service program.</p> <p>*SYSTEM The program or service program can be called by system-state programs.</p> <p>*USER The program can be called by user-state or system-state programs.</p>
EARLIEST_POSSIBLE_RELEASE	EARLY_REL	CHAR(6)	<p>The version, release, and modification level of the earliest release the program or service program is allowed to run on, in VvRrMm format. The target release (TGTRLS) parameter of the command can affect this value.</p>
RELEASE_CREATED_ON	CREATE_ON	CHAR(6) Nullable	<p>The version, release, and modification level of the operating system, in VvRrMm format, on which the program or service program was created.</p> <p>Contains the null value for OPM programs.</p>
TARGET_RELEASE	TGTRLS	CHAR(6) Nullable	<p>This is the release specified on the target release (TGTRLS) parameter of the command.</p> <p>Contains the null value for OPM programs.</p>

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MINIMUM_NUMBER_PARMS	MIN_PARMS	INTEGER Nullable	The minimum number of parameters that is to be received by the program when it is called. Contains the null value if this is a service program or if the information is not available.
MAXIMUM_NUMBER_PARMS	MAX_PARMS	INTEGER Nullable	The maximum number of parameters that may be received by the program when it is called. Contains the null value if this is a service program or if the information is not available.
ASSOCIATED_SPACE_SIZE	ASSOC_SIZE	INTEGER	The size (in bytes) of the associated space used by this program or service program.
PAGING_POOL	POOL	VARCHAR(8)	The paging pool used for the program or service program object. *BASE Use the base pool. *MACHINE Use the machine pool. *USER Use the user pool.
PAGING_AMOUNT	PAGEAMOUNT	VARCHAR(8)	The paging behavior for the program or service program. *BLOCK Page the program or service program in multiple-page blocks. *NOBLOCK Page the program or service program one page at a time.
ALLOW_RTVCLSRC	ALWRTVSRC	VARCHAR(4) Nullable	Value of the ALWRTVSRC parameter if this program was created using the Create CL Program (CRTCLPGM) command. *NO Source for the CL program is not saved with the program. *YES Source is saved. Contains the null value if this is not a CL program.
CONVERSION_REQUIRED	CONVREQ	VARCHAR(4)	Indicates whether the program or service program has been converted to the format required by the machine or if conversion is still required. *NO Conversion is not required. *YES Conversion is required.
CONVERSION_DETAIL	CONVDETAIL	VARCHAR(8)	Indicates details about the conversion status. *COMMON The object is compatible and requires only features common to all systems supported by this release. *COMPAT The object is compatible, but requires features not present on all systems supported by this release. *FEATURE The object is not compatible. The format is current, but the object requires features not present on the current machine implementation. *FORMAT The object is not compatible because it is in an older format.

The following columns apply to ILE programs and service programs. They will contain the null value for OPM programs.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PROGRAM_ENTRY_PROCEDURE_MODULE_LIBRARY	PEP_LIB	VARCHAR(10) Nullable	The library name that contained the module that contained the program entry procedure for this program when the bind was done. Contains the null value if this is a service program.
PROGRAM_ENTRY_PROCEDURE_MODULE	PEP_MODULE	VARCHAR(10) Nullable	The module name that contains the program entry procedure for this program. Contains the null value if this is a service program.
ACTIVATION_GROUP	ACTGRP	VARCHAR(10) Nullable	The name of the activation group in which this program or service program runs. If the activation group already exists when the program or service program is called, it runs in the existing activation group. If the activation group does not exist, a new activation group is created and the program or service program runs in it. Can contain the following special values: <p>*CALLER The program or service program runs in the activation group of the calling program.</p> <p>*DFACTGRP The program uses one of two existing activation groups created when the process is started. One default activation group is reserved for system-state programs. The other default activation group is reserved for user-state programs.</p> <p>*NEW A new unnamed activation group is created when this program is called. The program runs in this activation group.</p>
SHARED_ACTIVATION_GROUP	SHARED_AG	VARCHAR(4) Nullable	Whether the program or service program runs in a shared activation group. <p>*NO The activation group is not shared.</p> <p>*YES The activation group is shared.</p>
OBSERVABLE_INFO_COMPRESSED	OBS_COMP	VARCHAR(4) Nullable	Whether the observable information associated with the program or service program is compressed. <p>*NO The observable information is not compressed.</p> <p>*YES The observable information is compressed.</p>
RUNTIME_INFO_COMPRESSED	RUN_COMP	VARCHAR(4) Nullable	Whether the run-time information associated with the program or service program is compressed. <p>*NO The run-time information is not compressed.</p> <p>*YES The run-time information is compressed.</p>
ALLOW_UPDATE	ALWUPD	VARCHAR(4) Nullable	Whether the Update Program (UPDPGM) or Update Service Program (UPDSRVPGM) command is allowed on this program or service program. <p>*NO The update command cannot be run.</p> <p>*YES The update command can be run.</p>

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
ALLOW_BOUND_SRVPGM_LIBRARY_UPDATE	ALWLIBUPD	VARCHAR(4) Nullable	<p>Whether the Update Program (UPDPGM) or Update Service Program (UPDSRVPGM) command is allowed to change the bound *SRVPGM library names on this program or service program.</p> <p>*NO The update command cannot specify a library name for the SRVPGMLIB parameter.</p> <p>*YES The update command can specify a library name for the SRVPGMLIB parameter.</p>
ALL_CREATION_DATA	ALL_CREATE	VARCHAR(6) Nullable	<p>Whether the program or service program has all creation data and if that data is observable or unobservable. All observable creation data is needed to re-create the program or service program using the Change Program (CHGPGM) or Change Service Program (CHGSRVPGM) command. All creation data (either observable or unobservable) is needed to convert the program or service program during restore.</p> <p>*NO Not all of the creation data is present.</p> <p>*UNOBS All the creation data is present but not all of that data is observable.</p> <p>*YES All the creation data is present and all of that data is observable.</p>
PROFILING_DATA	PRF_DATA	VARCHAR(10) Nullable	<p>Specifies the profiling data attribute for this program or service program.</p> <p>*APYALL Block order and procedure order profiling data are applied to this program or service program. The collection of profiling data is no longer enabled.</p> <p>*APYBLKORD Block order profiling data has been applied to at least one module bound into this program or service program.</p> <p>*APYPRCORD Procedure order profiling data has been applied to this program or service program.</p> <p>*COL The collection of profiling data is enabled for at least one module bound into this program or service program. Any applied profiling data has been removed.</p> <p>*NOCOL The collection of profiling data is not enabled and profiling data is not applied.</p>
TERASPACE_STORAGE_ENABLED_MODULES	TERA_MOD	VARCHAR(5) Nullable	<p>The teraspace storage capability of the modules bound to this program or service program.</p> <p>*ALL All modules bound to this program or service program are teraspace storage enabled.</p> <p>*NONE No modules bound to this program or service program are teraspace storage enabled.</p> <p>*SOME One or more modules bound to this program or service program are teraspace storage enabled.</p>

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
TERASPACE_STORAGE_ENABLED_PEP	TERA_PEP	VARCHAR(4) Nullable	Indicates whether the Program Entry Procedure (PEP) is teraspace enabled. *NO The PEP module is not teraspace storage enabled. *YES The PEP module is teraspace storage enabled. Contains the null value if this is a service program.
STORAGE_MODEL	STGMDL	VARCHAR(10) Nullable	Where the automatic and static storage for this program or service program is allocated at run time. *INHERIT Automatic and static storage are allocated from either single-level storage or teraspace, depending on the activation. *SINGLVL Automatic and static storage are allocated from single-level storage. *TERASPACE Automatic and static storage are allocated from teraspace.
ARGUMENT_OPTIMIZATION	ARGOPT	VARCHAR(4) Nullable	Whether argument optimization was done during program or service program creation. *NO Argument optimization was not performed. *YES Argument optimization was performed.
NUMBER_OF_UNRESOLVED_REFERENCES	UNRESOLVED	INTEGER Nullable	The number of symbols that could not be resolved at Create Program (CRTPGM) or Create Service Program (CRTSRVPGM) command time.
ALLOW_STATIC_STORAGE_REINIT	ALWRINZ	VARCHAR(4) Nullable	Whether program or service program static storage can be reinitialized. *NO Static storage cannot be reinitialized. *YES Static storage can be reinitialized.
MINIMUM_STATIC_STORAGE_SIZE	MIN_STATIC	BIGINT Nullable	The minimum static storage size, in bytes, that this program or service program needs in order to run.
MAXIMUM_STATIC_STORAGE_SIZE	MAX_STATIC	BIGINT Nullable	The maximum static storage size, in bytes, that this program or service program may need in order to run.
AUXILIARY_STORAGE_SEGMENTS	NBR_AUX	INTEGER Nullable	The number of auxiliary storage segments in this program or service program.
MAXIMUM_AUXILIARY_STORAGE_SEGMENTS	MAX_AUX	INTEGER Nullable	The maximum number of auxiliary storage segments an ILE program or service program can have.
PROGRAM_SIZE	PGM_SIZE	INTEGER Nullable	The total size of the program or service program, in kilobytes.
MAXIMUM_PROGRAM_SIZE	MAXPGMSIZE	INTEGER Nullable	The maximum size that an ILE program or service program can be, in kilobytes.
MODULES	MODULES	INTEGER Nullable	The number of modules bound into this program or service program.
MAXIMUM_MODULES	MAXMODS	INTEGER Nullable	The maximum number of modules that can be bound into an ILE program or service program.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SERVICE_PROGRAMS	SRVPGMS	INTEGER Nullable	The number of service programs bound to this program or service program.
MAXIMUM_SERVICE_PROGRAMS	MAXSRVPGMS	INTEGER Nullable	The maximum number of service programs that can be bound to an ILE program or service program.
STRING_DIRECTORY_SIZE	STRDIRSIZE	INTEGER Nullable	The program or service program's string directory size.
MAXIMUM_STRING_DIRECTORY_SIZE	MAXSTRDIR	INTEGER Nullable	The maximum size that the string directory can be in an ILE program or service program.
COPYRIGHTS	COPYRIGHTS	INTEGER Nullable	The number of copyrights in this program or service program.
COPYRIGHT_STRING_SIZE	CPYRSTRSIZ	INTEGER Nullable	The program or service program's copyright string size.
MAXIMUM_COPYRIGHT_STRING_SIZE	MAXCPYRSTR	INTEGER Nullable	The maximum size of the copyright string in an ILE program or service program.
EXPORT_SOURCE_LIBRARY	EXP_SRCLIB	VARCHAR(10) Nullable	The name of the library that contains the export source file. Can contain the following special value: *LIBL The library list. Contains the null value if EXPORT_SOURCE_FILE is null.
EXPORT_SOURCE_FILE	EXP_SRCF	VARCHAR(10) Nullable	The name of the export source file that contains the export source file member. Contains the null value if this is not a service program, if there is no export source, if the export source was not specified using a source file, or if you are not authorized to this information.
EXPORT_SOURCE_FILE_MEMBER	EXP_SRCM	VARCHAR(10) Nullable	The name of the member in the export source file that was used to create this service program. Contains the null value if EXPORT_SOURCE_FILE is null.
EXPORT_SOURCE_STREAM_FILE	EXP_STMF	VARGRAPHIC(5000) CCSID 1200 Nullable	The path name of the stream file that contains the export source for this service program. Contains the null value if this is not a service program or if the export source was not specified using a stream file.
PROCEDURE_EXPORTS	PROCEXP	INTEGER Nullable	The number of procedures exported from this service program. Contains the null value if this is not a service program.
MAXIMUM_PROCEDURE_EXPORTS	MAXPROCEXP	INTEGER Nullable	The maximum number of procedures that can be exported by a service program. Contains the null value if this is not a service program.
DATA_EXPORTS	DATAEXP	INTEGER Nullable	The number of data items exported from this service program. Contains the null value if this is not a service program.
MAXIMUM_DATA_EXPORTS	MAXDATAEXP	INTEGER Nullable	The maximum number of data items that can be exported by a service program. Contains the null value if this is not a service program.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SIGNATURES	SIGNATURES	INTEGER Nullable	The number of signatures for this service program. This is the number of entries returned in EXPORT_SIGNATURES. Contains the null value if this is not a service program.
EXPORT_SIGNATURES	EXP_SIG	BLOB(557038) Nullable	The list of export signatures for this service program. Each signature is a binary string with a length of 16. A single hexadecimal 00 character separates values. The current export signature of this service program is the first one in the list. Contains the null value if this is not a service program.
MAXIMUM_SIGNATURES	MAXSIGS	INTEGER Nullable	The maximum number of signatures that can be in a service program. Contains the null value if this is not a service program.
The following columns apply to OPM programs. They will contain the null value for ILE programs and service programs.			
SOURCE_FILE_LIBRARY	SRCLIB	VARCHAR(10) Nullable	The name of the library that contains the source file used to create the program.. Contains the null value if no source file was used to create the program.
SOURCE_FILE	SRCFILE	VARCHAR(10) Nullable	The name of the source file used to create the program. Contains the null value if no source file was used to create the program.
SOURCE_FILE_MEMBER	SRCMBR	VARCHAR(10) Nullable	The name of the member in the source file. Contains the null value if no source file was used to create the program.
SOURCE_FILE_CHANGE_TIMESTAMP	SRCF_CHGTS	TIMESTAMP(0) Nullable	The timestamp of when the member in the source file was last updated. Contains the null value if no source file was used to create the program.
SORT_SEQUENCE_LIBRARY	SRTSEQ_LIB	VARCHAR(10) Nullable	The name of the library that contained the sort sequence table used when the module was compiled. This does not apply to SQL statements in the module. Can contain the following special values: <p>*LIBL The sort sequence table is found in the library list when PROGRAM_NAME runs this module.</p> <p>*CURLIB The sort sequence table is found in the current library when PROGRAM_NAME runs this module.</p> <p>Contains the null value if SORT_SEQUENCE contains a special value or is null.</p>

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SORT_SEQUENCE	SRTSEQ	VARCHAR(10) Nullable	<p>The name of the sort sequence table used when the program was compiled. This does not apply to SQL statements in the program. Can contain the following special values:</p> <p>*HEX No sort sequence is used.</p> <p>*JOBRUN The sort sequence value associated with the job at the time the program runs.</p> <p>*LANGIDSHR The shared sort sequence for the language identifier is used.</p> <p>*LANGIDUNQ The unique sort sequence for the language identifier is used.</p> <p>Contains the null value if the program does not contain any sort sequence information.</p>
LANGUAGE_ID	LANGID	VARCHAR(7) Nullable	<p>The language identifier used when the program was compiled. This does not apply to SQL statements in the program. Can contain the following special value:</p> <p>*JOBRUN The language identifier associated with the job at the time the program is run.</p> <p>Contains the null value if the program does not contain any language identification information.</p>
OBSERVABLE	OBSERVABLE	VARCHAR(6) Nullable	<p>Whether the OPM program contains creation data and if that data is observable or unobservable. All observable creation data is needed to re-create the program using CHGPGM. All creation data (either observable or unobservable) is needed to convert the program during restore.</p> <p>*ALL The program has all creation data and that data is observable.</p> <p>*NONE The program does not have all creation data.</p> <p>*UNOBS The program has all creation data but not all of that data is observable.</p>
OPTIMIZATION	OPTIMIZE	VARCHAR(11) Nullable	<p>Indicates what was specified on the OPTIMIZE parameter when the program was created or changed.</p> <p>*NOOPTIMIZE *NOOPTIMIZE was specified.</p> <p>*OPTIMIZE *OPTIMIZE was specified.</p>
LOG_COMMANDS	LOGCMD	VARCHAR(4) Nullable	<p>The value specified for the LOG parameter of the CRTCLPGM command.</p> <p>*JOB Logging of commands in a running CL program depends on the status of the job's logging flag.</p> <p>*NO Commands are not logged.</p> <p>*YES Commands are logged in all cases.</p> <p>Contains the null value if the program is not a CL program.</p>
FIX_DECIMAL_DATA	FIXDECDATA	VARCHAR(4) Nullable	<p>Whether decimal data that is not valid is corrected or an error is signaled.</p> <p>*NO An error is signaled to the program without correcting the data that is not valid.</p> <p>*YES Decimal data that is not valid is corrected.</p>

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
UPDATE_PASA	UPDPASA	VARCHAR(10) Nullable	The compiler may have allowed you to control this attribute through the GENOPT parameter of the command used to create the program. *NOUPDPASA Do not update internal program automatic storage area (PASA) stack information. *UPDPASA Update internal PASA stack information.
CLEAR_PASA	CLRPASA	VARCHAR(10) Nullable	The compiler may have allowed you to control this attribute through the GENOPT parameter of the command used to create the program. *CLRPASA Clear PASA storage. *NOCLRPASA Do not clear PASA storage.
COMPILER_ID	COMPILER	VARCHAR(14) Nullable	The licensed program identifier, version, release, and modification level of the compiler. The value has a ppppppbVvRrMm format, where: ppppppp The licensed program identifier. b A blank character. Vv The character V is followed by a 1-character version number. Rr The character R is followed by a 1-character release level. Mm The character M is followed by a 1-character modification level.
TERASPACE_STORAGE_ENABLED_PROGRAM	TS_PGM	VARCHAR(4) Nullable	The teraspace storage capability of the program. A program must be teraspace storage enabled to access teraspace storage. *NO This program is not teraspace storage enabled. *YES This program is teraspace storage enabled.
OPM_PROGRAM_SIZE	OPMPGMSIZE	INTEGER Nullable	The size (in bytes) of this program.
STATIC_STORAGE_SIZE	STATIC_STG	INTEGER Nullable	The size (in bytes) of the static storage used by the program.
AUTOMATIC_STORAGE_SIZE	AUTO_STG	INTEGER Nullable	The size (in bytes) of the automatic storage used by this program.
NUMBER_MI_INSTRUCTIONS	MI_INSTR	INTEGER Nullable	The number of machine interface (MI) instructions used by this program. Contains the null value if the program is not observable.
NUMBER_MI_ODT_ENTRIES	MI_ODT	INTEGER Nullable	The number of ODT (object definition table) entries for the program. There is a limit of 32767 ODT entries in a program. Contains the null value if the program is not observable.
SQL_STATEMENT_COUNT	NBRSTMTS	INTEGER Nullable	The number of SQL statements contained in the program. Contains 0 if there are no SQL statements in the program.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_RELATIONAL_DATABASE	RDB	VARCHAR(18) Nullable	The default relational database that was specified on the SQL precompile. Can contain the following special value: *LOCAL The program can only access data on the local system. Contains the null value if no package was created for the program by the SQL precompiler or if the program does not contain SQL statements.
SQL_COMMITMENT_CONTROL	ISOLATION	VARCHAR(5) Nullable	The level of commitment control that was specified on the SQL precompile. *ALL Read stability. *CHG Uncommitted read. *CS Cursor stability. *NONE No commit. *RR Repeatable read. Contains the null value if the program does not contain SQL statements.
SQL_NAMING	NAMING	VARCHAR(4) Nullable	The convention used for naming objects in SQL statements. *SQL The SQL naming convention is used. *SYS The system naming convention is used. Contains the null value if the program does not contain SQL statements.
SQL_DATE_FORMAT	DATFMT	VARCHAR(4) Nullable	The date format attribute. *DMY Day/month/year format (dd/mm/yy). *EUR European format (dd.mm.yyyy). *ISO International Standards Organization format (yyyy-mm-dd). *JIS Japanese Industrial Standard format (yyyy-mm-dd). *JUL Julian format (yy/dds). *MDY Month/day/year format (mm/dd/yy). *USA USA format (mm/dd/yyyy). *YMD Year/month/day format (yy/mm/dd). Contains the null value if the program does not contain SQL statements.
SQL_DATE_SEPARATOR	DATSEP	CHAR(1) Nullable	The date separator attribute. Contains the null value if the program does not contain SQL statements.
SQL_TIME_FORMAT	TIMFMT	VARCHAR(4) Nullable	The time format attribute. *EUR European format (hh.mm.ss). *HMS Hours/minutes/seconds format (hh:mm:ss). *ISO International Standards Organization format (hh.mm.ss). *JIS Japanese Industrial Standard format (hh.mm.ss). *USA USA format (hh:mm a.m. or p.m.). Contains the null value if the program does not contain SQL statements.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_TIME_SEPARATOR	TIMSEP	CHAR(1) Nullable	The time separator attribute. Contains the null value if the program does not contain SQL statements.
SQL_SORT_SEQUENCE_LIBRARY	SQL_SSEQLB	VARCHAR(10) Nullable	The name of the library that is used to locate the SQL sort sequence table. The following special values can be returned: *CURLIB The SQL sort sequence table is found by looking in the current library. *LIBL The SQL sort sequence table is found by looking in the library list. Contains the null value if SQL_SORT_SEQUENCE contains a special value or if the program does not contain SQL statements.
SQL_SORT_SEQUENCE	SQL_SRTSEQ	VARCHAR(10) Nullable	The sort sequence table name specified when the program was compiled. The following special values can be returned: *HEX No SQL sort sequence is used for the SQL statements. *JOBRUN The SQL sort sequence is the SRTSEQ value associated with the job at the time the SQL statements within the program are run. *LANGIDSHR The shared SQL sort sequence for the language identifier (LANGID) is used for the SQL statements. *LANGIDUNQ The unique SQL sort sequence for the language identifier (LANGID) is used for the SQL statements. Contains the null value if the program does not contain SQL statements.
SQL_LANGUAGE_ID	SQL_LANGID	VARCHAR(7) Nullable	The language identifier specified when the program was compiled. The following special value can be returned: *JOBRUN The language identifier is the LANGID associated with the job at the time the program is run. Contains the null value if the program does not contain SQL statements.
SQL_DEFAULT_SCHEMA	DFTRDBCOL	VARCHAR(10) Nullable	The schema name used for unqualified names of tables, views, indexes, and SQL packages in static statements. Contains the null value if there is no default schema name or if the program does not contain SQL statements.
SQL_PATH	SQLPATH	VARCHAR(3483) Nullable	The list of libraries used during resolution of functions, procedures, and data types within SQL statements. The list is in the form of repeating library names, each surrounded by double quotes and separated by commas. Contains the null value if the program does not contain SQL statements.

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_DYNAMIC_USER_PROFILE	DYNUSRPRF	VARCHAR(10) Nullable	<p>The user profile used for dynamic SQL statements. The following special values can be returned:</p> <p>*OWNER Local dynamic SQL statements are run under the profile of the program's owner. Distributed dynamic SQL statements are run under the profile of the SQL package's owner.</p> <p>*USER Local dynamic SQL statements are run under the profile of the job or thread. Distributed dynamic SQL statements are run under the profile of the application server job.</p> <p>Contains the null value if the program does not contain SQL statements.</p>
SQL_ALLOW_COPY_DATA	ALWCPYDTA	VARCHAR(9) Nullable	<p>Whether a copy of the data can be used in the implementation of an SQL query.</p> <p>*NO A copy of the data is not allowed.</p> <p>*OPTIMIZE A copy of the data is allowed whenever it might result is better performance.</p> <p>*YES A copy of the data is allowed, but only when necessary.</p> <p>Contains the null value if the program does not contain SQL statements.</p>
SQL_CLOSE_SQL_CURSOR	CLOSQCSR	VARCHAR(10) Nullable	<p>Specifies the CLOSQCSR attribute.</p> <p>*ENDJOB SQL cursors remain open between calls and can be fetched without running another OPEN statement. The programs higher on the call stack do not need to have run SQL statements. SQL cursors are left open, SQL prepared statements are preserved, and LOCK TABLE locks are held when the first SQL program on the call stack ends. SQL cursors are closed, SQL prepared statements are discarded, and LOCK TABLE locks are released when the job ends.</p> <p>*ENDPGM SQL cursors are closed and SQL prepared statements are discarded when the program ends. LOCK TABLE locks are released when the first SQL program on the call stack ends.</p> <p>*ENDSQL SQL cursors remain open between calls and rows can be fetched without running another OPEN statement. One of the programs higher on the call stack must have run at least one SQL statement. The SQL cursors are closed, SQL prepared statements are discarded, and LOCK TABLE locks are released when the first SQL program on the call stack ends. If you specify *ENDSQL for a program that is the first SQL program called (the first SQL program on the call stack), the program is treated as if *ENDPGM was specified.</p> <p>Contains the null value if the program does not contain SQL statements.</p>

Table 93. PROGRAM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SQL_DELAY_PREPARE	DLYPRP	VARCHAR(4) Nullable	Indicates the delay prepare attribute. *NO Dynamic statement validation is performed when the dynamic statements are prepared. *YES Dynamic statement validation is delayed until the dynamic statements are used. Contains the null value if the program does not contain SQL statements.
SQL_ALLOW_BLOCK	ALWBLK	VARCHAR(8) Nullable	Whether blocking is used to improve the performance of certain SQL statements. *ALLREAD Rows are blocked for read-only cursors. *NONE Rows are not blocked for retrieval of data for cursors. *READ Rows are blocked for read-only retrieval of data for cursors when: <ul style="list-style-type: none"> • The commitment control value is *NONE. • The cursor is declared with a FOR READ ONLY clause or there are no dynamic statements that could run a positioned UPDATE or DELETE statement for the cursor. Contains the null value if the program does not contain SQL statements.
SQL_PACKAGE_LIBRARY	SQLPKGLIB	VARCHAR(10) Nullable	The name of the library the SQL package is in. Contains the null value if the program is not distributed or if the program does not contain SQL statements.
SQL_PACKAGE	SQLPKG	VARCHAR(10) Nullable	The name of the SQL package created on the relational database specified on the RDB parameter of the command that created this program. Contains the null value if the program is not distributed or if it does not contain SQL statements.
SQL_RDB_CONNECTION_METHOD	RDBCNNMTH	VARCHAR(4) Nullable	Specifies the semantics used for CONNECT statements: *DUW CONNECT (Type 2) semantics are used to support distributed unit of work. *RUW CONNECT (Type 1) semantics are used to support remote unit of work Contains the null value if the program is not distributed or if it does not contain SQL statements.

Examples

- Summarize the activation group usage for all ILE programs in APPLIB.

```
SELECT ACTIVATION_GROUP, COUNT(*) AS ACTIVATION_GROUP_NAME_COUNT
FROM QSYS2.PROGRAM_INFO
WHERE PROGRAM_LIBRARY = 'APPLIB'
      AND PROGRAM_TYPE = 'ILE'
GROUP BY ACTIVATION_GROUP
ORDER BY 2 DESC;
```

- Examine the ownership of programs and service programs in APPLIB.

```
SELECT PROGRAM_OWNER, OBJECT_TYPE, COUNT(*) AS APPLICATION_OWNER_COUNT
FROM QSYS2.PROGRAM_INFO
WHERE PROGRAM_LIBRARY = 'APPLIB'
GROUP BY PROGRAM_OWNER, OBJECT_TYPE
ORDER BY 2, 3 DESC;
```

QCMDEXC procedure

The QCMDEXC procedure executes a CL command.

Authorization: Any authority requirements for the CL command apply to the use of this function.

►► QCMDEXC — (— *CL-command-string* —) ►►

The schema is QSYS2.

CL-command-string A character string expression containing a CL command.

The *CL-command-string* will be run as a CL command.

Examples

- Add a library to the library list.

```
CALL QSYS2.QCMDEXC('ADDLIB PRODLIB2');
```

- Add a library to the library list using an expression.

```
DECLARE V_LIBRARY_NAME VARCHAR(10);
SET V_LIBRARY_NAME = 'PRODLIB2';
CALL QSYS2/QCMDEXC('ADDLIB ' CONCAT V_LIBRARY_NAME);
```

QCMDEXC scalar function

The QCMDEXC scalar function executes a CL command.

Authorization: Any authority requirements for the CL command apply to the use of this function.

►► QCMDEXC — ( *command* —) ►►

The schema is QSYS2.

command A character string containing a CL command. The maximum length is 32000 characters.

The result of the function is an integer. If the command is successful, the function returns a value of 1. If the command execution fails, the function returns a value of -1.

Example

Hold any jobs that started running an SQL statement more than 2 hours ago.

```
SELECT JOB_NAME,
CASE WHEN QSYS2.QCMDEXC('HLDJOB ' CONCAT JOB_NAME) = 1 THEN 'Job Held'
ELSE 'Job not held'
END AS HLDJOB_RESULT
FROM TABLE (QSYS2.ACTIVE_JOB_INFO (DETAILED_INFO=> 'ALL'))
WHERE SQL_STATEMENT_START_TIMESTAMP < CURRENT_TIMESTAMP - 2 HOURS;
```

RECEIVE_DATA_QUEUE table function

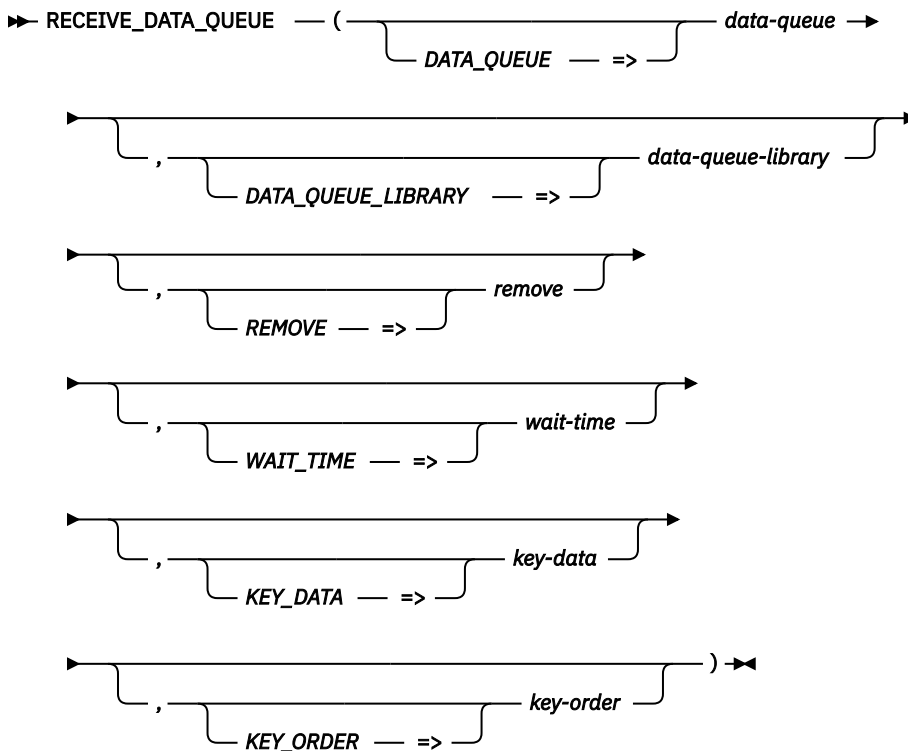
The RECEIVE_DATA_QUEUE table function returns a message from the specified data queue. The message data is returned as character, UTF-8, and binary data.

The MESSAGE_DATA, MESSAGE_DATA_UTF8, and MESSAGE_DATA_BINARY columns contain identical values. It is up to the user to determine which data type is most appropriate for working with the result of the receive data queue operation.

The values returned for the result columns of the table function are closely related to the values returned by the Receive Data Queue (QRCVDTAQ) API.

Authorization: The caller must have:

- *EXECUTE authority for the library, and
- *OBJOPR and *READ authority for the *DTAQ



The schema is QSYS2.

data-queue A character or graphic string containing the name of the data queue.

data-queue-library A character or graphic string containing the name of the library containing the data queue. Can be one of the following special values:

- *CURLIB** The job's current library is used.
- *LIBL** The library list is used. This is the default.

remove A character or graphic string indicating whether the message is to be removed from the data queue after it is received.

- NO** The message is not removed from the data queue.
- YES** The message is removed from the data queue. This is the default.

- wait-time** The amount of time to wait, in seconds and milliseconds, if no entries exist on the data queue. The value can be from -99999 to 99999.000.
- Any value less than zero indicates to wait forever.
 - A value of zero means to continue processing immediately. This is the default.
 - A positive value indicates the length of time to wait. The milliseconds part of the time is ignored for any value over 100 and for DDM data queues.

key-data A character string containing the data to use as the key for receiving a message from the data queue. This parameter is required for a keyed data queue. It must not be specified for a non-keyed data queue.

The length of *key-data* must be the length specified on the KEYLEN parameter on the Create Data Queue (CRTDTAQ) command. The KEY_LENGTH column of the QSYS2.DATA_QUEUE_INFO view contains this value.

When this parameter is specified, *key-order* must also be specified.

key-order The comparison criteria between the keys of messages on the data queue and the *key-data* parameter. When the system searches for the requested key, the entries are searched in ascending order from the lowest value key to the highest value key until a match is found. If there are entries with duplicate keys, the entry that was put on the queue first is received. Valid values are:

EQ	Equal
GE	Greater than or equal
GT	Greater than
LE	Less than or equal
LT	Less than
NE	Not equal

This parameter is ignored if *key-data* is not specified.

The result of the function is a table containing one row or no rows with the format shown in the following table. All the columns are nullable.

Table 94. RECEIVE_DATA_QUEUE table function

Column Name	Data Type	Description
MESSAGE_DATA	CLOB(64512)	The message received from the data queue as character data.
MESSAGE_DATA_UTF8	CLOB(64512) CCSID 1208	The message received from the data queue represented as character data in CCSID 1208.
MESSAGE_DATA_BINARY	BLOB(64512)	The message received from the data queue in binary form. This is the raw form of the data.
KEY_DATA	VARCHAR(256)	For a keyed data queue, the key value of the returned message. This is the actual key value, which could be different than the <i>key-data</i> parameter value. Contains the null value if this is not a keyed data queue.
SENDER_JOB_NAME	VARCHAR(28)	The qualified job name of the sender. Contains the null value if no sender information is available for the message.
SENDER_CURRENT_USER	VARCHAR(10)	The current user profile of the sender. Contains the null value if no sender information is available for the message.

Example

Get the message from data queue DQ1 in TESTLIB with key 456.

```
SELECT * FROM TABLE(QSYS2.RECEIVE_DATA_QUEUE(
                                DATA_QUEUE => 'DQ1',
```



```
DATA_QUEUE_LIBRARY => 'TESTLIB',
KEY_DATA => '456',
KEY_ORDER => 'EQ'));
```

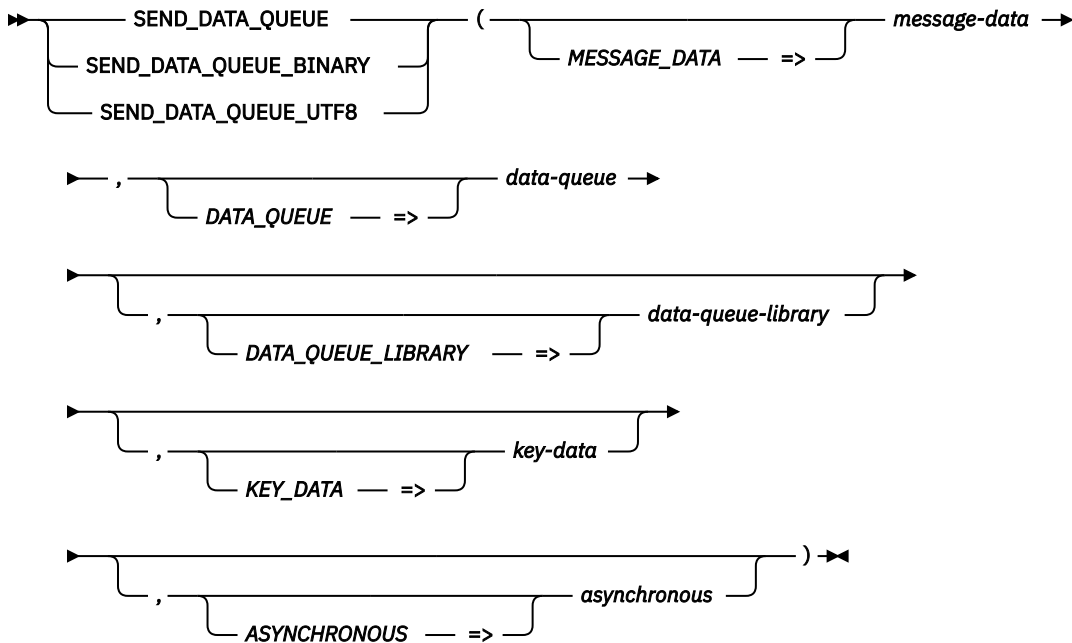
SEND_DATA_QUEUE, SEND_DATA_QUEUE_BINARY, and SEND_DATA_QUEUE_UTF8 procedures

The SEND_DATA_QUEUE, SEND_DATA_QUEUE_BINARY, and SEND_DATA_QUEUE_UTF8 procedures send a message to the specified data queue. The message data can be sent as character, UTF-8, or binary data.

This procedure provides function similar to the Send Data Queue (QSNDDTAQ) API.

Authorization: The caller must have:

- *EXECUTE authority for the library, and
- *OBJPR and *ADD authority for the *DTAQ



The schema is QSYS2.

- message-data** The data to send to the message queue. The string can be up to 64512 characters long.
- For SEND_DATA_QUEUE, the input value will be converted to a character string with the job CCSID.
 - For SEND_DATA_QUEUE_BINARY, the input value will be converted to a binary string.
 - For SEND_DATA_QUEUE_UTF8, the input value will be converted to a UTF-8 string.

data-queue A character or graphic string containing the name of the data queue.

data-queue-library A character or graphic string containing the name of the library containing the data queue. Can be one of the following special values:

- *CURLIB** The job's current library is used.
- *LIBL** The library list is used. This is the default.

key-data

A character string containing the key data to send to the data queue. This parameter is required for a keyed data queue. It must not be specified for a non-keyed data queue.

The length of the key must match the length specified on the KEYLEN parameter on the Create Data Queue (CRTDTAQ) command.

asynchronous

Indicates whether the send request to a DDM data queue should be processed asynchronously. Valid values are:

NO The request should not be processed asynchronously. This is the default.

YES The request should be processed asynchronously. This can only be specified for a DDM data queue.

If an error occurs during an asynchronous operation, the error will not be detected until the next time the data queue is accessed synchronously.

Example

Send a message to data queue DQ1 in TESTLIB with key 456.

```
CALL QSYS2.SEND_DATA_QUEUE(MESSAGE_DATA => 'This is the message data',
                           DATA_QUEUE => 'DQ1',
                           DATA_QUEUE_LIBRARY => 'TESTLIB',
                           KEY_DATA => '456');
```

SERVICES_INFO table

The SERVICES_INFO table returns information about system-supplied services.

Authorization: The caller must have:

- *OBJOPR and *READ authority to the QSYS2/SERV_INFO file, and
- *EXECUTE authority to the library containing the file.

The following table describes the columns in the table. The system name is SERV_INFO. The schema is QSYS2.

Table 95. SERVICES_INFO table

Column Name	System Column Name	Data Type	Description
SERVICE_CATEGORY	CATEGORY	VARCHAR(40)	Classification of the service. <ul style="list-style-type: none"> • APPLICATION • COMMUNICATION • DATABASE-APPLICATION • DATABASE-PERFORMANCE • DATABASE-PLAN CACHE • DATABASE-UTILITY • IFS • JAVA • JOURNAL • LIBRARIAN • MESSAGE HANDLING • PRODUCT • PTF • SECURITY • SPOOL • STORAGE • SYSTEM HEALTH • WORK MANAGEMENT

Table 95. SERVICES_INFO table (continued)

Column Name	System Column Name	Data Type	Description
SERVICE_SCHEMA_NAME	SYS_NAME	VARCHAR(128)	Name of the schema containing the service.
SERVICE_NAME	SERVNAME	VARCHAR(128)	Name of the service.
SQL_OBJECT_TYPE	SQLTYPE	VARCHAR(15)	The type of object. <ul style="list-style-type: none"> • PROCEDURE • SCALAR FUNCTION • TABLE • TABLE FUNCTION • VIEW
OBJECT_TYPE	OBJTYPE	VARCHAR(7) Nullable	The system object type of the service. <ul style="list-style-type: none"> • *FILE • *SRVPGM Contains null for procedures and functions implemented as external routines.
SYSTEM_OBJECT_NAME	SYS_ONAME	VARCHAR(10) Nullable	The system name of the service. <ul style="list-style-type: none"> • *FILE • *SRVPGM Contains null for procedures and functions implemented as external routines.
LATEST_DB2_GROUP_LEVEL	GROUPLVL	INTEGER Nullable	The Db2 for i PTF Group level which most recently changed this service. <ul style="list-style-type: none"> • *FILE • *SRVPGM Contains null if the service has not been enhanced in a PTF in this release.
INITIAL_DB2_GROUP_LEVEL	INITIALVL	INTEGER Nullable	The Db2 for i PTF Group level where this service was introduced. <ul style="list-style-type: none"> • *FILE • *SRVPGM Contains null if this service was available in the base for this release.
EARLIEST_POSSIBLE_RELEASE	MINRLS	VARCHAR(6)	The earliest release, in VxRxMx format, where a version of this service is available.
EXAMPLE	EXAMPLE	VARCHAR(5000)	An example SQL script that uses this service.

Example

Show all the available PTF services:

```
SELECT * FROM QSYS2.SERVICES_INFO
WHERE SERVICE_CATEGORY = 'PTF'
```

Db2 PTF Group dependencies

To complement the Db2 PTF Group level information provided by the SERVICES_INFO catalog table, you can determine the Db2 PTF Group dependency level for every static SQL statement within a module, program, or service program. The QSYS2.SYSPROGRAMSTMTSTAT catalog contains one row for every static SQL statement. The Db2 PTF Group dependency information is surfaced in two columns:

SQL_DB2_GROUP_LEVEL

Indicates the use of SQL language features. For example, new SQL statements or query clauses surface as dependencies upon having a certain Db2 PTF Group level (or higher) installed before the statement can be run.

This is an SQL syntax level and is an accurate indication of the dependency level.

SERVICES_DB2_GROUP_LEVEL

Indicates the consumption of IBM i Services. For example, queries that reference Db2 for i provided views, functions, procedures, or global variables can surface possible dependencies upon having a certain Db2 PTF Group level (or higher) installed before executing the

statement. If multiple services are used within a single SQL statement, the highest dependency level is returned.

The services that are instrumented are documented in [“IBM i Services” on page 391](#) and [Db2 for i Services](#). SQL built-in functions and built-in global variables are also tracked.

This is not an exact indication of the Db2 PTF Group that is needed. It depends on how the service is being used in your application. The information is provided based solely on the name of the service and the knowledge of when the latest enhancement was added for that service. If the name of an IBM-provided service matches an unqualified name in an SQL statement, it will be tracked as the IBM service. Based on the reported use of these services, you will need to determine whether the reported Db2 PTF Group is actually required.

To check all programs in APPLIB for potential SQL syntax and IBM i Service dependencies, execute the following query. Only programs created after the SERVICES_INFO table was introduced will report this information.

```
SELECT PROGRAM_NAME, SQL_DB2_GROUP_LEVEL, SERVICES_DB2_GROUP_LEVEL
FROM QSYS2.SYSPROGRAMSTMTSTAT
WHERE PROGRAM_SCHEMA = 'APPLIB' AND
      (SQL_DB2_GROUP_LEVEL IS NOT NULL OR
       SERVICES_DB2_GROUP_LEVEL IS NOT NULL);
```

To see more detailed information about which services are used in a program, including the name of each service and the Db2 PTF Group level required for the service, perform the following steps:

1. STRDBG UPDPROD(*YES)
2. Precompile your program or build your SQL procedure, function, or trigger.
 - To have informational messages written to the listing, add SET OPTION OUTPUT=*PRINT to your SQL routine or specify the OUTPUT(*PRINT) parameter on the CRTSQLxxx or RUNSQLSTM CL commands
3. For each reference to a service, message SQL7901 will be written to the joblog and, optionally, to the precompile listing.

If you precompile with a TGTRLS of 7.1 or later, a message will be issued for each of the earlier releases as well with an indication of the Db2 PTF Group level that is needed on that release. If the service is not supported for a release, message SQL795B will be issued.

This information can be used to determine whether your application contains any content that might require a certain level of Db2 PTF Group. If you need to deploy your application to a different partition or an earlier release, this feedback can alert you to potential dependencies.

After you have created one or more objects using the steps above, you can query your job log to see if any messages were issued that might need to be addressed.

```
SELECT MESSAGE_ID, MESSAGE_TEXT
FROM TABLE(QSYS2.JOBLOG_INFO('*')) X
WHERE MESSAGE_ID IN ('SQL7901', 'SQL795B')
ORDER BY ORDINAL_POSITION;
```

Here is one more query to help tie this information together. It will tell you the Db2 PTF Group level that is on a partition.

```
SELECT MAX(PTF_GROUP_LEVEL) AS DB2_PTF_LEVEL FROM QSYS2.GROUP_PTF_INFO
WHERE PTF_GROUP_NAME LIKE 'SF9970%' AND PTF_GROUP_STATUS = 'INSTALLED';
```


SET_PASE_SHELL_INFO procedure

The SET_PASE_SHELL_INFO procedure provides the ability to set the path to the PASE shell for the specified user or the path to the default shell returned for users that do not have a configured shell.

The path set by this procedure is returned in the pw_shell field in struct pw from PASE APIs such as Get User Information for User Name (getpwnam) and Get User Information for User ID (getpwuid). It is also returned by the QSYS2.USER_INFO view . If a user does not have a path set, the default shell path is returned; if the default shell is not set, an empty string is returned. The pw_shell field is used by PASE applications that need to execute shells for a user, such as the OpenSSH server. The OpenSSH server will start this application as the initial program when the user logs in. If it is not set it will use /QOpenSys/usr/bin/bsh instead.

Authorization:

- If AUTHORIZATION_NAME is *CURRENT or matches the caller of this procedure, no authorization is needed.
- Otherwise the user calling this procedure must have:
 - *SECADM special authority and
 - *OBJMGT and *USE to the user profile identified by AUTHORIZATION_NAME.



The schema is QSYS2.

authorization-name A character or graphic string expression that identifies an existing user profile name. Can also be one of the following special values:

- *CURRENT** Set the current user's shell.
- *DEFAULT** Set the PASE shell to be used by any user that does not have an explicit value set. The default does not apply to IBM supplied profiles. The default is saved in the QSYS user profile. This is equivalent to specifying 'QSYS' for *authorization-name*.

shell-path A character or graphic string expression that specifies the path to a PASE shell. The string must begin with a forward slash (/). If *shell-path* is blanks, the empty string, or NULL, the shell path is removed for the user. Once the path is removed, the value specified for *DEFAULT (if any) will apply to this user.

Examples

- Set the current user's shell to BASH shipped by 5733-OPS.

```
CALL QSYS2.SET_PASE_SHELL_INFO(*CURRENT',  
                               '/QOpenSys/QIBM/ProdData/OPS/tools/bin/bash');
```


- Set the default shell to be ksh for any users that do not have an explicit shell set.

```
CALL QSYS2.SET_PASE_SHELL_INFO(*DEFAULT', '/QOpenSys/usr/bin/ksh');
```

SPLIT table function

The SPLIT table function returns a result table that contains one row for each substring of *input-list* that is separated by *delimiter*.

Authorization: None required.



input-list An expression that contains a list to be deconstructed. The value is cast to CLOB(2G).

delimiter A character string expression that defines the separator between list elements. The value is cast to VARCHAR(32672).

escape A character string expression with a length of 1 that defines a character that is used to escape a delimiter sequence. If this parameter is provided, any delimiter sequence of characters immediately preceded by this character will not be interpreted as a delimiter. The escape character will be removed from the returned element value.

The table function returns one row for each substring of *input-list* containing the characters between *delimiter* strings. If *input-list* contains two adjacent *delimiter* strings, an empty string is returned to represent an element with no content. If a *delimiter* string starts at position 1 of *input-list*, a row containing a zero length string is returned as the first element. If a *delimiter* string ends at the last position of *input-list*, a row containing a zero length string is returned as the last element. Substrings of *input-list* that match *delimiter* are not included in the result.

If *input-list* is null, the result contains no rows. If *delimiter* is null, the empty string, or a string that is not found, *input-list* is returned.

The result of the function is a table containing a row for each substring of *input-list*. The columns of the result table are described in the following table. The result columns are nullable.

Table 96. SPLIT table function

Column Name	Data Type	Description
ORDINAL	INTEGER	The relative position of this element in the input string. The first row has a value of 1.
ELEMENT	CLOB(2G)	The value of the element.

Note

This function is provided in the SYSTOOLS schema as an example of how to break a string apart at a delimiting character by using an SQL table function. Creating customized versions of this table function to better suit a specific need is encouraged. Use the Insert Generated SQL feature in IBM i Access Client Solutions (ACS) to extract the source for this function. Then modify it and create a new procedure in a user-specified schema.

Example

- Return a list of all authorities collected for all users for objects in the APP1 schema. First, a common table expression breaks the detailed authority lists into separate rows for each authority using the SPLIT table function. Then, all the authorities for the object are recombined into a single list for each user, removing duplicate authorities and listing them alphabetically, using the LISTAGG aggregate function.

```

WITH EXPANDED_AUTHS AS (
SELECT AUTHORIZATION_NAME, OBJECT_NAME, VARCHAR(TRIM(ELEMENT), 10) AS AUTH
  
```

```


FROM QSYS2.AUTHORITY_COLLECTION,
     TABLE(SYSTOOLS.SPLIT(DETAILED_REQUIRED_AUTHORITY, ' '))
WHERE OBJECT_SCHEMA = 'APP1'
      AND CHECK_ANY_AUTHORITY = 0)
SELECT AUTHORIZATION_NAME, OBJECT_NAME,
       LISTAGG(DISTINCT AUTH, ' ') WITHIN GROUP(ORDER BY AUTH)
FROM EXPANDED_AUTHS
GROUP BY AUTHORIZATION_NAME, OBJECT_NAME
ORDER BY AUTHORIZATION_NAME, OBJECT_NAME;

```

STACK_INFO table function

The STACK_INFO table function returns one row for each entry in the call stack for either a specific thread or for every thread of the specified job. It returns information similar to what can be accessed through the Display Job (DSPJOB) CL command and the Retrieve Call Stack (QWVRCSTK) API.

Authorization: The authorization ID of the statement must be the same user profile that is running the specified *job-name*, or must have either *JOBCTL special authority or QIBM_SERVICE_THREAD function usage. If the authorization ID has *SERVICE special authority, the returned call stack information will include Licensed Internal Code (LIC) stack entries.



The schema is QSYS2.

job-name The qualified job name to return stack information. Can contain the following special value:

- * The current job name is used.

If *job-name* is not specified, the default is *.

thread-id A numeric expression indicating the thread identifier to return information for. Can contain one of the following special values:

- ALL** Information for all the threads in the job is returned.
- INITIAL** Information for the initial thread of the job is returned.

If *thread-id* is not specified:

- If *job-name* is *, the default is the value of the QSYS2.THREAD_ID global variable.
- Otherwise, the default is INITIAL.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 97. STACK_INFO table function

Column Name	Data Type	Description
THREAD_ID	BIGINT	The identifier for the specific thread.
THREAD_TYPE	VARCHAR(6)	Specifies how the thread was initiated. <ul style="list-style-type: none"> SYSTEM The thread was initiated by the operating system. USER The thread was initiated by a user process. Contains the null value unless ALL was specified for the <i>thread-id</i> input parameter.

Table 97. STACK_INFO table function (continued)

Column Name	Data Type	Description
ORDINAL_POSITION	INTEGER	A unique number for each row corresponding to a thread where 1 is the first invocation entry for this thread and the highest number is the most recent invocation entry for this thread.
ENTRY_TYPE	VARCHAR(4)	The type of stack entry. ILE This entry returns ILE program information. The columns specific to JAVA, PASE, and LIC contain the null value. JAVA This entry returns JAVA information. The columns specific to ILE and OPM, PASE, and LIC contain the null value. LIC This entry returns Licensed Internal Code (LIC) information. The columns specific to ILE and OPM, JAVA, and PASE contain the null value. OPM This entry returns OPM program information. The columns specific to JAVA, PASE, and LIC contain the null value. PASE This entry returns PASE information. The columns specific to ILE and OPM, JAVA, and LIC contain the null value.
--- ILE and OPM information -----		
PROGRAM_NAME	VARCHAR(10) Nullable	The name of the program or service program. Contains the null value if the program name is not available.
PROGRAM_LIBRARY_NAME	VARCHAR(10) Nullable	The name of the library in which the program is located. Contains the null value if the program is not located in a library or if the program library name is not available.
STATEMENT_IDENTIFIERS	VARCHAR(109) Nullable	The high-level language statement identifier. If this column contains the character representation of a number, the number is right-adjusted and padded on the left with zeros (for example, '0000000246'). If the call stack entry is for an integrated language environment (ILE) procedure, more than one statement identifier may exist. If more than one statement identifier is returned, each identifier will be ten characters long with a single blank between them. Up to ten identifiers will be returned. Returns the null value if a statement identifier cannot be determined.
REQUEST_LEVEL	INTEGER Nullable	The level of the request-processing program or procedure. Contains the null value if the program or procedure has not received a request message or incomplete information is available.
CONTROL_BOUNDARY	VARCHAR(3) Nullable	Whether a control boundary exists for a program or procedure. A control boundary is defined as any ILE call stack entry for which the immediately preceding call stack entry is for an ILE procedure or program object in a different activation group. NO No control boundary is active. YES A control boundary is active. Contains the null value if information is not available or incomplete information is available.
PROGRAM_ASP_NAME	VARCHAR(10) Nullable	The name of the auxiliary storage pool (ASP) device in which the program is located. Can contain the following special value: *SYSBAS The program is located in the system ASP or a basic user ASP Contains the null value if the name of the ASP cannot be determined.
PROGRAM_ASP_NUMBER	INTEGER Nullable	The numeric identifier of the ASP containing the program. 1 The program is in the system ASP. 2-32 The program is in a basic user ASP. 33-255 The program is in an independent ASP. Contains the null value if the ASP device cannot be determined.

Table 97. STACK_INFO table function (continued)

Column Name	Data Type	Description
MODULE_NAME	VARCHAR(10) Nullable	The module containing the integrated language environment (ILE) procedure. Contains the null value if this is not an ILE program or if the module name is not available.
MODULE_LIBRARY_NAME	VARCHAR(10) Nullable	The name of the library in which the module is located. Contains the null value if this is not an ILE program or if the module library name is not available.
PROCEDURE_NAME	VARCHAR(4096) Nullable	The name of the procedure at this level of the call stack. Returns the null value if this is not an ILE program or if the procedure name cannot be determined.
ACTIVATION_GROUP_NUMBER	DECIMAL(20,0) Nullable	The number of the activation group within which the program or procedure is running. This is an internal number that uniquely identifies the activation group within the job. Contains the null value if this is not an ILE program or incomplete information is available.
ACTIVATION_GROUP_NAME	VARCHAR(10) Nullable	The name of the activation group within which the program or procedure is running. Can contain the following special values: <p>*DFACTGRP The activation group does not have a specific name. The activation group is one of the default activation groups for the system.</p> <p>*NEW The activation group does not have a specific name. The activation group was created when the program was called.</p> <p>Contains the null value if this is not an ILE program or incomplete information is available.</p>
MI_INSTRUCTION_NUMBER	INTEGER Nullable	The current machine instruction number in the program. Contains the null value if this is not an OPM program.
--- JAVA information -----		
JAVA_LINE_NUMBER	INTEGER Nullable	The line number where the invocation was interrupted. Contains the null value if no line number can be determined.
JAVA_BYTE_CODE_OFFSET	INTEGER Nullable	The offset in bytes from the beginning of the Java method byte codes to the resume point for the invocation. Contains the null value if no Java byte code offset can be determined.
JAVA_METHOD_TYPE	VARCHAR(9) Nullable	The type of Java method. <p>DE The method is a direct execution Java method. The Java method has been precompiled by the Java Transformer.</p> <p>GLUE The invocation is a Java Virtual Machine glue frame used either to perform a call from the JVM to a Java method or perform a call to a Java native method.</p> <p>INTERPRET The method is an interpreted Java method. The Java method is being interpreted by the Java Interpreter.</p> <p>JIT The method is a JIT compiled Java method. The Java method has been compiled by the Java Just In Time Compiler.</p> <p>MMI The method is a MMI interpreted Java method. The Java method is being interpreted by the Mixed Mode Java Interpreter.</p> <p>Contains the null value if there is no information.</p>
JAVA_CLASS_NAME	DBCLOB(64000) CCSID 1200 Nullable	The name of the Java class at this level of the call stack. Returns the null value if the class name cannot be determined.
JAVA_METHOD_NAME	DBCLOB(64000) CCSID 1200 Nullable	The name of the Java method at this level of the call stack. Returns the null value if the method name cannot be determined.

Table 97. STACK_INFO table function (continued)

Column Name	Data Type	Description
JAVA_METHOD_SIGNATURE	DBCLOB(64000) CCSID 1200 Nullable	The signature of the Java method at this level of the call stack. Returns the null value if the signature cannot be determined.
JAVA_FILE_NAME	DBCLOB(64000) CCSID 1200 Nullable	The name of the Java file and directory that provides the location of where the Java class was loaded at this level of the call stack. If the Java class was loaded from a .jar or .zip file, then the location will be the path to and the name of the .jar or .zip file. If the class was loaded from a .class file, then the location will be the directory from which the class was loaded. Returns the null value if the file name cannot be determined.
JAVA_SOURCE_FILE_NAME	DBCLOB(64000) CCSID 1200 Nullable	The name of the Java source file at this level of the call stack. Returns the null value if the source file name cannot be determined.
--- PASE information -----		
PASE_LINE_NUMBER	BIGINT Nullable	The line number where the invocation was interrupted. Contains the null value if no line number can be determined.
PASE_INSTRUCTION_ADDRESS	DECIMAL(20,0) Nullable	The IBM PASE for i memory address for the instruction that will run when execution resumes for the invocation.
PASE_INSTRUCTION_OFFSET	BIGINT Nullable	The offset in bytes from the beginning of the start of the procedure to the instruction that is either the suspend point for the invocation or the resume point for the invocation.
PASE_KERNEL_CODE	VARCHAR(3) Nullable	Whether the invocation is running IBM PASE for i kernel code. NO The current invocation is not IBM PASE for i kernel code. YES The current invocation is IBM PASE for i kernel code. Contains the null value if there is no information.
PASE_BIT_CODE	INTEGER Nullable	Whether the invocation is running 32-bit or 64-bit IBM PASE for i code. 32 The invocation is running 32-bit IBM PASE for i code. 64 The invocation is running 64-bit IBM PASE for i code. Contains the null value if PASE_KERNEL_CODE is YES.
PASE_ALTERNATE_RESUME_POINT	VARCHAR(3) Nullable	Whether the current entry is a second entry for a given invocation. This flag is only used when the system can not reliably determine which of two possible resume points will be used when an invocation resumes execution. NO The current invocation does not have an alternate resume point. YES The current invocation has an alternate resume point. Contains the null value if there is no information.
PASE_PROCEDURE_NAME	DBCLOB(4000) CCSID 1200 Nullable	The name of the procedure at this level of the call stack. Returns the null value if the procedure name cannot be determined.
PASE_LOAD_MODULE_NAME	DBCLOB(1000) CCSID 1200 Nullable	The name of the load module at this level of the call stack. Returns the null value if the load module name cannot be determined.
PASE_LOAD_MODULE_PATH	DBCLOB(4000) CCSID 1200 Nullable	The path to the load module at this level of the call stack. Returns the null value if the load module path cannot be determined.
PASE_SOURCE_PATH_AND_FILE	DBCLOB(1000) CCSID 1200 Nullable	The path and name for the source file used to create the procedure. Returns the null value if the path and name for the source file cannot be determined.
--- LIC information -----		

Table 97. STACK_INFO table function (continued)

Column Name	Data Type	Description
LIC_INSTRUCTION_OFFSET	BIGINT Nullable	The offset in bytes from the beginning of the start of the procedure to the instruction that is either the suspend point for the invocation or the resume point for the invocation.
LIC_PROCEDURE_NAME	VARCHAR(4096) Nullable	The name of the procedure at this level of the call stack. Returns the null value if the procedure name cannot be determined.
LIC_LOAD_MODULE_NAME	VARCHAR(64) Nullable	The name of the load module at this level of the call stack. Returns the null value if the load module name cannot be determined.

Example

- Find out whether ILE program MYPGM is on the stack for the current thread.

```
SELECT * FROM TABLE(QSYS2.STACK_INFO('*')) A
WHERE PROGRAM_NAME = 'MYPGM';
```

- Create a table that contains the stack for all of the threads in a specific job.

```
CREATE TABLE STACK_DUMP AS (
SELECT * FROM TABLE(QSYS2.STACK_INFO('358788/QLIWISVR/ADMIN1', 'ALL')) AS X
) WITH DATA;
```


USER_INDEX_ENTRIES table function

The USER_INDEX_ENTRIES table function returns the entries of the specified user index (*USRIDX). The data is returned as character and binary data.

The values returned for the result columns of the table function are closely related to the values returned by the Retrieve User Index Entries (QUSRTVUI) API.

Authorization: The caller must have:

- *EXECUTE authority for the library, and
- *USE authority for the user index.



The schema is QSYS2.

user-index A character or graphic string containing the name of the user index.

user-index-library A character or graphic string containing the name of the library containing the user index. Can be one of the following special values:

***CURLIB** The job's current library is used.

***LIBL** The library list is used. This is the default.

The result of the function is a table containing one or more rows with the format shown in the following table. All the columns are nullable.

Table 98. USER_INDEX_ENTRIES table function

Column Name	Data Type	Description
ORDINAL_POSITION	INTEGER	The relative position of this row in the result data set.

Table 98. USER_INDEX_ENTRIES table function (continued)

Column Name	Data Type	Description
USER_INDEX_LIBRARY	VARCHAR(10)	The library in which the user index was found.
USER_INDEX	VARCHAR(10)	The name of the user index.
KEY	VARCHAR(2000)	The key for the index entry in character form. Contains the null value if the user index is not keyed.
KEY_BINARY	VARBINARY(2000)	The key for the index entry in binary form. This is the raw form of the data. Contains the null value if the user index is not keyed.
ENTRY	VARCHAR(2000)	The data for the index entry in character form. Contains the null value if the entry contains only a key value.
ENTRY_BINARY	VARBINARY(2000)	The data for the index entry in binary form. This is the raw form of the data. Contains the null value if the entry contains only a key value.

Example

Look at all the entries in user index IX1 in TESTLIB.

```
SELECT * FROM TABLE(QSYS2.USER_INDEX_ENTRIES(
                                USER_INDEX => 'IX1',
                                USER_INDEX_LIBRARY => 'TESTLIB'))
ORDER BY ORDINAL_POSITION;
```

USER_INDEX_INFO view

The USER_INDEX_INFO view returns the attributes of user indexes.

The values returned for the columns in the view are closely related to the values returned by the Retrieve User Index Attributes (QUSRUIAT) API.

Authorization: The caller must have:

- *EXECUTE authority to the library containing the user index, and
- *USE authority to the user index.

The following table describes the columns in the view. The system name is USRIDX_INF. The schema is QSYS2.

Table 99. USER_INDEX_INFO view

Column Name	System Column Name	Data Type	Description
USER_INDEX_LIBRARY	USRIDX_LIB	VARCHAR(10)	Library containing the user index.
USER_INDEX	USRIDX	VARCHAR(10)	Name of the user index.
ENTRY_TYPE	TYPE	VARCHAR(8)	The type of entries in the user index. FIXED Fixed-length entries VARIABLE Variable-length entries
ENTRY_LENGTH	LENGTH	INTEGER	When ENTRY_TYPE is FIXED, the length of each index entry. When ENTRY_TYPE is VARIABLE, the length of the longest entry that has ever been inserted into the index. Valid values are from 1 through 2000.
MAXIMUM_ENTRY_LENGTH	MAX_LENGTH	INTEGER	The maximum entry length any user index entry can have.

Table 99. USER_INDEX_INFO view (continued)

Column Name	System Column Name	Data Type	Description
INDEX_SIZE	SIZE	CHAR(4)	The maximum size of the user index. 4 GB The maximum size of the user index is 4 gigabytes. 1 TB The maximum size of the user index is 1 terabyte.
IMMEDIATE_UPDATE	IMMEDIATE	VARCHAR(3)	Whether updates to the index are written to auxiliary storage on each update to the index. NO No immediate update YES Immediate update
OPTIMIZATION	OPTIMIZE	VARCHAR(10)	The optimization method used for user index maintenance. RANDOM Random references SEQUENTIAL Sequential references
KEY_INSERTION	KEYED	VARCHAR(3)	Whether inserts into the index are by key. NO No insertion by key YES Insertion by key
KEY_LENGTH	KEY_LENGTH	INTEGER Nullable	The length of the key. Contains the null value when KEY_INSERTION is NO.
ENTRY_TOTAL	TOTAL	INTEGER	The number of entries currently in the index.
ENTRIES_ADDED	ADDED	INTEGER	The number of entries added to the user index.
ENTRIES_REMOVED	REMOVED	INTEGER	The number of entries removed from the user index.
OBJECT_DOMAIN	DOMAIN	VARCHAR(7)	The domain of the object. *SYSTEM The object is in the system domain. *USER The object is in the user domain.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The text description of the user index. Contains the null value if the user index has no description.

Example

- Return a list of attributes for all user indexes in MYLIB.

```
SELECT * FROM QSYS2.USER_INDEX_INFO
WHERE USER_INDEX_LIBRARY = 'MYLIB';
```


USER_SPACE table function

The USER_SPACE table function returns the contents of the specified user space (*USRSPC). The data is returned as character and binary data.

The values returned for the result columns of the table function are closely related to the values returned by the Retrieve User Space (QUSRTVUS) API.

Authorization: The caller must have:

- *EXECUTE authority for the library, and
- *USE authority for the user space.



The schema is QSYS2.

- user-space** A character or graphic string containing the name of the user space.
- user-space-library** A character or graphic string containing the name of the library containing the user space. Can be one of the following special values:
 - *CURLIB** The job's current library is used.
 - *LIBL** The library list is used. This is the default.

The result of the function is a table containing one row with the format shown in the following table. All the columns are nullable.

Table 100. USER_SPACE table function

Column Name	Data Type	Description
USER_SPACE_LIBRARY	VARCHAR(10)	The library in which the user space was found.
USER_SPACE	VARCHAR(10)	The name of the user space.
DATA	CLOB(16M)	The data from the user space as character data.
DATA_BINARY	BLOB(16M)	The data from the user space in binary form. This is the raw form of the data.

Example

Look at the untranslated contents of user space USERSPC1 in TESTLIB.

```
SELECT DATA_BINARY FROM TABLE(QSYS2.USER_SPACE(
    USER_SPACE => 'USRSPACE1',
    USER_SPACE_LIBRARY => 'TESTLIB'));
```

USER_SPACE_INFO view

The USER_SPACE_INFO view returns the attributes of user spaces.

The values returned for the columns in the view are closely related to the values returned by the Retrieve User Space Attributes (QUSRUSAT) API.

Authorization: The caller must have:

- *EXECUTE authority to the library containing the user space, and
- *USE authority to the user space.

The following table describes the columns in the view. The system name is USRSPC_INF. The schema is QSYS2.

Table 101. USER_SPACE_INFO view

Column Name	System Column Name	Data Type	Description
USER_SPACE_LIBRARY	USRSPC_LIB	VARCHAR(10)	Library containing the user space.
USER_SPACE	USRSPC	VARCHAR(10)	Name of the user space.
SIZE	SIZE	INTEGER	The size of the user space object in bytes.

Table 101. USER_SPACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
EXTENDABLE	EXTENDABLE	VARCHAR(3)	Whether the space is extended automatically by the system when the end of the space is encountered. NO Space is not automatically extendable YES Space is automatically extendable
INITIAL_VALUE	INITIAL	BINARY(1)	The initial value to which future extensions of the user space will be set.
OBJECT_DOMAIN	DOMAIN	VARCHAR(7)	The domain of the object. *SYSTEM The object is in the system domain. *USER The object is in the user domain.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The text description of the user space. Contains the null value if the user space has no description.

Example

- Return a list of user spaces in MYLIB.

```
SELECT USER_SPACE, SIZE, TEXT_DESCRIPTION FROM QSYS2.USER_SPACE_INFO
WHERE USER_SPACE_LIBRARY = 'MYLIB';
```

WATCH_DETAIL table function

The WATCH_DETAIL table function returns the details for watched messages, LIC logs, and PALs for a specific session identifier.

The values returned for the columns in the table function are closely related to the values returned by the WRKWCH (Work with Watches) CL command and by the Retrieve Watch Information (QSCRWCHI) API.

Authorization: The caller must have:

- *USE authority on the QSYS/QSCRWCHI program and
 - *SERVICE special authority or
 - authorized to the QIBM_SERVICE_WATCH and QIBM_SERVICE_TRACE function usage identifiers.

►► WATCH_DETAIL ((SESSION_ID =>) session-id) ►►

The schema is QSYS2.

session-id A character or graphic string expression that contains the session identifier that the watch details are returned for.

The result of the function is a table containing rows with the format shown in the following table. All columns are nullable.

Table 102. WATCH_DETAIL table function

Column Name	Data Type	Description
DETAIL_TYPE	VARCHAR(7)	The type of detail information returned by this row. LICLOG This row is for a watched LIC log. MESSAGE This row is for a watched message. PAL This row is for a watched PAL.

Table 102. WATCH_DETAIL table function (continued)

Column Name	Data Type	Description
COMPARISON_DATA_CCSSID	INTEGER	The CCSID that pertains to each of the xxx_COMPARISON_DATA columns. Contains the null value if this row has no comparison data.
The following columns apply when DETAIL_TYPE is MESSAGE. They will contain the null value when DETAIL_TYPE is LICLOG or PAL.		
MESSAGE_ID	VARCHAR(7)	The message to be watched. message-id The 7-character message identifier of the message to be watched. generic-name The generic name of the message to be watched. A generic name is a character string of one or more characters followed by an asterisk (*); for example, ABC*. The asterisk substitutes for any valid characters. A generic message name specifies all messages with identifiers that begin with the generic prefix. *ALL All messages are to be watched. This includes stored messages and immediate messages. *IMMED All immediate or impromptu messages will be watched.
MESSAGE_TYPE	VARCHAR(6)	The message type assigned to the message to be watched. ALL All the message types are to be watched. This includes completion, command, diagnostic, escape, informational, inquiry, notify, reply, request, sender's copy, scope and status message types. COMP A completion message is to be watched. DIAG A diagnostic message is to be watched. ESCAPE An escape message is to be watched. INFO An informational message is to be watched. INQ An inquiry message is to be watched. NOTIFY A notify message is to be watched. SCOPE A scope message is to be watched. STATUS A status message is to be watched.
MESSAGE_QUEUE_LIBRARY	VARCHAR(10)	The name of the library where the message queue is located. Contains the null value if MESSAGE_QUEUE is *JOBLOG.
MESSAGE_QUEUE	VARCHAR(10)	The name of the message queue to watch. Can contain the following special value: *JOBLOG Watch messages added to the job logs of the jobs specified for the watched job field.
MESSAGE_JOB_NAME	VARCHAR(10)	The job name of the job to be watched. Can contain the following special values: generic-name The generic name of the job to be watched. A generic name is a character string of one or more characters followed by an asterisk (*); for example, ABC*. The asterisk substitutes for any valid characters. A generic job name specifies all jobs with job names that begin with the generic prefix. *ALL All jobs with the specified job user name are watched. Contains the null value if MESSAGE_QUEUE is not *JOBLOG.

Table 102. WATCH_DETAIL table function (continued)

Column Name	Data Type	Description
MESSAGE_JOB_USER	VARCHAR(10)	<p>The user name of the job to be watched. Can contain the following special values:</p> <p>generic-name The generic name of the user name of the job to be watched. A generic name is a character string of one or more characters followed by an asterisk (*); for example, ABC*. The asterisk substitutes for any valid characters. A generic user name specifies all jobs with the specified job name and with user names that begin with the generic prefix.</p> <p>*ALL All jobs with the specified job name are watched.</p> <p>Contains the null value if MESSAGE_QUEUE is not *JOBLOG.</p>
MESSAGE_JOB_NUMBER	VARCHAR(6)	<p>The job number to further qualify the job name and user name. Can contain the following special value:</p> <p>*ALL All jobs with the specified job name and user name are watched.</p> <p>Contains the null value if a generic job name or a generic user name qualifier is specified, or if MESSAGE_QUEUE is not *JOBLOG.</p>
MESSAGE_SEVERITY	INTEGER	<p>The severity code, ranging from 0 through 99, of the message to be watched.</p>
MESSAGE_RELATIONAL_OPERATOR	CHAR(3)	<p>The relational operator against which the message severity code is compared.</p> <p>*EQ Equal</p> <p>*GE Greater than or equal</p> <p>*GT Greater than</p> <p>*LE Less than or equal</p> <p>*LT Less than</p>
MESSAGE_COMPARISON_DATA	VARCHAR(72) CCSID 65535	<p>The comparison data to be used if a message matching the specified message is added to the specified message queue or log. If the message data, the "From program" or the "To program" includes the specified text, the watched for condition is true.</p> <p>Contains the null value if no message comparison data was specified.</p>
MESSAGE_COMPARE_AGAINST	VARCHAR(8)	<p>The part of the message the data specified in MESSAGE_COMPARISON_DATA is to be compared against.</p> <p>*FROMPGM The message comparison data will be compared against the name of the program sending the message, or the name of the ILE program that contains the procedure sending the message.</p> <p>*MSGDTA The message comparison data will be compared against the message replacement data.</p> <p>*TOPGM The message comparison data will be compared against the name of the program the message was sent to, or the name of the ILE program that contains the procedure the message was sent to.</p> <p>Contains the null value if no message comparison data was specified.</p>
<p>The following columns apply when DETAIL_TYPE is LIC_LOG. They will contain the null value when DETAIL_TYPE is MESSAGE or PAL.</p>		
LIC_MAJOR_CODE	CHAR(4)	<p>The LIC log major code to be watched. A hexadecimal digit or a question mark can be specified for each character in the four-digit code. A question mark is a wildcard character that will match any digit in that position. Up to three wildcard characters can be specified. Can contain the following special value:</p> <p>*ALL Any LIC log entry major code will be considered to be a match.</p>

Table 102. WATCH_DETAIL table function (continued)

Column Name	Data Type	Description
LIC_MINOR_CODE	CHAR(4)	<p>The LIC log minor code to be watched. A hexadecimal digit or a question mark can be specified for each character in the four-digit code. A question mark is a wildcard character that will match any digit in that position. Up to three wildcard characters can be specified. Can contain the following special value:</p> <p>*ALL Any LIC log entry minor code will be considered to be a match.</p>
LIC_COMPARISON_DATA	VARCHAR(72) CCSID 65535	<p>The comparison data to be used if a log entry matching the specified major and minor codes is added to the licensed internal code (LIC) log. If this text is found in the LIC log entry data field specified by the LIC log compare against field, the watched for condition is true. This text is case sensitive. If *ALL is specified in the LIC log compare against field, the LIC log fields which will be compared are TDE number, task name, server type, job name, user ID, job number, thread ID, exception ID, LIC module compile binary timestamp, module offset, LIC module RU name, LIC module name, LIC module entry point name. The comparison data cannot be used to match across two fields, and can match an entire field or a substring of any field. The prefix MCH may be specified to compare only against the exception ID field and avoid possible substring matches with the other fields.</p> <p>Contains the null value if no LIC log comparison data was specified.</p>

Table 102. WATCH_DETAIL table function (continued)

Column Name	Data Type	Description
LIC_COMPARE_AGAINST	VARCHAR(10)	The part of the LIC log the data specified in LIC_COMPARISON_DATA is to be compared against.
		<p>*ALL The LIC log comparison data will be compared against all the fields described below.</p> <p>*EXCPID The LIC log comparison data will be compared against the exception that caused the LIC log entry to be requested. This is a 2-byte hexadecimal field formed by concatenating to the high-order 1-byte exception group number a low-order 1-byte exception subtype number. Exception identifier is binary zeros if the LIC log entry is not requested as a result of an exception.</p> <p>*JOBNAME The LIC log comparison data will be compared against the name of the job which requested the LIC log entry. LIC job name is blank (hex 40s) if the LIC log entry is not requested by a job.</p> <p>*JOBNBR The LIC log comparison data will be compared against the job number (000001-999999) to further qualify the job name and user name of the job which requested the LIC log entry. LIC job number is blank (hex 40s) if the LIC log entry is not requested by a job.</p> <p>*JOBUSR The LIC log comparison data will be compared against the user name of the job which requested the LIC log entry. LIC user name is blank (hex 40s) if the LIC log entry is not requested by a job.</p> <p>*MODEPNAME The LIC log comparison data will be compared against the name of the entry point which requested the LIC log entry. If the entry point name is greater than 128 characters, the LIC module entry point name is truncated to 128 characters.</p> <p>*MODNAME The LIC log comparison data will be compared against the LIC module name which requested the LIC log entry. If the module name is greater than 64 characters, the LIC module name is truncated to 64 characters.</p> <p>*MODOFFSET The LIC log comparison data will be compared against the byte offset into the LIC module text which requested the LIC log entry.</p> <p>*MODRUNAME The LIC log comparison data will be compared against the LIC module replacement unit name. LIC module RU name is always in upper case EBCDIC.</p> <p>*MODTSP The LIC log comparison data will be compared against the timestamp of when the LIC module was compiled. The format for this field is the system timestamp format.</p> <p>*SVRTYPE The LIC log comparison data will be compared against the type of server that requested the LIC log entry. Server type is blank (hex 40s) if the LIC log entry is not requested by a server.</p> <p>*TASKNAME The LIC log comparison data will be compared against the name of the task which requested the LIC log entry. Task name is blank (hex 40s) if the LIC log entry is not requested by a task.</p> <p>*TDENBR The LIC log comparison data will be compared against the number of the task dispatching element (TDE) which requested the LIC log entry.</p> <p>*THDID The LIC log comparison data will be compared against the thread which requested the LIC log entry. Thread identifier is binary zeros if the LIC log entry is not requested by a thread</p>
		Contains the null value if no LIC log comparison data was specified.

The following columns apply when **DETAIL_TYPE** is **PAL**. They will contain the null value when **DETAIL_TYPE** is **MESSAGE** or **LICLOG**.

Table 102. WATCH_DETAIL table function (continued)

Column Name	Data Type	Description
PAL_SYSTEM_REFERENCE_CODE	VARCHAR(8)	<p>The system reference code to be watched. A hexadecimal digit or a question mark can be specified for each character in the eight-digit code. A question mark is a wildcard character that will match any digit in that position. Up to seven wildcard characters can be specified. A generic name that is a character string of one or more characters followed by an asterisk (*) can also be specified; for example, ABC*. The asterisk substitutes for any valid characters. A generic name specifies all PALs with system reference codes that begin with the generic prefix. Can contain the following special value:</p> <p>*ALL Any PAL system reference code will be considered to be a match.</p>
PAL_COMPARISON_DATA	VARCHAR(10) CCSID 65535	<p>The comparison data to be used if a PAL entry matching the specified system reference code was created. If the data specified by PAL compare against field matches the specified text, the watched for condition is true. This text is case sensitive. Question marks (?) and asterisk (*) wildcard characters can be specified in the text string. A question mark is a single-character wildcard character and will match any character in the same position. For example, '??123' will match any value that is five characters long and ends with '123'. Multiple question mark wildcard characters can be specified for the comparison data value. An asterisk is a multiple-character wildcard character. A single asterisk wildcard character can be specified at the end of the comparison data value. For example, 'ABC*' will match any value that begins with the letters 'ABC'.</p> <p>Contains the null value if no PAL comparison data was specified.</p>
PAL_COMPARE_AGAINST	VARCHAR(9)	<p>The part of the PAL the data specified in PAL_COMPARISON_DATA is to be compared against.</p> <p>*RSCMODEL The PAL comparison data will be compared against the resource model.</p> <p>*RSCNAME The PAL comparison data will be compared against the resource name.</p> <p>*RSCTYPE The PAL comparison data will be compared against the resource type.</p> <p>Contains the null value if no PAL comparison data was specified.</p>

Example

- List all the watches that my user profile started with the STRWCH command including the detailed information.

```
SELECT *
FROM QSYS2.WATCH_INFO W, TABLE(QSYS2.WATCH_DETAIL(W.SESSION_ID)) D
WHERE W.WATCH_SESSION_TYPE = '*STRWCH' AND
      W.USER_ID = USER;
```

WATCH_INFO view

The WATCH_INFO view returns information about watch sessions on the system.

The values returned for the columns in the view are closely related to the values returned by the WRKWCH (Work with Watches) CL command and by the Retrieve Watch List (QSCRWCHL) and Retrieve Watch Information (QSCRWCHI) API.

Authorization: The caller must have:

- *USE authority on the QSYS/QSCRWCHL and QSYS/QSCRWCHI programs and
- *SERVICE special authority or
 - authorized to the QIBM_SERVICE_WATCH and QIBM_SERVICE_TRACE function usage identifiers.

The following table describes the columns in the view. The system name is WATCH_INFO. The schema is QSYS2.

Table 103. WATCH_INFO view

Column Name	System Column Name	Data Type	Description
SESSION_ID	SESSION_ID	VARCHAR(10)	The session identifier for the watch.
ORIGIN	ORIGIN	VARCHAR(9)	The name of the command or the API that started the watch. QSCSWCH Session started by Start Watch (QSCSWCH) API. STRCMNTRC Session started by Start Communications Trace (STRCMNTRC) command. STRTRC Session started by Start Trace (STRTRC) command. STRWCH Session started by Start Watch (STRWCH) command. TRCCNN Session started by Trace Connection (TRCCNN) command. TRCINT Session started by Trace Internal (TRCINT) command. TRTCPAPP Session started by Trace TCP/IP Application (TRTCPAPP) command.
ORIGIN_JOB	ORIGIN_JOB	VARCHAR(28)	The qualified job name that started the watch session.
START_TIMESTAMP	WCH_START	TIMESTAMP	The timestamp of when the watch was started.
USER_ID	USER_ID	VARCHAR(10)	The name of the user that started the watch session.
WATCH_SESSION_TYPE	WCH_TYPE	VARCHAR(7)	Identifies the type of watch according to its origin. *SRVMON Watch session started using the Service Monitor function of the operating system. *STRWCH Watch session started using the Start Watch (STRWCH) command or Start Watch (QSCSWCH) API. *TRCCMD Watch session started using Start Communications Trace (STRCMNTRC), Start Trace (STRTRC), Trace Internal (TRCINT), Trace Connection (TRCCNN), or Trace TCP/IP Application (TRTCPAPP) commands.
STATUS	STATUS	VARCHAR(6)	The status of the watch session. ACTIVE Session is in ACTIVE status ENDING Session is in ENDING status
JOB_RUN_PRIORITY	PRIORITY	INTEGER	The priority for the job where the watch session work will be run.
WATCHED_MESSAGE_COUNT	MSG_COUNT	INTEGER	The number of messages being watched.
WATCHED_LIC_LOG_COUNT	LIC_COUNT	INTEGER	The number of LIC logs being watched.
WATCHED_PAL_COUNT	PAL_COUNT	INTEGER	The number of PALs being watched.
WATCH_PROGRAM_LIBRARY	WCH_PGMLIB	VARCHAR(10) Nullable	The name of the library where the watch program is located. Contains the null value if WATCH_PROGRAM is null.
WATCH_PROGRAM	WCH_PGM	VARCHAR(10) Nullable	The name of the program called to notify that a specified watch event occurred. Contains the null value if a watch program is not specified or not found.

Table 103. WATCH_INFO view (continued)

Column Name	System Column Name	Data Type	Description
WATCH_PROGRAM_CALL_START	CALL_START	VARCHAR(3) Nullable	<p>Whether the watch program will be called before starting watching for any event.</p> <p>NO The watch program will not be called before starting watching for any event.</p> <p>YES The watch program will be called before starting watching for any event.</p> <p>Contains the null value if WATCH_PROGRAM is null.</p>
WATCH_PROGRAM_CALL_END	CALL_END	VARCHAR(3) Nullable	<p>Whether the watch program will be called when the watch session is ending.</p> <p>NO The watch program will not be called when the watch session is ending.</p> <p>YES The watch program will be called when the watch session is ending.</p> <p>Contains the null value if WATCH_PROGRAM is null.</p>
TIME_LIMIT	TIME_LIMIT	INTEGER Nullable	<p>The time limit, in minutes, for watching for a message, a LIC log entry, or a PAL entry. A non-null value is only returned for watch sessions started by a trace command.</p> <p>When the specified amount of time has elapsed, the program identified by the WATCH_PROGRAM column is called, the watch is ended, and message CPI3999 is sent to the history log.</p> <p>Contains the null value if the watch session was started by a trace command and did not specify a time limit or for sessions not started by a trace command.</p>
TIME_INTERVAL	INTERVAL	INTEGER Nullable	<p>The interval of time, in seconds, of how often the trace exit program is called.</p> <p>Contains the null value if the watch session was started by a trace command and did not specify an interval or for sessions not started by a trace command.</p>

Example

- List all the service monitor watches that are currently active.

```
SELECT *
FROM QSYS2.WATCH_INFO
WHERE WATCH_SESSION_TYPE = '*SRVMON' AND
STATUS = 'ACTIVE';
```

Backup, Recovery, and Media Services (BRMS) Services

These table functions and views provide information about BRMS.

[BRMS SQL Services](#)

Communication Services

These services provide communication information.

ACTIVE_DB_CONNECTIONS table function

The ACTIVE_DB_CONNECTIONS table function returns a result table that contains one row for each DRDA and DDM connection. For these connections, the specified job has either an application requester or application server role.

If a connection is no longer active on the remote system, it might still appear in the result of this table function. For this situation, the ERROR column will have a value of YES.

Authorization: To invoke this function, the caller must be the same user profile that is running the specified *job-name*, or the caller must have *JOBCTL user special authority, or QIBM_DB_SQLADM or QIBM_DB_SYSMON function usage authority.

►► ACTIVE_DB_CONNECTIONS ((JOB_NAME ⇒) job-name) ►►

The schema is QSYS2.

job-name An expression that contains the qualified job name for which connection information is to be returned. The special value of '*' indicates the current job.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 104. ACTIVE_DB_CONNECTIONS table function

Column Name	Data Type	Description
CONNECTION_USAGE	VARCHAR(21)	The connection usage of the local server (<i>job-name</i>) side of the connection. APPLICATION REQUESTER The remote job is the application requester, often called the client. APPLICATION SERVER The remote job is the application server, often called the server.
REMOTE_HOST_NAME	VARCHAR(255)	The name of the remote system for this connection.
REMOTE_JOB_NAME	VARCHAR(28)	The qualified job name of the remote job that is connected to <i>job-name</i> . Contains the null value if the remote server is not an IBM i.
REMOTE_USER	VARCHAR(255)	User name for the remote connection.
CLIENT_RDB	VARCHAR(18)	The RDB name of the application requester. Can contain the null value if the connection is only used for a DDM file operation.
SERVER_RDB	VARCHAR(18)	The RDB name of the application server. Can contain the null value if the connection is only used for a DDM file operation.
CONNECTION_TIME	TIMESTAMP	Timestamp for when the connection was established.
CONNECTION_TYPE	VARCHAR(11)	The type of connection established between <i>job-name</i> and REMOTE_JOB_NAME. OPTICONNECT This connection is an OptiConnect connection. SNA This connection is an SNA connection. TCP/IP This connection is a TCP/IP connection.
CLIENT_PORT	INTEGER	The client host port number.
CLIENT_ADDRESS	VARCHAR(45)	IP address of the client. This is shown in IPv6 address format.
SERVER_PORT	INTEGER	The server host port number.
SERVER_ADDRESS	VARCHAR(45)	IP address of the server. This is shown in IPv6 address format.

Table 104. ACTIVE_DB_CONNECTIONS table function (continued)

Column Name	Data Type	Description
SNA_DEVICE_NAME	VARCHAR(10)	The name of the device used to communicate with the remote location. Contains the null value if CONNECTION_TYPE is not SNA.
SNA_MODE_NAME	VARCHAR(8)	The name of the mode to be used to communicate between the local location and remote location. Contains the null value if CONNECTION_TYPE is not SNA.
SNA_LOCAL_LOCATION_NAME	VARCHAR(8)	Indicates the local location name by which this system is identified to the system on which the RDB is located. Contains the null value if CONNECTION_TYPE is not SNA.
SNA_REMOTE_LOCATION_NAME	VARCHAR(8)	The name of the remote location where the remote file exists. Contains the null value if CONNECTION_TYPE is not SNA.
SNA_REMOTE_NETWORK_ID	VARCHAR(8)	Indicates the remote network identifier of the system on which the RDB is located. Contains the null value if CONNECTION_TYPE is not SNA.
SNA_TRANSACTION_PROGRAM_NAME	VARCHAR(19)	Indicates the transaction program name (TPN) specified in the RDB directory entry. The value is returned as a hex string in the form 'X'aabbcc'. Contains the null value if CONNECTION_TYPE is not SNA.
SSL	VARCHAR(3)	Indicates whether this connection uses SSL. NO The connection does not use SSL. YES The connection uses SSL.
THREE_PART_NAMING	VARCHAR(3)	Indicates whether this connection was established using an SQL statement with 3-part naming. NO This connection was not established using an SQL statement with 3-part naming. YES This connection was established using an SQL statement with 3-part naming. Contains the null value if CONNECTION_USAGE is APPLICATION SERVER.
XA	VARCHAR(3)	Indicates whether this connection is used for XA. NO This is not an XA connection. YES This is an XA connection. Contains the null value if CONNECTION_USAGE is APPLICATION SERVER.
ACTIVATION_GROUP_NUMBER	BIGINT	Activation group number for job making the connection. Contains the null value if CONNECTION_USAGE is APPLICATION SERVER.
ACTIVATION_GROUP_NAME	VARCHAR(10)	Activation group name for job making the connection. Contains the null value if the activation group has no name or if CONNECTION_USAGE is APPLICATION SERVER.
THREAD_ID	BIGINT	The thread ID for the connection. Contains the null value if CONNECTION_USAGE is APPLICATION SERVER.
SCOPE	VARCHAR(16)	Scope of the connection. ACTIVATION GROUP The connection is scoped to the activation group. JOB The connection is scoped to the job. Contains the null value if CONNECTION_USAGE is APPLICATION SERVER.

Table 104. ACTIVE_DB_CONNECTIONS table function (continued)

Column Name	Data Type	Description
COMMIT_PROTOCOL	VARCHAR(9)	The commitment control protocol used for this connection. ONE-PHASE The connection uses one-phase commitment control. TWO-PHASE The connection uses two-phase commitment control. Contains the null value if CONNECTION_USAGE is APPLICATION SERVER.
REMOTE_CLASS	VARCHAR(255)	The class of the remote server. Some common values are: QAS Db2 for i QDB2 Db2 for z/OS QDB2/xxx Db2 for LUW
REMOTE_RELEASE	VARCHAR(255)	The release level of the remote server.
ERROR	VARCHAR(3)	Indicates whether the connection is currently in an error state. For example, a communication failure occurred. NO The connection is not in an error state. YES The connection is in an error state.

Example

- List all the connections that are active for a specific job.


```
SELECT * FROM TABLE (QSYS2.ACTIVE_DB_CONNECTIONS('900239/SKALSKY/QPADEV000F'));
```

DNS_LOOKUP table function

The DNS_LOOKUP table function queries a domain name server for IP address information about a hostname.

Authorization: If the HOSTALIASES environment variable is set, the caller must have:

- Execute (*X) data authority to each directory in the path of the host aliases file specified by the HOSTALIASES environment variable, and
- Read (*R) data authority to the host aliases file.



The schema is QSYS2.

search-name A character string containing the fully qualified domain name of the host to be resolved.

domain-server A character string containing either a host name or an IP address for the domain name server, or an IP address string for it.

If this parameter is omitted, the IPv4 or IPv6 configured name servers from CHGTCPDMN are used.

The result of the function is a table containing rows with the format shown in the following table. All columns are nullable.

Table 105. DNS_LOOKUP table function

Column Name	Data Type	Description
ADDRESS_SPACE_TYPE	CHAR(4)	The IP address space type for IP_ADDRESS IPV4 The address is specified using the IPv4 address space type. IPV6 The address is specified using the IPv6 address space type.
IP_ADDRESS	VARCHAR(45)	The IP address. When ADDRESS_SPACE_TYPE is IPV4, the address is in IPv4 format. When ADDRESS_SPACE_TYPE is IPV6, the address is in IPv6 format.
AUTHORITATIVE	VARCHAR(3)	Indicates whether the DNS response containing the IP address was an authoritative answer. NO The DNS response was a non-authoritative answer. YES The DNS response was an authoritative answer.
AUTHENTICATED_DATA	VARCHAR(3)	Whether the data was authenticated. NO Data was not DNSSEC authenticated YES Data was DNSSEC authenticated.

Example

- Determine the IP address for the ibm.com hostname.

```
SELECT * FROM TABLE(QSYS2.DNS_LOOKUP('ibm.com'));
```

ENV_SYS_INFO view

The ENV_SYS_INFO view contains information about the current server.

The following table describes the columns in the view. The system name is ENVSYSINFO. The schema is SYSIBMADM.

Table 106. ENV_SYS_INFO view

Column Name	System Column Name	Data Type	Description
OS_NAME	OS_NAME	VARCHAR(256) Nullable	Operating system name.
OS_VERSION	OS_VERSION	VARCHAR(256) Nullable	Operating system version.
OS_RELEASE	OS_RELEASE	VARCHAR(256) Nullable	Operating system release.
HOST_NAME	HOST_NAME	VARCHAR(256) Nullable	Name of the system.
TOTAL_CPUS	TOTAL_CPUS	INTEGER Nullable	The maximum number of virtual processors defined within the LPAR configuration.
CONFIGURED_CPUS	CONFIGCPUS	INTEGER Nullable	The number of virtual processors currently available to the partition.
CONFIGURED_MEMORY	CONFIGMEM	BIGINT Nullable	Total amount of configured memory on the system, in megabytes.
TOTAL_MEMORY	TOTAL_MEM	INTEGER Nullable	Total amount of memory on the system, in megabytes.

Example

Return information about the current server.

```
SELECT * FROM SYSIBMADM.ENV_SYS_INFO
```

HTTP_SERVER_INFO view

The HTTP_SERVER_INFO view returns information for HTTP Server for IBM i (powered by Apache). Only information for enabled and active functions is returned.

The values returned for the result columns of the table function are similar to the values shown by Real Time Server Statistics in the IBM Web Administration for i GUI.

Authorization: None required.

The following table describes the columns in the view. The system name is HTTP_SRVR. The schema is QSYS2.

Table 107. HTTP_SERVER_INFO view

Column Name	System Column Name	Data Type	Description
SERVER_NAME	SERVER	VARCHAR(10)	The name of the HTTP server instance.
JOB_NAME	JOB_NAME	VARCHAR(28)	Qualified job name for the server.
SERVER_START_TIME	START_T	TIMESTAMP(0)	Time the server was started.
SERVER_RESTART_TIME	RESTART_T	TIMESTAMP(0) Nullable	Time the server was restarted. Contains the null value if the server has never been restarted.
SERVER_CURRENT_TIME	CURRENT_T	TIMESTAMP(0)	Current time.
SERVER_NORMAL_CONNECTIONS	NORM_CONN	BIGINT	Total number of normal (non-secure) connections to the server.
SERVER_SSL_CONNECTIONS	SSL_CONN	BIGINT	Total number of SSL (secure) connections to the server.
SERVER_ACTIVE_THREADS	ACT_THRD	INTEGER	Number of active threads.
SERVER_IDLE_THREADS	IDLE_THRD	INTEGER	Number of idle threads.
SERVER_TOTAL_REQUESTS	TOT_REQ	BIGINT	Total number of requests to the server since the server was started.
SERVER_TOTAL_REQUESTS_REJECTED	TOT_REQREJ	BIGINT	Total number of rejected requests issued by the server since the server was started.
SERVER_TOTAL_RESPONSES	TOT_RESP	BIGINT	Total number of responses from the server since the server was started.

Table 107. HTTP_SERVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
HTTP_FUNCTION	HTTP_FUNC	VARCHAR(20)	HTTP server function or associated server for the remaining statistical information columns. A row is only returned for a function if the function is enabled. CGI Common Gateway Interface (CGI) CUSTOMER MODULE A customer or third-party module FRCA PROXY Fast Response Cache Accelerator (FRCA) proxy FRCA STATS Fast Response Cache Accelerator (FRCA) PROXY Proxy SERVER HANDLED Server transactions. This includes static HTML pages, HTML pages containing Server Side Includes (SSI), and images. SSL Secure Sockets Layer (SSL) WEBSPHERE WebSphere Application Server
REQUESTS	REQUESTS	BIGINT	Number of requests received.
RESPONSES	RESPONSES	BIGINT	Number of responses sent.
ERROR_RESPONSES	ERR_RESP	BIGINT	Number of error responses.
NONCACHE_RESPONSES	NC_RESP	BIGINT	Number of responses that did not come from cache.
BYTES_RECEIVED	BYTES_RCV	BIGINT	Total number of bytes received for all requests.
BYTES_SENT	BYTES_SENT	BIGINT	Total number of bytes sent for all requests.
NONCACHE_PROCESSING_TIME	NC_TIME	DECIMAL(20,3)	Processing time for non-cached requests, in seconds.
CACHE_PROCESSING_TIME	C_TIME	DECIMAL(20,3)	Processing time for cached requests, in seconds.

Example

- Examine information about the server functions associated with the ADMIN server.

```
SELECT * FROM QSYS2.HTTP_SERVER_INFO
WHERE SERVER_NAME = 'ADMIN';
```

NETSTAT_INFO view

The NETSTAT_INFO view returns information about IPv4 and IPv6 network connections.

The values returned for the columns in the view are closely related to the values returned by [List Network Connections API](#) and [Retrieve Network Connection Data API](#). Refer to the APIs for more detailed information.

The following table describes the columns in the view. The system name is NS_INFO. The schema is QSYS2.

Table 108. NETSTAT_INFO view

Column Name	System Column Name	Data Type	Description
CONNECTION_TYPE	CONN_TYPE	CHAR(4)	The type of connection. IPV4 The connection is an IPv4 connection. IPV6 The connection is an IPv6 connection.

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
REMOTE_ADDRESS	RMT_ADDR	VARCHAR(45)	The internet address of the remote host. For IPv4: <ul style="list-style-type: none"> The address is in IPv4 address format. A value of 0.0.0.0 indicates that either the system is waiting for a connection to open or that a UDP socket is being used. A value of 0 means that the connection is a listening or UDP socket so this field does not apply. For IPv6: <ul style="list-style-type: none"> The address is in IPv6 address format. A value of :: means that the connection is a listening socket so this field does not apply.
REMOTE_PORT	RMT_PORT	INTEGER	The remote host port number. A value of 0 means that the connection is a listening or UDP socket, so this field does not apply.
REMOTE_PORT_NAME	RMT_NAME	VARGRAPHIC(14) CCSID 1200 Nullable	The remote host well-known port name or the name from the service table entry. Contains null if there is no well-known port name.
LOCAL_ADDRESS	LOCAL_ADDR	VARCHAR(45)	The local address of this connection on this system. For IPv4: <ul style="list-style-type: none"> The address is in IPv4 address format. A value of 0.0.0.0 indicates that either the system is waiting for a connection to open or that a UDP socket is being used. For IPv6: <ul style="list-style-type: none"> The address is in IPv6 address format. A value of :: means the local application specified that any local internet address can be used.
LOCAL_PORT	LOCAL_PORT	INTEGER	The local system port number.
LOCAL_PORT_NAME	LOCAL_NAME	VARGRAPHIC(14) CCSID 1200 Nullable	The local system well-known port name or the name from the service table entry. Contains null if there is no well-known port name.
PROTOCOL	PROTOCOL	VARCHAR(3)	Identifies the type of connection protocol. TCP A Transmission Control Protocol (TCP) connection or socket. UDP A User Datagram Protocol (UDP) socket.

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
TCP_STATE	STATE	VARCHAR(12) Nullable	The state of the connection.
			CLOSED This connection has ended.
			CLOSE-WAIT Waiting for an end connection request from the local user.
			CLOSING Waiting for an end connection request acknowledgment from the remote host.
			ESTABLISHED The normal state in which data is transferred.
			FIN-WAIT-1 Waiting for the remote host to acknowledge the local system request to end the connection.
			FIN-WAIT-2 Waiting for the remote host request to end the connection.
			LAST-ACK Waiting for the remote host to acknowledge an end connection request.
			LISTEN Waiting for a connection request from any remote host.
			SYN-RECEIVED Waiting for a confirming connection request acknowledgment.
			SYN-SENT Waiting for a matching connection request after having sent a connection request.
			TIME-WAIT Waiting to allow the remote host enough time to receive the local system's acknowledgment to end the connection.
			Contains null if PROTOCOL is UDP.
IDLE_TIME	IDLE_TIME	DECIMAL(19,3)	The length of time, in seconds, since the last activity on this connection.
BIND_USER	BIND_USER	VARCHAR(10)	The user profile of the job on the local system which first performed a sockets API bind() of the socket.
BYTES_SENT_REMOTELY	BYTES_OUT	BIGINT	The number of bytes sent to the remote host.
BYTES_RECEIVED_LOCALLY	BYTES_IN	BIGINT	The number of bytes received from the remote host.
NETWORK_CONNECTION_TYPE	NET_TYPE	VARCHAR(4)	The type of connection or socket. *TCP Identifies a transmission control protocol (TCP) connection socket. *UDP Identifies a User Datagram Protocol (UDP) socket. For IPv4, the following additional value can be returned. *IPS Identifies an Internet Protocol (IP) over SNA connection or socket.
CONNECTION_OPEN_TYPE	OPN_TYPE	VARCHAR(7) Nullable	The type of open for the connection.
			ACTIVE The local system opens the connection. PASSIVE A remote host opens the connection.
			Contains null if PROTOCOL is UDP.
NUMBER_OF_ASSOCIATED_JOBS	NUM_JOBS	INTEGER	The number of jobs associated with this connection.
LINE_DESCRIPTION	LINE_DES	VARCHAR(10) Nullable	The local system line description associated with this connection. Contains null if this is an IPv4 connection or if the connection is not bound to a link local unicast interface.
VIRTUAL_LAN_ID	LAN_ID	VARCHAR(4) Nullable	The virtual LAN identifier associated with this connection. Can also contain the following special value:
			NONE No virtual LAN identifier is associated with this connection.
			Contains null if this is an IPv4 connection or if the connection is not bound to a link local unicast interface.

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
CONNECTION_TRANSPORT_LAYER	CNNTRANSPT	VARCHAR(5)	The transport that a connection is using. Values are: <ul style="list-style-type: none"> • IPS • TCPIP
IP_OPTIONS	IP_OPTIONS	BINARY(40) Nullable	The hex value of IP datagram options that may have been specified for a connection. Contains null if this is an IPv6 connection or if no IP datagram options have been specified.
ROUND_TRIP_TIME	ROUND_TRIP	BIGINT Nullable	The smoothed round-trip time interval in milliseconds. This is a measure of the time required for a segment on the connection to arrive at its destination, to be processed, and to return an acknowledgment to the client. Contains null if PROTOCOL is UDP.
ROUND_TRIP_VARIANCE	ROUND_VAR	BIGINT Nullable	The variance in milliseconds from the previous round-trip time. Contains null if PROTOCOL is UDP.
CURRENT_RETRANSMISSIONS	CT_RETRANS	BIGINT Nullable	The number of times the local system retransmitted the current segment without receiving an acknowledgment. Contains null if PROTOCOL is UDP.
TOTAL_RETRANSMISSIONS	TL_RETRANS	BIGINT Nullable	The total number of times the local system retransmitted a segment because an acknowledgement was not received. This is a cumulative count of all segments resent during the entire time the connection has been active. Contains null if PROTOCOL is UDP.
TCP_CONNECTIONS_CURRENTLY_ESTABLISHED	TCPCONN	BIGINT Nullable	The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT. Contains null if PROTOCOL is UDP.
TCP_ACTIVE_OPENS	TCPACTOPN	BIGINT Nullable	The number of times TCP connections have made a direct transition to the SYN-SENT state from the CLOSED state. This number is an indication of the number of times this local system opened a connection to a remote system. Contains null if PROTOCOL is UDP.
TCP_PASSIVE_OPENS	TCPPSVOPN	BIGINT Nullable	The number of times TCP connections have made a direct transition to the SYN-RECEIVED state from the LISTEN state. This number is an indication of the number of times a remote system opened a connection to this system. Contains null if PROTOCOL is UDP.
TCP_FAILED_OPENS	TCPFAILOPN	BIGINT Nullable	The total number of times TCP connections have made direct transitions to a CLOSED state from either the SYN-SENT state or the SYN-RECEIVED state and/or to LISTEN from SYN-RECEIVED. Contains null if PROTOCOL is UDP.
TCP_ESTABLISHED_AND_THEN_RESET	TCPESTRST	BIGINT Nullable	The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state. Contains null if PROTOCOL is UDP.
TCP_SEGMENTS_SENT	TCPSEGSNT	BIGINT Nullable	The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets. Contains null if PROTOCOL is UDP.
TCP_SEGMENTS_RETRANSMITTED	TCPSEGRTRN	BIGINT Nullable	The number of TCP segments transmitted containing one or more previously transmitted octets. Contains null if PROTOCOL is UDP.
TCP_SEGMENTS_RESET	TCPSEGRST	BIGINT Nullable	The number of TCP segments sent containing the RST flag. Contains null if PROTOCOL is UDP.

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
TCP_SEGMENTS_RECEIVED	TCPSEGRCV	BIGINT Nullable	The total number of segments received, including those received in error. This count includes segments received on currently established connections. Contains null if PROTOCOL is UDP.
TCP_SEGMENTS_RECEIVED_ERROR	TCPSEGRCVE	BIGINT Nullable	The total number of segments received in error (for example, bad TCP checksums). Contains null if PROTOCOL is UDP.
OUTGOING_BYTES_BUFFERED	BYTES_OUTB	BIGINT Nullable	The current number of bytes that an application has requested to send, but TCP has not yet sent. If TCP has sent the bytes to the remote system but has not yet received an acknowledgment, the bytes are considered 'not sent'. They are included in this count. Contains null if PROTOCOL is UDP.
USER_SEND_NEXT	USRSNDNXT	BIGINT Nullable	The sequence number of the next byte of data to be sent by the client application. Contains null if PROTOCOL is UDP.
SEND_NEXT	SEND_NEXT	BIGINT Nullable	The sequence number of the next byte of data that the local TCP application sends to the remote TCP application. Contains null if PROTOCOL is UDP.
SEND_UNACKNOWLEDGED	SNDUNACK	BIGINT Nullable	The sequence number of the last segment sent that was not acknowledged. This is the smallest sequence number of the send window. Contains null if PROTOCOL is UDP.
OUTGOING_PUSH_NUMBER	OUTPSHNBR	BIGINT Nullable	The sequence number of the last byte of push data in the outgoing stream. This value is zero if no push data is in the outgoing data stream. Contains null if PROTOCOL is UDP.
OUTGOING_URGENCY_NUMBER	OUTURGNBR	BIGINT Nullable	The sequence number of the last byte of urgent data in the outgoing data stream. This value is zero if no urgent data is in the outgoing data stream. Contains null if PROTOCOL is UDP.
OUTGOING_WINDOW_NUMBER	OUTWINNBR	BIGINT Nullable	The largest sequence number in the send window of the connection. The local TCP application cannot send data bytes with sequence numbers greater than the outgoing window number. Contains null if PROTOCOL is UDP.
INCOMING_BYTES_BUFFERED	BYTES_INB	BIGINT Nullable	The current number of bytes that are received and buffered by TCP. These bytes are available to be read by an application. Contains null if PROTOCOL is UDP.
RECEIVE_NEXT	RCVNEXT	BIGINT Nullable	The next sequence number the local TCP is expecting to receive. Contains null if PROTOCOL is UDP.
USER_RECEIVE_NEXT	USRRCVNXT	BIGINT Nullable	The sequence number of the next byte to be passed to the application by TCP. Contains null if PROTOCOL is UDP.
INCOMING_PUSH_NUMBER	INPSHNBR	BIGINT Nullable	The sequence number of the last byte of pushed data in the incoming data stream. This value is zero if no push data is in the incoming data stream. Contains null if PROTOCOL is UDP.
INCOMING_URGENCY_NUMBER	INURGNBR	BIGINT Nullable	The sequence number of the last byte of urgent data in the incoming data stream. This value is zero if no urgent data is in the incoming data stream. Contains null if PROTOCOL is UDP.
INCOMING_WINDOW_NUMBER	INWINNBR	BIGINT Nullable	The largest sequence number in the incoming window of this connection. Data bytes in the incoming stream having sequence numbers larger than this number are not accepted. Contains null if PROTOCOL is UDP.

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MAXIMUM_WINDOW_SIZE	MAXWINSIZ	BIGINT Nullable	The largest size of the send window, in bytes, during the entire time the connection has been active. Contains null if PROTOCOL is UDP.
CURRENT_WINDOW_SIZE	CURWINSIZ	BIGINT Nullable	The current send window size in bytes. Contains null if PROTOCOL is UDP.
LAST_UPDATE	LSTUPD	BIGINT Nullable	The sequence number of the incoming segment used for the last window update that occurred on the connection. Contains null if PROTOCOL is UDP.
LAST_UPDATE_ACKNOWLEDGED	LSTUPDACK	BIGINT Nullable	The acknowledgment number of the incoming segment used for the last window update that occurred on the connection. Contains null if PROTOCOL is UDP.
CONGESTION_WINDOW	CONGESTWIN	BIGINT Nullable	The number of segments that are sent on the next transmission. If an acknowledgment is received, the number is increased. If an acknowledgment is not received, the number is reset to the smallest allowable number. Contains null if PROTOCOL is UDP.
SLOW_START_THRESHOLD	SLWSTRTHR	BIGINT Nullable	The value of the slow-start threshold. Contains null if PROTOCOL is UDP.
MAXIMUM_SEGMENT_SIZE	MAXSEGSIZ	BIGINT Nullable	The size in bytes of the largest segment that may be transmitted on this connection. Contains null if PROTOCOL is UDP.
INITIAL_SEND_SEQUENCE_NUMBER	SNDSEQNBR	BIGINT Nullable	The first sequence number sent on this connection. Contains null if PROTOCOL is UDP.
INITIAL_RECEIVE_SEQUENCE_NUMBER	RCVSEQNBR	BIGINT Nullable	The first sequence number received on this connection. Contains null if PROTOCOL is UDP.
UDP_DATAGRAMS_SENT	UDPSENT	BIGINT Nullable	The total number of UDP datagrams sent from all connections since TCP/IP was started. Contains null if PROTOCOL is TCP.
UDP_DATAGRAMS_RECEIVED	UDPRCV	BIGINT Nullable	The total number of UDP datagrams received, including those received in error. This count includes datagrams received on currently established connections. Contains null if PROTOCOL is TCP.
UDP_DATAGRAMS_NOT_DELIVERED_PORT_NOT_FOUND	UDPNDPNF	BIGINT Nullable	The total number of received UDP datagrams for UDP users for which there was no application at the destination port. Contains null if PROTOCOL is TCP.
UDP_DATAGRAMS_NOT_DELIVERED_OTHER	UDPNDOTHER	BIGINT Nullable	The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port. Contains null if PROTOCOL is TCP.
SOCKET_STATE	SOCSTATE	VARCHAR(13)	The current state of the socket. Values are: <ul style="list-style-type: none"> • BOUND • CONNECTED • CONNECTING • DISCONNECTED • ERROR • LISTENING • UNBOUND • UNINITIALIZED

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SOCKET_BROADCAST	SOCBROAD	VARCHAR(3) Nullable	Indicates if messages can be sent to the broadcast address. NO Messages cannot be sent to the broadcast address. YES Messages can be sent to the broadcast address. Contains null if value is not specified or if socket is not an address family of AF_INET and type SOCK_DGRAM or SOCK_RAW.
SOCKET_BYPASS_ROUTE	SOCBYPASS	VARCHAR(3) Nullable	Indicates if the normal routing mechanism is being bypassed. NO The normal routing mechanism is being used. YES The normal routing mechanism is being bypassed. Contains null if value is not specified or if socket is not an address family of AF_INET or AF_INET6.
SOCKET_DEBUG	SOCDEBUG	VARCHAR(3) Nullable	Indicates if low-level debugging is active. NO Low-level debugging is not active. YES Low-level debugging is active. Contains null if value is not specified.
SOCKET_ERROR	SOCERROR	INTEGER Nullable	Indicates if there any pending errors in the socket. A value of zero indicates no pending errors. Otherwise, the value indicates the error number. Contains null if value is not specified.
SOCKET_KEEP_ALIVE	SOCALIVE	VARCHAR(3) Nullable	Indicates if the connection is being kept up by periodic transmissions. NO The connection is not being kept up by periodic transmissions. YES The connection is being kept up by periodic transmissions. Contains null if value is not specified or if socket is not an address family of AF_INET or AF_INET6 and type SOCK_STREAM.
SOCKET_LINGER	SOCLINGER	VARCHAR(3) Nullable	Indicates whether the system attempts to deliver any buffered data or if the system discards it when a close() is issued. NO The system attempts to send buffered data with an infinite wait time. YES The system attempts to send buffered data for SOCKET_LINGER_TIME seconds. If the data is not deliverable within that period of time, it is discarded. Contains null if value is not specified.
SOCKET_LINGER_TIME	SOCLTIME	BIGINT Nullable	The time, in seconds, the system will wait to send buffered data. Contains null if value is not specified or if SOCKET_LINGER is NO.
SOCKET_OUT_OF_BAND_DATA	SOCOUTBAND	VARCHAR(3) Nullable	Indicates if out-of-band data is received inline with normal data. NO Out-of-band data is not received inline with normal data. YES Out-of-band data is received inline with normal data. Contains null if value is not specified or if socket is not an address family of AF_INET or AF_INET6.
SOCKET_RECEIVE_BUFFER_SIZE	SOCRCVBUF	BIGINT Nullable	The size of the receive buffer. Contains null if value is not specified.
SOCKET_RECEIVE_LOW_WATER_MARK_SIZE	SOCRCVSZ	BIGINT Nullable	The size of the receive low-water mark. The default size is 1. Contains null if value is not specified or if socket is not type SOCK_STREAM.

Table 108. NETSTAT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SOCKET_REUSE_ADDRESS	SOCKREUSE	VARCHAR(3)	Indicates if the local socket address can be reused.
		Nullable	<p>NO The local socket address cannot be reused.</p> <p>YES The local socket address can be reused.</p> <p>Contains null if value is not specified or if socket is not an address family of AF_INET or AF_INET6 and type SOCK_STREAM or SOCK_DGRAM.</p>
SOCKET_SEND_BUFFER_SIZE	SOCKSENDUF	BIGINT	The size of the send buffer.
		Nullable	Contains null if value is not specified.
SOCKET_TYPE	SOCTYPE	VARCHAR(14)	The socket type. Values are:
		Nullable	<p>SOCK_DGRAM Datagram type.</p> <p>SOCK_RAW Raw type.</p> <p>SOCK_SEQPACKET Sequential packet type.</p> <p>SOCK_STREAM Stream type.</p> <p>Contains null if value is not specified.</p>
SOCKET_LOOPBACK	SOCKLOOPBK	VARCHAR(3)	Indicates if the loopback feature is being used.
		Nullable	<p>NO The loopback feature is not being used.</p> <p>YES The loopback feature is being used.</p> <p>Contains null if value is not specified.</p>
SOCKET_RECEIVE_TIMEOUT	SOCKRCVTO	BIGINT	The receive timeout value.
		Nullable	Contains null if value is not specified.
SOCKET_SEND_LOW_WATER_MARK_SIZE	SOCKSENDSZ	BIGINT	The size of the send low-water mark.
		Nullable	Contains null if value is not specified.
SOCKET_SEND_TIMEOUT	SOCKSENDTO	BIGINT	The send timeout value.
		Nullable	Contains null if value is not specified.

Example

Return information about all network connections for user QLWISVR.

```
SELECT * FROM QSYS2.NETSTAT_INFO
WHERE BIND_USER = 'QLWISVR'
```

Related information

[Internet Protocol version 6](#)

NETSTAT_INTERFACE_INFO view

The NETSTAT_INTERFACE_INFO view returns information about IPv4 and IPv6 interfaces.

The values returned for the columns in the view are closely related to the values returned by [List Network Interfaces API](#). Refer to the API for more detailed information.

The following table describes the columns in the view. The system name is NS_INTER. The schema is QSYS2.

Table 109. NETSTAT_INTERFACE_INFO view

Column Name	System Column Name	Data Type	Description
CONNECTION_TYPE	CONN_TYPE	CHAR(4)	The type of connection. IPV4 The connection is an IPv4 connection. IPV6 The connection is an IPv6 connection.
INTERNET_ADDRESS	INT_ADDR	VARCHAR(45)	The internet address of the interface. For IPv4: <ul style="list-style-type: none"> The address is in IPv4 address format. Can contain the special value: *IP4DHCP The interface has been configured to use DHCP to obtain an IPv4 address. For IPv6: <ul style="list-style-type: none"> The address is in IPv6 address format. Can contain the special value: *IP6SAC This interface will use Stateless Address Auto-configuration to obtain an IPv6 address.
NETWORK_ADDRESS	NET_ADDR	VARCHAR(45)	The internet address of the IP network or subnetwork to which the interface is attached. For IPv4: <ul style="list-style-type: none"> The address is in IPv4 address format. For IPv6: <ul style="list-style-type: none"> The address is in IPv6 address format.
SUBNET_MASK	SUBNET_MSK	VARCHAR(15) Nullable	The subnet mask for the network, subnet, and host address fields of the internet address that defines the subnetwork for an interface. Contains null if this is an IPv6 connection.
PREFIX_LENGTH	PRE_LEN	INTEGER Nullable	The prefix length defines how many bits of the IPv6 internet address are in the prefix. It specifies how many of the left-most bits of the address make up the prefix. The prefix length is used to generate network and host addresses. Contains null if this is an IPv4 connection.
LINE_DESCRIPTION	LINE_DES	VARCHAR(10)	The name of the communications line description that identifies the physical network associated with an interface. Can contain the following special values: *LOOPBACK This is the loopback interface. Processing associated with a loopback interface does not extend to a physical line. *VIRTUALIP The virtual interface is a circuitless interface. For IPv4, the following additional values can be returned. *IPS The interface is used by Internet Protocol (IP) over SNA. *OPC The interface is attached to the optical bus (OptiConnect).

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
INTERFACE_LINE_TYPE	LINE_TYPE	VARCHAR(6)	<p>The type of line used by the interface.</p> <p>ASYNC Asynchronous communications protocol.</p> <p>DDI Distributed Data Interface protocol.</p> <p>ELAN Ethernet local area network protocol.</p> <p>FR Frame relay network protocol.</p> <p>L2TP Layer Two Tunneling protocol.</p> <p>PPP Point-to-Point protocol.</p> <p>PPPOE Point-to-Point over Ethernet protocol.</p> <p>TDLC Twinaxial Datalink Control. Used for TCP/IP over Twinax.</p> <p>TRLAN Token-ring local area network protocol.</p> <p>VETH Virtual Ethernet protocol.</p> <p>WLS Wireless local area network protocol.</p> <p>X25 X.25 protocol.</p> <p>Can also contain one of the following special values:</p> <p>ERROR A system error other than those for NOTFND was received while trying to determine the link type for an interface.</p> <p>NONE Line is not defined. This value is used for the following interfaces: *LOOPBACK, *VIRTUALIP, *OPC. There is no line type value for this interface.</p> <p>NOTFND Not found. The line description object for this interface cannot be found.</p> <p>OTHER An Internet Protocol (IP) over SNA interface.</p>

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
INTERFACE_STATUS	STATUS	VARCHAR(12)	<p>The current status of the logical interface.</p> <p>ACTIVE The interface has been started and is running.</p> <p>DOD This interface is being used for Point-to-Point (PPP) Dial-on-Demand.</p> <p>ENDING The operating system is processing the request to end this interface.</p> <p>FAILED The line description associated with this interface has entered the failed state.</p> <p>FAILED (TCP) An error was detected in the IBM TCP/IP Licensed Internal Code.</p> <p>INACTIVE The interface has not been started.</p> <p>RCYCNL A hardware failure has occurred and the line description associated with this interface is in the recovery canceled (RCYCNL) state.</p> <p>RCYPND An error with the physical line associated with this interface was detected by the system. The line description associated with this interface is in the recovery pending (RCYPND) state.</p> <p>STARTING The operating system is processing the request to start this interface.</p> <p>For IPv4, the following additional value can be returned.</p> <p>ACQUIRING The operating system is attempting to obtain an IP address from a Dynamic Host Configuration Protocol (DHCP) server.</p> <p>For IPv6, the following additional value can be returned.</p> <p>DEPRECATED The interface address preferred lifetime has expired but the valid lifetime has not yet expired.</p>
INTERFACE_SOURCE	SOURCE	VARCHAR(9) Nullable	<p>Specifies how this interface was added to the protocol stack.</p> <p>LOOPBACK The interface was added by the protocol stack as the loopback address.</p> <p>STATELESS The interface was added by IPv6 address autoconfiguration.</p> <p>STATEFUL The interface was added by Dynamic Host Configuration Protocol version 6 (DHCPv6) configuration.</p> <p>MANUAL The interface was added by manual configuration.</p> <p>Contains null if this is an IPv4 connection.</p>

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SERVICE_TYPE	SRVC_TYPE	VARCHAR(9) Nullable	<p>The type of service that defines how the internet hosts and routers should make trade-offs between throughput, delay, reliability, and cost.</p> <p>MAXRLB A higher level of effort to ensure delivery is important for datagrams with the maximize reliability indication.</p> <p>MAXTHRPUT High data rate is important for datagrams with the maximize throughput indication.</p> <p>MINCOST Lower cost is important for datagrams with the minimize monetary cost indication.</p> <p>MINDELAY Prompt delivery is important for datagrams with the minimize delay indication.</p> <p>NORMAL Normal service is used for delivery of datagrams.</p> <p>OTHER An Internet Protocol (IP) over SNA interface.</p> <p>Contains null if this is an IPv6 connection.</p>
VIRTUAL_LAN_ID	LAN_ID	VARCHAR(4)	<p>The virtual LAN to which this interface belongs according to IEEE standard 802.1Q. Can also contain the following special value:</p> <p>NONE The interface does not belong to a virtual LAN.</p>
MAXIMUM_TRANSMISSION_UNIT	MTU	VARCHAR(10)	<p>The maximum transmission unit (MTU) value specified for this interface. Either an integer value or this special value:</p> <p>LIND The interface is not currently active and the MTU was specified as *LIND.</p> <p>For IPv4, the following additional value can be returned.</p> <p>OTHER An Internet Protocol (IP) over SNA interface.</p>
CONFIGURED_MAXIMUM_TRANSMISSION_UNIT	CFG_MTU	VARCHAR(10)	<p>The configured maximum transmission unit (MTU) value specified for this interface. Either an integer value or this special value:</p> <p>LIND The MTU was configured as *LIND, the maximum frame size found in the line description object associated with the interface.</p>
AUTOSTART	AUTOSTART	VARCHAR(3)	<p>Specifies whether the interface is automatically started when the protocol stack is activated.</p> <p>NO This interface is not automatically started.</p> <p>YES This interface is automatically started.</p>
DAD_MAX_TRANSMITS	DAD_MAX	BIGINT Nullable	<p>The maximum number of duplicate address detection (DAD) transmissions the protocol stack will send out on this interface.</p> <p>Contains null if this is an IPv4 connection.</p>

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
HOST_ADDRESS	HOST_ADDR	VARCHAR(45)	Host portion of the internet address. For IPv4: <ul style="list-style-type: none"> Host portion of the Internet address, in dotted decimal notation, as determined by the subnet mask specified for this interface. For IPv6: <ul style="list-style-type: none"> Host portion of the Internet address, in IPv6 address format, as determined by the prefix length configured for this interface.
DIRECTED_BROADCAST_ADDRESS	DIRBRDADR	VARCHAR(15) Nullable	The Internet address, in dotted decimal notation, used to broadcast to all systems attached to the same network or subnetwork as this interface. Contains null if this is an IPv6 connection or if interface is attached to a network that does not support a broadcast operation.
ASSOCIATED_LOCAL_INTERFACE	ASCLCLINT	VARCHAR(15) Nullable	The Internet address, in dotted decimal notation, of the local interface that has been associated with this interface. Contains null if this is an IPv6 connection or if no association has been made between this interface and another local interface.
CHANGE_STATUS	CHGSTS	VARCHAR(6) Nullable	The status of the most recent change to this interface in the dynamic tables used by the TCP/IP protocol stack. ADD Add interface request processed. CHANGE Change interface request processed. END End interface request processed. START Start interface request processed. Contains null if this is an IPv6 connection.
PACKET_RULES	PKT_RULES	VARCHAR(17) Nullable	The kind of packet rules data available for this line. NONE No filters and no NAT are loaded for this line. NAT NAT is enabled for this line. FILTERS Filters are defined for this line. NAT_FILTERS NAT enabled and filters defined for this line. FILTERS_IPSEC Filters and IPsec filters are defined for this line. NAT_FILTERS_IPSEC NAT enabled and Filters and IPsec filters defined for this line. Contains null if packet rules data is unknown.
INTERFACE_TYPE	TYPE	VARCHAR(12) Nullable	The interface type: BROADCAST Broadcast capable. NONBROADCAST Non-broadcast capable. UNNUMBERED Unnumbered network. Contains null if this is an IPv6 connection.
NETWORK_FULL_NAME	NET_FNAME	VARCHAR(24) Nullable	The complete name of the network that this interface is a part of. Contains null if this is an IPv6 connection or if there is no network name.

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
INTERFACE_FULL_NAME	FNAME	VARCHAR(24) Nullable	The complete interface name. Contains null if this is an IPv6 connection or if there is no interface name.
ALIAS_NAME	ALIAS_NAME	VARGRAPHIC(50) CCSID 1200 Nullable	Name given to the interface to use as an alternate to the IP address. Contains null if there is no alias name.
INTERFACE_TEXT	LABEL	VARGRAPHIC(50) CCSID 1200 Nullable	Description of the interface. Contains null if there is no interface description.
DHCP_CREATED	DHPCRT	VARCHAR(3) Nullable	Specifies whether this interface was created using Dynamic Host Configuration Protocol (DHCP). NO This interface was not created using DHCP. YES This interface was created using DHCP. Contains null if this is an IPv6 connection.
DHCP_DYNAMIC_DNS_UPDATES	DHCPDYDNS	VARCHAR(3) Nullable	Specifies whether dynamic updates to Domain Name System (DNS) tables are enabled or not. NO DNS updates are disabled. YES DNS updates are enabled. Contains null if this is an IPv6 connection or if interface was not created by DHCP.
DHCP_LEASE_EXPIRATION	DHCPLXP	TIMESTAMP(0) Nullable	The timestamp when the DHCP lease will expire. Contains null if this is an IPv6 connection or if interface was not created by DHCP.
DHCP_LEASE_OBTAINED	DHCPLOBT	TIMESTAMP(0) Nullable	The timestamp when the DHCP lease was obtained or renewed. Contains null if interface was not created by DHCP.
DHCP_USE_UNIQUE_ID	DHCPUSEUID	VARCHAR(3) Nullable	Whether the DHCP unique identifier (DUID) is used as the client identification for the Dynamic Host Configuration Protocol (DHCP). NO The hardware (MAC) address is used for the client ID. YES The DHCP unique identifier is used for the client ID or this is an IPv6 connection. Contains null if interface was not created by DHCP.
DHCP_SERVER_UNIQUE_ID	DHCPDRVUID	VARCHAR(30) Nullable	Specifies the DHCP unique identifier (DUID) of the DHCP server from which the IP address was obtained. Contains null if this is an IPv4 connection or if interface was not created by DHCP.
DHCP_SERVER_ADDRESS	DHCPDRVADD	VARCHAR(15) Nullable	The Internet address, in dotted decimal notation, of the DHCP server from which the DHCP lease was obtained or renewed. Contains null if this is an IPv6 connection or if interface was not created by DHCP.
PREFERRED_INTERFACE_DEFAULT_ROUTE	PREFDFTRTE	VARCHAR(3) Nullable	This field describes whether the preferred proxy interfaces are based on the system's default route. NO The default route is not used to determine the preferred interface. YES The default route is used to determine the preferred interface. Contains null if this is an IPv6 connection.

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PREFERRED_INTERFACE_LIST	PREFIFCLST	VARCHAR(159) Nullable	A list of up to 10 preferred interface internet addresses. Each internet address within the preferred interface list is given in dotted decimal notation. When there is more than one internet address, a single blank separates the addresses. Contains null if this is an IPv6 connection or if a preferred interface list is not being used.
PREFERRED_PHYSICAL_LINE_LIST	PREFLINLST	VARCHAR(159) Nullable	A list of up to 10 preferred physical line list entries. Each entry in the list is formatted as LINE_DESCRIPTION:VIRTUAL_LAN_ID. The line description can be up to 10 characters long. The virtual LAN ID can be up to 4 characters long. When there is more than one preferred physical line, a single blank separates entries. Contains null if this is an IPv4 connection.
ADDRESS_TYPE	ADDR_TYPE	VARCHAR(9) Nullable	The type of IPv6 address that is assigned to this network interface. ANYCAST An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance). MULTICAST An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address. UNICAST An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address. Contains null if this is an IPv4 connection.
ADDRESS_CLASS	ADDR_CLASS	VARCHAR(9) Nullable	The class of IPv6 address that is assigned to this network interface. PUBLIC The interface is a public one. TEMPORARY The interface is a temporary one used for privacy extensions. Contains null if this is an IPv4 connection.
ADDRESS_PREFERRED_LIFETIME	ADDPLIFE	BIGINT Nullable	The length of time that a "valid" address is preferred, in seconds. A negative value indicates that the address preferred lifetime expired that number of seconds ago. Contains null if this is an IPv4 connection or if this address has an infinite preferred lifetime.
ADDRESS_VALID_LIFETIME	ADDVLIFE	BIGINT Nullable	The length of time, in seconds, that an address remains in a "valid" state. A negative value indicates that the address valid lifetime expired that number of seconds ago. Contains null if this is an IPv4 connection or if this address has an infinite valid lifetime.
ADDRESS_PREFERRED_LIFETIME_EXPIRATION	ADDPFRLE	TIMESTAMP(0) Nullable	The timestamp when this address will no longer be in the preferred state. If the timestamp is in the future, the address is still preferred. If the timestamp is in the past, then this address is no longer preferred. Contains null if this is an IPv4 connection or if this address has an infinite preferred lifetime which never expires.

Table 109. NETSTAT_INTERFACE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
ADDRESS_VALID_LIFETIME_EXPIRATION	ADDVLDLE	TIMESTAMP(0) Nullable	The timestamp when this address will expire or did expire. If the timestamp is in the future, the address has not expired yet. If the timestamp is in the past, then this address has expired and is still being returned for a short period of time to indicate that the interface ceased to function because its valid lifetime expired. Contains null if this is an IPv4 connection or if this address has an infinite valid lifetime which never expires.
ON_LINK	ON_LINK	VARCHAR(3) Nullable	Whether this interface and all IPv6 addresses with the same prefix are on the same link. NO Addresses with the same prefix are not assumed to be on the same link. YES Addresses with the same prefix are assumed to be on the same link and directly reachable. Contains null if this is an IPv4 connection or if INTERNET_ADDRESS is *IP6SAC
PROXY_ARP_ENABLED	PRXARPENB	VARCHAR(3) Nullable	Indicates whether Proxy ARP is currently active for this interface. NO Proxy ARP not enabled. YES Proxy ARP enabled. Contains null if this is an IPv6 connection.
PROXY_ARP_ALLOWED	PRXARPALW	VARCHAR(3) Nullable	Indicates whether Proxy ARP has been configured to be allowed or not allowed. NO Proxy ARP not allowed. YES Proxy ARP allowed. Contains null if this is an IPv6 connection or if interface is not Opticonnect (*OPC) or Virtual Ethernet.
CURRENT_PROXY_AGENT_LINE	PRXAGT	VARCHAR(10) Nullable	Name of the communication line description that is used with the IPv6 interface for virtual IP address (VIPA) proxy Neighbor Discovery. Contains null if this is an IPv4 connection or if no association has been made between this interface and another physical interface.
CURRENT_PROXY_AGENT_LINE_VIRTUAL_LAN_ID	PRXAGTLAN	VARCHAR(4) Nullable	The virtual LAN to which the proxy agent line belongs according to IEEE standard 802.1Q. Can also contain the following special value: NONE There is no virtual LAN identifier associated with the current proxy agent line. Contains null if this is an IPv4 connection or if no association has been made between this interface and another physical interface.
LAST_CHANGE_TIMESTAMP	LASTCHG	TIMESTAMP(0) Nullable	The timestamp of the most recent change to this interface in the dynamic tables used by the protocol stack. Contains null if the interface has never been changed.

Example

Return information about all interfaces which are using Virtual Ethernet protocol..

```
SELECT * FROM QSYS2.NETSTAT_INTERFACE_INFO
WHERE INTERFACE_LINE_TYPE = 'VETH'
```

Related information

[Internet Protocol version 6](#)

NETSTAT_JOB_INFO view

The NETSTAT_JOB_INFO view returns information about jobs using IPv4 and IPv6 network connections.

The values returned for the columns in the view are closely related to the values returned by the Retrieve Network Connection Data (QtocRtvNetCnnDta) API.

Authorization: None required.

The following table describes the columns in the view. The system name is NS_JOB. The schema is QSYS2.

Table 110. NETSTAT_JOB_INFO view

Column Name	System Column Name	Data Type	Description
CONNECTION_TYPE	CONN_TYPE	CHAR(4)	The type of connection. IPV4 The connection is an IPv4 connection. IPV6 The connection is an IPv6 connection.
REMOTE_ADDRESS	RMT_ADDR	VARCHAR(45)	The internet address of the remote host. For IPv4: <ul style="list-style-type: none">The address is in IPv4 address format. A value of 0.0.0.0 indicates that either the system is waiting for a connection to open or that a UDP socket is being used. A value of 0 means that the connection is a listening or UDP socket so this field does not apply. For IPv6: <ul style="list-style-type: none">The address is in IPv6 address format. A value of :: means that the connection is a listening socket so this field does not apply.
REMOTE_PORT	RMT_PORT	INTEGER	The remote host port number. A value of 0 means that the connection is a listening or UDP socket, so this field does not apply.
REMOTE_PORT_NAME	RMT_NAME	VARGRAPHIC(14) CCSID 1200 Nullable	The remote host well-known port name or the name from the service table entry. Contains null if there is no well-known port name.
LOCAL_ADDRESS	LOCAL_ADDR	VARCHAR(45)	The local address of this connection on this system. For IPv4: <ul style="list-style-type: none">The address is in IPv4 address format. A value of 0.0.0.0 indicates that either the system is waiting for a connection to open or that a UDP socket is being used. For IPv6: <ul style="list-style-type: none">The address is in IPv6 address format. A value of :: means the local application specified that any local internet address can be used.
LOCAL_PORT	LOCAL_PORT	INTEGER	The local system port number.
LOCAL_PORT_NAME	LOCAL_NAME	VARGRAPHIC(14) CCSID 1200 Nullable	The local system well-known port name or the name from the service table entry. Contains null if there is no well-known port name.
AUTHORIZATION_NAME	USER_NAME	VARCHAR(10) Nullable	The effective user profile of the thread for which information is being retrieved. This name may differ from the user portion of the job name. Contains null when SLIC_TASK_NAME is not null or if JOB_NAME is the special value *SIGNON.

Table 110. NETSTAT_JOB_INFO view (continued)

Column Name	System Column Name	Data Type	Description
JOB_NAME	JOB_NAME	VARCHAR(28) Nullable	The qualified job name. Can also contain the following special value: *SIGNON This connection is a telnet connection and the system is performing sign-on processing or is displaying a sign-on prompt on it. Contains null when SLIC_TASK_NAME is not null.
SLIC_TASK_NAME	SLIC_TASK	VARCHAR(16) Nullable	The task name as identified to the system. Contains null when JOB_NAME is not null.
INTERNAL_JOB_ID	JOB_ID	BINARY(16) Nullable	A value that can be used by system APIs to speed the process of locating the job on the system. Contains null if JOB_NAME is the special value *SIGNON or if SLIC_TASK_NAME is not null.
JOB_TYPE	JOB_TYPE	VARCHAR(11) Nullable	The type of job: AUTOSTART The job is an autostart job. BATCH The job is a batch job. INTERACTIVE The job is an interactive job. MONITOR The job is a subsystem monitor job. READER The job is a spooled reader job. SCPF The job is the SCPF system job. SYSTEM The job is a system job. WRITER The job is a spooled writer job. Contains null if JOB_NAME is the special value *SIGNON or if SLIC_TASK_NAME is not null.

Example

Return information about all jobs using IPv4 network connections.

```
SELECT * FROM QSYS2.NETSTAT_JOB_INFO
WHERE CONNECTION_TYPE = 'IPV4'
```

Related information

[Internet Protocol version 6](#)

NETSTAT_ROUTE_INFO view

The NETSTAT_ROUTE_INFO view returns information about IPv4 and IPv6 routes.

The values returned for the columns in the view are closely related to the values returned by [List Network Routes API](#). Refer to the API for more detailed information.

The following table describes the columns in the view. The system name is NS_ROUTE. The schema is QSYS2.

Table 111. NETSTAT_ROUTE_INFO view

Column Name	System Column Name	Data Type	Description
CONNECTION_TYPE	CONN_TYPE	CHAR(4)	The type of connection. IPV4 The connection is an IPv4 connection. IPV6 The connection is an IPv6 connection.

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
ROUTE_DESTINATION	ROUTE_DEST	VARCHAR(45)	<p>The Internet Protocol address of the ultimate destination reached by this route.</p> <p>For IPv4:</p> <ul style="list-style-type: none"> The address is in IPv4 address format. When used in combination with the subnet mask and the type of service values, the route destination identifies a route to a network or system. A value of 0.0.0.0 means that the route destination is the default route. <p>For IPv6:</p> <ul style="list-style-type: none"> The address is in IPv6 address format. When used in combination with the prefix length, the route destination identifies a route to a network or host.
SUBNET_MASK	SUBNET_MSK	VARCHAR(15) Nullable	<p>The actual value of the subnet mask for the route destination in dotted-decimal notation. A value of 0.0.0.0 means no value is defined.</p> <p>Contains null if this is an IPv6 connection.</p>
NEXT_HOP	NEXT_HOP	VARCHAR(45)	<p>The internet address of the first system on the path from your system to the route destination.</p> <p>For IPv4:</p> <ul style="list-style-type: none"> The address is in IPv4 address format. <p>For IPv6:</p> <ul style="list-style-type: none"> The address is in IPv6 address format. <p>Can contain the following special value:</p> <p>*DIRECT This is the next hop value of a route that is automatically created.</p>
PREFIX_LENGTH	PRE_LEN	INTEGER Nullable	<p>The prefix length defines how many bits of the route destination IPv6 address are in the prefix. It specifies how many of the left-most bits of the address make up the prefix. The prefix length is used to generate network and host addresses.</p> <p>Contains null if this is an IPv4 connection.</p>
ROUTE_STATUS	ROUTE_STS	VARCHAR(10) Nullable	<p>The current state of the route.</p> <p>DOD This route is used for Point-to-Point (PPP) Dial-on-Demand. Currently, this Dial-on-Demand route is not available. The route will become available when a Dial-on-Demand session is initiated for the interface this route is associated with.</p> <p>For IPv4:</p> <p>YES The router specified by the next hop value for this interface is available for use.</p> <p>NO The router specified by the next hop value for this interface is not available for use.</p> <p>NO GATEWAY The router specified by the next hop value for this interface is not available for use, the router may be experiencing a problem.</p> <p>For IPv6:</p> <p>ACTIVE This route is currently active and is in the current route search path.</p> <p>INACTIVE This route is not in the route search path and is not being used.</p> <p>Contains null if the state is unknown.</p>

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
ROUTE_MAXIMUM_TRANSMISSION_UNIT	ROUTE_MTU	VARCHAR(10)	<p>The maximum transmission unit (MTU) value for this route in bytes. Can be either a number or one of the following special values:</p> <p>For IPv4:</p> <p>IFC The route is not currently active and the MTU was specified as *IFC.</p> <p>OTHER An Internet Protocol (IP) over SNA interface.</p> <p>For IPv6:</p> <p>*IP6LINMTU This route uses the MTU of the line it is bound to.</p>
CONFIGURED_ROUTE_MAXIMUM_TRANSMISSION_UNIT	CFG_RT_MTU	VARCHAR(10) Nullable	<p>A number representing the configured maximum transmission unit (MTU) value for this route, in bytes. Can be either a number or the following special value:</p> <p>*IP6LINMTU The route MTU was specified as *IP6LINMTU, the MTU value of the line to which this route is bound.</p> <p>Contains null if this is an IPv4 connection.</p>
ROUTE_TYPE	ROUTE_TYPE	VARCHAR(8) Nullable	<p>The type of route.</p> <p>DFTRROUTE A default route.</p> <p>DIRECT A route to a network or subnetwork to which this system has a direct physical connection.</p> <p>HOST A route to a specific remote host.</p> <p>NET An indirect route to a remote network.</p> <p>SUBNET An indirect route to a remote subnetwork. This option is only for IPv4 connections.</p> <p>Contains null if the type of route is unknown.</p>

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
ROUTE_SOURCE	ROUTE_SRC	VARCHAR(18) Nullable	<p>Specifies how this route was added to the routing table.</p> <p>For IPv4:</p> <p>CFG The route was added using the configuration commands of the local system.</p> <p>ICMP The route was added with the Internet Control Message Protocol (ICMP) redirect mechanism.</p> <p>OTHER The route was added with a sockets input/output control (IOctl) or other mechanism.</p> <p>RIP The route was added by the Routing Information Protocol (RIP).</p> <p>SNMP The route was added by the Simple Network Management Protocol (SNMP).</p> <p>For IPv6:</p> <p>AUTOCONFIG This route was added because of an interface added by stateless autoconfiguration.</p> <p>BGP This route was added by the Border Gateway Protocol (BGP).</p> <p>CFGIFC The route was added because of a manually configured interface.</p> <p>CFG RTE The route was manually configured.</p> <p>IDRP This route was added by the Inter-Domain Routing Protocol (IDRP).</p> <p>IGRP This route was added by the Interior Gateway Routing Protocol (IGRP).</p> <p>OSPF The route was added by the Open Shortest Path First (OSPF) protocol.</p> <p>RA_PREFIX_INFO This route was added because of the presence of a Prefix Information Option on a Router Advertisement packet received by the system.</p> <p>RA_ROUTE_INFO This route was added because of the presence of a Route Information Option on a Router Advertisement packet received by the system.</p> <p>RA_ROUTER_LIFETIME This route was added because of the presence of a non-zero value in the Router Lifetime field in a Router Advertisement packet received by the system.</p> <p>REDIRECT This route was added by the ICMPv6 redirect mechanism.</p> <p>RIP The route was added by the Routing Information Protocol (RIP).</p> <p>ROUTING This route was determined to be necessary and added by the TCP/IP stack on this system.</p> <p>SNMP This route was added by the Simple Network Management Protocol (SNMP).</p> <p>Contains null if the route source is not known.</p>

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SERVICE_TYPE	SRVC_TYPE	VARCHAR(9) Nullable	<p>The type of service that defines how the internet hosts and routers should make trade-offs between throughput, delay, reliability, and cost.</p> <p>MAXRLB A higher level of effort to ensure delivery is important for datagrams with the maximize reliability indication.</p> <p>MAXTHRPUT High data rate is important for datagrams with the maximize throughput indication.</p> <p>MINCOST Lower cost is important for datagrams with the minimize monetary cost indication.</p> <p>MINDELAY Prompt delivery is important for datagrams with the minimize delay indication.</p> <p>NORMAL Normal service is used for delivery of datagrams.</p> <p>OTHER An Internet Protocol (IP) over SNA interface.</p> <p>Contains null if this is an IPv6 connection.</p>
ROUTE_PROTOCOL	ROUTE_PTCL	VARCHAR(7) Nullable	<p>Specifies the protocol that was used to generate this route.</p> <p>BGP Border Gateway protocol.</p> <p>IDRP InterDomain Routing protocol.</p> <p>IGRP InterGateway Routing protocol.</p> <p>LOCAL Local configuration.</p> <p>NDISC Neighbor discovery.</p> <p>NETMGMT Network Management protocol.</p> <p>OSPF Open Shortest Path First protocol.</p> <p>OTHER None of the listed protocols.</p> <p>RIP Routing Information protocol.</p> <p>Contains null if this is an IPv4 connection.</p>
ROUTE_PREFERENCE	ROUTE_PREF	VARCHAR(6) Nullable	<p>The preference of this route during route selection.</p> <p>LOW This route has a low preference.</p> <p>MEDIUM This route has a medium preference.</p> <p>HIGH This route has a high preference.</p> <p>Contains null if this is an IPv4 connection.</p>
LOCAL_BINDING_TYPE	LOCALTYPE	VARCHAR(7) Nullable	<p>The type of line to which this route is bound.</p> <ul style="list-style-type: none"> • DYNAMIC • STATIC <p>Contains null if this is an IPv6 connection.</p>
LOCAL_BINDING_INTERFACE	LOCALIFC	VARCHAR(15) Nullable	<p>The IP interface to bind to this route.</p> <p>Contains null if this is an IPv6 connection.</p>

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCAL_BINDING_INTERFACE_STATUS	LOCALSTS	VARCHAR(12) Nullable	The current status of the logical interface.
			ACTIVE The interface has been started and is running.
			DOD This interface is being used for Point-to-Point (PPP) Dial-on-Demand.
			DUPLICATE Another host on the LAN responded to a packet destined for this logical interface.
			ENDING The operating system is processing the request to end this interface.
			FAILED The line description associated with this interface has entered the failed state.
			FAILED (TCP) An error was detected in the IBM TCP/IP Licensed Internal Code.
			INACTIVE The interface has not been started.
			RCYCNL A hardware failure has occurred and the line description associated with this interface is in the recovery canceled (RCYCNL) state.
			RCYPND An error with the physical line associated with this interface was detected by the system. The line description associated with this interface is in the recovery pending (RCYPND) state.
STARTING The operating system is processing the request to start this interface.			
			Contains null if this is an IPv6 connection.
LOCAL_BINDING_NETWORK_ADDRESS	LOCALADDR	VARCHAR(15) Nullable	The Internet address, in dotted decimal notation, of the IP network or subnetwork that the interface is attached to. Contains null if this is an IPv6 connection.
LOCAL_BINDING_SUBNET_MASK	LOCALMASK	VARCHAR(15) Nullable	The subnet mask for the network, subnet, and host address fields for the local binding network address, in dotted decimal notation, that defines the subnetwork for an interface. Contains null if this is an IPv6 connection.
LOCAL_BINDING_LINE_DESCRIPTION	LOCALLINE	VARCHAR(10)	The name of the communications line description or virtual line (L2TP) that identifies the network associated with an interface. Can contain the following special values:
			*LOOPBACK This is a loopback interface. Processing associated with the loopback interface does not extend to a physical line.
			*OPC This interface is attached to the optical bus (OptiConnect).
			*VIRTUALIP The virtual interface is a circuitless interface.
LOCAL_BINDING_LINE_STATUS	LOCALSTS	VARCHAR(8) Nullable	The current operational status of the communications line to which this route is bound. ACTIVE The line is operational. FAILED The desired state of the line is Active, but it is currently in the Inactive state. INACTIVE The line is not operational. Contains null if this is an IPv4 connection.

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCAL_BINDING_LINE_TYPE	LOCALLTYPE	VARCHAR(6)	<p>The type of line used by the interface.</p> <p>ASYNC Asynchronous communications protocol.</p> <p>DDI Distributed Data Interface protocol.</p> <p>ELAN Ethernet local area network protocol.</p> <p>FR Frame relay network protocol.</p> <p>L2TP Layer Two Tunneling protocol.</p> <p>PPP Point-to-Point protocol.</p> <p>PPPOE Point-to-Point over Ethernet protocol.</p> <p>TDLC Twinaxial Datalink Control. Used for TCP/IP over Twinax.</p> <p>TRLAN Token-ring local area network protocol.</p> <p>VETH Virtual Ethernet protocol.</p> <p>WLS Wireless local area network protocol.</p> <p>X25 X.25 protocol.</p> <p>Can also contain one of the following special values:</p> <p>ERROR A system error other than those for NOTFND was received while trying to determine the link type for an interface.</p> <p>NONE Line is not defined. This value is used for the following interfaces: *LOOPBACK, *VIRTUALIP, *OPC. There is no line type value for this interface.</p> <p>NOTFND Not found. The line description object for this interface cannot be found.</p> <p>OTHER An Internet Protocol (IP) over SNA interface.</p>
LOCAL_BINDING_VIRTUAL_LAN_ID	LOCALLAN	VARCHAR(4)	<p>The virtual LAN to which this route is bound. Can also contain the following special value:</p> <p>NONE No virtual LAN identifier is associated with the binding line.</p>
ROUTE_PRECEDENCE	ROUTE_PRCO	INTEGER Nullable	<p>Priority of route. Values are 1 to 10, with the lowest priority being 1.</p> <p>Contains null if this is an IPv6 connection.</p>
ROUTE_TEXT	LABEL	VARGRAPHIC(50) CCSID(1200) Nullable	<p>Text description associated with the route.</p> <p>Contains null if there is no description.</p>
DUPLICATE	DUPLICATE	VARCHAR(6) Nullable	<p>Indicates whether this route is a duplicate of another route in the routing table or not, and also whether there are any routes which are duplicates of this route.</p> <p>NO This route is not a duplicate of another route but it does have duplicates.</p> <p>UNIQUE This route is not a duplicate of another route and it does not have any duplicates.</p> <p>YES This route is a duplicate of another route.</p> <p>Contains null if this is an IPv4 connection.</p>
EXPIRATION	EXPIRATION	TIMESTAMP(0) Nullable	<p>The timestamp when this route will expire or did expire. If the timestamp is in the future, the route has not expired yet. If the timestamp is in the past, then this route has expired and is still being returned for a short period of time to indicate that the route ceased to function because its lifetime expired.</p> <p>Contains null if this is an IPv4 connection or if the route will never expire.</p>

Table 111. NETSTAT_ROUTE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PPP_CONFIGURATION_PROFILE	PPPCFGPRF	VARCHAR(10) Nullable	The name of the Point-to-Point Protocol (PPP) configuration profile associated with this route. Contains null if this is an IPv4 connection or if Point-to-Point Protocol is not being used with this route.
PPP_AUTHENTICATION_USER_ID	PPPAUTUSR	VARCHAR(24) Nullable	The Point-to-Point Protocol authentication user id associated with this route. Contains null if this is an IPv4 connection or if Point-to-Point Protocol is not being used with this route.
PPP_INTERNET_ADDRESS	PPPINTADD	VARCHAR(45) Nullable	The internet address, in IPv6 address format, to which this Point-to-Point route is bound. Contains null if this is an IPv4 connection or if Point-to-Point Protocol is not being used with this route.
PPP_DIAL_ON_DEMAND_PROFILE	PPPDODPRF	VARCHAR(10) Nullable	The name of the Dial-on-demand Remote Peer Enabled Point-to-Point profile associated with this route. Contains null if this is an IPv4 connection or if Point-to-Point Protocol is not being used with this route.
LAST_CHANGE_TIMESTAMP	LASTCHG	TIMESTAMP(0) Nullable	The timestamp of the most recent change to this route in the dynamic tables used by the protocol stack. Contains null if the interface has never been changed.

Example

Return information about all routes which are available for use.

```
SELECT * FROM QSYS2.NETSTAT_ROUTE_INFO
WHERE ROUTE_STATUS = 'YES' OR ROUTE_STATUS = 'ACTIVE'
```

Related information

[Internet Protocol version 6](#)

SERVER_SBS_CONFIGURATION view

The SERVER_SBS_CONFIGURATION view returns subsystem routing information for some IBM i servers. When a client attempts to use TCP/IP to form a connection to a server listed in this view, an attempt is made to attach to a prestart job in the subsystem configured for that server.

The information returned by this view is similar to the information shown by IBM Navigator for i in Network > Servers. For information about users who have alternate subsystem configurations for some IBM i servers, see [“SERVER_SBS_ROUTING view” on page 507](#).

The QSYS2.SET_SERVER_SBS_ROUTING procedure can be used to modify entries shown in this view.

Authorization: None required.

The following table describes the columns in the view. The system name is SERVER_CFG. The schema is QSYS2.

Table 112. SERVER_SBS_CONFIGURATION view

Column Name	System Column Name	Data Type	Description
SERVER_NAME	SERVER	VARCHAR(10)	The server name for this entry.
SERVER_SEARCH_ORDER	SEARCH_ORD	INTEGER Nullable	The search order for selecting this subsystem routing entry. The order starts with one for each SERVER_NAME value. Contains the null value for the default entry for each server. This is the entry that will be selected if no other specific entry is selected.
SUBSYSTEM	SUBSYSTEM	VARCHAR(10)	The subsystem name that incoming connections for this server will be rerouted to when this entry is selected.

Table 112. SERVER_SBS_CONFIGURATION view (continued)

Column Name	System Column Name	Data Type	Description
ALLOW_ROLLOVER	ROLLOVER	VARCHAR(3)	Indicates how incoming connection requests are handled if the subsystem is not active. NO Incoming connections will be rejected. YES Incoming connections will be routed to a batch immediate job in the subsystem where the server daemon job is active.
IP_ADDRESS_TYPE	ADDR_TYPE	CHAR(4) Nullable	The type of IP address for IP_ADDRESS_START and IP_ADDRESS_END. IPV4 The addresses are IPv4 addresses. IPV6 The addresses are IPv6 addresses. Contains the null value if no IP addresses are used for defining this row.
IP_ADDRESS_START	IP_START	VARCHAR(45) Nullable	The IP address for alternate routing or the starting IP address for a range of IP addresses. Contains the null value if no IP addresses are used for defining this row.
IP_ADDRESS_END	IP_END	VARCHAR(45) Nullable	The ending IP address for a range of IP addresses. Contains the null value if this row is not for a range of IP addresses.
SUBNET_MASK	SUBNET	VARCHAR(15) Nullable	The actual value of the subnet mask in dotted-decimal notation. Contains the null value if IP_ADDRESS_START is null, there is no subnet mask, or this is an IPv6 address.
PREFIX_LENGTH	PREFIX	INTEGER Nullable	The prefix length defines how many of the left-most bits of the IPv6 address make up the prefix. Contains the null value if IP_ADDRESS_START is null or an IPv4 address.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	Descriptive text for this entry. Contains the null value if there is no descriptive text.

Example

List all the servers that have alternate subsystems defined.

```
SELECT * FROM QSYS2.SERVER_SBS_CONFIGURATION
ORDER BY SERVER_NAME, SERVER_SEARCH_ORDER;
```

SERVER_SBS_ROUTING view

The SERVER_SBS_ROUTING view returns information about the users who have alternate subsystem configurations for some IBM i servers. When a user profile listed in this view attempts to use TCP/IP to form a connection to the server, an attempt is made to use the alternate subsystem instead of the default subsystem for that server.

For subsystem routing information for some IBM i servers, see [“SERVER_SBS_CONFIGURATION view” on page 506](#).

The QSYS2.SET_SERVER_SBS_ROUTING procedure can be used to modify the values shown in this view.

Authorization: You must have *OBJOPR and *READ authority to a *USRPRF or it will not be returned.

The following table describes the columns in the view. The system name is SRVR_RTG. The schema is QSYS2.

Table 113. SERVER_SBS_ROUTING view

Column Name	System Column Name	Data Type	Description
AUTHORIZATION_NAME	USER_NAME	VARCHAR(128) Nullable	The user profile that has an alternate subsystem configuration. Contains the value *ALL if this entry is for all users of the subsystem.
QRWTSRVR_SUBSYSTEM	DRDADDMSBS	VARCHAR(10) Nullable	The subsystem name that incoming DRDA or DDM connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QZDASOINIT_SUBSYSTEM	ZDASBS	VARCHAR(10) Nullable	The subsystem name that incoming database server connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QZRCRVS_SUBSYSTEM	ZRCSBS	VARCHAR(10) Nullable	The subsystem name that incoming remote command server connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QZHQSSRV_SUBSYSTEM	ZHQSBS	VARCHAR(10) Nullable	The subsystem name that incoming data queue server connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QZSCSRVS_SUBSYSTEM	ZSCSBS	VARCHAR(10) Nullable	The subsystem name that incoming central server connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QNPSERVS_SUBSYSTEM	NPSSBS	VARCHAR(10) Nullable	The subsystem name that incoming network print server connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QPWFSEVS_SO_SUBSYSTEM	PWFSBS	VARCHAR(10) Nullable	The subsystem name that incoming file server connections will be rerouted to. Contains the null value when an alternate subsystem for this server has not been configured for this user.
QRWTSRVR_ROLLOVER	DRDA_RO	VARCHAR(3)	Indicates how incoming DRDA or DDM connection requests are handled if the specified subsystem cannot be used. NO Incoming connections will be rejected. YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.
QZDASOINIT_ROLLOVER	ZDA_RO	VARCHAR(3)	Indicates how incoming database server connection requests are handled if the specified subsystem cannot be used. NO Incoming connections will be rejected. YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.
QZRCRVS_ROLLOVER	ZRC_RO	VARCHAR(3)	Indicates how incoming remote command connection requests are handled if the specified subsystem cannot be used. NO Incoming connections will be rejected. YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.

Table 113. SERVER_SBS_ROUTING view (continued)

Column Name	System Column Name	Data Type	Description
QZHQSSRV_ROLLOVER	ZHQ_RO	VARCHAR(3)	<p>Indicates how incoming data queue server connection requests are handled if the specified subsystem cannot be used.</p> <p>NO Incoming connections will be rejected.</p> <p>YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.</p>
QZSCSRVS_ROLLOVER	ZSC_RO	VARCHAR(3)	<p>Indicates how incoming central server connection requests are handled if the specified subsystem cannot be used.</p> <p>NO Incoming connections will be rejected.</p> <p>YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.</p>
QNPSERVS_ROLLOVER	NPS_RO	VARCHAR(3)	<p>Indicates how incoming network print server connection requests are handled if the specified subsystem cannot be used.</p> <p>NO Incoming connections will be rejected.</p> <p>YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.</p>
QPWFSEVS_ROLLOVER	PWF_RO	VARCHAR(3)	<p>Indicates how incoming file server connection requests are handled if the specified subsystem cannot be used.</p> <p>NO Incoming connections will be rejected.</p> <p>YES Incoming connections will be routed to the default subsystem. If an alternate subsystem is not configured for this server, YES is the default.</p>

The following table shows the servers that can have alternate subsystem configurations.

Server Description	Server Name
Central server	QZSCSRVS
Database server	QZDASOINIT
Data queue server	QZHQSSRV
DDM	QRWTSRVR
DRDA	QRWTSRVR
File server	QPWFSEVS
Network print server	QNPSERVS
Remote command server	QZRCSRVS

Example

Query subsystem routing information for all user profiles:

```
SELECT * FROM QSYS2.SERVER_SBS_ROUTING
```

SET_SERVER_SBS_ROUTING procedure

The SET_SERVER_SBS_ROUTING procedure provides the ability to configure a server to use an alternate subsystem. This can be for all users of the server, for a specific user profile, or by IP address.

This procedure allows an administrator to reposition certain connections to a server into an alternate, non-default, subsystem. When configured, new incoming TCP/IP server connections will use the alternate subsystem.

Connections to be re-routed can be defined three ways.

1. For all users of the server.
2. For an IP address or a range of IP addresses.
3. For a specific user profile. The user profile can be a group profile or a supplemental group profile. Some servers do not support routing by specific user profile; refer to [Table 115 on page 510](#) for details.

Routing will possibly be attempted for one entry from each of the three re-routing methods. First, it will consider an entry for all users of the server and attempt to route according to that entry. Then it will immediately attempt to route by IP address, if there is an appropriate entry. Finally, it will use a configuration option for a specific user profile and attempt to route to that subsystem. For example, if a server is configured to route to a user-specified subsystem by incoming TCP/IP address as well as by user profile, the job will first attempt to route to the subsystem configured for that TCP/IP address and then immediately attempt to route to the subsystem configured for the connecting user profile. Understanding this order of processing is key to recognizing failures that may occur.

By default, all users for the following servers use the listed default subsystem:

Table 115. Servers and default subsystems

Server Description	Server Name	Default subsystem	Supports routing by user	Supports routing by IP address
Central server	QZSCSRVS	QUSRWRK	Yes	Yes
Database server	QZDASOINIT	QUSRWRK	Yes	Yes
Data queue server	QZHQSSRV	QUSRWRK	Yes	Yes
DDM	QRWTSRVR	QUSRWRK	Yes	Yes
DRDA	QRWTSRVR	QUSRWRK	Yes	Yes
File server	QPWFSEVS	QSERVER	Yes	Yes
IBM i NetServer	QZLSFILE	QSERVER	No	Yes
Network print server	QNPSEVS	QUSRWRK	Yes	Yes
Remote command server	QZRCSRVS	QUSRWRK	Yes	Yes
Sign-on server	QZSOSIGN	QUSRWRK	No	Yes

For more information on these servers see [DRDA and DDM overview](#) and [Host servers by function](#).

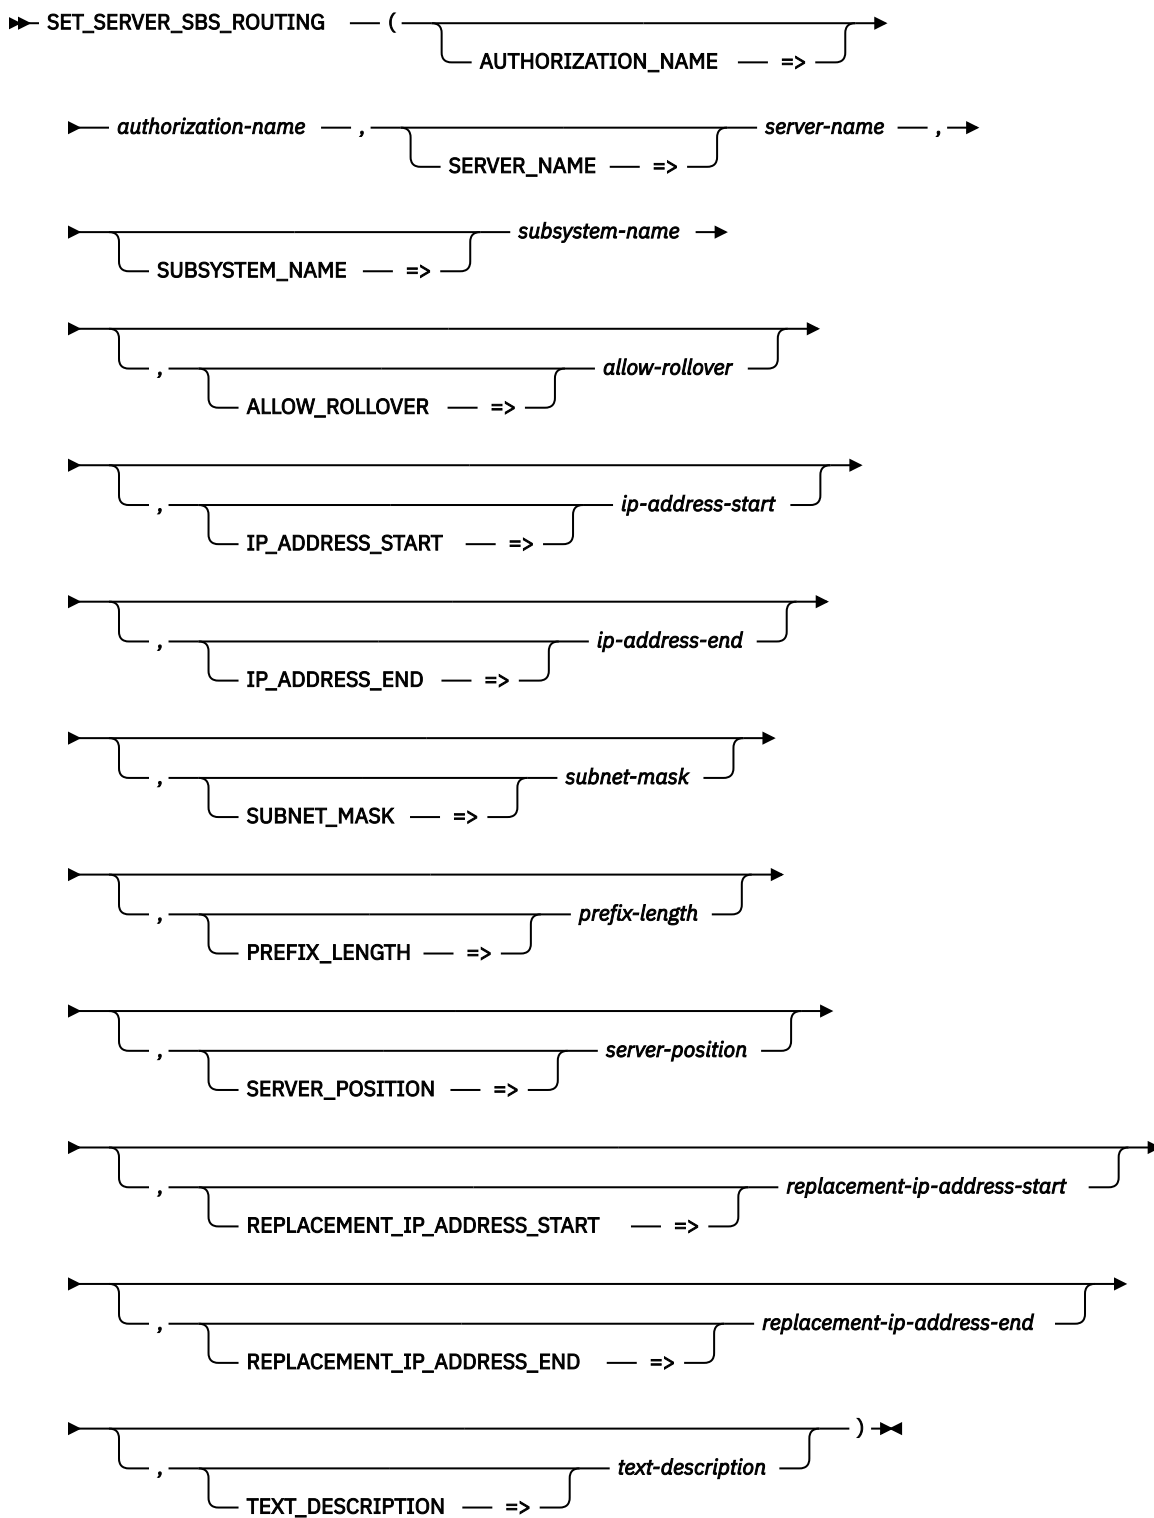
Routing information set by this procedure can be queried using the following views:

- For routing for a specific user, see [“SERVER_SBS_ROUTING view”](#) on page 507.
- For routing that applies to all users and IP addresses, see [“SERVER_SBS_CONFIGURATION view”](#) on page 506.

Authorization:

- When *authorization-name* specifies a user profile:

- The user calling this procedure must have *SECADM special authority.
- *OBJMGT and *USE authority is required to the user profile.
- When *authorization-name* is *ALL, *IOSYSCFG special authority is required



The schema is QSYS2.

authorization-name A character or graphic string expression that identifies an existing user or group profile name. Can contain the following special value:

***ALL** This configuration applies to all users of the server unless a specific alternate subsystem is defined for the user. To route using an IP address, *ALL must be specified.

server-name A character or graphic string expression that identifies the name of the server job that will be rerouted to *subsystem-name* for all users or for *authorization-name* whenever a connection is initiated to this server job. Valid *server-name* values are:

- QNPSERVS
- QPWFSERVSO
- QRWTSRVR
- QZDASOINIT
- QZHQSSRV
- QZLSFILE
- QZRCSRVS
- QZSCSRVS
- QZSOSIGN

The special value of *ALL can be used to indicate all of the *server-name* values that support routing by user, listed in [Table 115 on page 510](#). *ALL is not allowed if *authorization-name* is *ALL.

subsystem-name A character or graphic string expression that identifies the name of the subsystem that will be used, instead of the default subsystem, whenever a connection is initiated to the specified server job. No validation is done on *subsystem-name* to make sure it is a valid and active subsystem.

If *subsystem-name* is the null value:

- When *authorization-name* is not *ALL, the configuration entry for the *server-name* for this *authorization-name* will be cleared. The default subsystem will revert to the system supplied default as shown in [Table 115 on page 510](#).
- When *authorization-name* is *ALL:
 - If *ip-address-start* is not specified, the configuration entry for the *server-name* will be removed, reverting incoming connections to use the default subsystem.
 - If *ip-address-start* is specified, the configuration entry identified by the *server-name* and the specified *ip-address-start* will be removed.

allow-rollover A character or graphic string expression that indicates the action to take if the specified subsystem is not active. Valid values are:

NO If the alternate subsystem cannot be used, the connection request will fail.

YES If the alternate subsystem cannot be used, the connection request will succeed by using a batch immediate job in the default subsystem.

If this parameter is not specified, the default is YES.

When routing for all users of the server (*authorization-name* is *ALL), the following parameters are allowed:

ip-address-start A character or graphic string expression that identifies an IPv4 or IPv6 address for a single client.

- If *ip-address-end* is not specified, this is a specific IPv4 or IPv6 address.

- If *ip-address-end* is specified, this is the start value for a range of IPv4 or IPv6 addresses for a group of clients. When more than one range of addresses is defined for a server, the IP address ranges cannot overlap.

When *ip-address-start* is specified, either *subnet-mask* or *prefix-length* can be specified.

This parameter can only be specified when *authorization-name* is *ALL. If this parameter is not specified, the default is the null value, meaning that the connection will not be routed by IP address.

ip-address-end

A character or graphic string expression that identifies the end value for a range of IPv4 or IPv6 addresses for a group of clients. This parameter can only be specified if *ip-address-start* is specified.

This parameter can only be specified when *authorization-name* is *ALL. If this parameter is not specified, the default is the null value.

subnet-mask

A character or graphic string expression that identifies the IPv4 subnet mask.

This parameter can only be specified when *authorization-name* is *ALL. *subnet-mask* cannot be specified if *prefix-length* is specified.

prefix-length

An integer value that defines how many bits of the IPv6 address are in the prefix. When this parameter is specified with a non-zero value, the IP address parameters are treated as IPv6 addresses. It is required for an IPv6 address.

This parameter can only be specified when *authorization-name* is *ALL. *prefix-length* cannot be specified if *subnet-mask* is specified.

server-position

An integer value that indicates the position for this entry for *server-name* in the list returned by the QSYS2.SERVER_SBS_CONFIGURATION view in the SERVER_SEARCH_ORDER column. This value represents the search order to be used when determining which entry to use for routing.

This entry will be added at the specified position in the list for *server-name*. An existing entry at this position and all later entries will have their search order position incremented by one. If this value is greater than the current number of entries for the server, this entry will be added to the end of the server's list.

This parameter can only be specified when *authorization-name* is *ALL. If this parameter is not specified, the default is the null value. This means that a new entry will be added to the end of the list, and a change to an existing IP address entry will remain in its current position.

replacement-ip-address-start

A character or graphic string expression that identifies a replacement IPv4 or IPv6 address for the specified *ip-address-start* address. *ip-address-start* must already be configured for this server.

This parameter can only be specified when *authorization-name* is *ALL. If this parameter is not specified, the default is the null value.

replacement-ip-address-end

A character or graphic string expression that identifies a replacement IPv4 or IPv6 address for the specified *ip-address-end* address. *ip-address-end* must already be configured for this server.

This parameter can only be specified when *authorization-name* is *ALL. If this parameter is not specified, the default is the null value.

text-description

A character or graphic string expression describing this configuration.

This parameter can only be specified when *authorization-name* is *ALL. If this parameter is not specified, the default is the null value.

Notes

The prestart job entry must specify the *subsystem-name*.

When ALLOW_ROLLOVER is YES: if, for any reason, the alternate subsystem cannot be used to establish the connection, the connection will run in the default subsystem (or the last subsystem it was successfully routed to) as a batch immediate job. An example of this would be if the *authorization-name* does not have *USE authority to the subsystem description for *subsystem-name*.

If routing has been configured for a user profile, the user profile configuration will always be used, regardless of any group profile configuration or a subsystem configuration for *ALL. A group profile configuration will take precedence over any supplemental group profile configuration.

Examples

- Set new incoming DRDA and DDM TCP/IP server connections for user profile TIM to route to subsystem TIMSUBSYS.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING('TIM', 'QRWTSRVR', 'TIMSUBSYS')
```

- Reset incoming DRDA and DDM TCP/IP server connections for user profile TIM back to the original default subsystem.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING('TIM', 'QRWTSRVR', NULL)
```

- Configure group profile ADMIN to use an alternate subsystem for all of the servers supported by this procedure. Do not permit a connection request to rollover to use QUSRWRK.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING('ADMIN', '*ALL', 'ADHOC SBS', 'NO')
```

- Set new incoming Database server TCP/IP connections for user profile BOB to route to subsystem BOBSUBSYS.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING('BOB', 'QZDASOINIT', 'BOBSUBSYS')
```

- Construct a subsystem that will constrain the amount of system resources available to users who are known to execute expensive queries.

```
CRTSBS SBSD(QGPL/ADHOC SBS) POOLS((1 *BASE)) TEXT('Adhoc DRDA users SBS')
CRTCLS CLS(QGPL/ADHOCCLS) RUNPTY(55) TIMESLICE(100) TEXT('Adhoc DRDA users class')
ADDPJE SBSD(QGPL/ADHOC SBS) PGM(QSYS/QRWTSRVR) JOB(QGPL/QDFTSVR) CLS(QGPL/ADHOCCLS)
STRSBS SBSD(QGPL/ADHOC SBS)
CALL QSYS2.SET_SERVER_SBS_ROUTING('SLFUSER', 'QRWTSRVR', 'ADHOC SBS')
```

- Define a complete block of IP addresses to be routed to subsystem NEWSBS for QZDASOINIT jobs. The range includes all addresses in the address block of 192.168.1.0-255

```
CALL QSYS2.SET_SERVER_SBS_ROUTING(AUTHORIZATION_NAME => '*ALL',
                                  SERVER_NAME => 'QZDASOINIT',
                                  IP_ADDRESS_START => '192.168.1.0',
                                  SUBNET_MASK => '255.255.255.0',
                                  SUBSYSTEM_NAME => 'NEWSBS')
```

- Define a range of IP addresses for QZDASOINIT jobs to be routed to a subsystem NEWSBS.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING(AUTHORIZATION_NAME => '*ALL',
                                  SERVER_NAME => 'QZDASOINIT',
                                  IP_ADDRESS_START => '192.168.1.10',
                                  IP_ADDRESS_END => '192.168.1.30',
                                  SUBSYSTEM_NAME => 'NEWSBS')
```

- Change the range of IP addresses for QZDASOINIT jobs that are routed to a subsystem NEWSBS.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING(AUTHORIZATION_NAME => '*ALL',
                                  SERVER_NAME => 'QZDASOINIT',
                                  IP_ADDRESS_START => '192.168.1.10',
                                  IP_ADDRESS_END => '192.168.1.30',
                                  REPLACEMENT_IP_ADDRESS_START => '192.168.1.10',
```

```
REPLACEMENT_IP_ADDRESS_END => '192.168.1.100',
SUBSYSTEM_NAME => 'NEWSBS')
```

- Remove a range of IP addresses for QZRCRVS jobs that are routed to a subsystem NEWSBS.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING(AUTHORIZATION_NAME => '*ALL',
SERVER_NAME => 'QZRCRVS',
SUBSYSTEM_NAME => NULL,
IP_ADDRESS_START => '192.168.1.10',
IP_ADDRESS_END => '192.168.1.100')
```

- Add a new entry in the first position in the prioritized list of IP address routing entries for the QPWFSERVSO server.

```
CALL QSYS2.SET_SERVER_SBS_ROUTING(AUTHORIZATION_NAME => '*ALL',
SERVER_NAME => 'QPWFSEVSO',
SUBSYSTEM_NAME => 'SBS1',
IP_ADDRESS_START => '192.168.2.20',
SERVER_POSITION => 1,
TEXT_DESCRIPTION => 'Top rule for QPWFSERVSO')
```

Related information

[Use of prestart jobs](#)

TCPIP_INFO view

The TCPIP_INFO view contains TCP/IP information for the current host connection.

The following table describes the columns in the view. The schema is QSYS2.

Table 116. TCPIP_INFO view

Column Name	System Column Name	Data Type	Description
COLLECTED_TIME	COLLE00001	TIMESTAMP Nullable	Timestamp indicating when this row of information was collected.
LOCAL_HOST_NAME	LOCAL00001	VARCHAR(255) Nullable	TCP/IP host name of the local system.
CLIENT_IP_ADDRESS_TYPE	CLIEN00001	VARCHAR(10) Nullable	TCP/IP address version of the client.
CLIENT_IP_ADDRESS	CLIEN00002	VARCHAR(45) Nullable	TCP/IP address of the client.
CLIENT_PORT_NUMBER	CLIEN00003	INTEGER Nullable	TCP/IP port of the client.
SERVER_IP_ADDRESS_TYPE	SERVE00001	VARCHAR(10) Nullable	TCP/IP address version of the server.
SERVER_IP_ADDRESS	SERVE00002	VARCHAR(45) Nullable	TCP/IP address of the server.
SERVER_PORT_NUMBER	SERVE00003	INTEGER Nullable	TCP/IP port number of the server.
HOST_VERSION	HOST_00001	VARCHAR(10) Nullable	Operating system version.

Example

Return information about the current host connection.

```
SELECT * FROM QSYS2.TCPIP_INFO
```

IFS Services

These services provide information about the integrated file system.


IFS_JOB_INFO table function

The IFS_JOB_INFO table function returns a table that contains information about integrated file system references for a job.

This information is similar to what is returned by the Retrieve Referenced Objects (QPOLRRO) API.

The list of objects returned may be incomplete for objects residing in file systems other than the root (/), QOpenSys, and user-defined file systems. Objects in some of the other file systems can be locked with interfaces that do not use the integrated file system. Therefore, objects referenced by a job will only have references that were obtained as part of an integrated file system operation, or an operation that causes an integrated file system operation to occur.

Authorization: The user must be running with the same user profile as the job being retrieved or have *JOBCTL special authority.



job-name An expression that returns the qualified job name whose reference information is to be returned. Can contain the following special value:

- * Return information for the current job.

ignore-errors A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 117. IFS_JOB_INFO table function

Column Name	Data Type	Description
PATH_NAME	DBCLOB(16M) CCSID 1200	The path name of an integrated file system object that is referenced by the job. Contains the null value if the object is not linked to a path. This can occur in the root (/) file system, the QOpenSys file system, or a user-defined file system. Can also contain the null value if the object is in a remote file system or the optical file system (QOPT).

Table 117. IFS_JOB_INFO table function (continued)

Column Name	Data Type	Description
FILE_SYSTEM_TYPE	VARCHAR(15)	The file system for the object.
		NFS The Network File System (NFS).
		QDLS The Document Library Services (QDLS) file system.
		QFILSVR400 The QFileSvr.400 file system.
		QNTC The Windows NT® Server file system.
		QOPENSYS The QOpenSys file system.
		QOPT The optical file system (QOPT).
		QSYS The QSYS.LIB file system.
		QSYSIASP An independent ASP QSYS.LIB file system.
		ROOT The root (/) file system
UDFS A user-defined file system.		
UDFS MANAGEMENT A file system that manages the block special files (*BLKSF) for the user-defined file systems.		
FILE_IDENTIFIER_NUMBER	BIGINT	The file identifier number of the object. This number uniquely identifies the object with a file system. The file identifier number, generation identifier, and file system identifier used together uniquely identify the object on the system.
GENERATION_IDENTIFIER	BIGINT	The generation identifier associated with the object. The file identifier number, generation identifier, and file system identifier used together uniquely identify the object on the system.
FILE_SYSTEM_IDENTIFIER	BIGINT	The file system ID to which the object belongs. This number uniquely identifies the file system to which the object belongs. The file identifier number, generation identifier, and file system identifier used together uniquely identify the object on the system.
FILE_IDENTIFIER	BINARY(16)	An identifier associated with the referred to object.
REFERENCE_COUNT	INTEGER	Current number of references on the object for the specified job.
RO_SHARE_R_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with read and execute access intents only.
RO_SHARE_W_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with write access intents only.
RO_SHARE_RW_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
RO_SHARE_NONE_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with no other access intents.
WO_SHARE_R_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with read and execute access intents only.
WO_SHARE_W_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with write access intents only.
WO_SHARE_RW_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
WO_SHARE_NONE_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with no other access intents.
RW_SHARE_R_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with read and execute access intents only.
RW_SHARE_W_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with write access intents only.
RW_SHARE_RW_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
RW_SHARE_NONE_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with no other access intents.

Table 117. IFS_JOB_INFO table function (continued)

Column Name	Data Type	Description
XO_SHARE_R_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with read and execute access intents only.
XO_SHARE_W_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with write access intents only.
XO_SHARE_RW_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
XO_SHARE_NONE_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with no other access intents.
XR_SHARE_R_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with read and execute access intents only.
XR_SHARE_W_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with write access intents only.
XR_SHARE_RW_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
XR_SHARE_NONE_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with no other access intents.
CURRENT_DIRECTORY	VARCHAR(3)	The object is a directory that is being used as the current directory of the job. NO The object is not a directory that is being used as the current directory of the job. YES The object is the current directory of the job.
ROOT_DIRECTORY	VARCHAR(3)	The object is a directory that is being used as the root directory of the job. NO The object is not a directory that is being used as the root directory of the job. YES The object is the root directory of the job.
ATTRIBUTE_LOCK	VARCHAR(3)	Indicates whether attribute changes are prevented. NO Attribute changes are not prevented. YES Attribute changes are prevented.
SAVE_LOCK	VARCHAR(3)	Indicates whether the object is being referenced by an object save operation. NO Object is not being referenced by an object save operation. YES Object is being referenced by an object save operation.
INTERNAL_SAVE_LOCK	VARCHAR(3)	Indicates whether the object is being referenced internally during a save operation on a different object. NO The object is not being referenced internally during a save operation on a different object. YES The object is being referenced internally during a save operation on a different object.
LINK_CHANGES_LOCK	VARCHAR(3)	Indicates whether changes to links in the directory are prevented. NO Changes to links in the directory are not prevented. YES Changes to links in the directory are prevented.
CHECKED_OUT	VARCHAR(3)	Indicates whether the object is currently checked out. NO The object is not checked out. YES The object is checked out.
CHECKED_OUT_USER_NAME	VARCHAR(10)	The name of the user who has the object checked out. Contains the null value if CHECKED_OUT is NO.

Table 117. IFS_JOB_INFO table function (continued)

Column Name	Data Type	Description
FILE_SERVER_REFERENCE	VARCHAR(3)	<p>The File Server is holding a generic reference on the object on behalf of a client.</p> <p>NO The File Server is not holding a generic reference on the object on behalf of a client.</p> <p>YES The File Server is holding a generic reference on the object on behalf of a client.</p>
FILE_SERVER_WORKING_DIRECTORY	VARCHAR(3)	<p>The object is a directory, and the File Server is holding a working directory reference on it on behalf of a client.</p> <p>NO The object is not a directory or the File Server is not holding a working directory reference on it on behalf of a client.</p> <p>YES The object is a directory, and the File Server is holding a working directory reference on it on behalf of a client.</p>
NFS_SERVER_REFERENCE	VARCHAR(3)	<p>The Network File System (NFS) Version 4 server job is holding a generic reference on the object on behalf of a client.</p> <p>NO The NFS server job is not holding a generic reference on the object on behalf of a client.</p> <p>YES The NFS server job is holding a generic reference on the object on behalf of a client.</p>

Example

- List all the integrated file system references for the current job.

```
SELECT * FROM TABLE(QSYS2.IFS_JOB_INFO(JOB_NAME => '*'));
```

IFS_OBJECT_LOCK_INFO table function

The IFS_OBJECT_LOCK_INFO table function returns a result table that contains a row for each job that is known to be holding a reference, or lock, on the object.

This information is similar to what is returned by the Retrieve Object References (QPOLROR) API.

The list of object usages may be incomplete for objects residing in file systems other than the root (/), QOpenSys, and user-defined file systems. Objects in some of the other file systems can be locked with interfaces that do not use the integrated file system. Therefore, rows are only returned for references that were obtained as part of an integrated file system operation, or an operation that cause the integrated file system operation to occur.

Authorization: The user must have:

- Execute (*X) data authority to each directory preceding the object whose references are to be obtained and
- Read (*R) data authority to the object whose references are to be obtained.

```

>> IFS_OBJECT_LOCK_INFO ( ( path-name )
                          { PATH_NAME => }
,
  { IGNORE_ERRORS => } ignore-errors ) <<

```

path-name

An expression that defines the path name to the object whose reference information is to be returned. If the last element of the path is a symbolic link, the reference information will be for the symbolic link itself. If an absolute path name is not specified, the current working directory is used in combination with the relative path name to resolve to the object.

ignore-errors

A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 118. IFS_OBJECT_LOCK_INFO table function

Column Name	Data Type	Description
PATH_NAME	DBCLOB(16M) CCSID 1200	The full path name of the object.
JOB_NAME	VARCHAR(28)	The qualified job name holding the reference.
RO_COUNT	INTEGER	Total number of read only access references for this job.
WO_COUNT	INTEGER	Total number of write only access references for this job.
RW_COUNT	INTEGER	Total number of read and write access references for this job.
XO_COUNT	INTEGER	Total number of execute only access references for this job.
SHARE_R_COUNT	INTEGER	Total number of references for this job where the sharing mode allows sharing with read and execute access intents only.
SHARE_W_COUNT	INTEGER	Total number of references for this job where the sharing mode allows sharing with write access intents only.
SHARE_RW_COUNT	INTEGER	Total number of references for this job where the sharing mode allows sharing with read, execute, and write access intents.
SHARE_NONE_COUNT	INTEGER	Total number of references for this job where the sharing mode allows sharing with no other access intents.
ATTRIBUTE_LOCK	VARCHAR(3)	Indicates whether attribute changes are prevented for this job. NO Attribute changes are not prevented. YES Attribute changes are prevented.
SAVE_LOCK	VARCHAR(3)	Indicates whether the object is being referenced by an object save operation by this job. NO Object is not being referenced by an object save operation. YES Object is being referenced by an object save operation.
INTERNAL_SAVE_LOCK	VARCHAR(3)	Indicates whether the object is being referenced internally during a save operation on a different object by this job. NO The object is not being referenced internally during a save operation on a different object. YES The object is being referenced internally during a save operation on a different object.
LINK_CHANGES_LOCK	VARCHAR(3)	Indicates whether changes to links in the directory are prevented for this job. NO Changes to links in the directory are not prevented. YES Changes to links in the directory are prevented.
CHECKED_OUT	VARCHAR(3)	Indicates whether the object is currently checked out by this job. NO The object is not checked out. YES The object is checked out.
CHECKED_OUT_USER_NAME	VARCHAR(10)	The name of the user who has the object checked out. Contains the null value if CHECKED_OUT is NO.

Table 118. IFS_OBJECT_LOCK_INFO table function (continued)

Column Name	Data Type	Description
RO_SHARE_R_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with read and execute access intents only.
RO_SHARE_W_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with write access intents only.
RO_SHARE_RW_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
RO_SHARE_NONE_COUNT	INTEGER	Number of read only accesses for the job where the sharing mode allows sharing with no other access intents.
WO_SHARE_R_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with read and execute access intents only.
WO_SHARE_W_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with write access intents only.
WO_SHARE_RW_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
WO_SHARE_NONE_COUNT	INTEGER	Number of write only accesses for the job where the sharing mode allows sharing with no other access intents.
RW_SHARE_R_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with read and execute access intents only.
RW_SHARE_W_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with write access intents only.
RW_SHARE_RW_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
RW_SHARE_NONE_COUNT	INTEGER	Number of read and write accesses for the job where the sharing mode allows sharing with no other access intents.
XO_SHARE_R_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with read and execute access intents only.
XO_SHARE_W_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with write access intents only.
XO_SHARE_RW_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
XO_SHARE_NONE_COUNT	INTEGER	Number of execute only accesses for the job where the sharing mode allows sharing with no other access intents.
XR_SHARE_R_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with read and execute access intents only.
XR_SHARE_W_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with write access intents only.
XR_SHARE_RW_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with read, execute, and write access intents.
XR_SHARE_NONE_COUNT	INTEGER	Number of execute and read accesses for the job where the sharing mode allows sharing with no other access intents.
CURRENT_DIRECTORY	VARCHAR(3)	Indicates whether the object is a directory that is being used as the current directory of the job. NO The object is not a directory that is being used as the current directory of the job. YES The object is a directory that is being used as the current directory of the job.
ROOT_DIRECTORY	VARCHAR(3)	Indicates whether the object is a directory that is being used as the root directory of the job. NO The object is not a directory that is being used as the root directory of the job. YES The object is a directory that is being used as the root directory of the job.

Table 118. IFS_OBJECT_LOCK_INFO table function (continued)

Column Name	Data Type	Description
FILE_SERVER_REFERENCE	VARCHAR(3)	<p>Indicates whether the File Server is holding a generic reference on the object on behalf of a client for this job.</p> <p>NO The File Server is not holding a generic reference on the object on behalf of a client.</p> <p>YES The File Server is holding a generic reference on the object on behalf of a client.</p>
FILE_SERVER_WORKING_DIRECTORY	VARCHAR(3)	<p>Indicates whether the object is a directory, and the File Server is holding a working directory reference on it on behalf of a client for this job.</p> <p>NO The object is not a directory or the File Server is not holding a working directory reference on it on behalf of a client.</p> <p>YES The object is a directory, and the File Server is holding a working directory reference on it on behalf of a client.</p>
NFS_SERVER_REFERENCE	VARCHAR(3)	<p>Indicates whether the Network File System (NFS) Version 4 server job is holding a generic reference on the object on behalf of a client for this job.</p> <p>NO The NFS server job is not holding a generic reference on the object on behalf of a client.</p> <p>YES The NFS server job is holding a generic reference on the object on behalf of a client.</p>
NETSERVER_SESSION	CLOB(1M)	<p>This column contains a list of NetServer sessions, with each entry formatted as follows:</p> <ol style="list-style-type: none"> 1. 20 characters containing the session ID number 2. 10 character user name 3. 15 character workstation name 4. 45 character workstation address 5. 1 semicolon (;) <p>Contains the null value if there are no NetServer sessions.</p>

Example

- List all the jobs that have a lock on /usr/test1

```
SELECT * FROM TABLE(QSYS2.IFS_OBJECT_LOCK_INFO(PATH_NAME => '/usr/test1'))
```


IFS_OBJECT_PRIVILEGES table function

The IFS_OBJECT_PRIVILEGES table function returns a row for every user authorized to the object identified by the path name, along with their associated object and data authorities.

This information is similar to the information available through the Display Authority (DSPAUT) CL command and the Qp0lGetAttr()--Get Attributes API.

Authorization: The user needs either *ALLOBJ authority or the following authorities:

- For objects not in the QSYS.LIB file system:
 - For each directory included in the path name prior to the object name, *X
 - For the object, *OBJMGT
- For objects in the QSYS.LIB file system:
 - For each directory included in the path name prior to the object name, *X
 - For a *MBR object, *RX and *OBJMGT
 - For all other object types, *OBJMGT



path-name An expression that returns the path name identifying the object. A relative path name is relative to the current directory. If an absolute path name is not specified, the current working directory is used in combination with the relative path name to resolve to the object. If the last element of the path is a symbolic link, the privilege information will be for the symbolic link itself.

ignore-errors A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 119. IFS_OBJECT_PRIVILEGES table function

Column Name	Data Type	Description
PATH_NAME	DBCLOB(16M) CCSID 1200	The full path name of the object.
OBJECT_TYPE	VARCHAR(8)	The type of the object.
OWNER	VARCHAR(10)	The user profile that owns the object. Contains the null value if no owner is available.
PRIMARY_GROUP	VARCHAR(10)	The name of the user profile that is the primary group of the object. Can contain the following special value: *NOUSRPRF This special value is used by the Network File System to indicate that there is no user profile on the local server on the IBM i with a group ID (GID) matching the GID of the remote object. Contains the null value if the object has no primary group.
AUTHORIZATION_LIST	VARCHAR(10)	The name of the authorization list if the object is secured by an authorization list. Contains the null value if the object is not secured by an authorization list.
AUTHORIZATION_NAME	VARCHAR(10)	User profile name for this row. Can contain the following special values: *NOUSRPRF The authorities of either the owner or the primary group of the object for which the profile name could not be determined. This value is used by the Network File System only. It indicates that the user ID (UID) or the group ID (GID) for the remote object does not match any profile on the local server on the IBM i with that UID or GID. *PUBLIC This row contains the public authority for the object.

Table 119. IFS_OBJECT_PRIVILEGES table function (continued)

Column Name	Data Type	Description
DATA_AUTHORITY	VARCHAR(12)	<p>The operation, use, or access that AUTHORIZATION_NAME has to the object. Contains one of the following special values:</p> <p>*AUTL The public authority specified in the authorization list used by this object is used.</p> <p>*EXCLUDE All operations on the object are prohibited.</p> <p>*NONE The user does not have any data authorities.</p> <p>*R Allows access to the object attributes.</p> <p>*RW Allows access to the object attributes and allows the object to be changed. The user cannot use the object.</p> <p>*RWX Allows all operations on the object except those that are limited to the owner or controlled by the object rights.</p> <p>*RX Allows access to the object attributes and use of the object. The user cannot change the object.</p> <p>*W Allows the object to be changed.</p> <p>*WX Allows use of the object and allows the object to be changed. The user cannot access the object attributes.</p> <p>*X Allows the use of the object.</p> <p>USER DEFINED The specific data authorities do not match any of the predefined authority levels.</p>
OBJECT_OPERATIONAL	VARCHAR(3)	<p>Indicates the object operational authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_MANAGEMENT	VARCHAR(3)	<p>The object management authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_EXISTENCE	VARCHAR(3)	<p>The object existence authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_ALTER	VARCHAR(3)	<p>The object alter authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_REFERENCE	VARCHAR(3)	<p>The object reference authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
DATA_READ	VARCHAR(3)	<p>The data read authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
DATA_ADD	VARCHAR(3)	<p>The data add authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
DATA_UPDATE	VARCHAR(3)	<p>The data update authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>

Table 119. IFS_OBJECT_PRIVILEGES table function (continued)

Column Name	Data Type	Description	
DATA_DELETE	VARCHAR(3)	The data delete authority for AUTHORIZATION_NAME.	
		NO	The user does not have this authority.
		YES	The user has this authority.
DATA_EXECUTE	VARCHAR(3)	The data execute authority for AUTHORIZATION_NAME.	
		NO	The user does not have this authority.
		YES	The user has this authority.

Example

- List all the authorities for all objects in the /usr directory.

```
WITH OBJS AS (SELECT PATH_NAME
              FROM TABLE (QSYS2.IFS_OBJECT_STATISTICS(START_PATH_NAME => '/usr'))
              SELECT * FROM OBJS, TABLE(QSYS2.IFS_OBJECT_PRIVILEGES(PATH_NAME));
```

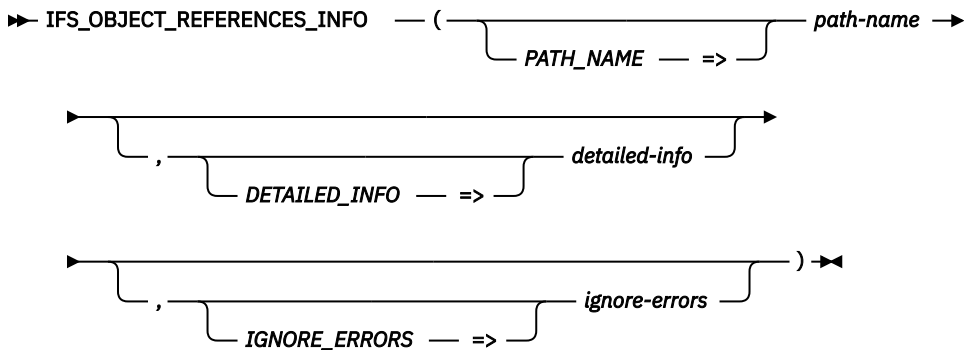
IFS_OBJECT_REFERENCES_INFO table function

The IFS_OBJECT_REFERENCES_INFO table function returns a single row result table that contains information about integrated file system references on an object.

This information is similar to what is returned by the Retrieve Object References (QPOLROR) API. The information may not be complete for objects residing in file systems other than the root (/), QOpenSys, and user-defined file systems.

Authorization: The user must have:

- Execute (*X) data authority to each directory preceding the object whose references are to be obtained and
- Read (*R) data authority to the object whose references are to be obtained.



path-name An expression that defines the path name to the object whose reference information is to be returned. If the last element of the path is a symbolic link, the reference information will be for the symbolic link itself. If an absolute path name is not specified, the current working directory is used in combination with the relative path name to resolve to the object.

detailed-info A character or graphic string expression that indicates the type of information to be returned.

- NO** Only basic information is returned. Values for columns through the `CHECKED_OUT_USER_NAME` are returned; NULL is returned for the remaining columns. This is the default.

YES Values are returned for all columns.

ignore-errors

A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing a single row with the format shown in the following table. All the columns are nullable.

Table 120. IFS_OBJECT_REFERENCES_INFO table function

Column Name	Data Type	Description
PATH_NAME	DBCLOB(16M) CCSID 1200	The full path name of the object.
REFERENCE_COUNT	INTEGER	Current number of references on the object. This may be 0 even though IN_USE has a value of YES.
IN_USE	VARCHAR(3)	Whether the object is currently in use. NO The object is not in use and all of the reference type fields are 0. YES The object is in use. At least one of the reference type fields is greater than 0. This condition may occur even if REFERENCE_COUNT is 0.
RO_COUNT	INTEGER	Total number of read only access references.
WO_COUNT	INTEGER	Total number of write only access references.
RW_COUNT	INTEGER	Total number of read and write access references.
XO_COUNT	INTEGER	Total number of execute only access references.
SHARE_R_COUNT	INTEGER	Total number of references where the sharing mode allows sharing with read and execute access intents only.
SHARE_W_COUNT	INTEGER	Total number of references where the sharing mode allows sharing with write access intents only.
SHARE_RW_COUNT	INTEGER	Total number of references where the sharing mode allows sharing with read, execute, and write access intents.
SHARE_NONE_COUNT	INTEGER	Total number of references where the sharing mode allows sharing with no other access intents.
ATTRIBUTE_LOCK	VARCHAR(3)	Indicates whether attribute changes are prevented. NO Attribute changes are not prevented. YES Attribute changes are prevented.
SAVE_LOCK	VARCHAR(3)	Indicates whether the object is being referenced by an object save operation. NO Object is not being referenced by an object save operation. YES Object is being referenced by an object save operation.
INTERNAL_SAVE_LOCK	VARCHAR(3)	Indicates whether the object is being referenced internally during a save operation on a different object. NO The object is not being referenced internally during a save operation on a different object. YES The object is being referenced internally during a save operation on a different object.

Table 120. IFS_OBJECT_REFERENCES_INFO table function (continued)

Column Name	Data Type	Description
LINK_CHANGES_LOCK	VARCHAR(3)	Indicates whether changes to links in the directory are prevented. NO Changes to links in the directory are not prevented. YES Changes to links in the directory are prevented.
CHECKED_OUT	VARCHAR(3)	Indicates whether the object is currently checked out. NO The object is not checked out. YES The object is checked out.
CHECKED_OUT_USER_NAME	VARCHAR(10)	The name of the user who has the object checked out. Contains the null value if CHECKED_OUT is NO.
Values are returned for the following columns when the DETAILED_INFO parameter is YES. They will contain the null value if DETAILED_INFO is NO.		
RO_SHARE_R_COUNT	INTEGER	Total number of read only access references. The sharing mode allows sharing with read and execute access intents only.
RO_SHARE_W_COUNT	INTEGER	Total number of read only access references. The sharing mode allows sharing with write access intents only.
RO_SHARE_RW_COUNT	INTEGER	Total number of read only access references. The sharing mode allows sharing with read, execute, and write access intents.
RO_SHARE_NONE_COUNT	INTEGER	Total number of read only access references. The sharing mode allows sharing with no other access intents.
WO_SHARE_R_COUNT	INTEGER	Total number of write only access references. The sharing mode allows sharing with read and execute access intents only.
WO_SHARE_W_COUNT	INTEGER	Total number of write only access references. The sharing mode allows sharing with write access intents only.
WO_SHARE_RW_COUNT	INTEGER	Total number of write only access references. The sharing mode allows sharing with read, execute, and write access intents.
WO_SHARE_NONE_COUNT	INTEGER	Total number of write only access references. The sharing mode allows sharing with no other access intents.
RW_SHARE_R_COUNT	INTEGER	Total number of read and write access references. The sharing mode allows sharing with read and execute access intents only.
RW_SHARE_W_COUNT	INTEGER	Total number of read and write access references. The sharing mode allows sharing with write access intents only.
RW_SHARE_RW_COUNT	INTEGER	Total number of read and write access references. The sharing mode allows sharing with read, execute, and write access intents.
RW_SHARE_NONE_COUNT	INTEGER	Total number of read and write access references. The sharing mode allows sharing with no other access intents.
XO_SHARE_R_COUNT	INTEGER	Total number of execute only access references. The sharing mode allows sharing with read and execute access intents only.
XO_SHARE_W_COUNT	INTEGER	Total number of execute only access references. The sharing mode allows sharing with write access intents only.
XO_SHARE_RW_COUNT	INTEGER	Total number of execute only access references. The sharing mode allows sharing with read, execute, and write access intents.
XO_SHARE_NONE_COUNT	INTEGER	Total number of execute only access references. The sharing mode allows sharing with no other access intents.
XR_SHARE_R_COUNT	INTEGER	Total number of execute and read access references. The sharing mode allows sharing with read and execute access intents only.
XR_SHARE_W_COUNT	INTEGER	Total number of execute and read access references. The sharing mode allows sharing with write access intents only.
XR_SHARE_RW_COUNT	INTEGER	Total number of execute and read access references. The sharing mode allows sharing with read, execute, and write access intents.
XR_SHARE_NONE_COUNT	INTEGER	Total number of execute and read access references. The sharing mode allows sharing with no other access intents.
CURRENT_DIRECTORY_COUNT	INTEGER	Total number of jobs where the object is a directory that is being used as the current directory of the job.

Table 120. IFS_OBJECT_REFERENCES_INFO table function (continued)

Column Name	Data Type	Description
ROOT_DIRECTORY_COUNT	INTEGER	Total number of jobs where the object is a directory that is being used as the root directory of the job.
FILE_SERVER_REFERENCE_COUNT	INTEGER	Total number of jobs where the File Server is holding a generic reference on the object on behalf of a client.
FILE_SERVER_WORKING_DIRECTORY_COUNT	INTEGER	Total number of jobs where the object is a directory, and the File Server is holding a working directory reference on it on behalf of a client.
NFS_SERVER_REFERENCE_COUNT	INTEGER	Total number of jobs where the Network File System (NFS) Version 4 server job is holding a generic reference on the object on behalf of a client.

Example

- Determine how /usr/test is currently being used.

```
SELECT * FROM TABLE (QSYS2.IFS_OBJECT_REFERENCES_INFO(PATH_NAME => '/usr/test'));
```

IFS_OBJECT_STATISTICS table function

The IFS_OBJECT_STATISTICS table function returns a table of objects contained in the starting path name or accessible from the starting path name.

This information is similar to what is returned by the Retrieve Directory Information (RTVDIRINF) command or the Qp0lGetAttr()--Get Attributes API.

No rows are returned for remote file system objects. This means that for the QNTC file system, only a row for /QNTC is returned. For the Network File System (NFS) and QFileSvr.400 file systems, no rows are returned.

Authorization: The user needs the following authorities:

For file systems other than QDLS and QSYS:

- For each directory included in the path name used to start the search, *X
- For each directory processed recursively, *RX and *OBJMGT
- For each object returned, *OBJMGT

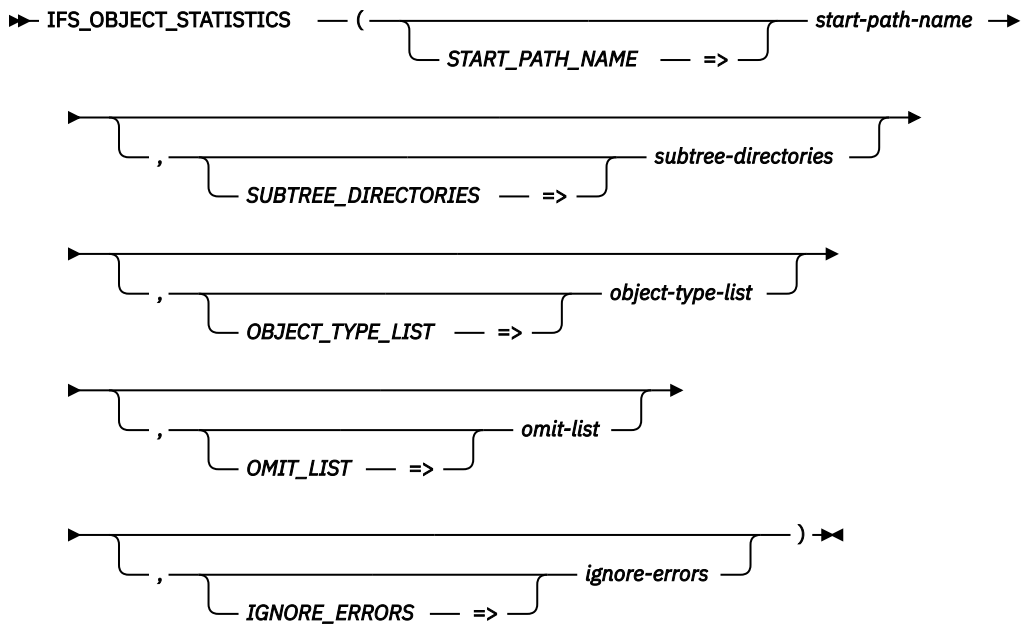
For the QDLS file system:

- For each directory included in the path name, except for QDLS, used to start the search, *X
- For every object being returned or processed recursively, *RWX and *OBJEXIST *OBJMGT *OBJALTER *OBJREF

For the QSYS file system:

- For the library or object included in the path name used to start the search, *USE
- For each library or object processed recursively by the service, *USE and *OBJMGT
- For each object returned, *OBJMGT

To return values for OBJECT_AUDIT and OBJECT_AUDIT_CREATE, the user must have *ALLOBJ or *AUDIT special authority.



start-path-name An expression that returns a path name for starting the search. A relative path name is relative to the current directory. If an absolute path name is not specified, the current working directory is used in combination with the relative path name to resolve to the object.

subtree-directories An expression that indicates whether all subdirectories should be recursively processed or not.

NO Only objects in the directory identified by *start-path-name* are processed.

YES All subdirectories for the directory identified by *start-path-name* are processed. This is the default.

object-type-list A list of one or more object types that should be returned. One or more blanks separate multiple values. The default is the empty string, indicating all objects are returned. Values include all the standard system object types, for example *PGM or *STMF, and the following special values:

***ALLDIR** Select all directory object types. This includes *LIB, *DIR, *FLR, *FILE, and *DDIR object types.

***ALLSTMF** Select all stream file object types. This includes *MBR, *DOC, *STMF, *DSTMF, and *USRSPC object types.

***MBR** Select all database file member types.

***NOQDLS** Exclude all QDLS file system object types.

***NOQOPT** Exclude all QOPT and QNTC file system object types.

***NOQSYS** Exclude all QSYS.LIB object types. This includes all objects in the QSYS.LIB file system and all independent ASP QSYS.LIB file systems which are available when the function is invoked.

omit-list A character string containing one or more path names to exclude from processing. All objects and sub directories accessed from this path are excluded as well. The default is the empty string, indicating no path names are excluded from processing. The root path ('/') must not be included in the list. If more than one path is provided, the values are separated by one or more blanks. If a blank exists in a path name, the path

name must be enclosed in either apostrophes (') or quotes("). When delimiting a path name and there is an apostrophe or quote (matching the enclosing character) in the path name, the embedded character needs to be doubled. For example, the following are valid path strings.

- One path name. No delimiter needed:
/dir1/dir2/content
- Two path names. No delimiter needed:
/dir1/content /dir1/dir3
- Blank in path name, path delimited with apostrophes:
'/dir1/content string'
- Two path names. Blank and apostrophe in first path. First path name is delimited with apostrophes and the apostrophe in the name is doubled. Second path does not require delimiters:
'/dir1/dir2/my file"s content' /dir4/test'123
- Two path names. Blank and apostrophe in first path. First path name is delimited with quotes. Second path is delimited even though not required:
"/dir1/dir2/my file's content" "/dir4/test'123"

ignore-errors A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 121. IFS_OBJECT_STATISTICS table function

Column Name	Data Type	Description
PATH_NAME	DBCLOB(16M) CCSID 1200	The path name of an integrated file system object.
OBJECT_TYPE	VARCHAR(10)	The object type.
SYMBOLIC_LINK	DBCLOB(16M) CCSID 1200	The symbolic link for this object. Contains the null value if OBJECT_TYPE is not *SYMLNK.
ASP_NUMBER	INTEGER	The auxiliary storage pool (ASP) in which the object is stored. Contains the null value if the object does not reside on this IBM i.
TEXT_DESCRIPTION	VARGRAPHIC(50) CCSID 1200	The text string associated with the object. Contains the null value if there is no text string.
FILE_IDENTIFIER_NUMBER	BIGINT	The file identifier number of the object. This number uniquely identifies the object with a file system. The file identifier number, generation identifier, and file system identifier used together uniquely identify the object on the system.
GENERATION_IDENTIFIER	BIGINT	The generation identifier associated with the object. The file identifier number, generation identifier, and file system identifier used together uniquely identify the object on the system.
FILE_SYSTEM_IDENTIFIER	BIGINT	The file system ID to which the object belongs. This number uniquely identifies the file system to which the object belongs. The file identifier number, generation identifier, and file system identifier used together uniquely identify the object on the system.

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
FILE_IDENTIFIER	BINARY(16)	An identifier associated with the referred to object.
FILE_ACCESS	BINARY(4)	Bit string containing information including the file type and file mode
CREATE_TIMESTAMP	TIMESTAMP(0)	The time the object was created.
ACCESS_TIMESTAMP	TIMESTAMP(0)	The time that the object's data was last accessed. Contains the null value if the object has never been accessed.
DATA_CHANGE_TIMESTAMP	TIMESTAMP(0)	The time that the object's data was last changed. Contains the null value if the object's data has never been changed.
OBJECT_CHANGE_TIMESTAMP	TIMESTAMP(0)	The time that the object's data or attributes were last changed. Contains the null value if the object's data or attributes have never been changed.
LAST_USED_TIMESTAMP	TIMESTAMP(0)	The date the object was last used. Contains the null value if the object has never been used or if usage data is not maintained for the IBM i type or the file system to which an object belongs.
DAYS_USED_COUNT	INTEGER	The number of days an object has been used. Usage has different meanings according to the specific file system and according to the individual object types supported within a file system. Usage can indicate the opening or closing of a file or can refer to adding links, renaming, restoring, or checking out an object. This count is incremented once each day that an object is used and can be reset.
LAST_RESET_TIMESTAMP	TIMESTAMP(0)	The timestamp when the days used count was last reset to zero. Contains the null value if the days used count has never been reset.
ALLOCATED_SIZE	BIGINT	The number of bytes that have been allocated for this object.
DATA_SIZE	BIGINT	The size, in bytes, of the data in this object. This size does not include object headers or the size of extended attributes associated with the object.
CCSID	INTEGER	The CCSID of the data and extended attributes of the object. Contains the null value if there is no CCSID.
CODE_PAGE	INTEGER	The code page derived from the coded character set identifier (CCSID) used for the data in the file or the extended attributes of the directory. Contains the null value if there is more than one code page or if the CCSID is not a supported CCSID.
EXTENDED_ATTRIBUTE_COUNT	BIGINT	Number of extended attributes associated with this object.
CRITICAL_EXTENDED_ATTRIBUTE_COUNT	BIGINT	Number of critical extended attributes associated with this object.
EXTENDED_ATTRIBUTE_SIZE	BIGINT	The total number of extended attribute bytes.
HARD_LINK_COUNT	INTEGER	The number of hard links to the object.
OBJECT_READ_ONLY	VARCHAR(3)	Whether the object can be written to or deleted, have its extended attributes changed or deleted, or have its size changed. NO The object can be changed. YES The object can only be read.
OBJECT_HIDDEN	VARCHAR(3)	Whether the object can be displayed using an ordinary directory listing. NO The object is not hidden. YES The object is hidden.
TEMPORARY_OBJECT	VARCHAR(3)	Whether the object is a temporary object. NO The object is a permanent object. YES The object is a temporary object.

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
SYSTEM_FILE	VARCHAR(3)	Whether the object is a system file and is excluded from normal directory searches. NO The object is not a system file. YES The object is a system file.
SYSTEM_USAGE	VARCHAR(6)	Whether the file has a special use by the system. NONE The file is a generic stream file. NWSSTG The file is a network server storage space. VRTVOL The file is a virtual volume. Examples include tape and optical virtual volumes. Contains the null value if OBJECT_TYPE is not *STMF.
DEVICE_SPECIAL_FILE	BIGINT	It the object is a device special file, the real device it represents. Contains the null value if this is not a device special file.
OBJECT_OWNER	VARCHAR(10)	The name of the user profile that is the owner of the object. Can contain the following special value: *NOUSRPRF This special value is used by the Network File System to indicate that there is no user profile on the local server on the IBM i with a user ID (UID) matching the UID of the remote object. Contains the null value if there is no owner.
USER_ID_NUMBER	BIGINT	User ID (UID) number of the owner of the object. Contains the null value if there is no owner.
PRIMARY_GROUP	VARCHAR(10)	The name of the user profile that is the primary group of the object. Can contain the following special value: *NOUSRPRF This special value is used by the Network File System to indicate that there is no user profile on the local server on the IBM i with a group ID (GID) matching the GID of the remote object. Contains the null value if the object has no primary group.
GROUP_ID_NUMBER	BIGINT	Group ID (GID) number of the user profile that is the primary group of the object. Contains the null value if the object has no primary group.
AUTHORIZATION_LIST	VARCHAR(10)	The name of the authorization list that is used to secure the named object. Contains the null value if no authorization list is used in determining authority to the object.
SET_EFFECTIVE_USER_ID	VARCHAR(3)	Set effective user ID (UID) at execution time. NO The user ID (UID) is not set at execution time. YES The object owner is the effective user ID (UID) at execution time. Contains the null value if OBJECT_TYPE is *DIR.
SET_EFFECTIVE_GROUP_ID	VARCHAR(3)	Set effective group ID (GID) at execution time. NO If the object is a file, the group ID (GID) is not set at execution time. If the object is a directory in the "root" (/), QOpenSys, or user-defined file systems, the group ID (GID) of objects created in the directory is set to the effective GID of the thread creating the object. This value cannot be set for other file systems. YES If the object is a file, the group ID (GID) is set at execution time. If the object is a directory, the group ID (GID) of objects created in the directory is set to the GID of the parent directory.

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
OBJECT_AUDIT	VARCHAR(7)	<p>The auditing value associated with the object.</p> <p>*ALL Audit all access to this object by all users on the system. All access is defined as a read or change operation.</p> <p>*CHANGE Audit all change access to this object by all users on the system.</p> <p>*NONE No auditing occurs for this object when it is read or changed regardless of the user who is accessing the object.</p> <p>*USRPRF Audit this object only if the current user is being audited. The current user is tested to determine if auditing should be done for this object. The user profile can specify if only change access is audited or if both read and change accesses are audited for this object.</p> <p>Contains the null value if the user is not allowed to retrieve the current auditing value.</p>
OBJECT_AUDIT_CREATE	VARCHAR(7)	<p>The create object auditing value associated with the directory. This is the auditing value given to any objects created in the directory.</p> <p>*ALL Audit all access to this object by all users on the system. All access is defined as a read or change operation.</p> <p>*CHANGE Audit all change access to this object by all users on the system.</p> <p>*NONE No auditing occurs for this object when it is read or changed regardless of the user who is accessing the object.</p> <p>*SYSVAL The object auditing value for the objects created in the directory is determined by the system auditing value (QCRTOBJAUD).</p> <p>*USRPRF Audit this object only if the current user is being audited. The current user is tested to determine if auditing should be done for this object. The user profile can specify if only change access is audited or if both read and change accesses are audited for this object. The OBJAUD parameter of the Change User Auditing (CHGUSRAUD) command is used to change the auditing for a specific user.</p> <p>Contains the null value if the user is not allowed to retrieve the current create object auditing value.</p>
JOURNALED	VARCHAR(3)	<p>Current journaling status of the object.</p> <p>NO The object is not currently being journaled.</p> <p>YES The object is currently being journaled.</p>
JOURNAL_LIBRARY	VARCHAR(10)	<p>The name of the library containing the journal currently being used .</p> <p>Contains the null value if the object is not journaled.</p>
JOURNAL_NAME	VARCHAR(10)	<p>The name of the journal currently being used.</p> <p>Contains the null value if the object is not journaled.</p>
JOURNAL_BEFORE_IMAGE	VARCHAR(3)	<p>Indicates whether the image of the object before a change is journaled when journaling is active.</p> <p>NO The image of the object before a change is not journaled.</p> <p>YES The image of the object before a change is journaled.</p> <p>Contains the null value if the object has never been journaled.</p>
JOURNAL_AFTER_IMAGE	VARCHAR(3)	<p>Indicates whether the image of the object after a change is journaled when journaling is active.</p> <p>NO The image of the object after a change is not journaled.</p> <p>YES The image of the object after a change is journaled.</p> <p>Contains the null value if the object has never been journaled.</p>
JOURNAL_IDENTIFIER	VARCHAR(10)	<p>The journal identifier (JID) for this object.</p> <p>Contains the null value if the object has never been journaled.</p>

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
JOURNAL_START_TIMESTAMP	TIMESTAMP(0)	The timestamp when the object had most recently had journaling started for it. Contains the null value if the object has never been journaled.
JOURNAL_OPTIONAL_ENTRIES	VARCHAR(3)	When journaling is active, entries that are considered optional are journaled. The list of optional journal entries varies for each object type. See Integrated file system for information regarding these optional entries for various objects. NO Optional entries for this object are not journaled. YES Optional entries for this object are journaled. Contains the null value if the object has never been journaled.
JOURNAL_SUBTREE	VARCHAR(3)	Indicates whether this object is a directory or library with inherit journal semantics. NO This object does not use inherit journal semantics. YES If this object is a directory, new objects created or linked within this directory will inherit the journal options and state from this directory. If this object is a library, new objects created or linked within this library will inherit the journal options and state from this library according to the journal inherit rules for this library. Contains the null value if the object has never been journaled.
PARTIAL_TRANSACTION	CHAR(1)	Indicates whether the object contains a partial transaction: N The object does not contain a partial transaction. R A rollback abnormally ended prior to completion. It is recommended that the object be restored as it can not be used. As a last resort, the Change Journaled Object (CHGJRNOBJ) command can be used to allow the object to be used. Doing this, however, may leave the object in an inconsistent state. Y The object was saved while active with a partial transaction. A subsequent restore of the object contains the partial transaction. The user should apply changes from the journal to complete the transaction. Contains the null value if the object has never been journaled.
APPLY_STARTING_RECEIVER_LIBRARY	VARCHAR(10)	The name of the library that contains the journal receiver. Contains the null value if the object has never been journaled.
APPLY_STARTING_RECEIVER	VARCHAR(10)	The oldest journal receiver needed to successfully Apply Journaled Changes (APYJRNCHG). If PARTIAL_TRANSACTION has a value of Y, the journal receiver contains the journal entries representing the start of the partial transaction. Otherwise, the journal receiver contains the journal entries representing the start-of-the-save operation. Contains the null value if PARTIAL_TRANSACTION has a value of R. Also contains the null value if the object has never been journaled.
APPLY_STARTING_RECEIVER_ASP	VARCHAR(10)	The name of the ASP for the library that contains the starting journal receiver. Can contain the special value *SYSBAS. Contains the null value if the object has never been journaled.
OBJECT_SIGNED	VARCHAR(3)	Whether an object has a digital signature. NO The object does not have a digital signature. YES The object does have a digital signature. Contains the null value if OBJECT_TYPE is not *STMF.

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
SYSTEM_TRUSTED_SOURCE	VARCHAR(3)	<p>Whether the object was signed by a source that is trusted by the system.</p> <p>NO None of the signatures came from a source that is trusted by the system.</p> <p>YES The object was signed by a source that is trusted by the system. If the object has multiple signatures, at least one of the signatures came from a source that is trusted by the system.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF or if OBJECT_SIGNED is NO.</p>
MULTIPLE_SIGNATURES	VARCHAR(3)	<p>Whether an object has more than one digital signature.</p> <p>NO The object has only one digital signature.</p> <p>YES The object has more than one digital signature. If SYSTEM_TRUSTED_SOURCE is YES, at least one of the signatures is from a source trusted by the system.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF or if OBJECT_SIGNED is NO.</p>
OBJECT_DOMAIN	VARCHAR(7)	<p>The domain of the object.</p> <p>*SYSTEM The object exists in system domain.</p> <p>*USER The object exists in user domain.</p>
BLOCK_SIZE	INTEGER	The block size of the object.
AUX_STORAGE_ALLOCATION	VARCHAR(8)	<p>Determines how auxiliary storage is allocated by the system for the specified object.</p> <p>DYNAMIC The system will dynamically determine the optimal auxiliary storage allocation for the object, balancing space used versus disk I/O operations.</p> <p>MINIMIZE The auxiliary storage will be allocated to minimize the space used by the object. That is, as additional auxiliary storage is required, it will be allocated in small sized extents to accommodate the current space requirement. Accessing an object composed of many small extents may increase the number of disk I/O operations for that object.</p> <p>NORMAL The auxiliary storage will be allocated normally. That is, as additional auxiliary storage is required, it will be allocated in logically sized extents to accommodate the current space requirement, and anticipated future requirements, while minimizing the number of disk I/O operations.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF.</p>
AUX_STORAGE_OVERFLOW	VARCHAR(3)	<p>Whether the object has overflowed the auxiliary storage pool it resides in.</p> <p>NO The auxiliary storage pool is not overflowed.</p> <p>YES The auxiliary storage pool is overflowed.</p>

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
MAIN_STORAGE_ALLOCATION	VARCHAR(8)	<p>Determines how main storage is allocated by the system for the specified object.</p> <p>DYNAMIC The system will dynamically determine the optimal main storage allocation for the object depending on other system activity and main storage contention. That is, when there is little main storage contention, as much storage as possible will be allocated and used to minimize the number of disk I/O operations. When there is significant main storage contention, less main storage will be allocated and used to minimize the main storage contention.</p> <p>MINIMIZE The main storage will be allocated to minimize the space used by the object. That is, as little main storage as possible will be allocated and used. This minimizes main storage usage while increasing the number of disk I/O operations since less information is cached in main storage.</p> <p>NORMAL The main storage will be allocated normally. That is, as much main storage as possible will be allocated and used. This minimizes the number of disk I/O operations since the information is cached in main storage.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF.</p>
STORAGE_FREED	VARCHAR(3)	<p>Whether the object's data has been moved offline, freeing its online storage.</p> <p>NO The object's data is not offline.</p> <p>YES The object's data is offline.</p>
STORED_LOCAL	VARCHAR(3)	<p>Indicates whether an object is stored locally or stored on a remote system. The decision of whether a file is local or remote varies according to the respective file system rules. Objects in file systems that do not carry either a local or remote indicator are treated as remote.</p> <p>NO The object's data is on a remote system.</p> <p>YES The object's data is stored locally.</p>
VIRTUAL_DISK_STORAGE	VARCHAR(3)	<p>Whether the object is the storage which was allocated for Integrated xSeries servers to use as virtual disk drives for the xSeries servers.</p> <p>NO Object is not virtual disk storage.</p> <p>YES Object is virtual disk storage.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF.</p>
DIRECTORY_FORMAT	CHAR(6)	<p>The format of the specified directory object.</p> <p>*TYPE1 The directory has the original directory format. The Convert Directory (CVTDIR) command may be used to convert from the *TYPE1 format to the *TYPE2 format.</p> <p>*TYPE2 The directory is optimized for performance, size, and reliability compared to directories having the *TYPE1 format.</p> <p>Contains the null value if OBJECT_TYPE is not *DIR.</p>
STREAM_FILE_FORMAT	CHAR(6)	<p>The format of the stream file.</p> <p>*TYPE1 The object has the same format as *STMF objects created on releases prior to V4R4. It has a minimum object size of 4096 bytes and a maximum object size of approximately 128 gigabytes.</p> <p>*TYPE2 This format was introduced in V4R4. It has a minimum object size of 4096 bytes and a maximum object size of approximately one terabyte in the "root" (/), QOpenSys and user-defined file systems. Otherwise, the maximum is approximately 256 gigabytes.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF.</p>

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
UDFS_FILE_FORMAT	CHAR(6)	<p>The default file format of stream files (*STMF) created in the user-defined file system.</p> <p>*TYPE1 The stream file (*STMF) has the same format as *STMFs created on releases prior to V4R4. It has a minimum object size of 4096 bytes and a maximum object size of approximately 128 gigabytes.</p> <p>*TYPE2 This format was introduced in V4R4. It has a minimum object size of 4096 bytes and a maximum object size of approximately one terabyte in the "root" (/), QOpenSys and user-defined file systems. Otherwise, the maximum is approximately 256 gigabytes.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF.</p>
UDFS_PREFERRED_STORAGE	CHAR(3)	<p>The preferred storage media for the objects in the UDFS.</p> <p>ANY No storage media is preferred. Storage will be allocated from any available storage media.</p> <p>SSD Solid state drive storage media is preferred. Storage should be allocated from solid state drive storage media, if available.</p> <p>Contains the null value if the preferred storage media is not known.</p>
UDFS_TEMPORARY_OBJECT	VARCHAR(3)	<p>Whether the objects in the UDFS are temporary objects.</p> <p>NO The objects in the UDFS are permanent objects.</p> <p>YES The objects in the UDFS are temporary objects.</p>
CASE_SENSITIVE_FILE_SYSTEM	VARCHAR(3)	<p>The case sensitivity of the file system that contains this object.</p> <p>NO The file system is not case sensitive.</p> <p>YES The file system is case sensitive.</p>
RESTRICT_RENAME_AND_UNLINK	VARCHAR(3)	<p>Restricted renames and unlinks for objects within a directory. Objects can be linked into a directory that has this attribute set on, but cannot be renamed or unlinked from it unless one or more of the following are true for the user performing the operation:</p> <ul style="list-style-type: none"> • The user is the owner of the object. • The user is the owner of the directory. • The user has *ALLOBJ special authority. <p>NO No additional restrictions for rename and unlink operations.</p> <p>YES Additional restrictions for rename and unlink operations.</p>
PC_ARCHIVE	VARCHAR(3)	<p>Whether the object has changed since the last time it was saved on the PC.</p> <p>NO The object has not changed.</p> <p>YES The object has changed.</p>
SYSTEM_ARCHIVE	VARCHAR(3)	<p>Whether the object has changed and needs to be saved on the IBM i. The value will be YES when an object's change time is updated, and set to NO when the object has been saved.</p> <p>NO The object has not changed and does not need to be saved.</p> <p>YES The object has changed and needs to be saved.</p>
ALLOW_SAVE	VARCHAR(3)	<p>Whether the object can be saved or not.</p> <p>NO This object will not be saved when using the Save Object (SAV) command or the QsrSave() API.</p> <p>Additionally, if this object is a directory, none of the objects in the directory's subtree will be saved unless they were explicitly specified as an object to be saved.</p> <p>YES This object will be saved when using the Save Object (SAV) command or the QsrSave() API.</p>

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
SYSTEM_RESTRICT_SAVE	VARCHAR(3)	<p>Whether the system prevents the object from being saved.</p> <p>NO The system does not prevent the object from being saved.</p> <p>YES The system has determined that the object cannot be saved because of system restrictions.</p>
INHERIT_ALLOW_CHECKPOINT_WRITER	VARCHAR(3)	<p>Whether new objects created within a directory should inherit the save-while-active checkpoint processing options of its parent.</p> <p>NO New directory objects created within this directory will have the QPOL_ATTR_INHERIT_ALWCKPWRT attribute set to NO. New objects created within this directory will have the QPOL_ATTR_ALWCKPWRT attribute set to QPOL_NOT_ALWCKPWRT.</p> <p>YES New directory objects created within this directory will have the QPOL_ATTR_INHERIT_ALWCKPWRT attribute set to YES. New objects created within this directory will have the QPOL_ATTR_ALWCKPWRT attribute set to QPOL_ALWCKPWRT.</p>
ALLOW_WRITE_DURING_SAVE	VARCHAR(3)	<p>If the object is a stream file, indicates whether the stream file can be shared with readers and writers during save-while-active checkpoint processing. If the object is a directory, indicates whether links can be added, removed, or renamed in the directory during a save-while-active operation.</p> <p>NO If the object is a stream file, it can be shared with readers only during save-while-active checkpoint processing. If the object is a directory, links can not be added, removed, or renamed in the directory during a save-while-active operation.</p> <p>YES If the object is a stream file, it can be shared with readers and writers during save-while-active checkpoint processing. If the object is a directory, links can be added, removed, or renamed in the directory during a save-while-active operation.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF or *DIR.</p>
EXIT_PROGRAM_SCAN	VARCHAR(11)	<p>Whether the object will be scanned when exit programs are registered with any of the integrated file system scan-related exit points.</p> <p>CHANGE ONLY The object will be scanned according to the rules described in the scan-related exit programs only if the object has been modified since the last time the object was scanned. It will not be scanned if the scanning software has been updated. This attribute only takes effect if the Scan file systems control (QSCANFCTL) system value has *USEOCOATR specified. Otherwise, it will be treated the same as YES.</p> <p>NO The object will not be scanned according to the rules described in the scan-related exit programs.</p> <p>YES The object will be scanned according to the rules described in the scan-related exit programs if the object has been modified or if the scanning software has been updated since the last time the object was scanned.</p> <p>Contains the null value if OBJECT_TYPE is not *STMF.</p>

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
EXIT_PROGRAM_SCAN_DIRECTORY	VARCHAR(11)	<p>Whether the objects created in a directory will be scanned when exit programs are registered with any of the integrated file system scan-related exit points.</p> <p>CHANGE ONLY After an object is created in the directory, the object will be scanned according to the rules described in the scan-related exit programs only if the object has been modified since the last time the object was scanned. It will not be scanned if the scanning software has been updated. This attribute only takes effect if the Scan file systems control (QSCANFCTL) system value has *USEOCOATR specified. Otherwise, it will be treated the same as YES.</p> <p>NO After an object is created in the directory, the object will not be scanned according to the rules described in the scan-related exit programs.</p> <p>YES After an object is created in the directory, the object will be scanned according to the rules described in the scan-related exit programs if the object has been modified or if the scanning software has been updated since the last time the object was scanned.</p> <p>Contains the null value if OBJECT_TYPE is not *DIR.</p>
SCAN_STATUS	VARCHAR(12)	<p>The scan status associated with this object.</p> <p>FAILURE The object has been scanned by a scan-related exit program, and at the time of that last scan request, the object failed the scan and the operation did not complete. Once an object has been marked as a failure, it will not be scanned again until the object's scan signature is different than the global scan key signature or independent ASP group scan key signature as appropriate. Therefore, subsequent requests to work with the object will fail with a scan failure indication if that access meets the criteria for when an object is to be scanned.</p> <p>NOT REQUIRED The object does not require any scanning because the object is marked to not be scanned.</p> <p>PENDING The object is in a file system that has not completely converted to the *TYPE2 directory format, and therefore will not be scanned until the file system is completely converted.</p> <p>REQUIRED A scan is required for the object either because it has not yet been scanned by the scan-related exit programs, or because the object data or CCSID has been modified since it was last scanned.</p> <p>SUCCESS The object has been scanned by a scan-related exit program, and at the time of that last scan request, the object did not fail the scan.</p>
CCSID_SCAN	INTEGER	<p>A CCSID value that the object has been scanned in if it was previously scanned in a CCSID. If SCAN_STATUS is SUCCESS, the object was successfully scanned in this CCSID. If SCAN_STATUS is FAILURE, the object failed the scan in this CCSID.</p> <p>Contains the null value if no CCSID applies.</p>
CCSID_SCAN_SUCCESS	INTEGER	<p>If SCAN_STATUS is SUCCESS, the object was successfully scanned in this CCSID.</p> <p>Contains the null value if the SCAN_STATUS is FAILURE or if CCSID_SCAN is the null value.</p>
SCAN_SIGNATURES_DIFFERENT	VARCHAR(3)	<p>When an object is in an independent ASP group, the object scan signature is compared to the associated independent ASP group scan signature. When an object is not in an independent ASP group, the object scan signature is compared to the global scan signature value.</p> <p>NO The compared signatures are not different.</p> <p>YES The compared signatures are different.</p>

Table 121. IFS_OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
BINARY_SCAN	VARCHAR(3)	The object was scanned in binary mode when it was previously scanned.
		NO The object was not scanned in binary mode.
		YES The object was scanned in binary mode.
CHECKED_OUT	VARCHAR(3)	Whether the object is checked out.
		NO The object is not checked out.
		YES The object is checked out.
CHECKED_OUT_TIMESTAMP	TIMESTAMP(0)	The time the object was checked out.
		Contains the null value if the object is not checked out.
CHECKED_OUT_USER	VARCHAR(10)	The user who has the object checked out.
		Contains the null value if the object is not checked out.

Example

- List basic information for all the objects in directory /usr.

```
SELECT PATH_NAME, OBJECT_TYPE, DATA_SIZE, OBJECT_OWNER
FROM TABLE (QSYS2.IFS_OBJECT_STATISTICS(START_PATH_NAME => '/usr',
                                         SUBTREE_DIRECTORIES => 'NO'));
```

- List basic information for all the objects in /usr, processing all subdirectories as well.

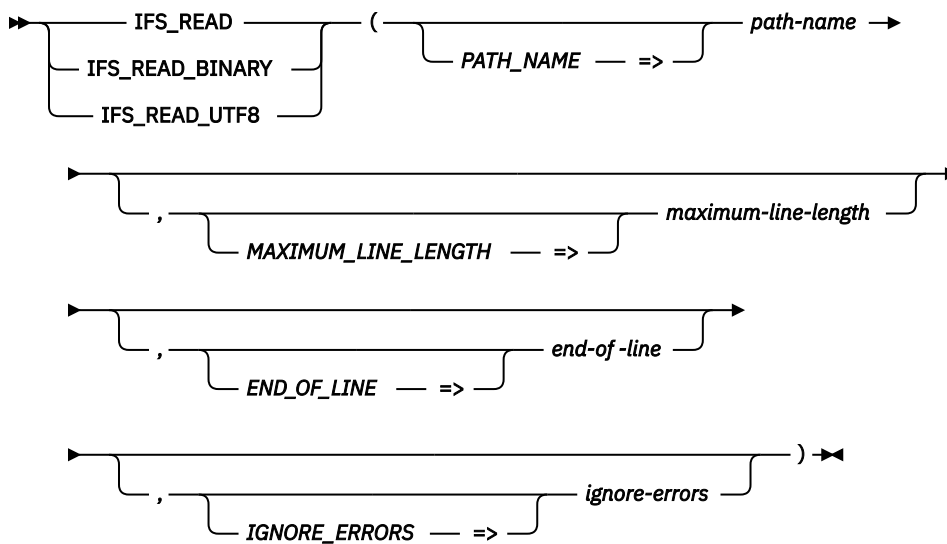
```
SELECT PATH_NAME, OBJECT_TYPE, DATA_SIZE, OBJECT_OWNER
FROM TABLE (QSYS2.IFS_OBJECT_STATISTICS(START_PATH_NAME => '/usr',
                                         SUBTREE_DIRECTORIES => 'YES'));
```

IFS_READ, IFS_READ_BINARY, and IFS_READ_UTF8 table functions

The IFS_READ, IFS_READ_BINARY, and IFS_READ_UTF8 table functions read an integrated file system stream file identified by *path-name*. The file's data is returned as character, binary, or UTF-8 data. It can be returned as one string of data, or it can be broken into multiple lines using a specified length or end of line characters.

Authorization: The caller must have:

- For objects not in the QSYS.LIB file system:
 - Execute (*X) data authority to each directory preceding the stream file being read and
 - Read (*R) data authority to the stream file
- For objects in the QSYS.LIB file system:
 - Execute (*X) data authority to each directory preceding the object being read and
 - For a *SAVF object, Read, Write, and Execute (*RWX) data authority to the object
 - For all other object types, Read (*R) data authority to the object



The schema is QSYS2.

- path-name** An expression that returns the path name identifying the stream file to read. A relative path name is relative to the current directory. If an absolute path name is not specified, the current working directory is used in combination with the relative path name to resolve to the object. If the object is not a stream file, an error is issued.
- maximum-line-length** An integer value that specifies the maximum number of characters returned for each line. It must be greater than 0. The default is 2 gigabytes.
 If *end-of-line* is NONE or if no end of line sequence is found in the stream file, the number of characters returned for each line will be limited to this value.
 If an end of line sequence is encountered before this length is reached, the line will end at that point. The next line returned will start with the character directly after the end of line sequence.
- end-of-line** A character or graphic string that specifies the end of line characters to be recognized in the stream file. Each occurrence of an end of line sequence determines a line which is returned. The end of line character sequence is not returned with the line. When using IFS_READ_BINARY, end of line characters are never processed, so this parameter must have a value of NONE.
 The carriage-return character is always X'0D'. Based on the CCSID of the stream file being read, the line feed character is X'25' for an EBCDIC CCSID and X'0A' for ASCII and UTF-8 CCSIDs.
- ANY** Any of the four end of line sequences indicate the end of a line. This is the default for IFS_READ and IFS_READ_UTF8.
- CR** A carriage return indicates the end of a line.
- CRLF** A carriage return and line feed indicate the end of a line.
- LF** A line feed indicates the end of a line.
- LFCR** A line feed and carriage return indicate the end of a line.
- NONE** No end of line characters are recognized. *maximum-line-length* determines the number of characters to be returned. This is the default for IFS_READ_BINARY.
- ignore-errors** A character or graphic string expression that identifies what to do when an error is encountered.

- NO** An error is returned.
- YES** A warning is returned.
No row is returned when an error is encountered. This is the default.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 122. IFS_READ table function

Column Name	Data Type	Description
LINE_NUMBER	INTEGER	Relative position of this line in the stream file.
LINE	For IFS_READ: CLOB(2G) For IFS_READ_BINARY: BLOB(2G) For IFS_READ_UTF8: CLOB(2G) CCSID 1208	The data for this line. <ul style="list-style-type: none"> • For IFS_READ, the data from the stream file will be returned in the job CCSID. • For IFS_READ_BINARY, the data from the stream file will be returned exactly as it is stored without conversion. • For IFS_READ_UTF8, the data from the stream file will be returned in CCSID 1208.

Example

- Read the data from stream file /usr/file1. Break lines when a carriage return/line feed sequence is encountered. The result will be in the job's CCSID.

```
SELECT * FROM TABLE(QSYS2.IFS_READ(PATH_NAME => '/usr/file1',
                                   END_OF_LINE => 'CRLF'));
```

- Read the data from stream file /usr/file2. If the file size is less than 2 gigabytes, the result will be a single row containing the entire file. If the file size is greater than 2 gigabytes, multiple rows will be returned.

```
SELECT * FROM TABLE(QSYS2.IFS_READ_BINARY(PATH_NAME => '/usr/file2'));
```

IFS_WRITE, IFS_WRITE_BINARY, and IFS_WRITE_UTF8 procedures

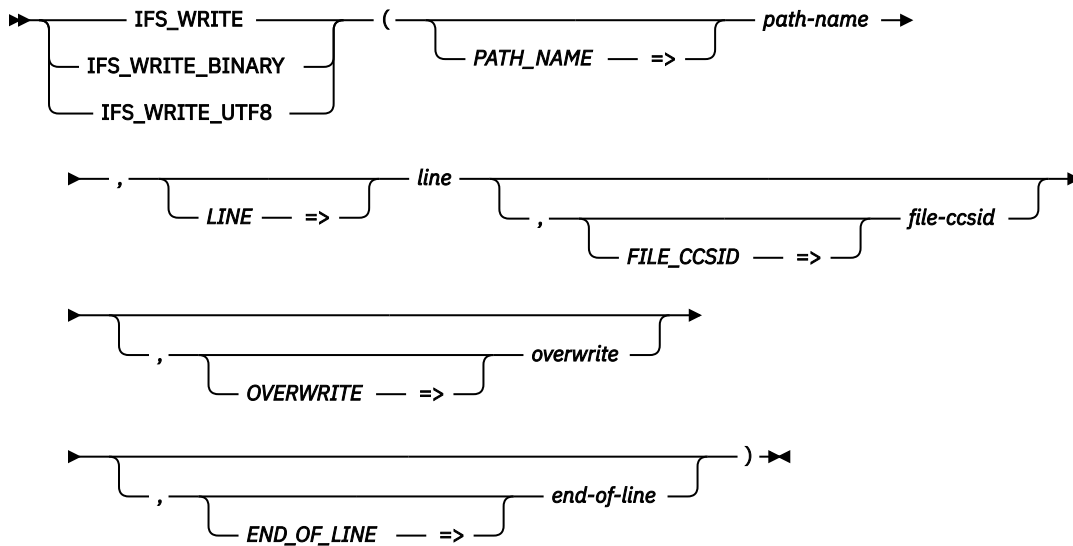
The IFS_WRITE, IFS_WRITE_BINARY, and IFS_WRITE_UTF8 procedures write data to an integrated file system stream file. The data can be written as character, binary, or UTF-8 data. Data can be either replaced or appended for an existing file, or a new file can be created.

Up to 2 gigabytes of data can be written by one call to this procedure. It can contain embedded end of line characters, or the procedure can append end of line characters to the input data.

Authorization: The caller must have:

- For objects not in the QSYS.LIB file system when the stream file exists and is not being replaced:
 - Execute (*X) data authority to each directory preceding the stream file being written and
 - Write (*W) data authority to the stream file
- For objects not in the QSYS.LIB file system when the stream file does not exist or is being replaced:
 - Execute (*X) data authority to each directory preceding the stream file being written and
 - Write and Execute (*WX) authority to the parent directory of the stream file
- For objects in the QSYS.LIB file system when the object exists and is not being replaced:
 - Execute (*X) data authority to each directory preceding the object being written and
 - For a *SAVF object, Read, Write, and Execute (*RWX) data authority to the object
 - For all other object types, Write (*W) data authority to the object
- For objects in the QSYS.LIB file system when the object does not exist or is being replaced:

- Execute (*X) data authority to each directory preceding the object being written and
 - For a *SAVF object, Read and Execute and Add (*RX and *ADD) data authority to the parent directory of the object
 - For a physical file member, Add (*ADD) data authority to the parent directory of the object
 - For all other object types, *OBJMGT or *OBJALTER authority to the parent directory of the object



The schema is QSYS2.

path-name A character or graphic string that defines the path name for the file to be written. If an absolute path name is not specified, the current working directory is used in combination with the relative path name to resolve to the object.

If *path-name* identifies an existing object, the object must be a stream file. Otherwise, a stream file will be created.

When a stream file is created, the default authority for the parent directory is used.

line A character or graphic string containing the data to be written to the stream file at *path-name*. It can be up to 2 gigabytes long.

- For IFS_WRITE, *line* is a character string in the job CCSID. If the stream file is not in the job CCSID, *line* will be converted to the stream file's CCSID as it is written.
- For IFS_WRITE_BINARY, *line* is a binary string. The data will not be converted when writing to the stream file.
- For IFS_WRITE_UTF8, *line* is a UTF-8 string. If the stream file is not UTF-8, *line* will be converted to the stream file's CCSID as it is written.

file-ccsid An integer value that specifies the CCSID to be used when creating a new stream file. This parameter is ignored when appending to an existing file.

If this parameter is not specified, the default is 1208 for IFS_WRITE_UTF8 and 0 for IFS_WRITE and IFS_WRITE_BINARY.

When *file-ccsid* is 0, the job CCSID is used when creating a new stream file.

overwrite A character or graphic string that specifies whether the write operation appends to the stream file, replaces the stream file, or fails when a stream file with the specified name already exists.

APPEND The data in *line* is added to the end of the existing stream file. If the stream file does not exist, it is created. This is the default.

NONE The write operation fails if the stream file exists.

REPLACE The data in *line* replaces the existing stream file if it exists. An existing stream file is deleted and a new stream file is created. The CCSID of the stream file might change. If the stream file does not exist, it is created.

end-of-line A character or graphic string that specifies the end of line characters to write to the stream file after *line* is written. When using IFS_WRITE_BINARY, end of line characters are never appended, so this parameter must have a value of NONE.
The carriage-return character is always X'0D'. Based on the CCSID of the stream file being written, the line feed character is X'25' for an EBCDIC CCSID and X'0A' for ASCII and UTF-8 CCSIDs.

CR A carriage return is appended.

CRLF A carriage return and line feed are appended. This is the default for IFS_WRITE and IFS_WRITE_UTF8.

LF A line feed is appended.

LFCR A line feed and carriage return are appended.

NONE No end of line characters are appended. This is the default for IFS_WRITE_BINARY.

Examples

- Create a stream file which contains a list of all of the libraries on the system, one per line. The file is created in job CCSID.

```
BEGIN
-- Make sure output file is empty to start
CALL QSYS2.IFS_WRITE(PATH_NAME => '/tmp/library_names',
                    LINE => '',
                    OVERWRITE => 'REPLACE',
                    END_OF_LINE => 'NONE');

-- Add lines to the output file
FOR SELECT OBJNAME AS LIBNAME FROM TABLE(QSYS2.OBJECT_STATISTICS('*ALLSIMPLE', 'LIB')) DO
CALL QSYS2.IFS_WRITE(PATH_NAME => '/tmp/library_names',
                    LINE => LIBNAME);

END FOR;
END;
```

The QSYS2.IFS_READ table function can be used to read the contents of the generated stream file.

```
SELECT * FROM TABLE(QSYS2.IFS_READ('/tmp/library_names'));
```

- Create a UTF-8 (CCSID 1208) stream file which contains a UTF-8 Byte Order Mark (BOM) and the string 'Hello'.

```
CALL QSYS2.IFS_WRITE_BINARY(PATH_NAME => '/usr/utf8file',
                            LINE => BLOB(X'EFBBBF48656C66'),
                            FILE_CCSID => 1208,
                            OVERWRITE => 'REPLACE'
                            );
```

SERVER_SHARE_INFO view

The SERVER_SHARE_INFO view returns information about IBM i NetServer shares.

This information is similar to what is returned by the List Server Information (QZLSLSTI) and Open List of Server Information (QZLSOLST) APIs.

Authorization: None required.

The following table describes the columns in the view. The system name is SHARE_INFO. The schema is QSYS2.

Table 123. SERVER_SHARE_INFO view

Column Name	System Column Name	Data Type	Description
SERVER_SHARE_NAME	SHARE	VARCHAR(12)	The network name of the resource.
SHARE_TYPE	SHARE_TYPE	VARCHAR(5)	The type of share. FILE This is a file share PRINT This is a print share
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	An optional comment about the shared resource or computer. Contains the null value if there is no text description.
The following columns can contain values when SHARE_TYPE is FILE. They will contain the null value when SHARE_TYPE is PRINT.			
PATH_NAME	PATH_NAME	DBCLOB(16M) CCSID 1200 Nullable	The file share path in the integrated file system. Contains the null value when SHARE_TYPE is PRINT.
PERMISSIONS	PERMISSION	VARCHAR(3) Nullable	Permissions to be applied against the file for sharing. *R Read only *RW Read/write Contains the null value when SHARE_TYPE is PRINT.
MAXIMUM_CONNECTIONS	MAX_CONN	INTEGER Nullable	The maximum number of concurrent connections that the shared file resource can accommodate. Contains the null value if the number is unlimited or SHARE_TYPE is PRINT.
CURRENT_CONNECTIONS	CUR_CONN	INTEGER Nullable	The number of connections that are currently made to the resource. Contains the null value if the value could not be returned or SHARE_TYPE is PRINT.
TEXT_CONVERSION_ENABLED	TEXT_CONV	VARCHAR(5) Nullable	Whether the server enables text file data conversion for this file share. MIXED Text conversion enabled and mixed data is allowed NO Text conversion not enabled YES Text conversion enabled Contains the null value when SHARE_TYPE is PRINT.
TEXT_CONVERSION_CCSID	TEXT_CCSID	INTEGER Nullable	The CCSID that is used for text file data conversion. If the value is 0, no CCSID was specified so the IBM i NetServer's default CCSID will be used. Contains the null value if TEXT_CONVERSION_ENABLED is NO or when SHARE_TYPE is PRINT.
FILE_EXTENSION_COUNT	EXT_COUNT	INTEGER Nullable	The number of file extension entries returned in FILE_EXTENSIONS. Contains the null value when SHARE_TYPE is PRINT.

Table 123. SERVER_SHARE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FILE_EXTENSIONS	EXTENSIONS	VARBINARY(5000) Nullable	<p>A string containing a list of file extensions. The format of each entry is a 2 byte integer length followed by that number of characters followed by a single blank. If there are more file extensions than what fit in this column, the last file extension in the list will have a length of 3 with a value of +++ to indicate the list was truncated.</p> <p>Examples of extensions are:</p> <ul style="list-style-type: none"> * The server will convert all files. . The server will convert all files without an extension. TXT, .TXT The server will convert all files ending with .TXT (that is, a.TXT, a.b.TXT). ..TXT, ...TXT,TXT Extensions with more than one leading period will have no effect on the server. No translation will be done. T*T The server will convert all files ending with an extension that substitutes any number of characters for the * wild card (that is, a.T123T, b.TXT, c.TEST). T?T The server will convert all files ending with an extension that substitutes any one character for the ? wild card (that is, a.T1T, b.TXT). <p>Contains the null value when SHARE_TYPE is PRINT.</p>
<p>The following columns can contain values when SHARE_TYPE is PRINT. They will contain the null value when SHARE_TYPE is FILE.</p>			
SPOOLED_FILE_TYPE	SPOOL_TYPE	VARCHAR(9) Nullable	<p>The type of spooled files that will be created using this print share.</p> <ul style="list-style-type: none"> AFP Advanced Function Presentation AUTOSENSE Automatic type sensing SCS Simplified Character Set USERASCII User ASCII <p>Contains the null value when SHARE_TYPE is FILE.</p>
OUTPUT_QUEUE_LIBRARY	OUTQLIB	VARCHAR(10) Nullable	<p>The library containing the output queue.</p> <p>Contains the null value when SHARE_TYPE is FILE.</p>
OUTPUT_QUEUE	OUTQ	VARCHAR(10) Nullable	<p>The name of the output queue.</p> <p>Contains the null value when SHARE_TYPE is FILE.</p>

Table 123. SERVER_SHARE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PRINT_DRIVER	PRT_DRIVER	VARCHAR(50) Nullable	The text string that identifies the print driver appropriate for this print share. When personal computers connect to this shared printer, this identifies the print driver they should use. This text should match the name of a print driver known to the personal computer operating system. Contains the null value when there is no print driver or when SHARE_TYPE is FILE.
PRINTER_FILE_LIBRARY	PRTF_LIB	VARCHAR(10) Nullable	The library containing the printer file. Contains the null value when there is no printer file or when SHARE_TYPE is FILE.
PRINTER_FILE	PRTF	VARCHAR(10) Nullable	The name of the printer file. Contains the null value when there is no printer file or when SHARE_TYPE is FILE.
PUBLISH_PRINT_SHARE	PUBLISH	VARCHAR(3) Nullable	Whether the print share is to be published NO Print share is not published YES Print share is published Contains the null value when SHARE_TYPE is FILE.

Example

- List all the file shares.

```
SELECT * FROM QSYS2.SERVER_SHARE_INFO
WHERE SHARE_TYPE = 'FILE';
```

Java Services

This view and procedure provide Java information and JVM management options.

JVM_INFO view

The JVM_INFO view returns information about active Java Virtual Machine (JVM) jobs. The information is a subset of what can be found interactively using the Work with JVM Jobs (WRKJVMJOB) command.

Authorization: The caller must have *JOBCTL special authority to see values for columns other than JOB_NAME and PROCESS_ID.

The following table describes the columns in the view. The schema is QSYS2.

Table 124. JVM_INFO view

Column Name	System Column Name	Data Type	Description
JOB_NAME	JOB_NAME	VARCHAR(28)	The qualified job name for the active JVM.
PROCESS_ID	PROCESS_ID	INTEGER	The process identifier used by the kernel to uniquely identify the process.
START_TIME	START_TIME	TIMESTAMP	The current time when the JVM was started.
INITIAL_THREAD_TASKCOUNT	INITTHDNUM	BIGINT	The taskcount or TDE number of the JVM's initial thread. The taskcount or TDE number is a unique identifier assigned to each job, thread, and task running in the system.
JAVA_THREAD_COUNT	JAVATHDNUM	BIGINT	The current number of java threads within the JVM job.
TOTAL_GC_TIME	ACCUMTIME	BIGINT	Total time spent performing garbage collection tasks in milliseconds.
GC_CYCLE_NUMBER	GC_CYCLE	INTEGER	The current or last garbage collection cycle performed.

Table 124. JVM_INFO view (continued)

Column Name	System Column Name	Data Type	Description
GC_POLICY_NAME	GCPOLICY	VARGRAPHIC(16) CCSID 1200	The name of the garbage collection policy in use.
JAVA_HOME	JAVA_HOME	VARGRAPHIC(1024) CCSID 1200	The java.home environment variable value in effect for this JVM. This value indicates the JDK that is used when running a Java application. The location of the Java tools and utilities is in one of two directories, either <JAVA_HOME>/jre/bin or <JAVA_HOME>/bin, where <JAVA_HOME> is the value of the JAVA_HOME environment variable. For example, if JAVA_HOME is set to /QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit, indicating that IBM Technology for Java 6 32-bit is to be used, then the Java tools and utilities directories would be: /QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit/bin /QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit/jre/bin
USER_DIRECTORY	USER_DIR	VARGRAPHIC(1024) CCSID 1200	The user working directory for the JVM. This also indicates the location where diagnostic detail will be dumped for the JVM.
NUM_CURRENT_PROPERTIES	NUMPROP	INTEGER	Total number of Java system properties currently present.
INITIAL_HEAP_SIZE	INTHEAP	BIGINT	The initial heap size available to the JVM code, in kilobytes.
CURRENT_HEAP_SIZE	CURHEAP	BIGINT	The amount of memory, in kilobytes, currently allocated for heap space.
IN_USE_HEAP_SIZE	INUSEHEAP	BIGINT	The amount of memory, in kilobytes, currently in use by the heap.
MAX_HEAP_SIZE	MAXHEAP	BIGINT	The maximum heap size available to the JVM code, in kilobytes.
MALLOC_MEMORY_SIZE	MALLOCSIZE	BIGINT	The amount of memory, in kilobytes, that has been allocated with malloc().
INTERNAL_MEMORY_SIZE	INTMEM	BIGINT	The amount of memory, in kilobytes, that the JVM is using for internal operations.
JIT_MEMORY_SIZE	JITSIZE	BIGINT	The size of the memory space, in kilobytes, that is used by the JIT (Just in Time) compiler.
SHARED_CLASS_SIZE	SHAREDSIZE	BIGINT	The amount of memory, in kilobytes, that the JVM is using for shared classes.
BIT_MODE	BIT_MODE	INTEGER	The Java version of this job. 32 32 bit Java job 64 64 bit Java job

Example

Examine the active JVM jobs, ordered by top heap space consumption.

```
SELECT * FROM QSYS2.JVM_INFO
ORDER BY CURRENT_HEAP_SIZE DESC
```

SET_JVM procedure

The SET_JVM procedure can be used to manage specific JVM jobs.

This actions provided by this Db2 for i procedure can also be accomplished interactively using the Work with JVM Jobs (WRKJVMJOB) command.

➤➤ SET_JVM — (— *job_name* — , — *action* —) ➤➤

The schema is QSYS2.

job_name A character or graphic string expression that identifies the qualified job name of the job to change.

action A character or graphic string expression that specifies that action to perform. Supported actions are:

GC_ENABLE_VERBOSE Enable verbose garbage collection detail.

GC_DISABLE_VERBOSE Disable verbose garbage collection detail.

GENERATE_HEAP_DUMP Generates information about the JVM's heap. Generates a dump of all the heap space allocations which have not yet been freed.

GENERATE_SYSTEM_DUMP Generates system detail for the JVM. Generates a binary format raw memory image of the job that was running when the dump was initiated.

GENERATE_JAVA_DUMP Generates Java detail for the JVM. Generates multiple files that contain diagnostic information for the JVM and the Java applications running within the JVM.

Example

- Change a specific web admin JVM to provide verbose garbage collection details:

```
CALL QSYS2.SET_JVM('121376/QWEBADMIN/ADMIN4', 'GC_ENABLE_VERBOSE') ;
```

Journal Services

These services provide information about audit journals and data journals.

ASSOCIATE_JOURNAL_RECEIVER table function

The ASSOCIATE_JOURNAL_RECEIVER table function associates one or more journal receivers with a journal if the journal receiver was originally associated with the journal.

Authorization: The user must have:

- *OBJMGT authority to the journal receiver
- *EXECUTE authority to the library containing the journal receiver
- *OBJMGT and *OBJOPR authority to the journal
- *EXECUTE authority to the library containing the journal

```
➤ ASSOCIATE_JOURNAL_RECEIVER ( journal-library ➤
    JOURNAL_LIBRARY =>
    , journal-name , ➤
    JOURNAL_NAME =>
    , receiver-library , ➤
    RECEIVER_LIBRARY =>
    , receiver-name ) ➤
    RECEIVER_NAME =>
```

The schema is QSYS2.

- journal-library** A character or graphic string expression that identifies the library that contains *journal-name*. Can contain the special values *LIBL or *CURLIB.
- journal-name** A character or graphic string expression that identifies the name of the journal.
- receiver-library** A character or graphic string expression that identifies the name of the library containing *receiver-name*. Can contain the special values *LIBL or *CURLIB.
- receiver-name** A character or graphic string expression that identifies the name of the journal receiver. A generic name can be specified to select specific journal receivers. If the last character in the name is an asterisk, the name is a generic name. For example, to select all journal receivers that start with 'JR', specify a *receiver-name* value of 'JR*'.
The special value of *ALL can be used to select all journal receivers in *receiver-library*. *ALL is not allowed if *receiver-library* is *LIBL.

The result of the function is a table containing rows with the format shown in the following table. All columns are nullable.

Table 125. ASSOCIATE_JOURNAL_RECEIVER table function

Column name	Data type	Description
JOURNAL_LIBRARY	VARCHAR(10)	The library containing the journal that the journal receiver attempted to associate with.
JOURNAL_NAME	VARCHAR(10)	The journal that the journal receiver attempted to associate with.
RECEIVER_LIBRARY	VARCHAR(10)	The library containing the journal receiver.
RECEIVER_NAME	VARCHAR(10)	The journal receiver.
RESULT	VARCHAR(7)	Result type. FAILURE The operation did not complete. SUCCESS The operation completed successfully. WARNING The operation completed but a warning was issued.
RESULT_DETAIL	VARGRAPHIC(100) CCSID 1200	Descriptive text that corresponds to RESULT.

Table 125. ASSOCIATE_JOURNAL_RECEIVER table function (continued)

Column name	Data type	Description
RETURN_CODE	INTEGER	<p>The return code that corresponds to the result.</p> <ul style="list-style-type: none"> 0 The specified journal receiver was successfully associated with the specified journal. 1 The specified journal receiver was already associated with the specified journal. 2 The specified journal receiver could not be associated with the specified journal because the journal receiver was already associated with a different journal. 3 The specified journal receiver could not be associated with the specified journal because the journal receiver was never associated with that journal or could not be associated with that journal in a remote journal network. 4 The specified journal receiver could not be associated with the specified journal because the journal receiver's library could not be found. 5 The specified journal receiver could not be associated with the specified journal because the journal receiver could not be found. 6 The specified journal receiver could not be associated with the specified journal because the journal's library could not be found. 7 The specified journal receiver could not be associated with the specified journal because the journal could not be found. 8 The specified journal receiver could not be associated with the specified journal because the journal receiver could not be locked. 9 The specified journal receiver could not be associated with the specified journal because the journal could not be locked. 10 The specified journal receiver could not be associated with the specified journal because the user is not authorized to the journal receiver's library. 11 The specified journal receiver could not be associated with the specified journal because the user is not authorized to the journal receiver. 12 The specified journal receiver could not be associated with the specified journal because the user is not authorized to the journal's library. 13 The specified journal receiver could not be associated with the specified journal because the user is not authorized to the journal. 14 The specified journal receiver could not be associated with the specified journal because the journal receiver is damaged. 15 The specified journal receiver could not be associated with the specified journal because the journal is damaged. 16 The specified journal receiver could not be associated with the specified journal because the journal receiver and journal are not both in the system or user auxiliary storage pools (ASPs) or the journal receiver and journal are in different independent auxiliary storage pool (IASP) groups.

Example

Re-establish the association of journal APPLIB/APPJRN with journal receiver RCVRLIB/RCV001.

```
SELECT *
FROM TABLE(QSYS2.ASSOCIATE_JOURNAL_RECEIVER('APPLIB', 'APPJRN', 'RCVRLIB', 'RCV001'));
```

Audit journal entry services

These table functions provide detailed information for audit journal entries.

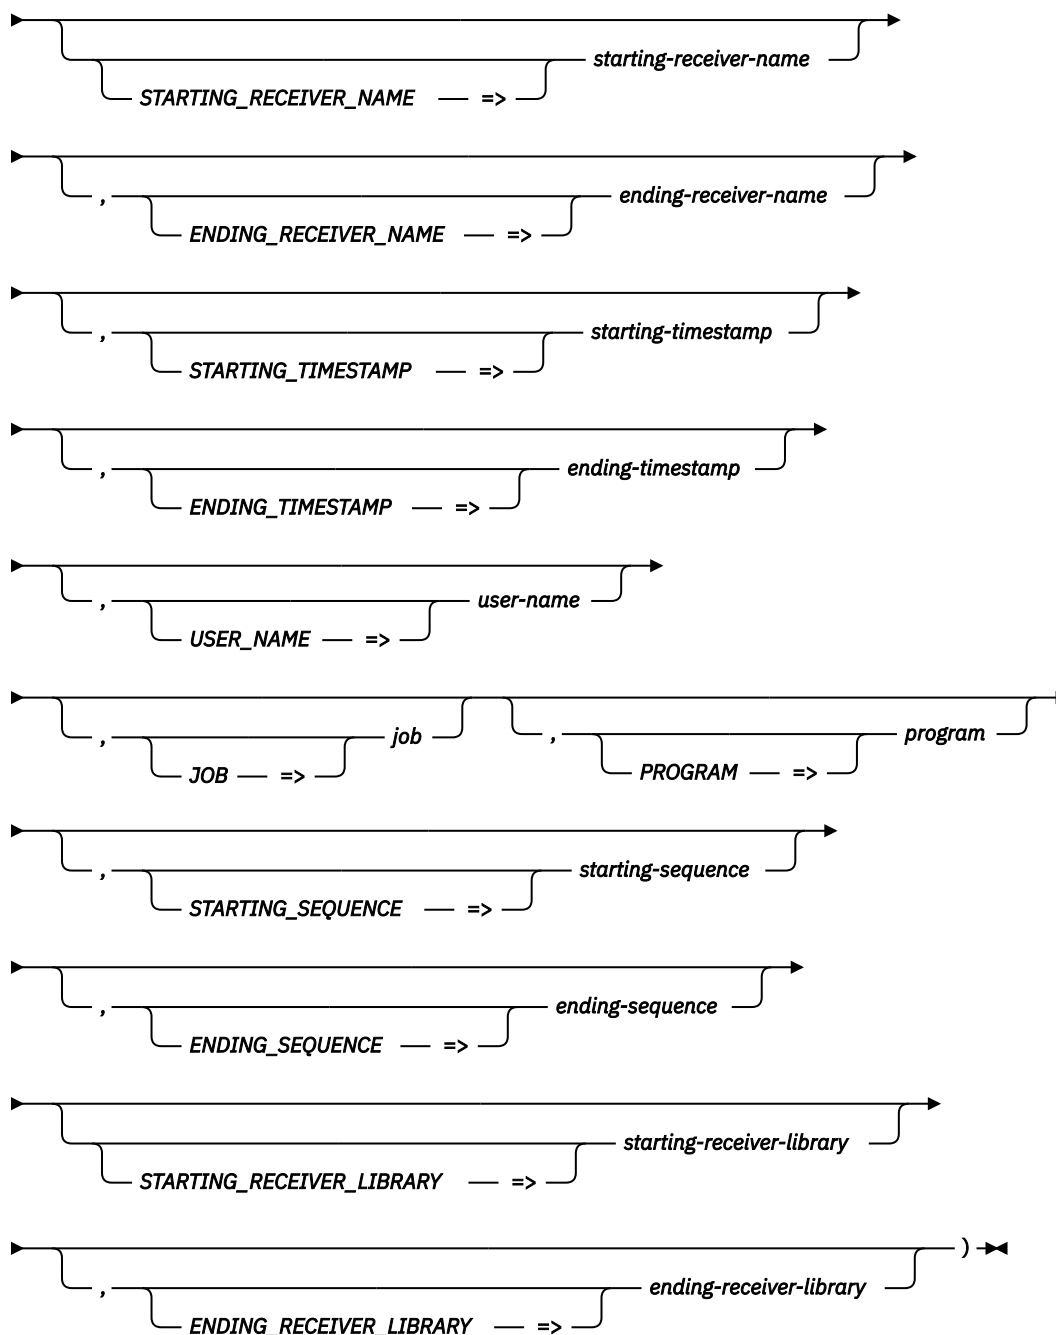
AUDIT JOURNAL table function common information

Each of the AUDIT_JOURNAL_xx table functions returns information from the audit journal specific to an individual entry type. The invocation of every function and the first set of columns returned by the function are identical. This common information is described in this section.

Authorization: The caller must have:

- *USE authority to the audit journal and to all requested journal receivers, and
- *OBJEXIST authority to the audit journal

►► AUDIT_JOURNAL_xx — (→



The schema is SYSTOOLS.

The function name, identified as AUDIT_JOURNAL_xx in the syntax diagram, represents any of the audit journal functions. For example, to invoke the AUDIT_JOURNAL_PW function, replace the xx with PW.

starting-receiver-name	<p>A character or graphic string expression that identifies the name of the starting journal receiver. The special values *CURRENT, *CURCHAIN, *CURAVLCHN, and *CURSEQCHN are supported. Otherwise, <i>starting-receiver-name</i> must identify a valid journal receiver.</p> <p>If no journal receiver is specified, *CURAVLCHN is used.</p>
ending-receiver-name	<p>A character or graphic string expression that identifies the name of the ending journal receiver. The special value *CURRENT is supported. Otherwise, <i>ending-receiver-name</i> must identify a valid journal receiver.</p> <p>If <i>ending-receiver-name</i> is not specified, *CURRENT is used.</p>
starting-timestamp	<p>A timestamp value that specifies the starting timestamp to use¹.</p> <p>A value cannot be specified for both <i>starting-timestamp</i> and <i>starting-sequence</i>.</p> <p>If no starting timestamp is specified, CURRENT DATE - 1 DAY is used.</p>
ending-timestamp	<p>A timestamp value that specifies the ending timestamp to use¹.</p> <p>A value cannot be specified for both <i>ending-timestamp</i> and <i>ending-sequence</i>.</p> <p>If no ending timestamp is specified, CURRENT TIMESTAMP is used.</p>
user-name	<p>A character or graphic string expression that identifies the user profile name for the current user of the job. If <i>user-name</i> is not specified, *ALL is used.</p>
job	<p>A character or graphic string expression that identifies the name of a job. Two forms of a job name are supported:</p> <ol style="list-style-type: none">1. A fully qualified job name in the form <i>job-number/job-user/job-name</i>.2. The first 10 characters are the job name, the second 10 characters are the user name, and the last 6 characters are the job number. <p>If <i>job</i> is not specified, *ALL is used.</p>
program	<p>A character or graphic string expression that identifies the name of a program. If <i>program</i> is not specified, *ALL is used.</p>
starting-sequence	<p>A decimal expression that identifies the starting sequence number to use. This is the value shown in the SEQUENCE_NUMBER column in the QSYS2.DISPLAY_JOURNAL results. If the <i>starting-sequence</i> value is not found in the receiver range, an error is returned.</p> <p>A value cannot be specified for both <i>starting-timestamp</i> and <i>starting-sequence</i>.</p> <p>If no starting sequence is specified, *FIRST is used.</p>
ending-sequence	<p>A decimal expression that identifies the ending sequence number to use. If the <i>ending-sequence</i> value is not found in the receiver range, an error is returned.</p> <p>A value cannot be specified for both <i>ending-timestamp</i> and <i>ending-sequence</i>.</p> <p>If no ending sequence is specified, *LAST is used.</p> <p>When <i>ending-sequence</i> is used, the query results will end when the first ending sequence value is encountered. If the journal has had its sequence numbers reset, <i>ending-sequence</i> will only return results through the first match of <i>ending-sequence</i>.</p>

¹ The accuracy of the entry timestamp stored in journal receivers is only accurate to 16 microseconds. Hence, a value passed as a starting-timestamp and ending-timestamp will be truncated such that the actual timestamps being searched for may be from 0 to 15 microseconds less than the specified value.

starting-receiver-library A character or graphic string expression that identifies the library that contains *starting-receiver-name*.

ending-receiver-library A character or graphic string expression that identifies the library that contains *ending-receiver-name*.

The special values supported for the function arguments are the same as for the Display Journal (DSPJRN) CL command.

The first set of columns returned by each of the audit journal functions are information from the audit journal entry's common header. They have the format shown in the following table. All the columns are nullable.

Table 126. Common columns returned from the audit journal entry header

Column Name	Data Type	Description
ENTRY_TIMESTAMP	TIMESTAMP	The system date and time when the journal entry was added to the journal receiver.
SEQUENCE_NUMBER	DECIMAL(21,0)	A number assigned by the system to each journal entry.
USER_NAME	VARCHAR(10)	The name of the effective user profile under which the thread was running when the entry was created. Contains *NONE if the effective user profile is not available.
QUALIFIED_JOB_NAME	VARCHAR(28)	The qualified job name that added the entry. Contains the null value if the system job is running in a task rather than a process.
JOB_NAME	VARCHAR(10)	The name of the job that added the entry.
JOB_USER	VARCHAR(10)	The user profile that started the job. Contains the null value if the system job is running in a task rather than a process.
JOB_NUMBER	VARCHAR(6)	The job number of the job that added the entry. Contains the null value if the system job is running in a task rather than a process.
THREAD	BIGINT	Identifies the thread within the process that added the journal entry. Contains the null value if the system job is running in a task rather than a process.
PROGRAM_LIBRARY	VARCHAR(10)	The name of the library that contains PROGRAM_NAME. Contains *NONE if PROGRAM_NAME returns *NONE.
PROGRAM_NAME	VARCHAR(10)	The name of the program that caused the journal entry to be added. This can also be the name of a service program or the partial name of a class file used in a compiled Java program. If an application program or CL program did not cause the entry, contains the name of a system-supplied program such as QCMD. Contains *NONE if one of the following conditions is true: <ul style="list-style-type: none"> The program name does not apply to this entry type. The program name was not available.
PROGRAM_LIBRARY_ASP_DEVICE	VARCHAR(10)	The name of the ASP device that contains PROGRAM_NAME.
PROGRAM_LIBRARY_ASP_NUMBER	INTEGER	The number for the auxiliary storage pool that contains PROGRAM_NAME.
REMOTE_PORT	INTEGER	The port number of the remote address associated with this journal entry. Contains the null value if a remote port is not available.
REMOTE_ADDRESS	VARCHAR(46)	The remote address associated with the journal entry in IPv4 or IPv6 format. Contains the null value if a remote address is not available.
SYSTEM_NAME	VARCHAR(8)	The name of the system on which the entry is being retrieved.
SYSTEM_SEQUENCE_NUMBER	DECIMAL(21,0)	The system sequence number which indicates the relative sequence of when this journal entry was deposited into the journal.

Table 126. Common columns returned from the audit journal entry header (continued)

Column Name	Data Type	Description
RECEIVER_LIBRARY	VARCHAR(10)	The name of the library containing RECEIVER_NAME.
RECEIVER_NAME	VARCHAR(10)	The name of the receiver holding the journal entry.
RECEIVER_ASP_DEVICE	VARCHAR(10)	The name of the ASP device containing RECEIVER_NAME.
RECEIVER_ASP_NUMBER	INTEGER	The number for the auxiliary storage pool containing RECEIVER_NAME.
ARM_NUMBER	INTEGER	The number of the disk arm that contains the journal entry.

AUDIT_JOURNAL_AF table function

The AUDIT_JOURNAL_AF table function returns rows from the audit journal that contain information from the AF (Authority Failure) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 127. AUDIT_JOURNAL_AF table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the AF audit journal entry.		
VIOLATION_TYPE	CHAR(1)	The type of authority violation. <ul style="list-style-type: none"> A Not authorized to object B Restricted instruction C Validation failure. See VALIDATION_ERROR_ACTION for additional details. D Use of unsupported interface, object domain failure E Hardware storage protection error, program constant space violation H Scan exit program action. See VALIDATION_ERROR_ACTION for additional details. I System Java inheritance not allowed J Submit job profile error K Special authority violation N Profile token not a regenerable token O Optical Object Authority Failure P Profile swap error R Hardware protection error S Default sign-on attempt T Not authorized to TCP/IP port U User permission request not valid V Profile token not valid for generating new profile token W Profile token not valid for swap X System violation. See OPERATION_VIOLATION_CODE for additional details. Y Not authorized to the current JUID field during a clear JUID operation. Z Not authorized to the current JUID field during a set JUID operation.
VIOLATION_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the violation type.

Table 127. AUDIT_JOURNAL_AF table function (continued)

Column Name	Data Type	Description
VALIDATION_ERROR_ACTION	CHAR(1)	<p>Action taken after validation error detected, set when VIOLATION_TYPE is C or H.</p> <p>A The translation of the object was not attempted or it failed. The QALWBJRST system value setting allowed the object to be restored. The user doing the restore did not have *ALLOBJ special authority and the system security level is set to 10, 20, or 30. Therefore, all authorities to the object were retained.</p> <p>B The translation of the object was not attempted or it failed. The QALWBJRST system value setting allowed the object to be restored. The user doing the restore did not have *ALLOBJ special authority and the system security level is set to 40 or above. Therefore, all authorities to the object were revoked.</p> <p>C The translation of the object was successful. The translated copy was restored on the system.</p> <p>D The translation of the object was not attempted or it failed. The QALWBJRST system value setting allowed the object to be restored. The user doing the restore had *ALLOBJ special authority. Therefore, all authorities to the object were retained.</p> <p>E System install time error detected.</p> <p>F The object was not restored because the signature is not IBM i format.</p> <p>G Unsigned system or inherit state object found when checking system.</p> <p>H Unsigned user state object found when checking system.</p> <p>I Mismatch between object and its signature found when checking system.</p> <p>J IBM certificate not found when checking system.</p> <p>K Invalid signature format found when checking system.</p> <p>M Scan exit program modified the object that was scanned.</p> <p>X Scan exit program wanted object marked as having a scan failure.</p> <p>Contains the null value if VIOLATION_TYPE is not C or H.</p>
VALIDATION_ERROR_ACTION_DETAIL	VARCHAR(200)	<p>Descriptive text that corresponds to the violation error action.</p> <p>Contains the null value if VIOLATION_TYPE is not C or H.</p>
OPERATION_VIOLATION_CODE	CHAR(3)	<p>The type of operation violation that occurred, set when VIOLATION_TYPE is X.</p> <p>AAC Not authorized to use SST Advanced Analysis Command.</p> <p>HCA Service tool user profile not authorized to perform hardware configuration operation (QYHCHCOP).</p> <p>LIC LIC indicates that a Licensed Internal Code fix was not applied because of a signature violation.</p> <p>SFA Not authorized to activate the environment attribute for system file access.</p> <p>CMD An attempt was made to use a command that has been disabled by a system administrator.</p> <p>Contains the null value if the VIOLATION_TYPE is not X.</p>
OBJECT_LIBRARY	VARCHAR(10)	<p>The name of the library containing the object.</p> <p>When VIOLATION_TYPE is K, contains the name of the program's library or the command's library that detected the error.</p> <p>Contains the null value if there is no library name.</p>

Table 127. AUDIT_JOURNAL_AF table function (continued)

Column Name	Data Type	Description
OBJECT_NAME	VARCHAR(10)	<p>The name of the object.</p> <p>When VIOLATION_TYPE is K, contains the name of the command or program that detected the error. If the command has several alternative names, the command name in the audit record might not match the specific command name used but will be one of the equivalent alternatives. A special value of *INSTR indicates that a machine instruction detected the error.</p> <p>When OBJECT_TYPE is *LIC, contains a Licensed Internal Code replacement unit (Ru) name.</p> <p>Contains the null value if there is no object name.</p>
OBJECT_TYPE	VARCHAR(7)	<p>The type of the object.</p> <p>When VIOLATION_TYPE is K, contains the object type of the command or program that detected the error.</p> <p>When VIOLATION_TYPE is G, contains the name of the *SRVPGM that contained the exit that detected the error.</p> <p>Contains the null value if there is no object type.</p>
OBJECT_ASP_NAME	VARCHAR(10)	<p>The name of the auxiliary storage pool (ASP) in which the object resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.</p>
OBJECT_ASP_NUMBER	INTEGER	<p>The number of the ASP device.</p>
FIELD_NAME	VARCHAR(10)	<p>The system name of the column.</p> <p>Contains the null value if the authority is not related to a column or the column name is not available.</p>
TCPIP_PORT	INTEGER	<p>The TCP/IP port the user is not authorized to use, when VIOLATION_TYPE is T.</p> <p>Contains the null value if VIOLATION_TYPE is not T.</p>
API_NAME	VARCHAR(20)	<p>The full API name of the API or exit point name that detected the error, when VIOLATION_TYPE is K.</p> <p>Contains the null value when the VIOLATION_TYPE is not K or if there is no API or exit point information.</p>
PTF_NUMBER	CHAR(7)	<p>The PTF number that failed to apply because of a signature violation when the VIOLATION_TYPE is X and OPERATION_VIOLATION_CODE is LIC.</p> <p>Contains the null value if VIOLATION_TYPE is not X with an OPERATION_VIOLATION_CODE of LIC.</p>
AAC_NAME	VARCHAR(30)	<p>The Advanced Analysis Command name, when the VIOLATION_TYPE is X and the OPERATION_VIOLATION_CODE is AAC.</p> <p>Contains the null value if VIOLATION_TYPE is not X with an OPERATION_VIOLATION_CODE of AAC.</p>
USER_PROFILE_NAME	VARCHAR(10)	<p>The name of the user that caused the authority failure.</p> <p>Contains the null value if the user name is not available.</p>
WORKSTATION_NAME	VARCHAR(10)	<p>The name of the workstation or workstation type.</p> <p>Contains the null value if the workstation name is not available.</p>
PROGRAM_INSTRUCTION	INTEGER	<p>The instruction number of the program.</p> <p>Contains the null value if the instruction number is not available.</p>
PATH_NAME	VARGRAPHIC(5000) CCSID 1200	<p>The path name of the object.</p> <p>Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
PATH_NAME_INDICATOR	VARCHAR(3)	<p>Path name indicator.</p> <p>NO The PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name.</p> <p>YES The PATH_NAME column contains complete absolute path name for the object.</p> <p>Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>

Table 127. AUDIT_JOURNAL_AF table function (continued)

Column Name	Data Type	Description
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	When PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PATH_NAME column. Contains the null value when PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	The name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OBJECT_FILE_ID	BINARY(16)	The file ID of the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PARENT_FILE_ID	BINARY(16)	The file ID of the parent directory. Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OFFICE_USER	VARCHAR(10)	The name of the office user. Contains the null value if there is no office user.
OFFICE_ON_BEHALF_OF_USER	VARCHAR(10)	User working on behalf of another user. Contains the null value if the user name is not available.
DLO_NAME	VARCHAR(12)	The name of the document library object. Contains the null value if there is no document library object.
FOLDER_PATH	VARCHAR(63)	The path of the folder. Contains the null value if there is no folder path.

Example

- Find any authority failures for Integrated File System (IFS) objects in the past 24 hours.

```
SELECT * FROM TABLE(
  SYSTOOLS.AUDIT_JOURNAL_AF(
    STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 1 DAY
  )
) WHERE PATH_NAME IS NOT NULL;
```

- Determine the number of 'Not authorized to object' authority failures for user BOB in the last week.

```
SELECT COUNT(*) FROM TABLE(
  SYSTOOLS.AUDIT_JOURNAL_AF(
    STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 7 DAYS,
    USER_NAME => 'BOB'
  )
) WHERE VIOLATION_TYPE = 'A';
```

AUDIT_JOURNAL_CA table function

The AUDIT_JOURNAL_CA table function returns rows from the audit journal that contain information from the CA (Authority Changes) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 128. AUDIT_JOURNAL_CA table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the CA audit journal entry.		

Table 128. AUDIT_JOURNAL_CA table function (continued)

Column Name	Data Type	Description								
OBJECT_LIBRARY	VARCHAR(10)	The name of the library containing the object. Contains the null value if there is no library name.								
OBJECT_NAME	VARCHAR(10)	The name of the object. Contains the null value if there is no object name.								
OBJECT_TYPE	VARCHAR(7)	The type of the object. Contains the null value if there is no object type.								
OBJECT_ATTRIBUTE	VARCHAR(10)	The attribute for the object type. Contains the null value if there is no object attribute.								
OBJECT_ASP_NAME	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which OBJECT_NAME resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.								
OBJECT_ASP_NUMBER	INTEGER	The number of the auxiliary storage pool to which storage for OBJECT_NAME is allocated.								
FIELD_NAME	VARCHAR(10)	The system name of the column. Contains the null value if the authority is not related to a column or the column name is not available.								
COMMAND_TYPE	VARCHAR(9)	The type of command used. <table border="0"> <tr> <td>GRANT</td> <td>Grant</td> </tr> <tr> <td>GRTUSRAUT</td> <td>GRTUSRAUT operation</td> </tr> <tr> <td>REPLACE</td> <td>Grant with replace</td> </tr> <tr> <td>REVOKE</td> <td>Revoke</td> </tr> </table>	GRANT	Grant	GRTUSRAUT	GRTUSRAUT operation	REPLACE	Grant with replace	REVOKE	Revoke
GRANT	Grant									
GRTUSRAUT	GRTUSRAUT operation									
REPLACE	Grant with replace									
REVOKE	Revoke									
USER_PROFILE_NAME	VARCHAR(10)	The name of the user profile whose authority is being granted or revoked. Contains the null value if a user profile is not being changed.								
AUTHORIZATION_LIST_MANAGEMENT	VARCHAR(3)	Indicates whether authorization list management (*AUTLMGT) authority is granted or revoked. <table border="0"> <tr> <td>NO</td> <td>The authority is not changed.</td> </tr> <tr> <td>YES</td> <td>The authority might have changed.</td> </tr> </table>	NO	The authority is not changed.	YES	The authority might have changed.				
NO	The authority is not changed.									
YES	The authority might have changed.									
OBJECT_EXCLUDE	VARCHAR(3)	Indicates whether exclude (*EXCLUDE) authority is granted or revoked. <table border="0"> <tr> <td>NO</td> <td>The authority is not changed.</td> </tr> <tr> <td>YES</td> <td>The authority might have changed.</td> </tr> </table>	NO	The authority is not changed.	YES	The authority might have changed.				
NO	The authority is not changed.									
YES	The authority might have changed.									
OBJECT_OPERATIONAL	VARCHAR(3)	Indicates whether object operational (*OBJOPR) authority is granted or revoked. <table border="0"> <tr> <td>NO</td> <td>The authority is not changed.</td> </tr> <tr> <td>YES</td> <td>The authority might have changed.</td> </tr> </table>	NO	The authority is not changed.	YES	The authority might have changed.				
NO	The authority is not changed.									
YES	The authority might have changed.									
OBJECT_MANAGEMENT	VARCHAR(3)	Indicates whether object management (*OBJMGT) authority is granted or revoked. <table border="0"> <tr> <td>NO</td> <td>The authority is not changed.</td> </tr> <tr> <td>YES</td> <td>The authority might have changed.</td> </tr> </table>	NO	The authority is not changed.	YES	The authority might have changed.				
NO	The authority is not changed.									
YES	The authority might have changed.									
OBJECT_EXISTENCE	VARCHAR(3)	Indicates whether object existence (*OBJEXIST) authority is granted or revoked. <table border="0"> <tr> <td>NO</td> <td>The authority is not changed.</td> </tr> <tr> <td>YES</td> <td>The authority might have changed.</td> </tr> </table>	NO	The authority is not changed.	YES	The authority might have changed.				
NO	The authority is not changed.									
YES	The authority might have changed.									
OBJECT_ALTER	VARCHAR(3)	Indicates whether object alter (*OBJALTER) authority is granted or revoked. <table border="0"> <tr> <td>NO</td> <td>The authority is not changed.</td> </tr> <tr> <td>YES</td> <td>The authority might have changed.</td> </tr> </table>	NO	The authority is not changed.	YES	The authority might have changed.				
NO	The authority is not changed.									
YES	The authority might have changed.									

Table 128. AUDIT_JOURNAL_CA table function (continued)

Column Name	Data Type	Description
OBJECT_REFERENCE	VARCHAR(3)	Indicates whether object reference (*OBJREF) authority is granted or revoked. NO The authority is not changed. YES The authority might have changed.
DATA_READ	VARCHAR(3)	Indicates whether data read (*READ) authority is granted or revoked. NO The authority is not changed. YES The authority might have changed.
DATA_ADD	VARCHAR(3)	Indicates whether data add (*ADD) authority is granted or revoked. NO The authority is not changed. YES The authority might have changed.
DATA_UPDATE	VARCHAR(3)	Indicates whether data update (*UPD) authority is granted or revoked. NO The authority is not changed. YES The authority might have changed.
DATA_DELETE	VARCHAR(3)	Indicates whether data delete (*DLT) authority is granted or revoked. NO The authority is not changed. YES The authority might have changed.
DATA_EXECUTE	VARCHAR(3)	Indicates whether data execute (*EXECUTE) authority is granted or revoked. NO The authority is not changed. YES The authority might have changed.
PREV_AUTHORIZATION_LIST_MANAGEMENT	VARCHAR(3)	Indicates whether the user previously had authorization list management (*AUTLMGT) authority. NO The user did not have this authority. YES The user had this authority. Contains the null value if the AUTHORIZATION_LIST_MANAGEMENT column is NO.
PREV_OBJECT_EXCLUDE	VARCHAR(3)	Indicates whether the user previously had exclude (*EXCLUDE) authority. NO The user did not have this authority. YES The user had this authority. Contains the null value if the OBJECT_EXCLUDE column is NO.
PREV_OBJECT_OPERATIONAL	VARCHAR(3)	Indicates whether the user previously had object operational (*OBJOPR) authority. NO The user did not have this authority. YES The user had this authority. Contains the null value if the OBJECT_OPERATIONAL column is NO.
PREV_OBJECT_MANAGEMENT	VARCHAR(3)	Indicates whether the user previously had object management (*OBJMGT) authority. NO The user did not have this authority. YES The user had this authority. Contains the null value if the OBJECT_MANAGEMENT column is NO.

Table 128. AUDIT_JOURNAL_CA table function (continued)

Column Name	Data Type	Description
PREV_OBJECT_EXISTENCE	VARCHAR(3)	<p>Indicates whether the user previously had object existence (*OBJEXIST) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the OBJECT_EXISTENCE column is NO.</p>
PREV_OBJECT_ALTER	VARCHAR(3)	<p>Indicates whether the user previously had object alter (*OBJALTER) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the OBJECT_ALTER column is NO.</p>
PREV_OBJECT_REFERENCE	VARCHAR(3)	<p>Indicates whether the user previously had object reference (*OBJREF) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the OBJECT_REFERENCE column is NO.</p>
PREV_DATA_READ	VARCHAR(3)	<p>Indicates whether the user previously had data read (*READ) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the DATA_READ column is NO.</p>
PREV_DATA_ADD	VARCHAR(3)	<p>Indicates whether the user previously had data add (*ADD) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the DATA_ADD column is NO.</p>
PREV_DATA_UPDATE	VARCHAR(3)	<p>Indicates whether the user previously had data update (*UPD) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the DATA_UPDATE column is NO.</p>
PREV_DATA_DELETE	VARCHAR(3)	<p>Indicates whether the user previously had data delete (*DLT) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the DATA_DELETE column is NO.</p>
PREV_DATA_EXECUTE	VARCHAR(3)	<p>Indicates whether the user previously had data execute (*EXECUTE) authority.</p> <p>NO The user did not have this authority.</p> <p>YES The user had this authority.</p> <p>Contains the null value if the DATA_EXECUTE column is NO.</p>
AUTHORIZATION_LIST	VARCHAR(10)	<p>The name of the authorization list that is being modified.</p> <p>Contains the null value if an authorization list is not being changed.</p>
AUTHORIZATION_LIST_PUBLIC	VARCHAR(3)	<p>Indicates whether authorization list (*AUTL public authority) authority has been granted or revoked.</p> <p>NO The authority is not changed.</p> <p>YES The authority might have changed.</p>
PREV_AUTHORIZATION_LIST	VARCHAR(10)	<p>The name of the previous authorization list.</p> <p>Contains the null value if an authorization list is not being changed.</p>

Table 128. AUDIT_JOURNAL_CA table function (continued)

Column Name	Data Type	Description
PREV_AUTHORIZATION_LIST_PUBLIC	VARCHAR(3)	Indicates whether the user previously had authorization list (*AUTL public authority) authority. NO The user did not have this authority. YES The user had this authority. Contains the null value if the AUTHORIZATION_LIST_PUBLIC column is NO.
PERSONAL_STATUS_CHANGED	VARCHAR(3)	The personal status changed. NO The status did not changed. YES The status changed.
ACCESS_CODE_CHANGED	VARCHAR(6)	Whether the access code changed. ADD The access code was added. REMOVE The access code was removed. Contains the null value if the access code was not changed.
ACCESS_CODE	VARCHAR(4)	Access code. Contains the null value if the access code was not changed.
PATH_NAME	VARGRAPHIC(5000) CCSID 1200	The path name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PATH_NAME_INDICATOR	VARCHAR(3)	Path name indicator. NO The PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name. YES The PATH_NAME column contains complete absolute path name for the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	When PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PATH_NAME column. Contains the null value when PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	The name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OBJECT_FILE_ID	BINARY(16)	The file ID of the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PARENT_FILE_ID	BINARY(16)	The file ID of the parent directory. Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OFFICE_USER	VARCHAR(10)	The name of the office user. Contains the null value if there is no office user.
OFFICE_ON_BEHALF_OF_USER	VARCHAR(10)	User working on behalf of another user. Contains the null value if the user name is not available.
DLO_NAME	VARCHAR(12)	The name of the document library object. Contains the null value if there is no document library object.
FOLDER_PATH	VARCHAR(63)	The path of the folder. Contains the null value if there is no folder path.

Example

- List any files in APPLIB1 that had the ability to change data granted to them in the last week.

```
SELECT OBJECT_NAME, USER_PROFILE_NAME FROM TABLE(
  SYSTOOLS.AUDIT_JOURNAL_CA(
    STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 7 DAYS)
)
WHERE COMMAND_TYPE = 'GRANT' AND
      OBJECT_LIBRARY = 'APPLIB1' AND
      OBJECT_TYPE = '*FILE' AND
      (DATA_ADD = 'YES' OR DATA_UPDATE = 'YES' OR DATA_DELETE = 'YES');
```

AUDIT_JOURNAL_CD table function

The AUDIT_JOURNAL_CD table function returns rows from the audit journal that contain information from the CD (Command String) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 129. AUDIT_JOURNAL_CD table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the CD audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. <ul style="list-style-type: none"> C Command run L OCL statement O Operator control command P S/36 procedure S Command run after command substitution took place U Utility control statement X Proxy command
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
OBJECT_LIBRARY	VARCHAR(10)	The name of the library where the object is stored.
OBJECT_NAME	VARCHAR(10)	The name of the object.
OBJECT_TYPE	VARCHAR(7)	The type of object.
OBJECT_ASP_NAME	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which the object resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.
OBJECT_ASP_NUMBER	INTEGER	The number of the ASP device.
WHERE_RUN	CHAR(1)	Where the CL command was run. <ul style="list-style-type: none"> B In a batch job but not for any of the reason listed under Y, R, or E. Typical case would be that the CL command was run using STRDBRDR or SBMDBJOB command or was specified on the CMD parameter of the SBMJOB command. E The command string was passed as a parameter to one of the Command Analyzer APIs: QCMDEXC, QCAPCMD, or QCAEXEC N Interactively from a command line or by choosing a menu option that runs a CL command R From a REXX procedure Y From a compiled OPM CL program or an ILE CL program
WHERE_RUN_DETAIL	VARCHAR(200)	Descriptive text that corresponds to where the CL command was run.

Table 129. AUDIT_JOURNAL_CD table function (continued)

Column Name	Data Type	Description
COMMAND_STRING	VARCHAR(6000)	The command that was run, with parameters. Contains the null value if no command was run.

Example

- List any change commands that were run this week.

```
SELECT *
FROM TABLE (
  SYSTOOLS.AUDIT_JOURNAL_CD (STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 7 DAYS)
)
WHERE ENTRY_TYPE = 'C' AND
      OBJECT_LIBRARY = 'QSYS' AND
      OBJECT_NAME LIKE 'CHG%' ;
```

AUDIT_JOURNAL_CO table function

The AUDIT_JOURNAL_CO table function returns rows from the audit journal that contain information from the CO (Create Object) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 130. AUDIT_JOURNAL_CO table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the CO audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. N Create of new object R Replacement of existing object
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
OBJECT_LIBRARY	VARCHAR(10)	The name of the library containing the object. Contains the null value if there is no library name.
OBJECT_NAME	VARCHAR(10)	The name of the object. Contains the null value if there is no object name.
OBJECT_TYPE	VARCHAR(7)	The type of the object. Contains the null value if there is no object type.
OBJECT_ATTRIBUTE	VARCHAR(10)	The attribute for the object type. Contains the null value if there is no object attribute.
OBJECT_ASP_NAME	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which the object resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.
OBJECT_ASP_NUMBER	INTEGER	The number of the ASP device.
PATH_NAME	VARGRAPHIC(5000) CCSID 1200	The path name of the object. Contains the null value if the path name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.

Table 130. AUDIT_JOURNAL_CO table function (continued)

Column Name	Data Type	Description
PATH_NAME_INDICATOR	VARCHAR(3)	Path name indicator. NO The PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name. YES The PATH_NAME column contains complete absolute path name for the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	When PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PATH_NAME column. Contains the null value when PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	The name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OBJECT_FILE_ID	BINARY(16)	The file ID of the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PARENT_FILE_ID	BINARY(16)	The file ID of the parent directory. Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OFFICE_USER	VARCHAR(10)	The name of the office user. Contains the null value if there is no office user.
OFFICE_ON_BEHALF_OF_USER	VARCHAR(10)	User working on behalf of another user. Contains the null value if the user name is not available.
DLO_NAME	VARCHAR(12)	The name of the document library object. Contains the null value if there is no document library object.
FOLDER_PATH	VARCHAR(63)	The path of the folder. Contains the null value if there is no folder path.

Example

- Find the objects created in APPLIB in the last 2 months.

```
SELECT OBJECT_LIBRARY, OBJECT_NAME, OBJECT_TYPE FROM TABLE(
  SYSTOOLS.AUDIT_JOURNAL_CO(
    STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 2 MONTHS
  )
) WHERE OBJECT_LIBRARY = 'APPLIB';
```

AUDIT_JOURNAL_CP table function

The AUDIT_JOURNAL_CP table function returns rows from the audit journal that contain information from the CP (User Profile Changes) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 131. AUDIT_JOURNAL_CP table function

Column Name	Data Type	Description
<p>The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the CP audit journal entry.</p>		
ENTRY_TYPE	CHAR(1)	The type of entry. A Change to a user profile
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
USER_PROFILE	VARCHAR(10)	The user profile that was changed.
COMMAND_TYPE	CHAR(3)	The type of command used. CHG Change User Profile (CHGUSRPRF) or Change Expiration Scd Entry (CHGEXPSCDE) commands CRT Create User Profile (CRTUSRPRF) command DST QSECOFR password reset using DST RPA Reset Profile Attributes (QSYRESPA) API RST Restore User Profile (RSTUSRPRF) command SQL QSYS2/SET_SERVER_SBS_ROUTING() Db2® for i procedure
STATUS	VARCHAR(9)	User profile status. *DISABLED The user profile cannot be used *ENABLED The user profile is valid Contains the null value if the value was not changed.
PASSWORD_CHANGED	VARCHAR(3)	Whether the password changed. YES Password changed Contains the null value if the value was not changed.
NO_PASSWORD_INDICATOR	VARCHAR(3)	Whether the password is *NONE. YES Password is *NONE Contains the null value if the value was not changed.
PASSWORD_EXPIRED	VARCHAR(4)	Whether the password is expired. *NO Password is not expired *YES Password is expired Contains the null value if the value was not changed.
LOCAL_PASSWORD_MANAGEMENT	VARCHAR(4)	Specifies whether the user profile password should be managed locally. *NO The password is not managed on the local system. *YES The password is managed on the local system. Contains the null value if the value was not changed.
PASSWORD_CONFORMANCE	VARCHAR(8)	Indicates whether the new password conforms to the password composition rules. *EXITPGM Checked but does not conform because of an exit program response. *NOCHECK Not checked; password was changed. *NONE Not checked; *NONE was specified for the new password. *PASSED Checked and conforms. *SYSVAL Checked but does not conform because of a system value based rule. Contains the null value if PASSWORD_CHANGED is not YES.

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
BLOCK_PASSWORD_CHANGE	VARCHAR(7)	<p>Specifies the value that the block password change has been changed to.</p> <p>1-99 Blocked hours.</p> <p>*NONE No block period.</p> <p>*SYSVAL The system value QPWDCHGBLK is used.</p> <p>Contains the null value if the value was not changed.</p>
PASSWORD_EXPIRATION_INTERVAL	VARCHAR(7)	<p>Specifies the value that the password expiration interval has been changed to.</p> <p>1-366 The size of the expiration interval in days.</p> <p>*NOMAX No expiration interval.</p> <p>*SYSVAL The system value QPWDEXPITV is used.</p> <p>Contains the null value if the value was not changed.</p>
USER_EXPIRATION_DATE	VARCHAR(10)	<p>Specifies the date when the user profile expires. The user profile is automatically disabled or deleted on this date. Can contain the special value *NONE to indicate the user profile does not expire.</p> <p>Contains the null value if the value was not changed.</p>
USER_EXPIRATION_ACTION	CHAR(3)	<p>The action performed on the profile when it expires. This value is always DSB when using the CRTUSRPRF and CHGUSRPRF commands. When using the CHGEXPSCDE command, this value is one of the following values.</p> <p>DLT The profile is deleted when it expires.</p> <p>DSB The profile is disabled when it expires.</p> <p>Contains the null value if the value was not changed.</p>
OWNED_OBJECT_OPTION	VARCHAR(7)	<p>The type of operation performed on the objects owned by the expiring profile. The owned object option value is specified on the OWNBOBJOPT parameter of the CHGEXPSCDE ACTION(*DELETE) command.</p> <p>*CHGOWN The owned objects for the user profile have ownership transferred to the new owner user profile. The user profile is deleted if the transfer of all owned objects is successful.</p> <p>*DLT The owned objects for the user profile are deleted. The user profile is deleted if the deletion of all owned objects is successful.</p> <p>*NODLT The owned objects for the user profile are not changed, and the user profile is not deleted if the user owns any objects.</p> <p>Contains the null value if USER_EXPIRATION_ACTION is not DLT or if the value was not changed.</p>
OWNED_OBJECT_OPTION_NEW_OWNER	VARCHAR(10)	<p>The profile that will own all of the objects owned by the expiring profile.</p> <p>Contains the null value if OWNED_OBJECT_OPTION is not *CHGOWN or if the value was not changed.</p>
PRIMARY_GROUP_OPTION	VARCHAR(7)	<p>The type of operation performed on the objects that have the expiring user profile as their primary group. The primary group option value is specified on the PGPOPT parameter of the CHGEXPSCDE ACTION(*DELETE) command.</p> <p>*CHGPGP The objects the user profile is the primary group for are transferred to the new primary group user profile. The user profile is deleted if the transfer of all objects is successful.</p> <p>*NOCHG The objects the user profile is the primary group for do not change, and the user profile is not deleted if the user is the primary group for any objects.</p> <p>Contains the null value if USER_EXPIRATION_ACTION is not DLT or if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
NEW_PRIMARY_GROUP	VARCHAR(10)	<p>The profile that will become the new primary group of the objects for which the expiring profile is the primary group. Can contain the following special value:</p> <p>*NONE All of the objects for which the expiring user is the primary group will no longer have a primary group.</p> <p>Contains the null value when PRIMARY_GROUP_OPTION is not *CHGPGP or if the value was not changed.</p>
NEW_PRIMARY_GROUP_AUTHORITY	VARCHAR(8)	<p>The authority the new primary group has to the object.</p> <p>*ALL The new primary group has *ALL authority to the object.</p> <p>*CHANGE The new primary group has *CHANGE authority to the object.</p> <p>*EXCLUDE The new primary group has *EXCLUDE authority to the object.</p> <p>*OLDPGP The new primary group has the same authority to the object as the old primary group.</p> <p>*PRIVATE The new primary group has the same authority to the object as its private authority to the object was.</p> <p>*USE The new primary group has *USE authority to the object.</p> <p>Contains the null value when PRIMARY_GROUP_OPTION is not *CHGPGP or NEW_PRIMARY_GROUP is *NONE or if the value was not changed.</p>
USER_CLASS_NAME	VARCHAR(7)	<p>The user class of the user.</p> <p>*PGMR Programmer</p> <p>*SECADM Security administrator</p> <p>*SECOFR Security officer</p> <p>*SYSOPR System operator</p> <p>*USER User</p> <p>Contains the null value if the value was not changed.</p>
SPECIAL_AUTHORITIES	VARCHAR(72)	<p>Current list of special authorities. This is a single string containing the list of special authorities separated by one blank.</p> <p>Contains the null value if the value was not changed.</p>
ALLOBJ	VARCHAR(3)	<p>Current *ALLOBJ special authority</p> <p>YES Profile has *ALLOBJ special authority</p> <p>Contains the null value if the value was not changed.</p>
JOBCTL	VARCHAR(3)	<p>Current *JOBCTL special authority</p> <p>YES Profile has *JOBCTL special authority</p> <p>Contains the null value if the value was not changed.</p>
SAVSYS	VARCHAR(3)	<p>Current *SAVSYS special authority</p> <p>YES Profile has *SAVSYS special authority</p> <p>Contains the null value if the value was not changed.</p>
SECADM	VARCHAR(3)	<p>Current *SECADM special authority</p> <p>YES Profile has *SECADM special authority</p> <p>Contains the null value if the value was not changed.</p>
SPLCTL	VARCHAR(3)	<p>Current *SPLCTL special authority</p> <p>YES Profile has *SPLCTL special authority</p> <p>Contains the null value if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
SERVICE	VARCHAR(3)	Current *SERVICE special authority YES Profile has *SERVICE special authority Contains the null value if the value was not changed.
AUDIT	VARCHAR(3)	Current *AUDIT special authority YES Profile has *AUDIT special authority Contains the null value if the value was not changed.
IOSYSCFG	VARCHAR(3)	Current *IOSYSCFG special authority YES Profile has *IOSYSCFG special authority Contains the null value if the value was not changed.
PREVIOUS_SPECIAL_AUTHORITIES	VARCHAR(72)	Previous list of special authorities. This is a single string containing the list of special authorities separated by one blank. Contains the null value if the value was not changed.
PREVIOUS_ALLOBJ	VARCHAR(3)	Previous *ALLOBJ special authority YES Profile had *ALLOBJ special authority Contains the null value if the value was not changed.
PREVIOUS_JOBCTL	VARCHAR(3)	Previous *JOBCTL special authority YES Profile had *JOBCTL special authority Contains the null value if the value was not changed.
PREVIOUS_SAVSYS	VARCHAR(3)	Previous *SAVSYS special authority YES Profile had *SAVSYS special authority Contains the null value if the value was not changed.
PREVIOUS_SECADM	VARCHAR(3)	Previous *SECADM special authority YES Profile had *SECADM special authority Contains the null value if the value was not changed.
PREVIOUS_SPLCTL	VARCHAR(3)	Previous *SPLCTL special authority YES Profile had *SPLCTL special authority Contains the null value if the value was not changed.
PREVIOUS_SERVICE	VARCHAR(3)	Previous *SERVICE special authority YES Profile had *SERVICE special authority Contains the null value if the value was not changed.
PREVIOUS_AUDIT	VARCHAR(3)	Previous *AUDIT special authority YES Profile had *AUDIT special authority Contains the null value if the value was not changed.
PREVIOUS_IOSYSCFG	VARCHAR(3)	Previous *IOSYSCFG special authority YES Profile had *IOSYSCFG special authority Contains the null value if the value was not changed.
GROUP_PROFILE_NAME	VARCHAR(10)	The user's group profile. Can contain the special value *NONE to indicate the user is not a member of any group profiles. Contains the null value if the value was not changed.

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
GROUP_OWNER	VARCHAR(7)	Owner of objects created as a member of a group profile. *GRPPRF The user's group profile becomes the owner *USRPRF The user's user profile becomes the owner Contains the null value if the value was not changed.
GROUP_AUTHORITY	VARCHAR(8)	The authority the user's group profile has to objects the user creates. *ALL ALL authority *CHANGE CHANGE authority *EXCLUDE The group is denied access *NONE No authority is granted to the group profile *USE USE authority Contains the null value if the value was not changed.
GROUP_AUTHORITY_TYPE	VARCHAR(8)	The value of the group authority type parameter. *PGP The group becomes the object's primary group and is given the authority specified in GROUP_AUTHORITY. *PRIVATE The authority defined in the GROUP_AUTHORITY is assigned as private authority to the group profile. Contains the null value if the value was not changed.
SUPPLEMENTAL_GROUP_LIST	VARCHAR(150)	The names of up to 15 supplemental group profiles for the user. Can contain the special value *NONE to indicate the user has no supplemental groups. Each entry except for the last one is padded with blanks to fill 10 characters. Contains the null value if the value was not changed.
INITIAL_PROGRAM_LIBRARY	VARCHAR(10)	The library where the initial program is found. Can contain the special value *LIBL. Contains the null value if INITIAL_PROGRAM is *NONE or if the value was not changed.
INITIAL_PROGRAM	VARCHAR(10)	The user's initial program. Can contain the special value *NONE to indicate the user has no initial program. Contains the null value if the value was not changed.
INITIAL_MENU_LIBRARY	VARCHAR(10)	The library where the initial menu is found. Can contain the special value *LIBL. Contains the null value if INITIAL_MENU is *SIGNOFF or if the value was not changed.
INITIAL_MENU	VARCHAR(10)	The user's initial menu. Can contain the special value *SIGNOFF to indicate the user is limited to running the initial program specified for this profile. Contains the null value if the value was not changed.
CURRENT_LIBRARY_NAME	VARCHAR(10)	The user's current library. Can contain the special value *CRTDFT to indicate the user has no current library. Contains the null value if the value was not changed.
HOME_DIRECTORY	VARGRAPHIC(5000) CCSID 1200	Path name of the home directory or the following special value: *USRPRF The home directory assigned to the user will be /home/USRPRF, where USRPRF is the name of the user profile. Contains the null value if the value was not changed.

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
LOCALE_PATH_NAME	VARGRAPHIC(5000) CCSID 1200	<p>Path name of the locale or one of the following special values:</p> <ul style="list-style-type: none"> *C The C locale path name is assigned to this user. *NONE No locale path name is assigned to this user. *POSIX The POSIX locale path name is assigned to this user. *SYSVAL The system value, QLOCALE, is used to determine the locale path name to be assigned to this user. <p>Contains the null value if the value was not changed.</p>
LIMIT_CAPABILITIES	VARCHAR(8)	<p>The value of limited capabilities parameter.</p> <ul style="list-style-type: none"> *NO No limitations. *PARTIAL The user can change the initial menu, but cannot change the initial program, current library, or attention key handling program. The user can run commands from command lines. *YES The user cannot change the initial program, initial menu, current library, and attention key handling programs. The user cannot run commands from command lines. <p>Contains the null value if the value was not changed.</p>
ASSISTANCE_LEVEL	VARCHAR(9)	<p>The user interface that will be used.</p> <ul style="list-style-type: none"> *ADVANCED The expert system interface is used. *BASIC The Operational Assistant user interface is used. *INTERMED The system interface is used. *SYSVAL The system value, QASTLVL, is used to determine the user interface that will be used. <p>Contains the null value if the value was not changed.</p>
USER_OPTIONS	VARCHAR(70)	<p>The level of help information detail to be shown and the default function of the Page Up and Page Down keys. This column can contain up to seven values, each ten characters long.</p> <ul style="list-style-type: none"> *CLKWD Parameter keywords are shown instead of the possible parameter values when a control language (CL) command is prompted. *EXPERT More detailed information is shown when the user is performing display and edit options to define or change the system. *HLPFULL Help text is shown on a full display rather than in a window. *NONE Detailed information is not shown. *NOSTMSG Status messages are not displayed when sent to the user. *PRTMSG A message is sent to this user's message queue when a spooled file for this user is printed or held by the printer writer. *ROLLKEY The actions of the Page Up and Page Down keys are reversed. *STMSG Status messages are displayed when sent to the user. <p>Contains the null value if the value was not changed.</p>
SPECIAL_ENVIRONMENT	VARCHAR(7)	<p>The special environment in which the user operates after signing on.</p> <ul style="list-style-type: none"> *NONE The user operates in the IBM i system environment after signing on the system. *S36 The user operates in the System/36 environment after signing on the system. *SYSVAL The system value, QSPCENV, is used to determine the system environment in which the user operates after signing on the system. <p>Contains the null value if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
DISPLAY_SIGNON_INFORMATION	VARCHAR(7)	<p>Indicates if the sign-on information display is shown.</p> <p>*NO The sign-on information display is not shown.</p> <p>*SYSVAL The system value, QDSPSGNINF, is used to determine whether the sign-on information display is shown.</p> <p>*YES The sign-on information display is shown.</p> <p>Contains the null value if the value was not changed.</p>
LIMIT_DEVICE_SESSIONS	VARCHAR(7)	<p>The number of device sessions allowed for a user is limited.</p> <p>0 The user is not limited to a specific number of device sessions. This value has the same meaning as *NO.</p> <p>1 The user is limited to a single device session. This value has the same meaning as *YES.</p> <p>2-9 The user is limited to the specified number of device sessions.</p> <p>*NO The user is not limited to a specific number of device sessions.</p> <p>*SYSVAL The system value, QLMTDEVSSN, is used to determine whether the user is limited to a specific number of device sessions.</p> <p>*YES The user is limited to a single device session.</p> <p>Contains the null value if the value was not changed.</p>
KEYBOARD_BUFFERING	VARCHAR(10)	<p>The keyboard buffering value to be used when a job is initialized for this user profile.</p> <p>*NO The type-ahead feature and attention key buffering option are not active.</p> <p>*SYSVAL The system value, QKBDBUF, is used to determine the keyboard buffering value.</p> <p>*TYPEAHEAD The type-ahead feature is active, but the attention key buffering option is not.</p> <p>*YES The type-ahead feature and attention key buffering option are active.</p> <p>Contains the null value if the value was not changed.</p>
MAXIMUM_ALLOWED_STORAGE	BIGINT	<p>The amount of auxiliary storage (in kilobytes) assigned to store permanent objects owned by this user profile in the system auxiliary storage pool (ASP) and on all the basic ASPs combined. In addition, this value also controls the maximum amount of auxiliary storage that can be used to store permanent objects owned by this user profile on each Independent ASP (IASP). A value of -1 indicates *NOMAX.</p> <p>Contains the null value if the value was not changed.</p>
PRIORITY_LIMIT	INTEGER	<p>The value of the priority limit parameter. Values are 0 to 9.</p> <p>Contains the null value if the value was not changed.</p>
JOB_DESCRIPTION_LIBRARY	VARCHAR(10)	<p>The library where the job description is found. Can contain the special value *LIBL.</p> <p>Contains the null value if the value was not changed.</p>
JOB_DESCRIPTION	VARCHAR(10)	<p>The job description used for jobs that start through subsystem work station entries whose job description parameter values indicate the user JOBD(*USRPRF).</p> <p>Contains the null value if the value was not changed.</p>
ALTERNATE_SUBSYSTEM_NAME	VARCHAR(10)	<p>The alternative subsystem that will be used for this user, instead of the default subsystem, whenever a connection is initiated to the server job specified in ALTERNATE_SERVER_JOB_NAME.</p> <p>Contains the null value when COMMAND_TYPE is not SQL or if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
ALTERNATE_SERVER_JOB_NAME	VARCHAR(10)	<p>When a connection to this server is initiated for this user it will be routed to the subsystem specified in the ALTERNATIVE_SUBSYSTEM_NAME column. Can contain the special value *ALL.</p> <p>To understand the Server Job Name mapping to server names and the default subsystem use, see Server table.</p> <p>Contains the null value when COMMAND_TYPE is not SQL or if the value was not changed.</p>
ACCOUNTING_CODE	VARCHAR(15)	<p>The accounting code that is associated with this user profile or the following special value:</p> <p>*BLANK An accounting code of 15 blanks is assigned to this user profile.</p> <p>Contains the null value if the value was not changed.</p>
MESSAGE_QUEUE_LIBRARY	VARCHAR(10)	<p>The library where the message queue is found. Can contain the special value *LIBL.</p> <p>Contains the null value if the value was not changed.</p>
MESSAGE_QUEUE	VARCHAR(10)	<p>The message queue to which messages are sent or the following special value:</p> <p>*USRPRF A message queue with the same name as the user profile is used as the message queue for this user. The message queue is located in the QUSRSYS library.</p> <p>Contains the null value if the value was not changed.</p>
MESSAGE_QUEUE_DELIVERY_METHOD	VARCHAR(7)	<p>How messages sent to the message queue for this user are to be delivered.</p> <p>*BREAK The job to which the message queue is assigned is interrupted when a message arrives at the message queue.</p> <p>*DFT The default reply to the inquiry message is sent.</p> <p>*HOLD The messages are held in the message queue until they are requested by the user or program.</p> <p>*NOTIFY The job to which the message queue is assigned is notified when a message arrives at the message queue.</p> <p>Contains the null value if the value was not changed.</p>
MESSAGE_QUEUE_SEVERITY	INTEGER	<p>The lowest severity code that a message can have and still be delivered to a user in break or notify mode. This is a value from 0 to 99.</p> <p>Contains the null value if the value was not changed.</p>
PRINT_DEVICE	VARCHAR(10)	<p>The default printer device for this user or one of the following special values:</p> <p>*SYSVAL The system value, QPRTDEV, is used to determine the printer device.</p> <p>*WRKSTN The printer assigned to the user's work station is used.</p> <p>Contains the null value if the value was not changed.</p>
OUTPUT_QUEUE_LIBRARY	VARCHAR(10)	<p>The library where the output queue is found. Can contain the special value *LIBL.</p> <p>Contains the null value if the value was not changed.</p>
OUTPUT_QUEUE	VARCHAR(10)	<p>The output queue to be used by this user profile or one of the following special values:</p> <p>*DEV The output queue associated with the printer specified for the Printer Device is used.</p> <p>*WRKSTN The output queue assigned to the user's work station is used.</p> <p>Contains the null value if the value was not changed.</p>
ATTENTION_KEY_HANDLING_PROGRAM_LIBRARY	VARCHAR(10)	<p>The library where the ATTN program is found. Can contain the special value *LIBL.</p> <p>Contains the null value if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
ATTENTION_KEY_HANDLING_PROGRAM	VARCHAR(10)	<p>The program to be used as the Attention (ATTN) key handling program for this user or one of the following special values:</p> <p>*ASSIST The Operational Assistant ATTN key handling program, QEZMAIN, is used.</p> <p>*NONE No ATTN key handling program is used by this user.</p> <p>*SYSVAL The system value, QATNPGM, is used to determine the ATTN key handling program.</p> <p>Contains the null value if the value was not changed.</p>
SORT_SEQUENCE_TABLE_LIBRARY	VARCHAR(10)	<p>The library where the sort sequence table is found. Can contain the special value *LIBL.</p> <p>Contains the null value if the value was not changed.</p>
SORT_SEQUENCE_TABLE	VARCHAR(10)	<p>The sort sequence table to be used for string comparisons for this user profile or one of the following special values:</p> <p>*HEX A sort sequence table is not used. The hexadecimal values of the characters are used to determine the sort sequence.</p> <p>*LANGIDSHR A shared-weight sort table is used.</p> <p>*LANGIDUNQ A unique-weight sort table is used.</p> <p>*SYSVAL The system value, QSRTSEQ, is used to determine the sort sequence table.</p> <p>Contains the null value if the value was not changed.</p>
LANGUAGE_ID	VARCHAR(7)	<p>The language identifier to be used for this user profile or the following special value:</p> <p>*SYSVAL The system value, QLANGID, is used to determine the language identifier.</p> <p>Contains the null value if the value was not changed.</p>
COUNTRY_OR_REGION_ID	VARCHAR(7)	<p>The country or region identifier to be used for this user profile or the following special value:</p> <p>*SYSVAL The system value, QCNTYID, is used to determine the country or region ID.</p> <p>Contains the null value if the value was not changed.</p>
CCSID	VARCHAR(7)	<p>The coded character set identifier (CCSID) to be used for this user profile.</p> <p>*SYSVAL The system value, QCCSID, is used to determine the CCSID.</p> <p>Contains the null value if the value was not changed.</p>
CHARACTER_IDENTIFIER_CONTROL	VARCHAR(9)	<p>The character identifier control (CHRIDCTL) for the job.</p> <p>*DEVD Performs the same function as it does on the CHRID parameter for display files, printer files, and panel groups.</p> <p>*JOBCCSID Performs the same function as it does on the CHRID parameter for display files, printer files, and panel groups.</p> <p>*SYSVAL The system value, QCHRIDCTL, is used to determine the CHRIDCTL for the job.</p> <p>Contains the null value if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
LOCALE_JOB_ATTRIBUTES	VARCHAR(60)	<p>The job attributes that are to be taken from the locale when the job is initiated. This column can contain up to six values, each ten characters long.</p> <p>*CCSID The coded character set identifier from the locale is used.</p> <p>*DATFMT The date format from the locale is used.</p> <p>*DATSEP The date separator from the locale is used.</p> <p>*DECfmt The decimal format from the locale is used.</p> <p>*NONE No job attributes are taken from the locale.</p> <p>*SRTSEQ The sort sequence from the locale is used.</p> <p>*SYSVAL The system value, QSETJOBATR, is used to determine which job attributes are taken from the locale.</p> <p>*TIMSEP The time separator from the locale is used.</p> <p>Contains the null value if the value was not changed.</p>
DOCUMENT_PASSWORD_CHANGED	VARCHAR(3)	<p>Indicates if the document password has been changed.</p> <p>YES Document password changed.</p> <p>Contains the null value if the value was not changed.</p>
DOCUMENT_PASSWORD_NONE	VARCHAR(3)	<p>Indicates if the document password is *NONE.</p> <p>YES Document password is *NONE.</p> <p>Contains the null value if the value was not changed.</p>
EIM_ID	VARCHAR(128)	<p>Enterprise Identity Mapping (EIM) identifier name or the following special value:</p> <p>*USRPRF The EIM identifier is the same name as the user profile.</p> <p>Contains the null value if no EIM values were changed.</p>
EIM_ASSOCIATION_TYPE	VARCHAR(7)	<p>EIM association type.</p> <p>*ADMIN Administrative association.</p> <p>*ALL All association types.</p> <p>*SOURCE Source association.</p> <p>*TARGET Target association.</p> <p>*TGTSRC Target and source associations.</p> <p>Contains the null value if no EIM values were changed.</p>
EIM_ASSOCIATION_ACTION	VARCHAR(8)	<p>EIM association action.</p> <p>*ADD Add an association.</p> <p>*REMOVE Remove an association.</p> <p>*REPLACE Associations of the specified type will be removed from all EIM identifiers that have an association for this user profile and local EIM registry. A new association will be added to the specified EIM identifier.</p> <p>Contains the null value if no EIM values were changed.</p>
CREATE_EIM_ID	VARCHAR(11)	<p>Indicates whether the EIM identifier should be created if it does not exist.</p> <p>*CRTEIMID EIM identifier gets created if it does not exist.</p> <p>*NOCRTEIMID EIM identifier does not get created.</p> <p>Contains the null value if no EIM values were changed.</p>
USER_ID_NUMBER	VARCHAR(10)	<p>The UID for the user. Can contain the special value *GEN to indicate the uid number is generated for the user.</p> <p>Contains the null value if the value was not changed.</p>

Table 131. AUDIT_JOURNAL_CP table function (continued)

Column Name	Data Type	Description
GROUP_ID_NUMBER	VARCHAR(10)	The GID for the user. Can contain the special values *GEN to indicate the gid number is generated for the user or *NONE to indicate the user profile does not have a group profile. Contains the null value if the value was not changed.

Example

- List any user profiles that were created in the last 6 months.

```
SELECT USER_PROFILE FROM TABLE(
  SYSTOOLS.AUDIT_JOURNAL_CP(
    STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 6 MONTHS
  )
)
WHERE COMMAND_TYPE = 'CRT';
```

AUDIT_JOURNAL_DO table function

The AUDIT_JOURNAL_DO table function returns rows from the audit journal that contain information from the DO (Delete Operation) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 132. AUDIT_JOURNAL_DO table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the DO audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. A Object was deleted (not under commitment control) C A pending object delete was committed D A pending object create was rolled back P The object delete is pending (the delete was performed under commitment control) R A pending object delete was rolled back
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
OBJECT_LIBRARY	VARCHAR(10)	The name of the library containing the object. Contains the null value if there is no library name.
OBJECT_NAME	VARCHAR(10)	The name of the object. Contains the null value if there is no object name.
OBJECT_TYPE	VARCHAR(7)	The type of the object. Contains the null value if there is no object type.
OBJECT_ATTRIBUTE	VARCHAR(10)	The attribute for the object type. Contains the null value if there is no object attribute.
OBJECT_ASP_NAME	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which the object resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.
OBJECT_ASP_NUMBER	INTEGER	The number of the ASP device.
PATH_NAME	VARGRAPHIC(5000) CCSID 1200	The path name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.

Table 132. AUDIT_JOURNAL_DO table function (continued)

Column Name	Data Type	Description
PATH_NAME_INDICATOR	VARCHAR(3)	<p>Path name indicator.</p> <p>NO The PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name.</p> <p>YES The PATH_NAME column contains complete absolute path name for the object.</p> <p>Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	<p>When PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PATH_NAME column.</p> <p>Contains the null value when PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	<p>The name of the object.</p> <p>Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
OBJECT_FILE_ID	BINARY(16)	<p>The file ID of the object.</p> <p>Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
PARENT_FILE_ID	BINARY(16)	<p>The file ID of the parent directory.</p> <p>Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
OFFICE_USER	VARCHAR(10)	<p>The name of the office user.</p> <p>Contains the null value if there is no office user.</p>
OFFICE_ON_BEHALF_OF_USER	VARCHAR(10)	<p>User working on behalf of another user.</p> <p>Contains the null value if the user name is not available.</p>
DLO_NAME	VARCHAR(12)	<p>The name of the document library object.</p> <p>Contains the null value if there is no document library object.</p>
FOLDER_PATH	VARCHAR(63)	<p>The path of the folder.</p> <p>Contains the null value if there is no folder path.</p>

Example

- List any *FILE objects that were deleted this week. Return the user profile in effect when the delete occurred.

```
SELECT OBJECT_LIBRARY, OBJECT_NAME, USER_NAME FROM TABLE(
  SYSTOOLS.AUDIT_JOURNAL_DO(
    STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 7 DAYS
  )
) WHERE OBJECT_TYPE = '*FILE';
```

AUDIT_JOURNAL_EV table function

The AUDIT_JOURNAL_EV table function returns rows from the audit journal that contain information from the EV (Environment Variable) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information” on page 552](#).

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 133. AUDIT_JOURNAL_EV table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the EV audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. A Add C Change D Delete I Initialize environment variable space
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
ENVIRONMENT_VARIABLE_NAME	VARGRAPHIC(1000) CCSID 1200	The name of the environment variable. Contains the null value when ENTRY_TYPE is I.
NAME_TRUNCATED	VARCHAR(3)	Indicates whether ENVIRONMENT_VARIABLE_NAME is truncated. NO Environment variable name is not truncated. YES Environment variable name is truncated. Contains the null value when ENVIRONMENT_VARIABLE_NAME is null.
ENVIRONMENT_VARIABLE_VALUE	VARGRAPHIC(1000) CCSID 1200	The value of the environment variable. Contains the null value when no environment variable value is available.
VALUE_TRUNCATED	VARCHAR(3)	Indicates whether ENVIRONMENT_VARIABLE_VALUE is truncated. NO Environment variable value is not truncated. YES Environment variable value is truncated. Contains the null value when ENVIRONMENT_VARIABLE_VALUE is null.

Example

- List any environment variables that have been changed in the last month.

```
SELECT *
  FROM TABLE (
    SYSTOOLS.AUDIT_JOURNAL_EV (STARTING_TIMESTAMP => CURRENT DATE - 1 MONTH))
 WHERE ENTRY_TYPE = 'C';
```

AUDIT_JOURNAL_GR table function

The AUDIT_JOURNAL_GR table function returns rows from the audit journal that contain information from the GR (Generic Record) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 134. AUDIT_JOURNAL_GR table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the GR audit journal entry.		

Table 134. AUDIT_JOURNAL_GR table function (continued)

Column Name	Data Type	Description
ENTRY_TYPE	CHAR(1)	The type of entry. A Exit program added D Exit program removed F Function registration operations O ObjectConnect operations R Exit program replaced
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
ACTION	CHAR(2)	The action performed. ZC Change ZR Read When ENTRY_TYPE is O: SV Save RS Restore
ACTION_DETAIL	VARCHAR(200)	Descriptive text that describes the action.
EXIT_POINT_NAME	VARCHAR(20)	When ENTRY_TYPE is A, D, or R, the exit point name. Contains the null value when ENTRY_TYPE is F, or O.
EXIT_POINT_FORMAT	CHAR(8)	When ENTRY_TYPE is A, D, or R, the exit point format name associated with the exit point Contains the null value when ENTRY_TYPE is F, or O.
EXIT_PROGRAM_NUMBER	INTEGER	When ENTRY_TYPE is A, D, or R, the exit program number associated with the exit program. Contains the null value when ENTRY_TYPE is F, or O.
EXIT_PROGRAM_LIBRARY	VARCHAR(10)	When ENTRY_TYPE is A, D, or R, the library in which EXIT_PROGRAM resides. Contains the null value when ENTRY_TYPE is F, or O.
EXIT_PROGRAM	VARCHAR(10)	When ENTRY_TYPE is A, D, or R, the name of the exit program.. Contains the null value when ENTRY_TYPE is F, or O.
USER_PROFILE_NAME	VARCHAR(10)	User profile name When ENTRY_TYPE is F, contains the name of the user the function registration operation was performed against. When a REGISTER or REREGISTER operation is being performed that includes the default usage, contains *DEFAULT. When ENTRY_TYPE is O, contains the name of the user performing the ObjectConnect operation. Contains the null value when ENTRY_TYPE is A, D, or R, or if a user profile name does not apply to the entry.

Table 134. AUDIT_JOURNAL_GR table function (continued)

Column Name	Data Type	Description
FUNCTION_REGISTRATION_OPERATION	VARCHAR(13)	<p>When ENTRY_TYPE is F, contains the description of the function registration operation that was performed. The possible values are:</p> <p>CHANGE USAGE Function usage information has been changed, such as which user profiles are allowed to use a function.</p> <p>CHECK USAGE Function usage information was checked for a user and determined the user is allowed to use the specified function.</p> <p>DEREGISTER Function and all associated usage information has been removed from the registration facility.</p> <p>REGISTER Function has been registered with the registration facility.</p> <p>REREGISTER Function has been updated and replaced within the registration facility.</p> <p>USAGE FAILURE Function usage information was checked for a user and determined the user is not allowed to use the specified function..</p> <p>Contains the null value when ENTRY_TYPE is not F.</p>
FUNCTION_NAME	VARCHAR(30)	<p>When ENTRY_TYPE is F, contains the name of the function that was operated on.</p> <p>Contains the null value when ENTRY_TYPE is not F.</p>
USAGE_SETTING	VARCHAR(7)	<p>When ENTRY_TYPE is F, contains the usage value for the user specified in USER_PROFILE_NAME. If USER_PROFILE_NAME is *DEFAULT, this is the default usage.</p> <p>ALLOWED The user is allowed to use the function.</p> <p>DENIED The user is not allowed to use the function.</p> <p>REMOVED The user's previous setting is removed.</p> <p>UNKNOWN The user did not previously have a usage value.</p> <p>Contains the null value when ENTRY_TYPE is not F or if no value is available.</p>
PREVIOUS_USAGE	VARCHAR(7)	<p>When ENTRY_TYPE is F, contains the previous usage value for a user. If USER_PROFILE_NAME is *DEFAULT, this is the previous default usage.</p> <p>ALLOWED The user is allowed to use the function.</p> <p>DENIED The user is not allowed to use the function.</p> <p>REMOVED The user's previous setting is removed.</p> <p>UNKNOWN The user did not previously have a usage value.</p> <p>Contains the null value when ENTRY_TYPE is not F or if no value is available.</p>
FUNCTION_ALLOBJ	VARCHAR(8)	<p>When ENTRY_TYPE is F and FUNCTION_REGISTRATION_OPERATION is REGISTER or REREGISTER, contains the setting that indicates whether all object (*ALLOBJ) special authority may be used to give a user access to the function.</p> <p>NOT USED For a user with *ALLOBJ special authority to use the function, the usage information specified for the function must indicate that the user is allowed to use the function</p> <p>USED A user with *ALLOBJ special authority is always allowed to use the function.</p> <p>Contains the null value when ENTRY_TYPE is not F or if no value is available.</p>

Table 134. AUDIT_JOURNAL_GR table function (continued)

Column Name	Data Type	Description
PREVIOUS_ALLOBJ	VARCHAR(8)	<p>When ENTRY_TYPE is F and FUNCTION_REGISTRATION_OPERATION is REREGISTER, contains the previous setting that indicates whether all object (*ALLOBJ) special authority may be used to give a user access to the function.</p> <p>NOT USED For a user with *ALLOBJ special authority to use the function, the usage information specified for the function must indicate that the user is allowed to use the function</p> <p>USED A user with *ALLOBJ special authority is always allowed to use the function.</p> <p>Contains the null value when ENTRY_TYPE is not F or if no value is available.</p>
OBJECTCONNECT_COMMAND	VARCHAR(9)	<p>When ENTRY_TYPE is O, contains the ObjectConnect CL command.</p> <p>SAVRST Save/Restore Integrated File System Object</p> <p>SAVRSTCFG Save/Restore Configuration</p> <p>SAVRSTCHG Save/Restore Changed Object</p> <p>SAVRSTDLO Save/Restore Document Library Object</p> <p>SAVRSTLIB Save/Restore Library</p> <p>SAVRSTOBJ Save/Restore Object</p> <p>Contains the null value when ENTRY_TYPE is not O.</p>
SAVE_SYSTEM	VARCHAR(15)	<p>When ENTRY_TYPE is O and ACTION is RS, contains the name of the system on which the objects are saved.</p> <p>When ENTRY_TYPE is O and ACTION is SV, contains the name of the system on which the objects are restored.</p> <p>Contains the null value when ENTRY_TYPE is not O.</p>
SAVE_ASP	VARCHAR(10)	<p>When ENTRY_TYPE is O and OBJECTCONNECT_COMMAND is SAVRST, SAVRSTCHG, SAVRSTLIB, or SAVRSTOBJ, the ASP device name. Can contain one of the following special values:</p> <ul style="list-style-type: none"> • * • *ALLAVL • *ANY • *CURASPGRP • *SYSBAS • 1-32 <p>Contains the null value when ENTRY_TYPE is not O or there is no ASP device.</p>
SAVE_LIBRARY	VARCHAR(10)	<p>When ENTRY_TYPE is O and OBJECTCONNECT_COMMAND is SAVRST, SAVRSTCHG, SAVRSTLIB, or SAVRSTOBJ, the library name is set when processing objects from the QSYS file system. It is the name of the saved library or the library from which the objects were saved.</p> <p>Contains the null value when ENTRY_TYPE is not O or there is no saved library.</p>
RESTORE_ASP_DEVICE	VARCHAR(10)	<p>When ENTRY_TYPE is O and OBJECTCONNECT_COMMAND is SAVRST, SAVRSTCHG, SAVRSTLIB, or SAVRSTOBJ, the ASP device name. Can contain the following special value:</p> <ul style="list-style-type: none"> • *SAVASPDEV <p>Contains the null value when ENTRY_TYPE is not O or RESTORE_ASP_NUMBER is set.</p>
RESTORE_ASP_NUMBER	VARCHAR(10)	<p>When ENTRY_TYPE is O and OBJECTCONNECT_COMMAND is SAVRST, SAVRSTCHG, SAVRSTLIB, or SAVRSTOBJ, the ASP number, 1-32. Can contain the following special value:</p> <ul style="list-style-type: none"> • *SAVASP <p>Contains the null value when ENTRY_TYPE is not O or RESTORE_ASP_DEVICE is set.</p>

Table 134. AUDIT_JOURNAL_GR table function (continued)

Column Name	Data Type	Description
RESTORE_LIBRARY	VARCHAR(10)	When ENTRY_TYPE is O and OBJECTCONNECT_COMMAND is SAVRST, SAVRSTCHG, SAVRSTLIB, or SAVRSTOBJ, the library name is set when processing objects from the QSYS file system otherwise it is blank. It is the name of the restored library or the library to which the objects were restored. Contains the null value when ENTRY_TYPE is not O or there is no restore library name.
OBJECTCONNECT_UUID	VARCHAR(16)	When ENTRY_TYPE is O, contains the Universal Unique Identifier (UUID) of the ObjectConnect operation. Contains the null value when ENTRY_TYPE is not O.
RESTORE_USER	VARCHAR(10)	When ENTRY_TYPE is O and ACTION is SV, contains the name of the user under which the restore will be performed. Can contain the following special values: *CURRENT *KERBEROS *NONE Contains the null value when ENTRY_TYPE is not O with an ACTION of SV.

Example

- List any changes to function usage definitions from the last two weeks

```
SELECT *
FROM TABLE (
  SYSTOOLS.AUDIT_JOURNAL_GR (STARTING_TIMESTAMP => CURRENT DATE - 14 DAYS))
WHERE ENTRY_TYPE = 'F' AND
      FUNCTION_REGISTRATION_OPERATION = 'CHANGE USAGE';
```

AUDIT_JOURNAL_JS table function

The AUDIT_JOURNAL_JS table function returns rows from the audit journal that contain information from the JS (Job Change) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 135. AUDIT_JOURNAL_JS table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the JS audit journal entry.		

Table 135. AUDIT_JOURNAL_JS table function (continued)

Column Name	Data Type	Description
ENTRY_TYPE	CHAR(1)	<p>The type of entry.</p> <p>A ENDJOBABN command</p> <p>B Submit</p> <p>C Change</p> <p>E End</p> <p>H Hold</p> <p>I Disconnect</p> <p>J The current job is attempting to interrupt another job</p> <p>K The current job is about to be interrupted</p> <p>L The interruption of the current job has completed</p> <p>M Change profile or group profile</p> <p>N ENDJOB command</p> <p>P Attach prestart or batch immediate job</p> <p>Q Change query attributes</p> <p>R Release</p> <p>S Start</p> <p>T Change profile or group profile using a profile token</p> <p>U CHGUSRTRC</p> <p>V Virtual device changed by QWSACCD5 API</p>
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
JOB_TYPE	VARCHAR(3)	<p>The job type. This value is based on the values of the JOB_TYPE_BASIC and JOB_SUBTYPE columns.</p> <p>ASJ Autostart</p> <p>BCH Batch</p> <p>BCI Batch Immediate</p> <p>EVK Started by a procedure start request</p> <p>INT Interactive</p> <p>M36 Advanced 36 server job</p> <p>MRT Multiple requester terminal</p> <p>PDJ Print driver job</p> <p>PJ Prestart job</p> <p>RDR Spool reader</p> <p>SBS Subsystem monitor</p> <p>SYS System</p> <p>WTR Spool writer</p> <p>Contains the null value if no value is available.</p>

Table 135. AUDIT_JOURNAL_JS table function (continued)

Column Name	Data Type	Description
JOB_TYPE_BASIC	CHAR(1)	The type of job. A Autostart B Batch I Interactive M Subsystem monitor R Reader S System W Writer X SCPF Contains the null value if no value is available.
JOB_SUBTYPE	CHAR(1)	The subtype of the job. D Batch immediate E Procedure start request J Prestart P Print device driver Q Query T MRT U Alternate spool user Contains the null value if there is no subtype for the job.
TARGET_QUALIFIED_JOB_NAME	VARCHAR(28)	The qualified job name of the job that is being operated on.
TARGET_JOB_NAME	VARCHAR(10)	The name of the job that is being operated on.
TARGET_JOB_USER	VARCHAR(10)	The user profile that of the job that is being operated on.
TARGET_JOB_NUMBER	VARCHAR(6)	The job number of the job that is being operated on.
DEVICE_NAME	VARCHAR(10)	The name of the device. Contains the null value if no value is available.
EFFECTIVE_USER_PROFILE	VARCHAR(10)	The name of the effective user profile for the thread. When this audit record is generated because one job performs an operation on another job this column contains data from the initial thread of the job that is being operated on. In all other cases, the column contains data from the thread that performed the operation. Contains the null value if no value is available.
EFFECTIVE_GROUP_PROFILE	VARCHAR(10)	The name of the effective group profile for the thread. When this audit record is generated because one job performs an operation on another job, this column contains data from the initial thread of the job that is being operated on. In all other cases, the column contains data from the thread that performed the operation. Contains the null value if no value is available.
SUPPLEMENTAL_GROUP_PROFILES	VARCHAR(150)	The names of the supplemental group profiles for the thread. When this audit record is generated because one job performs an operation on another job, this column contains data from the initial thread of the job that is being operated on. In all other cases, the column contains data from the thread that performed the operation. Contains the null value if no value is available.
REAL_USER_PROFILE	VARCHAR(10)	The name of the real (initial) user profile for the thread. Contains the null value if no value is available.
SAVED_USER_PROFILE	VARCHAR(10)	The name of the saved user profile for the thread. Contains the null value if no value is available.
REAL_GROUP_PROFILE	VARCHAR(10)	The name of the real (initial) group profile for the thread. Contains the null value if no value is available.

Table 135. AUDIT_JOURNAL_JS table function (continued)

Column Name	Data Type	Description
SAVED_GROUP_PROFILE	VARCHAR(10)	The name of the saved group profile for the thread. Contains the null value if no value is available.
REAL_USER_CHANGED	VARCHAR(3)	The real user profile was changed. NO The real user profile was not changed. YES The real user profile was changed. Contains the null value when ENTRY_TYPE is not M or T.
EFFECTIVE_USER_CHANGED	VARCHAR(3)	The effective user profile was changed. NO The effective user profile was not changed. YES The effective user profile was changed. Contains the null value when ENTRY_TYPE is not M or T.
SAVED_USER_CHANGED	VARCHAR(3)	The saved user profile was changed. NO The saved user profile was not changed. YES The saved user profile was changed. Contains the null value when ENTRY_TYPE is not M or T.
REAL_GROUP_CHANGED	VARCHAR(3)	The real group profile was changed. NO The real group profile was not changed. YES The real group profile was changed. Contains the null value when ENTRY_TYPE is not M or T.
EFFECTIVE_GROUP_CHANGED	VARCHAR(3)	The effective group profile was changed. NO The effective group profile was not changed. YES The effective group profile was changed. Contains the null value when ENTRY_TYPE is not M or T.
SAVED_GROUP_CHANGED	VARCHAR(3)	The saved group profile was changed. NO The saved group profile was not changed. YES The saved group profile was changed. Contains the null value when ENTRY_TYPE is not M or T.
SUPPLEMENTAL_GROUPS_CHANGED	VARCHAR(3)	The supplemental group profiles were changed. NO The supplemental group profiles were not changed. YES The supplemental group profiles were changed. Contains the null value when ENTRY_TYPE is not M or T.
JOB_DESCRIPTION_LIBRARY	VARCHAR(10)	The name of the library for the job description. Contains the null value if no value is available.
JOB_DESCRIPTION	VARCHAR(10)	The name of the job description for the job. Contains the null value if no value is available.
JOB_DESCRIPTION_ASP_NAME	VARCHAR(10)	ASP name for JOB_DESCRIPTION_LIBRARY. Contains the null value if no value is available.
JOB_DESCRIPTION_ASP_NUMBER	INTEGER	ASP number for JOB_DESCRIPTION_LIBRARY. Contains the null value if no value is available.
JOB_QUEUE_LIBRARY	VARCHAR(10)	The name of the library for the job queue. Contains the null value if no value is available.
JOB_QUEUE_NAME	VARCHAR(10)	The name of the job queue for the job. Contains the null value if no value is available.

Table 135. AUDIT_JOURNAL_JS table function (continued)

Column Name	Data Type	Description
JOB_QUEUE_ASP_NAME	VARCHAR(10)	ASP name for JOB_QUEUE_LIBRARY. Contains the null value if no value is available.
JOB_QUEUE_ASP_NUMBER	INTEGER	ASP number of JOB_QUEUE_LIBRARY. Contains the null value if no value is available.
OUTPUT_QUEUE_LIBRARY	VARCHAR(10)	The name of the library for the output queue. Contains the null value if no value is available.
OUTPUT_QUEUE_NAME	VARCHAR(10)	The name of the output queue for the job. Contains the null value if no value is available.
PRINTER_DEVICE	VARCHAR(10)	The name of the printer device for the job. Contains the null value if no value is available.
TIME_ZONE_DESCRIPTION_NAME	VARCHAR(10)	The timezone description name. Contains the null value if no value is available.
THREAD_ASP_NAME	VARCHAR(10)	The ASP group for the current thread. Contains the null value if no value is available.
LIBRARY_LIST_COUNT	INTEGER	The number of libraries in LIBRARY_LIST.
LIBRARY_LIST	VARCHAR(2680)	The library list for the job. The list is an array of entries, each ten characters long. When this audit record is generated because one job performs an operation on another job, this column contains data from the initial thread of the job that is being operated on. In all other cases, the column contains data from the thread that performed the operation. Contains the null value when LIBRARY_LIST_COUNT is 0.
JOB_USER_IDENTITY_DESCRIPTION	VARCHAR(5)	Describes the meaning of JOB_USER_IDENTITY. CLEAR The clear JUID API was called. JOB_USER_IDENTITY contains the new value. JOB JOB_USER_IDENTITY contains the value for the JOB. SET The set JUID API was called. JOB_USER_IDENTITY contains the new value.
JOB_USER_IDENTITY	VARCHAR(10)	The job user identity value. Contains the null value if no value is available.
WORKLOAD_GROUP	VARCHAR(10)	The name of the workload group associated with the job. Contains the null value when ENTRY_TYPE is not C, E, or S.
EXIT_QUALIFIED_JOB_NAME	VARCHAR(28)	<ul style="list-style-type: none"> When ENTRY_TYPE is J, the qualified job name of the job that was interrupted by the current job. When ENTRY_TYPE is K or L, the qualified job name of the job that requested the interruption of the current job. Contains the null value when ENTRY_TYPE is not J, K, or L.
EXIT_JOB_NAME	VARCHAR(10)	<ul style="list-style-type: none"> When ENTRY_TYPE is J, the name of the job that was interrupted by the current job. When ENTRY_TYPE is K or L, the name of the job that requested the interruption of the current job. Contains the null value when ENTRY_TYPE is not J, K, or L.
EXIT_JOB_USER	VARCHAR(10)	<ul style="list-style-type: none"> When ENTRY_TYPE is J, the user of the job that was interrupted by the current job. When ENTRY_TYPE is K or L, the user of the job that requested the interruption of the current job. Contains the null value when ENTRY_TYPE is not J, K, or L.

Table 135. AUDIT_JOURNAL_JS table function (continued)

Column Name	Data Type	Description
EXIT_JOB_NUMBER	VARCHAR(6)	<ul style="list-style-type: none"> When ENTRY_TYPE is J, the job number of the job that was interrupted by the current job. When ENTRY_TYPE is K or L, the number of the job that requested the interruption of the current job. Contains the null value when ENTRY_TYPE is not J, K, or L.
EXIT_PROGRAM_LIBRARY	VARCHAR(10)	The library name of the exit program used to interrupt the job. Contains the null value when ENTRY_TYPE is not J, K, or L.
EXIT_PROGRAM	VARCHAR(10)	The exit program used to interrupt the job. Contains the null value when ENTRY_TYPE is not J, K, or L.

Example

- List jobs that were ended today with the ENDJOB or ENDJOBABN CL commands.

```
SELECT *
FROM TABLE (
  SYSTOOLS.AUDIT_JOURNAL_JS (STARTING_TIMESTAMP => CURRENT DATE)
)
WHERE ENTRY_TYPE IN ('A', 'N');
```

AUDIT_JOURNAL_OM table function

The AUDIT_JOURNAL_OM table function returns rows from the audit journal that contain information from the OM (Object Management Change) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in “AUDIT JOURNAL table function common information” on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 136. AUDIT_JOURNAL_OM table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the OM audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. M Object moved to a different library. R Object renamed.
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
LIBRARY_NAME	VARCHAR(10)	The name of the library to which the object was moved. Contains the null value if there is no new library name.
OBJECT_NAME	VARCHAR(10)	The new name of the object which was moved. Contains the null value if there is no new object name.
OBJECT_TYPE	VARCHAR(7)	The type of object. Contains the null value if there is no object type.
OBJECT_ATTRIBUTE	VARCHAR(10)	The attribute of the object. Contains the null value if there is no object attribute.
PREV_LIBRARY_NAME	VARCHAR(10)	The name of the library in which the previous object resides. Contains the null value if there is no previous library name.
PREV_OBJECT_NAME	VARCHAR(10)	The previous name of the object. Contains the null value if there is no previous object name.

Table 136. AUDIT_JOURNAL_OM table function (continued)

Column Name	Data Type	Description
OBJECT_ASP_NAME	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which the object resides. If the new object is in a library, this is the ASP information of the object's library. If the new object is not in a library, this is the ASP information of the object. A value of *SYSBAS indicates the system ASP and all basic user ASPs. Contains the null value if there is no ASP information.
OBJECT_ASP_NUMBER	INTEGER	The number of the ASP device. If the new object is in a library, this is the ASP information of the object's library. If the new object is not in a library, this is the ASP information of the object. Contains the null value if there is no ASP information.
PREV_OBJECT_ASP_NAME	VARCHAR(10)	The previous name of the auxiliary storage pool (ASP) in which the object resides. If the previous object is in a library, this is the ASP information of the object's library. If the previous object is not in a library, this is the ASP information of the object. A value of *SYSBAS indicates the system ASP and all basic user ASPs. Contains the null value if there is no ASP information.
PREV_OBJECT_ASP_NUMBER	INTEGER	The previous number of the ASP device. If the previous object is in a library, this is the ASP information of the object's library. If the previous object is not in a library, this is the ASP information of the object. Contains the null value if there is no ASP information.
PATH_NAME	VARGRAPHIC(5000) CCSID 1200	The new path name of the object. Contains the null value if the path name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PATH_NAME_INDICATOR	VARCHAR(3)	New path name indicator: NO The PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name. YES The PATH_NAME column contains the absolute path name for the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	When PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PATH_NAME column. Contains the null value when PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	The name of the new object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OBJECT_FILE_ID	BINARY(16)	The file ID of the new object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PARENT_FILE_ID	BINARY(16)	The file ID of the new parent directory. Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PREV_PATH_NAME	VARGRAPHIC(5000) CCSID 1200	The previous path name of the object. Contains the null value if the path name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.

Table 136. AUDIT_JOURNAL_OM table function (continued)

Column Name	Data Type	Description
PREV_PATH_NAME_INDICATOR	VARCHAR(3)	<p>Previous path name indicator:</p> <p>NO The PREV_PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The PREV_RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name.</p> <p>YES The PREV_PATH_NAME column contains the absolute path name for the object.</p> <p>Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
PREV_RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	<p>When PREV_PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PREV_PATH_NAME column.</p> <p>Contains the null value when PREV_PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
PREV_IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	<p>The name of the previous object.</p> <p>Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
PREV_OBJECT_FILE_ID	BINARY(16)	<p>The file ID of the previous object.</p> <p>Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
PREV_PARENT_FILE_ID	BINARY(16)	<p>The file ID of the previous parent directory.</p> <p>Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.</p>
OFFICE_USER	VARCHAR(10)	<p>The name of the office user.</p> <p>Contains the null value if there is no office user.</p>
OFFICE_ON_BEHALF_OF_USER	VARCHAR(10)	<p>User working on behalf of another user.</p> <p>Contains the null value if there is no user name.</p>
DLO_NAME	VARCHAR(12)	<p>The new name of the folder or document.</p> <p>Contains the null value if there is no new folder or document.</p>
FOLDER_PATH	VARCHAR(63)	<p>The new path of the folder.</p> <p>Contains the null value if there is no new folder path.</p>
PREV_DLO_NAME	VARCHAR(12)	<p>The previous name of the folder or document.</p> <p>Contains the null value if there is no previous folder or document.</p>
PREV_FOLDER_PATH	VARCHAR(63)	<p>The previous path of the folder.</p> <p>Contains the null value if there is no previous folder path.</p>

Example

- List any objects that were moved or renamed in the APPLIB library this week.

```
SELECT LIBRARY_NAME, OBJECT_NAME, PREV_LIBRARY_NAME, PREV_OBJECT_NAME, OBJECT_TYPE
FROM TABLE (
    SYSTOOLS.AUDIT_JOURNAL_OM (STARTING_TIMESTAMP => CURRENT DATE - 7 DAYS)
)
WHERE PREV_LIBRARY_NAME = 'APPLIB';
```

AUDIT_JOURNAL_OW table function

The AUDIT_JOURNAL_OW table function returns rows from the audit journal that contain information from the OW (Ownership Change) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 137. AUDIT_JOURNAL_OW table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the OW audit journal entry.		
OBJECT_LIBRARY	VARCHAR(10)	The name of the library where the object is stored.
OBJECT_NAME	VARCHAR(10)	The name of the object.
OBJECT_TYPE	VARCHAR(7)	The type of object.
OBJECT_ASP_NAME	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which the object resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.
OBJECT_ASP_NUMBER	INTEGER	The number of the ASP device.
PREVIOUS_OWNER	VARCHAR(10)	Previous owner of the object.
NEW_OWNER	VARCHAR(10)	New owner of the object.
PATH_NAME	VARGRAPHIC(5000) CCSID 1200	The path name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PATH_NAME_INDICATOR	VARCHAR(3)	Path name indicator. NO The PATH_NAME column does not contain an absolute path name for the object, instead it contains a relative path name. The RELATIVE_DIRECTORY_FILE_ID can be used to form an absolute path name with this relative path name. YES The PATH_NAME column contains complete absolute path name for the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	When PATH_NAME_INDICATOR is NO, contains the file ID of the directory that contains the object identified in the PATH_NAME column. Contains the null value when PATH_NAME_INDICATOR is YES, or if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
IFS_OBJECT_NAME	VARGRAPHIC(512) CCSID 1200	The name of the object. Contains the null value if the object name is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OBJECT_FILE_ID	BINARY(16)	The file ID of the object. Contains the null value if the object is not in the "root" (/), QOpenSys, or user-defined file systems.
PARENT_FILE_ID	BINARY(16)	The file ID of the parent directory. Contains the null value if the file ID is not available or the object is not in the "root" (/), QOpenSys, or user-defined file systems.
OFFICE_USER	VARCHAR(10)	The name of the office user.
OFFICE_ON_BEHALF_OF_USER	VARCHAR(10)	User working on behalf of another user.
DLO_NAME	VARCHAR(12)	The name of the document library object.
FOLDER_PATH	VARCHAR(63)	The path of the folder.

Example

- List the objects in APPLIB1 or APPLIB2 that had object ownership changes in the last 3 months.

```
SELECT *
FROM TABLE (
    SYSTOOLS.AUDIT_JOURNAL_OW (STARTING_TIMESTAMP => CURRENT_TIMESTAMP - 3 MONTHS)
)
WHERE OBJECT_LIBRARY IN ('APPLIB1', 'APPLIB2');
```


AUDIT_JOURNAL_PW table function

The AUDIT_JOURNAL_PW table function returns rows from the audit journal that contain information from the PW (Password) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 138. AUDIT_JOURNAL_PW table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the PW audit journal entry.		
VIOLATION_TYPE	CHAR(1)	The type of violation. <ul style="list-style-type: none"> A APPC bind failure. C User authentication with the CHKPWD command failed. D Service tools user ID name not valid (QSYCHGDS API, CRTSSTUSR, CHGSSTUSR, DLTSSSTUSR commands). E Service tools user ID password not valid (QSYCHGDS API, CRTSSTUSR, CHGSSTUSR, DLTSSSTUSR commands). P Password not valid. Q Attempted signon (user authentication) failed because user profile is disabled. R Attempted signon (user authentication) failed because password was expired. S SQL Decryption password is not valid. U User name not valid. X Service tools user ID is disabled. Y Service tools user ID not valid (service tools interface). Z Service tools user ID password not valid (service tools interface).
VIOLATION_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the violation type.
AUDIT_USER_NAME	VARCHAR(10)	The job user name or the service tools user ID name. Contains the null value if no user name is available.
DEVICE_NAME	VARCHAR(40)	The name of the device or communications device on which the password or user ID was entered. Contains the null value when VIOLATION_TYPE is D, E, X, Y, or Z or if the device name is not available.
INTERFACE_NAME	VARCHAR(40)	The name of the interface being used. Contains the null value when VIOLATION_TYPE is not D, E, X, Y, or Z or if the interface name is not available.
REMOTE_LOCATION	VARCHAR(8)	Name of the remote location for the APPC bind. Contains the null value when VIOLATION_TYPE is not A.
LOCAL_LOCATION	VARCHAR(8)	Name of the local location for the APPC bind. Contains the null value when VIOLATION_TYPE is not A.
NETWORK_ID	VARCHAR(8)	Network ID for the APPC bind. Contains the null value when VIOLATION_TYPE is not A.
DECRYPT_HOST_VARIABLE	VARCHAR(3)	Whether the user attempted to decrypt data in a host variable. <ul style="list-style-type: none"> NO The user did not attempt to decrypt data in a host variable. YES The user attempted to decrypt data in a host variable. Contains the null value if VIOLATION_TYPE is not S.

Table 138. AUDIT_JOURNAL_PW table function (continued)

Column Name	Data Type	Description
DECRYPT_OBJECT_LIBRARY	VARCHAR(10)	The library that contains OBJECT_NAME. Contains the null value if VIOLATION_TYPE is not S or if there is no library name.
DECRYPT_OBJECT_NAME	VARCHAR(10)	The name of the object being decrypted. Contains the null value if VIOLATION_TYPE is not S or if there is no object name.
DECRYPT_OBJECT_TYPE	VARCHAR(8)	The type of the object. Contains the null value if VIOLATION_TYPE is not S or if there is no object.
DECRYPT_OBJECT_ASP_NAME	VARCHAR(10)	The name of the ASP device where OBJECT_NAME resides. Contains the null value if VIOLATION_TYPE is not S or if there is no object name.
DECRYPT_OBJECT_ASP_NUMBER	INTEGER	The number of the ASP device where OBJECT_NAME resides. Contains the null value if VIOLATION_TYPE is not S or if there is no object name.

Example

- For all the password audit journal entries from yesterday and today, list the number of each type of audit violation.

```
SELECT VIOLATION_TYPE CONCAT ' - ' CONCAT VIOLATION_TYPE_DETAIL,
       COUNT(*) AS VIOLATION_COUNT
FROM TABLE (
  SYSTOOLS.AUDIT_JOURNAL_PW ( )
)
GROUP BY VIOLATION_TYPE CONCAT ' - ' CONCAT VIOLATION_TYPE_DETAIL
ORDER BY 2 DESC;
```

AUDIT_JOURNAL_ST table function

The AUDIT_JOURNAL_ST table function returns rows from the audit journal that contain information from the ST (Service Tools Action) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 139. AUDIT_JOURNAL_ST table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the ST audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. A Service record
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.

Table 139. AUDIT_JOURNAL_ST table function (continued)

Column Name	Data Type	Description
SERVICE_TOOL	CHAR(2)	The service tool identifier. AR ARM diagnostic trace (see ARMSRV QShell command) AS Storage altered by Display/Alter/Dump service tool or by a remote service tool debugger CD QTACTLDV, QTADMPDV CE QWTCTLTR CS STRCPYSCN CT DMPCLUTRC DC DLTCMNTRC DD DMPDLO DF QWTDMPFR, QWTDMPFL DI QSCDIRD DM DMPMEMINF DO DMPOBJ DS DMPYSOBB, QTADMPTS, QTADMPDV, QWTDMPFL DU DMPUSRPRF DW STRDW, ENDDW, ADDDWDFN, RMVDWDFN EC ENDCMNTRC ER ENDRMTSPT FF FFDC (First Failure Data Capture) GS QSMGSSTD HD QYHCHCOP (DASD) HL QYHCHCOP (LPAR) JW STRJW, ENDJW, ADDJWDFN, RMVJWDFN MC QWTMAINT (change) MD QWTMAINT (dump) MP End system job MQ Restart system job OP Operations console PC PRTC MNTRC PE PRTERLOG, QTADMPDV

Table 139. AUDIT_JOURNAL_ST table function (continued)

Column Name	Data Type	Description
SERVICE_TOOL (continued)		PI PRTINTDTA, QTADMPDV PS QPOFPTOS SC STRCMNTRC, QSCCHGCT SE QWTSETTR SF QWCCDSIC, QWVRCSTK (Display internal stack entry) SJ STRSRVJOB SN QPZSYNC SR STRRMTSPT SS QFPHPSF ST STRSST SV QSRSRV TA TRCTCPAPP TC TRCCNN (*FORMAT specified) TE ENDTRC, ENDPEX, TRCJOB(*OFF or *END specified) TI TRCINT, or TRCCNN with SET(*ON), SET(*OFF), or SET(*END) TO QTOBSRV TQ QWCTMQTM TS STRTRC, STRPEX, TRCJOB(*ON specified) UD QTAUPDDV WE ENDWCH, QSCEWCH WS STRWCH, QSCSWCH WT WRKTRC WW WRKWCH, QSCRWCHI, QSCRWCHL
SERVICE_TOOL_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the service tool.
OBJECT_LIBRARY	VARCHAR(10)	Name of the library for OBJECT_NAME. Contains the null value if there is no library name.
OBJECT_NAME	VARCHAR(10)	Name of the object accessed. Contains the null value if there is no object name.
OBJECT_TYPE	VARCHAR(7)	Type of object. Contains the null value if there is no object type.
OBJECT_ASP_NAME	VARCHAR(10)	ASP name for object library. Contains the null value if there is no object ASP name.
OBJECT_ASP_NUMBER	INTEGER	ASP number for object library. Contains the null value if there is no object ASP number.
DLO_NAME	VARCHAR(12)	Name of the document library object. Contains the null value if there is no document library object.
FOLDER_PATH	VARCHAR(63)	The folder containing the document library object Contains the null value if there is no folder.
SOURCE_NODE_ID	VARCHAR(8)	Source node ID. Contains the null value if there is no source node ID.
SOURCE_USER	VARCHAR(10)	Source user. Contains the null value if there is no source user.
TARGET_QUALIFIED_JOB_NAME	VARCHAR(28)	The qualified job name of the target job. Contains the null value if there is no target job.

Table 139. AUDIT_JOURNAL_ST table function (continued)

Column Name	Data Type	Description
TARGET_JOB_NAME	VARCHAR(10)	The job name of the target job. Contains the null value if there is no target job.
TARGET_JOB_USER	VARCHAR(10)	The job user of the target job. Contains the null value if there is no target job.
TARGET_JOB_NUMBER	VARCHAR(6)	The job number of the target job. Contains the null value if there is no target job.
TARGET_JOB_USER_IDENTITY	VARCHAR(10)	The job user identity value of the target job. Contains the null value if there is no target job.
DMPYSOBJ_LIBRARY	VARCHAR(30)	Name of the library for the object for DMPYSOBJ. Contains the null value when SERVICE_TOOL is not DS.
DMPYSOBJ_OBJECT	VARCHAR(30)	Name of the object for DMPYSOBJ. Contains the null value when SERVICE_TOOL is not DS.
DMPYSOBJ_TYPE	VARCHAR(7)	Type of the object for DMPYSOBJ. Contains the null value when SERVICE_TOOL is not DS.
DMPYSOBJ_ASP_NAME	VARCHAR(10)	ASP name for DMPYSOBJ object library. Contains the null value when SERVICE_TOOL is not DS.
DMPYSOBJ_ASP_NUMBER	INTEGER	ASP number for DMPYSOBJ object library. Contains the null value when SERVICE_TOOL is not DS.
AA_COMMAND_NAME	VARCHAR(30)	Advanced Analysis Command name. See QMGTTOOLS: Run AA Macros Contains the null value when SERVICE_TOOL is not GS.
EARLY_TRACE_ACTION	VARCHAR(6)	The action requested for early job tracing *OFF Early tracing turned off *ON Early tracing turned on *RESET Early tracing turned off and trace information deleted. Contains the null value when SERVICE_TOOL is not CE.
ARM_TRACE	VARCHAR(10)	ARM diagnostic trace. ACTIVATE Activate DEACTIVATE Deactivate Contains the null value when SERVICE_TOOL is not AR.
TRCTCPAPP_OPTION	VARCHAR(7)	The trace option specified on TRCTCPAPP. ENDED Collection of trace information ended and all trace information purged (no output created) STARTED Collection of trace information started STOPPED Collection of trace information stopped and trace information written to spooled file Contains the null value when SERVICE_TOOL is not TA.
APPLICATION_TRACED	VARCHAR(10)	The name of the application being traced. Contains the null value when SERVICE_TOOL is not AR or TA.
SERVICE_TOOLS_PROFILE	VARCHAR(10)	The name of the service tools profile used for Start System Service Tools (STRSST). Contains the null value when SERVICE_TOOL is not ST or OP.

Table 139. AUDIT_JOURNAL_ST table function (continued)

Column Name	Data Type	Description
CONSOLE_TYPE	VARCHAR(7)	The console type. *DIRECT *HMC *LAN Contains the null value when SERVICE_TOOL is not OP.
CONSOLE_ACTION	VARCHAR(9)	The console action. *RECOVERY *TAKEOVER Contains the null value when SERVICE_TOOL is not OP.
ADDRESS_FAMILY	VARCHAR(5)	The address family. *IPV4 *IPV6 Contains the null value when SERVICE_TOOL is not OP or CONSOLE_TYPE is not *LAN.
CURRENT_IP_ADDRESS	VARCHAR(45)	The IP address of the current console device for *LAN. Contains the null value when SERVICE_TOOL is not OP or CONSOLE_TYPE is not *LAN.
CURRENT_DEVICE_ID	VARCHAR(10)	The service tools device ID of the current console device for *LAN. Contains the null value when SERVICE_TOOL is not OP or CONSOLE_TYPE is not *LAN.
PREVIOUS_IP_ADDRESS	VARCHAR(45)	The IP address of the previous console device for *LAN. Contains the null value when SERVICE_TOOL is not OP or CONSOLE_TYPE is not *LAN or there is no previous console device.
PREVIOUS_DEVICE_ID	VARCHAR(10)	The service tools device ID of the previous console device for *LAN. Contains the null value when SERVICE_TOOL is not OP or CONSOLE_TYPE is not *LAN or there is no previous console device.
WATCH_SESSION	VARCHAR(10)	Watch session ID. Contains the null value when SERVICE_TOOL is not WE or WS.
SERVICE_TOOL_USERID	VARCHAR(10)	Service tools user ID if storage was altered from DST or *DEBUG if storage was altered by a remote service tool debugger. Contains the null value when SERVICE_TOOL is not AS.
USER_PROFILE	VARCHAR(10)	User profile name if storage was altered from SST. Contains the null value when SERVICE_TOOL is not AS.
LIC_RU_NAME	VARCHAR(8)	The Licensed Internal Code Replaceable Unit name. Contains the null value when SERVICE_TOOL is not AS.
ADDRESS_OF_ALTERED_STORAGE	BINARY(8)	Address of storage that was altered. Contains the null value when SERVICE_TOOL is not AS.
SEGMENT_TYPE	BINARY(2)	Type of segment that was altered. Contains the null value when SERVICE_TOOL is not AS.
NUMBER_OF_ALTERED_BYTES	INTEGER	The number of bytes of storage in ALTERED_STORAGE_VALUE that were changed. Contains the null value when SERVICE_TOOL is not AS.
ALTERED_STORAGE_VALUE	BINARY(16)	Altered storage value. Contains the null value when SERVICE_TOOL is not AS.
PREVIOUS_STORAGE_VALUE	BINARY(16)	Original storage value. Contains the null value when SERVICE_TOOL is not AS.

Example

- Check whether anyone has changed storage using STRSST today.

```
SELECT *
FROM TABLE (
    SYSTOOLS.AUDIT_JOURNAL_ST (STARTING_TIMESTAMP => CURRENT DATE)
)
WHERE SERVICE_TOOL = 'ST' AND NUMBER_OF_ALTERED_BYTES > 0;
```

AUDIT_JOURNAL_SV table function

The AUDIT_JOURNAL_SV table function returns rows from the audit journal that contain information from the SV (Action to System Value) journal entries.

Every audit journal table function shares a common authorization requirement and a common set of parameters. These are described in [“AUDIT JOURNAL table function common information”](#) on page 552.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 140. AUDIT_JOURNAL_SV table function

Column Name	Data Type	Description
The first columns returned by this table function are from the common audit journal entry header. See Common columns returned from the audit journal entry header for the column definitions. After the common columns are the following columns that describe the entry specific data for the SV audit journal entry.		
ENTRY_TYPE	CHAR(1)	The type of entry. A Change to system values B Change to service attributes C Change to system clock D Adjustment to Coordinated Universal Time (UTC) E Change to option F Change to system-wide journal attribute
ENTRY_TYPE_DETAIL	VARCHAR(200)	Descriptive text that corresponds to the entry type.
SYSTEM_VALUE	VARCHAR(10)	The name of the system value or service attribute that was changed. Can contain one of the following special values. CACHEWAIT Changed journal maximum cache wait time JRNRCYCNT Changed journal recovery count value QINPIDCO Change the current install disk configuration option with QINPIDCO API.
OLD_VALUE	VARCHAR(1500)	The value of the system value or service attribute before it was changed.
NEW_VALUE	VARCHAR(1500)	The new value of the system value or service attribute.

Example

- List any system value changes that have been made in the current calendar year.

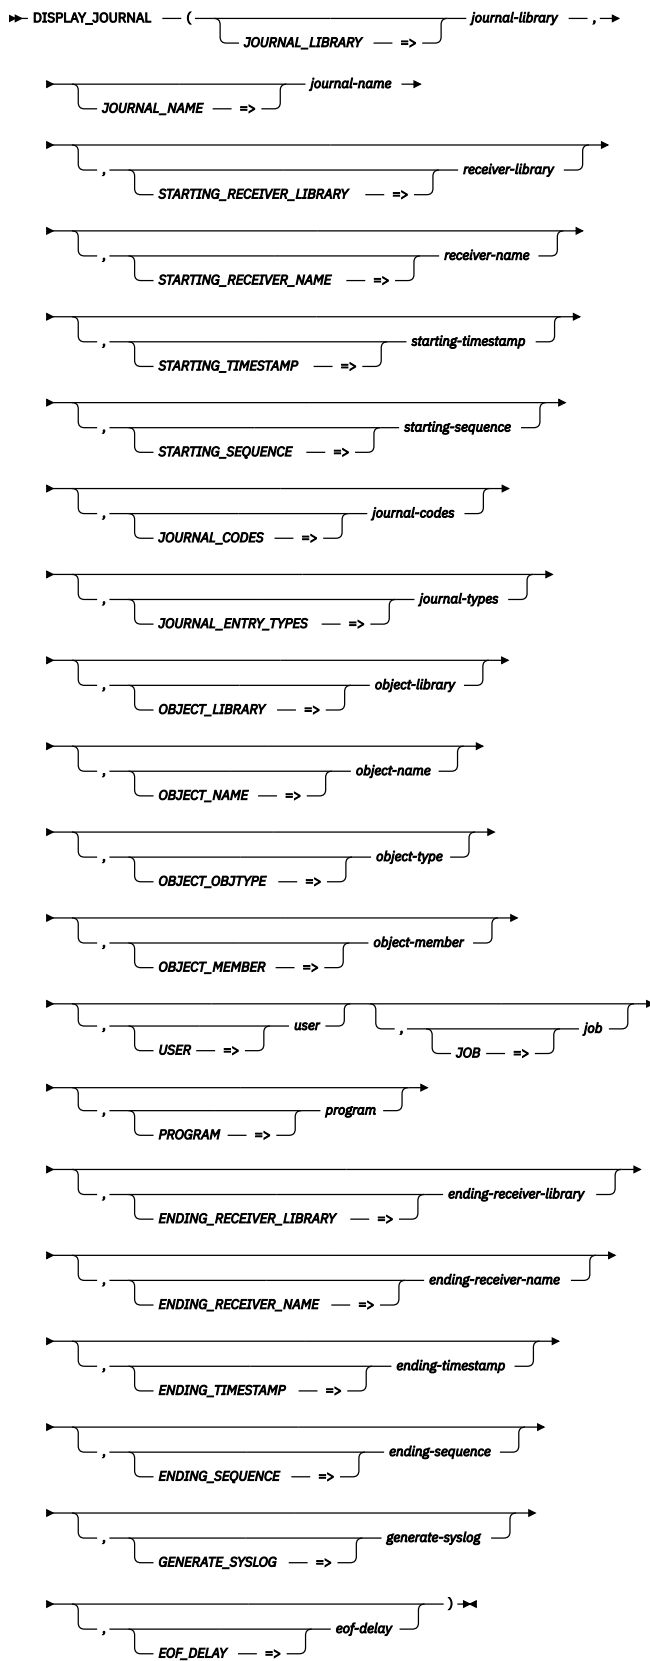
```
SELECT *
FROM TABLE (
    SYSTOOLS.AUDIT_JOURNAL_SV (STARTING_TIMESTAMP => TRUNC_TIMESTAMP(CURRENT DATE, 'YEAR'))
)
WHERE ENTRY_TYPE = 'A';
```

DISPLAY_JOURNAL table function

The DISPLAY_JOURNAL table function returns information about journal entries. It returns information similar to what is returned by the Display Journal (DSPJRN) CL command and the Retrieve Journal Entries (QjoRetrieveJournalEntries) API.

Authorization:

- The caller must have *USE authority to the journal and to all requested journal receivers.
- *OBJEXIST authority is required to the journal if *object-name* is omitted or if *object-name* specifies an object that no longer exists.
- If *object-name* is *ALL, the caller must be authorized to every object associated with a journal entry.
 - For objects in the QSYS file system, the caller must have:
 - *USE authority to the object, and
 - *EXECUTE to the library containing the object.
 - For objects in the integrated file system, the caller must have:
 - *R for the object
 - *X for every directory in the object's path.



The schema is QSYS2.

journal-library A character or graphic string expression that identifies the name of the library containing the journal. The name cannot be *LIBL or *CURLIB.

journal-name	A character or graphic string expression that identifies the name of the journal.
receiver-library	A character or graphic string expression that identifies the name of the starting journal receiver library. The name can be *LIBL or *CURLIB.
receiver-name	A character or graphic string expression that identifies the name of the starting journal receiver. If one of the special values is specified, the <i>receiver-library</i> value will be ignored. Otherwise, the <i>receiver-name</i> and <i>receiver-library</i> must identify a valid journal receiver. If no journal receiver is specified, *CURRENT is used.
*CURRENT	The journal receiver that is currently attached when starting to convert journal entries is used.
*CURCHAIN	The journal receiver chain that includes the journal receiver that is currently attached when starting to convert journal entries is used. This receiver chain does not cross a break in the chain. If there is a break in the chain, the receiver range is from the most recent break in the chain through the receiver that is attached when starting to convert journal entries.
*CURAVLCHN	The journal receiver chain that includes the journal receiver that is attached when starting to convert journal entries is used. This receiver chain does not cross a break in the chain. If there is a break in the chain, the receiver range is from the most recent break in the chain through the receiver that is attached when starting to convert journal entries. If journal receivers exist in the receiver chain that are not available because they were saved with the storage freed option, those journal receivers will be ignored and entries will be converted starting with the first available journal receiver in the chain.
*CURSEQCHN	The journal receiver chain that includes the journal receiver that is attached when starting to convert journal entries is used, starting with the most recent journal receiver in which the journal sequence number was reset if the journal sequence number was reset in the receiver chain. This receiver chain does not cross a break in the chain. If there is a break in the chain, the receiver range is from the most recent break in the chain through the receiver that is attached when starting to convert journal entries. If journal receivers exist in the receiver chain that are not available because they were saved with the storage freed option, those journal receivers will be ignored and entries will be converted starting with the first available journal receiver in the chain.
starting-timestamp	A timestamp value that specifies the starting timestamp to use ² . A value cannot be specified for both <i>starting-timestamp</i> and <i>starting-sequence</i> . If no starting timestamp is specified, *FIRST is used.
starting-sequence	A decimal expression that identifies the starting sequence number to use. If the <i>starting-sequence</i> value is not found in the receiver range, an error is returned. A value cannot be specified for both <i>starting-timestamp</i> and <i>starting-sequence</i> . If no starting sequence is specified, *FIRST is used.
journal-codes	A character or graphic string expression that lists the journal codes to return. The string can contain the special values of *ALL or *CTL, or it can be a list of one or more

² The accuracy of the entry timestamp stored in journal receivers is only accurate to 16 microseconds. Hence, a value passed as a *starting_timestamp* and *ending_timestamp* will be truncated such that the actual timestamps being searched for may be from 0 to 15 microseconds less than the specified value.

journal codes. Journal codes in the string can be separated by one or more separators. Separators are blank and comma. For example, a valid string can be 'RJ' or 'R J' or 'R,J' or 'R, J'.

If no string is provided, *ALL is used.

journal-types

A character or graphic string expression that lists the journal entry types to return. The string can contain the special values of *ALL or *RCD, or it can be a list of one or more journal entry types. Journal entry types in the string can be separated by one or more separators. Separators are blank and comma. For example, a valid string can be 'JFCT' or 'JF CT' or 'JF,CT' or 'JF, CT'.

If no string is provided, *ALL is used.

object-library

A character or graphic string expression that identifies the name of an object library. The values *LIBL and *CURLIB are allowed.

object-name

A character or graphic string expression that contains up to 300 object names. Multiple names can be separated by one or more separators. Separators are blank and comma.

If *object-name* is the special value of *ALL, *object-library* must contain a library name and *object-type* must contain a valid object type. Otherwise, each object name must identify a valid object using the same *object-library*, *object-type*, and *object-member* parameters.

If no object name is provided, a value of *ALLFILE is used for the journaled file name on the API interface.

object-type

A character or graphic string expression that identifies the system object type for the object. The value must be *DTAARA, *DTAQ, *FILE, or *LIB.

object-member

A character or graphic string expression that identifies the name of a member. It can be a special value of *FIRST, *ALL, or *NONE or a valid member name. If the object type is not *FILE, the member name is ignored.

user

A character or graphic string expression that identifies the user profile name for the current user of the job. If *user* is not specified, *ALL is used.

job

A character or graphic string expression that identifies the name of a job. Two forms of a job name are supported.

1. A fully qualified job name in the form *job-number/job-user/job-name*.

2. The first 10 characters are the job name, the second 10 characters are the user name, and the last 6 characters are the job number.

If *job* is not specified, *ALL is used.

program

A character or graphic string expression that identifies the name of a program. If *program* is not specified, *ALL is used.

ending-receiver-library

A character or graphic string expression that identifies the name of the ending journal receiver library. The name can be *LIBL or *CURLIB. If *ending-receiver-name* is not *CURRENT, a value for *ending-receiver-library* must be specified.

The value of this parameter is ignored if *eof-delay* is greater than zero.

ending-receiver-name

A character or graphic string expression that identifies the name of the ending journal receiver. If the special value *CURRENT is specified, the *ending-receiver-library* value will be ignored. Otherwise, the *ending-receiver-name* and *ending-receiver-library* must identify a valid journal receiver.

If *ending-receiver-name* is not specified, *CURRENT is used.

The value of this parameter is ignored if *eof-delay* is greater than zero.

ending-timestamp

A timestamp value that specifies the ending timestamp to use².

A value cannot be specified for both *ending-timestamp* and *ending-sequence*. This parameter cannot be specified if *eof-delay* is greater than zero.

If no ending timestamp is specified, *LAST is used.

ending-sequence

A decimal expression that identifies the ending sequence number to use. If the *ending-sequence* value is not found in the receiver range, an error is returned.

A value cannot be specified for both *ending-timestamp* and *ending-sequence*. This parameter cannot be specified if *eof-delay* is greater than zero.

If no ending sequence is specified, *LAST is used.

When *ending-sequence* is used, the query results will end when the first ending sequence value is encountered. If the journal has had its sequence numbers reset, *ending-sequence* will only return results through the first match of *ending-sequence*.

generate-syslog

A character or graphic string expression that indicates whether to transform journal entries into syslog formatted detail. Values are:

NO No syslog information will be returned. The SYSLOG_EVENT, SYSLOG_FACILITY, SYSLOG_SEVERITY, and SYSLOG_PRIORITY columns will contain the null value.

RFC3164 Values will be returned for the SYSLOG_EVENT, SYSLOG_FACILITY, SYSLOG_SEVERITY, and SYSLOG_PRIORITY columns if syslog information is defined for the journal entry. The SYSLOG_EVENT column will contain a syslog header that matches the RFC3164 format as described by the Internet Engineering Task Force (IETF) Request For Comments (RFC) 3164.

RFC5424 Values will be returned for the SYSLOG_EVENT, SYSLOG_FACILITY, SYSLOG_SEVERITY, and SYSLOG_PRIORITY columns if syslog information is defined for the journal entry. The SYSLOG_EVENT column will contain a syslog header that matches the RFC5424 format as described by the Internet Engineering Task Force (IETF) Request For Comments (RFC) 5424.

DISPLAY_JOURNAL only returns syslog information for the audit journal. If RFC3164 or RFC5424 is specified, *journal-library* must be QSYS and *journal-name* must be QAUDJRN.

If *generate-syslog* is not specified or is the null value, NO is used.

eof-delay

An integer expression that specifies the number of seconds to sleep when all audit journal entries have been read. This delay allows the caller to establish a polling service that will continually return rows, sleeping for the specified interval whenever all entries have been processed.

A value of zero indicates no delay is used and a finite set of rows will be returned. A value greater than zero indicates that the table function will sleep, as needed, to wait for new audit journal entries and never end. If *eof-delay* is not specified or is the null value, zero is used.

If this parameter has a value greater than zero, the *generate-syslog* parameter must be RFC3164 or RFC5424, the *ending-receiver-library* and *ending-receiver-name* are ignored, and the *ending-timestamp* and *ending-sequence* parameters cannot be specified with a value other than their default values.

When using a non-zero *eof-delay* parameter, avoid using query clauses that depend on returning a finite number of rows. For example, using the FETCH FIRST n ROWS clause can cause the query to end when the requested number of rows has been satisfied. A query using the DISPLAY_JOURNAL function with a non-zero *eof-delay* parameter does not allow data to be copied (ALWCOPYDTA(*NO)). This means that a query requiring a

copy of data, such as one using an ORDER BY clause or UNION DISTINCT, will issue an error and not be allowed. When using *eof-delay*, consider using a simple query to avoid blocking of rows. When rows are blocked for data transport efficiency, rows won't be returned until the block is full. Therefore, you should decide whether you favor data transport efficiency or moving events as soon as they occur.

The special values supported for the function arguments are the same as for the Display Journal (DSPJRN) CL command.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 141. DISPLAY_JOURNAL table function

Column Name	Data Type	Description
ENTRY_TIMESTAMP	TIMESTAMP	The system date and time when the journal entry was added to the journal receiver ² .
SEQUENCE_NUMBER	DECIMAL(21,0)	A number assigned by the system to each journal entry.
JOURNAL_CODE	CHAR(1)	The primary category of the journal entry.
JOURNAL_ENTRY_TYPE	CHAR(2)	Further identifies the type of user-created or system-created entry.
COUNT_OR_RRN	BIGINT	Contains either the relative record number (RRN) of the record that caused the journal entry or a count that is pertinent to the specific type of journal entry.
ENTRY_DATA	BLOB(2G)	The entry specific data returned for this journal entry. See Notes section for row and column access control considerations.
NULL_VALUE_INDICATORS	VARCHAR(8000)	The null value indicators returned for this journal entry.
OBJECT	VARCHAR(30)	The name of the object for which the journal entry was added.
OBJECT_TYPE	VARCHAR(10)	The type of object in the entry.
OBJECT_TYPE_INDICATOR	CHAR(1)	An indicator with respect to the information in the object field.
FILE_TYPE_INDICATOR	CHAR(1)	Identifies whether or not this journal entry is associated with a logical file.
JOURNAL_IDENTIFIER	VARCHAR(10)	The journal identifier (JID) for the object.
USER_NAME	VARCHAR(10)	The name of the effective user profile under which the job was running when the entry was created. This value is identical to what is returned in the CURRENT_USER column.
JOB_NAME	VARCHAR(10)	The name of the job that added the entry.
JOB_USER	VARCHAR(10)	The user profile name of the user that started the job.
JOB_NUMBER	VARCHAR(6)	The job number of the job that added the entry.
THREAD	BIGINT	Identifies the thread within the process that added the journal entry.
PROGRAM_NAME	VARCHAR(10)	The name of the program that added the entry.
PROGRAM_LIBRARY	VARCHAR(10)	The name of the library that contains the program that added the journal entry.
PROGRAM_LIBRARY_ASP_DEVICE	VARCHAR(10)	The name of the ASP device that contains the program.
PROGRAM_LIBRARY_ASP_NUMBER	INTEGER	The number for the auxiliary storage pool that contains the program that added the journal entry.
COMMIT_CYCLE	DECIMAL(21,0)	A number that identifies the commit cycle.
NESTED_COMMIT_LEVEL	BIGINT	Indicates the nesting level of the commit cycle that was open when a journal entry representing an object level change was deposited.
XID	VARCHAR(140)	The transaction identifier, as defined by the Open Group's XA specification, for commit cycles related to an XA transaction branch.
LUW	VARCHAR(39)	The logical unit of work identifies entries to be associated with a given unit of work.
REMOTE_PORT	INTEGER	The port number of the remote address associated with this journal entry.
REMOTE_ADDRESS	VARCHAR(46)	The remote address associated with the journal entry.
SYSTEM_NAME	VARCHAR(8)	The name of the system on which the entry is being retrieved.

Table 141. DISPLAY_JOURNAL table function (continued)

Column Name	Data Type	Description
SYSTEM_SEQUENCE_NUMBER	DECIMAL(21,0)	The system sequence number indicates the relative sequence of when this journal entry was deposited into the journal.
REFERENTIAL_CONSTRAINT	CHAR(1)	Whether this entry was recorded for actions that occurred on records that are part of a referential constraint.
TRIGGER	CHAR(1)	Whether this entry was created as result of a trigger program.
IGNORE_ON_APPLY	CHAR(1)	Whether this entry is ignored during an Apply Journalized Changes (APYJRNCHG) or Remove Journalized Changed (RMVJRNCHG) command.
MINIMIZED_ENTRY_DATA	CHAR(1)	Whether this entry has minimized entry specific data as a result of the journal having specified MINENTDTA for the object type of the entry.
MINIMIZED_ON_FIELD_BOUNDARY	CHAR(1)	Whether this entry has minimized entry specific data on field boundaries as a result of the journal having been specified with MINENTDTA(*FLDBDY).
INDICATOR_FLAG	CHAR(1)	An indicator for the operation.
RECEIVER_NAME	VARCHAR(10)	The name of the receiver holding the journal entry.
RECEIVER_LIBRARY	VARCHAR(10)	The name of the library containing the receiver holding the journal entry.
RECEIVER_ASP_DEVICE	VARCHAR(10)	The name of the ASP device containing the receiver holding the journal entry.
RECEIVER_ASP_NUMBER	INTEGER	The number for the auxiliary storage pool containing the receiver holding the journal entry.
ARM_NUMBER	INTEGER	The number of the disk arm that contains the journal entry.
OBJECT_ASP_DEVICE	VARCHAR(10)	ASP device name.
OBJECT_ASP_NUMBER	INTEGER	ASP number.
PARENT_FILE_ID	BINARY(16)	File ID for parent directory.
OBJECT_FILE_ID	BINARY(16)	File ID for object.
RELATIVE_DIRECTORY_FILE_ID	BINARY(16)	File ID of directory containing object in PATH_NAME.
OBJECT_FILE_NAME	VARGRAPHIC(2002) CCSID 1200	Object name.
PATH_NAME	DBCLOB(16M) CCSID 1200	Name of IFS path.
DLO_NAME	VARCHAR(12)	DLO name.
FOLDER_PATH	VARCHAR(63)	DLO folder path.
CURRENT_USER	VARCHAR(10)	The name of the effective user profile under which the job was running when the entry was created. This value is identical to what is returned in the USER_NAME column.
SYSLOG_EVENT	VARGRAPHIC(2048) CCSID 1200	The Common Event Format (CEF) syslog event for the journal entry preceded with a header of the requested type. If a header-type of RFC3164 is requested, the maximum length is 1024 characters. If a header-type of RFC5424 is requested, the maximum length is 2048 characters. The string will be truncated with no warning if it exceeds the maximum length. The audit journal entry types that generate syslog information and the key names returned for journal entries are listed in the Notes section. Contains the null value if there is no syslog event defined for the journal entry or if NO was specified for the GENERATE_SYSLOG parameter.
SYSLOG_FACILITY	INTEGER	The syslog facility assigned to the event. 4 Security/authorization messages. This value is returned for all T and U audit journal entries. Contains the null value if there is no syslog event defined for the journal entry or if NO was specified for the GENERATE_SYSLOG parameter.

Table 141. DISPLAY_JOURNAL table function (continued)

Column Name	Data Type	Description
SYSLOG_SEVERITY	INTEGER	<p>The syslog severity assigned to the event.</p> <p>2 Critical condition</p> <p>4 Warning condition</p> <p>5 Notice: A normal but significant condition</p> <p>6 Informational message</p> <p>The severity assigned to each journal entry is listed in the Notes section.</p> <p>Contains the null value if there is no syslog event defined for the journal entry or if NO was specified for the GENERATE_SYSLOG parameter.</p>
SYSLOG_PRIORITY	INTEGER	<p>The syslog priority number assigned to the event.</p> <p>Contains the null value if there is no syslog event defined for the journal entry or if NO was specified for the GENERATE_SYSLOG parameter.</p>

Notes

Row and column access control: This table function recognizes whether ROW ACCESS CONTROL or COLUMN ACCESS CONTROL exists and is activated for the target table. If any row or column access control is active for the table, the rule text logic defined for the row permissions and/or column masks is applied before returning the value in ENTRY_DATA. When the rule text for a row permission determines that the user invoking the function should not see the row, the ENTRY_DATA column contains the text NOT AUTHORIZED. If the user is allowed to see the row and a column mask exists, the rule text for the column mask determines the value returned for ENTRY_DATA.

LOB data considerations: For journal code R (any entry type) and code F (IZ entry type), when a LOB data type is encountered that is not null or a zero-length string, 16 'Q's are placed in the Entry Specific Data, followed by the LOB data.

Filtering considerations: When using DISPLAY_JOURNAL to review journal activity for specific objects, there are some considerations to ensure complete results are returned.

- When querying the audit journal (JOURNAL_LIBRARY => 'QSYS' and JOURNAL_NAME => 'QAUDJRN'), do not use the object filters because they will result in no entries being returned. For an audit journal, use a WHERE clause to limit the rows returned.
- When querying a data journal, consider whether it's possible that the Journal Identifier (JID) has changed for the object. The object filters use the JID found in the object and will not return any entries with a different JID. If the JID for the object might have changed, avoid using the object filters. In this case, use a WHERE clause to limit the rows returned.

Syslog information: Syslog information is returned for all audit journal entries with T and U journal codes. Syslog information is also available for history log messages. See [HISTORY_LOG_INFO table function](#) for more details.

The following audit journal entries with T journal codes generate syslog information:

AD	Auditing changes
AF	Authority failure
AP	Obtaining adopted authority
AU	Attribute changes
AX	Row and column access control
CA	Authority changes
CD	Command string audit
CO	Create object
CP	User profile changed, created, or restored

CQ	Change of *CRQD object
CU	Cluster operations
CV	Connection verification
CY	Cryptographic configuration
DI	Directory server
DO	Delete object
DS	DST security password reset
EV	System environment variables
GR	Generic record
GS	Socket description was given to another job
IM	Intrusion monitor
IP	Interprocess communication
IR	IP rules actions
IS	Internet security management
JD	Change to user parameter of a job description
JS	Actions that affect jobs
KF	Key ring file
LD	Link, unlink, or look up directory entry
ML	Office services mail actions
NA	Network attribute changed
ND	APPN directory search filter violation
NE	APPN end point filter violation
OM	Object move or rename
OR	Object restore
OW	Object ownership changed
O1	Optical access
O2	Optical access
O3	Optical access
PA	Program changed to adopt authority
PF	PTF operations
PG	Change of an object's primary group
PO	Printed output
PS	Profile swap
PU	PTF object changes
PW	Invalid password
RA	Authority change during restore
RJ	Restoring job description with user profile specified
RO	Change of object owner during restore

RP	Restoring adopted authority program
RQ	Restoring a *CRQD object
RU	Restoring user profile authority
RZ	Changing a primary group during restore
SD	Changes to system distribution directory
SE	Subsystem routing entry changed
SF	Actions to spooled files
SG	Asynchronous signals
SK	Sockets connections
SM	Systems management changes
SO	Server security user information actions
ST	Use of service tools
SV	System value changed
VO	Validation list actions
VP	Network password error
XD	Directory server extension
XO	Network authentication
X1	Identity token
X2	Query manager profile changes
YC	DLO object accessed (change)
YR	DLO object accessed (read)
ZC	Object accessed (change)
ZR	Object accessed (read)

Audit journal entries with the U journal code are assigned a SYSLOG_SEVERITY of 6. The audit journal entries with T journal codes are assigned a SYSLOG_SEVERITY value in the following way:

- Severity 2 Critical condition
 - SV - System value when QAUDCTL is changed to *NONE
- Severity 4 Warning condition
 - AF - Authority failure
 - DI - Directory server (operation type 'AF')
 - GR - Generic record, when function usage was checked and failed for a function name with a prefix of QIBM_DB_
 - IM - Intrusion monitor
 - IP - Interprocess communication (entry type 'F')
- Severity 5 Notice: A normal but significant condition
 - AD - Auditing changes
 - AX - Row and column access control
 - CA - Authority changes
 - CP - User profile changed, created, or restored
 - DI - Directory server (operation types 'AD', 'CA', 'CP', 'OM', 'OW', and 'PW')

- DS - DST security password reset
- IP - Interprocess communication (entry type 'A')
- JD - Change to user parameter of a job description
- JS - Actions that affect jobs (entry types 'M' and 'T')
- OM - Object move or rename
- OW - Object ownership changed
- O3 - Optical access (entry type 'L')
- PG - Change of an object's primary group
- PS - Profile swap
- PW - Invalid password
- RA - Authority change during restore
- RO - Change of object owner during restore
- RU - Restoring user profile authority
- RZ - Change a primary group during restore
- SO - Server security user information actions
- VO - Validation list actions (entry type 'U')
- VP - Network password error
- X0 - Network authentication (entry types '2' - '6', '8', '9' and 'A' - 'F')
- X1 - Identity token (entry types 'F' and 'U')
- X2 - Query manager profile changes
- Severity 6 Informational message
 - AP - Obtaining adopted authority
 - AU - Attribute changes
 - CD - Command string audit
 - CO - Create object
 - CQ - Change of *CRQD object
 - CU - Cluster operations
 - CV - Connection verification
 - CY - Cryptographic configuration
 - DI - Directory server (all operation types other than 'AD', 'AF', 'CA', 'CP', 'OM', 'OW', and 'PW')
 - DO - Delete object
 - EV - System environment variables
 - GR - Generic record, except for the Severity 4 case where function usage was checked and failed
 - GS - Socket description was given to another job
 - IP - Interprocess communication (all entry types other than 'A' and 'F')
 - IR - IP rules actions
 - IS - Internet security management
 - JS - Actions that affect jobs (all entry types other than 'M' and 'T')
 - KF - Key ring file
 - LD - Link, unlink, or look up directory entry
 - ML - Office services mail actions
 - NA - Network attribute changed
 - ND - APPN directory search filter violation

- NE - APPN end point filter violation
- OR - Object restore
- O1 - Optical access
- O2 - Optical access
- O3 - Optical access (all entry types other than 'L')
- PA - Program changed to adopt authority
- PF - PTF operations
- PO - Printed output
- PU - PTF object changes
- RJ - Restoring job description with user profile specified
- RP - Restoring adopted authority program
- RQ - Restoring a *CRQD object
- SD - Changes to system distribution directory
- SE - Subsystem routing entry changed
- SF - Actions to spooled files
- SG - Asynchronous signals
- SK - Sockets connections
- SM - Systems management changes
- ST - Use of service tools
- SV - System value changed, except for QAUDCTL severity 2 case
- VO - Validation list actions (all entry types other than 'U')
- XD - Directory server extension
- X0 - Network authentication (all entry types other than '2' - '6', '8', '9' and 'A' - 'F')
- X1 - Identity token (all entry types other than 'F' and 'U')
- YC - DLO object accessed (change)
- YR - DLO object accessed (read)
- ZC - Object accessed (change)
- ZR - Object accessed (read)

The Common Event Format key names that are generated within the SYSLOG_EVENT column are:

<i>Table 142. Common Event Format key names</i>	
Common Event Format key name	Description
attrName	Attribute name (extracted from ENTRY_DATA column)
attrValue	Attribute value (extracted from ENTRY_DATA column)
deviceExternalId	Device name (extracted from ENTRY_DATA column)
dloName	Document Library Object name (DLO_NAME column)
dloPath	Document Library Object folder path (FOLDER_PATH column)
dproc	Destination job (process) name (extracted from ENTRY_DATA column)

Table 142. Common Event Format key names (continued)

Common Event Format key name	Description
dpt	Destination port number (extracted from ENTRY_DATA column)
dst	Destination IP address (extracted from ENTRY_DATA column)
duser	Destination user name (extracted from ENTRY_DATA column)
filePath	IFS stream file path (PATH_NAME column)
fileType	Object type (OBJECT_TYPE column)
fname	IFS stream file name (OBJECT_FILE_NAME column)
msg	Additional information from the audit record not included in other keys (extracted from ENTRY_DATA column)
objName	Object name (OBJECT column)
oldAttrValue	Attribute value (before change) (extracted from ENTRY_DATA column)
oldDloName	Document Library Object name (before rename) (extracted from ENTRY_DATA column)
oldDloPath	Document Library Object folder path (before rename) (extracted from ENTRY_DATA column)
oldFileName	IFS stream file name (before rename) (extracted from ENTRY_DATA column)
oldFilePath	IFS stream file path (before rename) (extracted from ENTRY_DATA column)
oldObjName	Object name (before rename) (extracted from ENTRY_DATA column)
reason	Text description of the audit journal entry
shost	Source system (host) name (SYSTEM_NAME column)
sproc	Source job (process) name (JOB_NAME, JOB_USER, JOB_NUMBER columns)
spt	Source port number (REMOTE_PORT column)
src	Source IP address (REMOTE_ADDRESS column)
suser	Source user name (USER_NAME column)

Examples

- Select all entries from the *CURRENT receiver of journal TESTLIB/QSQJRN.

```
SELECT * FROM TABLE (
    QSYS2.DISPLAY_JOURNAL( 'TESTLIB', 'QSQJRN')) AS JT;
```

- Find all changes made by SUPERUSER against the PRODDATA/SALES table. The first two arguments are passed without names since they correspond with the first two parameters for the function. The

other four arguments are passed using the parameter name syntax to avoid specifying a value for the parameters that are not needed.

```
SELECT journal_code, journal_entry_type, object, object_type, X.*
FROM TABLE (
  QSYS2.Display_Journal(
    'PRODDATA', 'QSQJRN', -- Journal library and name
    OBJECT_LIBRARY=>'PRODDATA', OBJECT_NAME=>'SALES',
    OBJECT_OBJTYPE=> '*FILE', OBJECT_MEMBER=>'SALES'
  ) ) AS X
WHERE journal_entry_type in ('DL', 'PT', 'PX', 'UP') AND "CURRENT_USER" = 'SUPERUSER'
ORDER BY entry_timestamp DESC;
```

- Review audit journal entries for the REQUESTS file in MYCO library. For an audit journal, a predicate is used to designate the object name.

```
SELECT journal_code, journal_entry_type, object, object_type, X.*
FROM TABLE (QSYS2.Display_Journal('QSYS', 'QAUDJRN') ) AS X
WHERE LEFT(OBJECT,20) = CHAR('REQUESTS', 10) CONCAT CHAR('MYCO', 10)
ORDER BY entry_timestamp DESC;
```

- Review changes from the last hour for the REQUESTS file in MYCO library using the current JID. This query will only find entries where the JID of MYCO/REQUESTS *FILE is an exact match to the entry. It filters for the following three journal codes:

- D** Database File Operation
- F** Database File Member Operation
- R** Operation on Specific Record

```
SELECT journal_code, journal_entry_type, object, object_type, X.*
FROM TABLE (QSYS2.Display_Journal('MYCO', 'QSQJRN',
  JOURNAL_CODES => 'D,F,R',
  STARTING_RECEIVER_NAME => '*CURCHAIN',
  OBJECT_OBJTYPE=> '*FILE',
  OBJECT_LIBRARY=> 'MYCO',
  OBJECT_NAME=> 'REQUESTS',
  OBJECT_MEMBER=> '*ALL'
) ) AS X
WHERE entry_timestamp > CURRENT_TIMESTAMP - 1 HOUR
ORDER BY entry_timestamp DESC ;
```

- Review all changes from the last hour for the REQUESTS file in MYCO library, including any that might have a different Journal ID (JID). To see entries for all JIDs, a predicate is used to designate the object name.

```
SELECT journal_code, journal_entry_type, hex( journal_identifier ),
  object, object_type, X.*
FROM TABLE (QSYS2.Display_Journal('MYCO', 'QSQJRN',
  JOURNAL_CODES => 'D,F,R',
  STARTING_RECEIVER_NAME => '*CURCHAIN'
) ) AS X
WHERE LEFT(OBJECT,20) = CHAR('REQUESTS', 10) CONCAT CHAR('MYCO', 10)
  and entry_timestamp > CURRENT_TIMESTAMP - 1 HOUR
ORDER BY entry_timestamp DESC ;
```

- Select entries from the audit journal that return syslog information and format them with an RFC5424 header.

```
SELECT syslog_facility, syslog_severity, syslog_event
FROM TABLE (QSYS2.DISPLAY_JOURNAL('QSYS', 'QAUDJRN',
  GENERATE_SYSLOG => 'RFC5424'
) ) AS X
WHERE syslog_event IS NOT NULL;
```

JOURNAL_INFO view

The JOURNAL_INFO view contains information about journals, including remote journals.

The values returned for the columns in the view are closely related to the values returned by [Retrieve Journal Information API](#). Refer to the API for more detailed information.

The following table describes the columns in the view. The system name is JRNINFO. The schema is QSYS2.

Table 143. JOURNAL_INFO view

Column Name	System Column Name	Data Type	Description
JOURNAL_NAME	JRNNAME	VARCHAR(10)	The name of the journal.
JOURNAL_LIBRARY	SYS_DNAME	VARCHAR(10)	The name of the library that contains the journal.
ASP_NUMBER	ASPNUMBER	INTEGER	The number of the auxiliary storage pool to which storage for the journal is allocated.
JOURNAL_ASPPGRP	JRNASPPGRP	VARCHAR(10)	The name of the auxiliary storage pool (ASP) in which the journal resides. A value of *SYSBAS indicates the system ASP and all basic user ASPs.
ATTACHED_JOURNAL_RECEIVER_NAME	ATTRCVNAME	VARCHAR(10) Nullable	The name of the journal receiver that is currently attached to this journal. Contains the null value when there is no attached receiver.
ATTACHED_JOURNAL_RECEIVER_LIBRARY	ATTRCVLIB	VARCHAR(10) Nullable	The name of the library that contains the attached journal receiver. Contains the null value when there is no attached receiver.
MESSAGE_QUEUE	MSGQNAME	VARCHAR(10)	The name of the message queue that is associated with this journal.
MESSAGE_QUEUE_LIBRARY	MSGQLIB	VARCHAR(10)	The name of the library that contains the message queue.
DELETE_RECEIVER_OPTION	DLTRCOPT	VARCHAR(3)	Indicates whether the system deletes detached journal receivers that are associated with this journal when they are no longer needed for IPL or IASP vary on recovery. NO Detached journal receivers that are associated with this journal are not deleted when they are no longer needed for IPL or IASP vary on recovery. YES Detached journal receivers that are associated with this journal are deleted when they are no longer needed for IPL or IASP vary on recovery.
DELETE_RECEIVER_DELAY	DLTRCVDLY	INTEGER Nullable	The delay time (in minutes) between attempts to delete journal receivers associated with this journal. Contains the null value when DELETE_RECEIVER_OPTION is NO.
JOURNAL_TYPE	TYPE	VARCHAR(10)	The scope of the journal and some of its characteristics. *LOCAL This is a local journal. *REMOTE This is a remote journal.

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
JOURNAL_STATE	STATE	VARCHAR(10)	<p>An indication as to whether journal entries are currently being sent to a journal. For a remote journal, this is whether the journal is actively receiving journal entries from the source system journal.</p> <p>*ACTIVE If this is a local journal, this means journal entries can be deposited to this journal. If this is a remote journal, this means journal entries can be received from a source journal.</p> <p>*CTLINACT The remote journal is in the process of a controlled inactivate.</p> <p>*FAILED If this is a remote journal, this means journal entries cannot be received from a source journal due to a remote journal function failure. Does not apply to local journals.</p> <p>*INACTIVE If this is a remote journal, this means journal entries cannot be received from a source journal.</p> <p>*INACTPEND If this is a remote journal, this means a request is being processed to set the journal state to *INACTIVE. Does not apply to local journals.</p> <p>*PENDING The remote journal is transitioning from an *INACTIVE state to an *ACTIVE state.</p> <p>*STANDBY If this is a local journal, this means that most journal entries are not deposited into the journal and there will be no errors indicating that the entry was not deposited. Does not apply to remote journals.</p>
NUMBER_JOURNAL_RECEIVERS	NUMJRNRCV	INTEGER	The total number of journal receivers that are associated with the journal.
TOTAL_SIZE_JOURNAL_RECEIVERS	SIZJRNRCV	BIGINT	The total size of the journal receivers (in kilobytes) associated with the journal.
NUMBER_REMOTE_JOURNALS	RMTJRNS	INTEGER	The total number of remote journals that are directly downstream of this journal.
REDIRECTED_RECEIVER_LIBRARY	RDRRCVLIB	VARCHAR(10) Nullable	<p>For a local or *TYPE1 remote journal, the redirected receiver library name that is currently in place on this journal's local journal for any downstream journal receivers associated with *TYPE1 remote journals.</p> <p>Contains *NONE if no *TYPE1 remote journals have been added or if no receiver library redirection was specified when *TYPE1 remote journals were added.</p> <p>Contains the redirected receiver library name that is currently in place on this remote journal if the specified journal is a *TYPE2 remote journal.</p> <p>Contains the null value if no *TYPE1 remote journals have been added.</p>
MAXIMUM_REMOTE_JOURNALS_ENTRIES_BEHIND	MAXRMTENTB	INTEGER Nullable	<p>The maximum number of entries that are waiting to be sent to the target system for any remote journal.</p> <p>Contains the null value if NUMBER_REMOTE_JOURNALS is 0 or if no attached remote journals are active with async delivery mode.</p>

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MAXIMUM_REMOTE_JOURNALS_TIME_BEHIND	MAXRMTSECB	BIGINT Nullable	The maximum value (in hundredths of seconds) that the source journal is behind in sending journal entries to the target system for any remote journal. Contains the null value if NUMBER_REMOTE_JOURNALS is 0 or if no attached remote journals are active with async delivery mode.
MAXIMUM_REMOTE_JOURNALS_RETRANSMISSIONS	MAXRMTRETR	BIGINT Nullable	The maximum value for any remote journal of the total number of times the local system retransmitted a segment because an acknowledgement was not received. Contains the null value if NUMBER_REMOTE_JOURNALS is 0 or if no attached remote journals are active using TCP/IP.
JOURNAL_TEXT	TEXT	VARCHAR(50) Nullable	The text description of the journal. Contains the null value if the journal has no text.
MANAGE_RECEIVER_OPTION	MNGRCVOPT	VARCHAR(10) Nullable	Indicates whether the system or user manages the changing of journal receivers. *SYSTEM The system manages the changing of journal receivers. *USER The user manages the changing of journal receivers. Contains the null value for a remote journal.
MANAGE_RECEIVER_DELAY	MNGRCVDLY	INTEGER Nullable	The delay time (in minutes) between attempts to attach new journal receivers to this journal. Contains the null value when MANAGE_RECEIVER_OPTION is *USER or the null value.
REMOVE_INTERNAL_ENTRIES	RMVINTENT	VARCHAR(3) Nullable	Handling of internal system entries. NO The internal system entries are not removed. YES The size of the attached receivers is reduced by automatic removal of the internal system entries. Contains the null value for a remote journal.
REMOVE_FIXED_LENGTH_DETAIL	MINFIXLEN	VARCHAR(3) Nullable	Handling of fixed length details. NO Fixed length data is not removed. YES The size of the journal entries that are deposited into the attached journal receivers is reduced by the automatic removal of all fixed length data such as job name, machine sequence number, and so on. Contains the null value for a remote journal.

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
RECEIVER_MAXIMUM_SIZE	MAXOPT	VARCHAR(10) Nullable	<p>The receiver size option that applies to this journal receiver.</p> <p>*MAXOPT1 The journal receivers attached to the journal can have a maximum receiver size of approximately one terabyte (1,099,511,627,776 bytes) and a maximum sequence number of 9,999,999,999. Additionally, the maximum size of the journal entry that can be deposited is 15,761,440 bytes.</p> <p>*MAXOPT2 The journal receivers attached to the journal can have a maximum receiver size of approximately one terabyte (1,099,511,627,776 bytes) and a maximum sequence number of 9,999,999,999. Additionally, the maximum size of the journal entry which can be deposited is 4,000,000,000 bytes.</p> <p>*MAXOPT3 The journal receivers attached to the journal can have a maximum receiver size of approximately one terabyte (1,099,511,627,776 bytes) and a maximum sequence number of 18,446,744,073,709,551,600. Additionally, the maximum size of the journal entry which can be deposited is 4,000,000,000 bytes.</p> <p>*NONE The journal receivers attached to the journal can have a maximum journal receiver size of approximately 1.9 gigabytes and a maximum sequence number of 2,147,483,136.</p> <p>Contains the null value for a remote journal.</p>
MINIMIZE_ESD_FOR_DATA_AREAS	MINDTAARA	VARCHAR(3) Nullable	<p>Indicates whether journal entries for data areas may have minimized entry specific data.</p> <p>NO Journal entries for data areas do not have minimized entry specific data.</p> <p>YES Journal entries for data areas have minimized entry specific data.</p> <p>Contains the null value for a remote journal.</p>
MINIMIZE_ESD_FOR_FILES	MINFILE	VARCHAR(19) Nullable	<p>Indicates whether journal entries for files may have minimized entry specific data.</p> <p>MINIMIZED Journal entries for files may have minimized entry specific data. The minimizing does not occur on field boundaries. Therefore, the entry specific data may not be viewable and may not be used for auditing purposes.</p> <p>MINIMIZED FOR AUDIT Journal entries for files may have minimized entry specific data. The minimizing occurs on field boundaries. Therefore, the entry specific data will be viewable and may be used for auditing purposes.</p> <p>NO Journal entries for files will have complete entry specific data.</p> <p>Contains the null value for a remote journal.</p>

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
JOURNAL_CACHE	JRNCACHE	VARCHAR(3) Nullable	Specifies whether journal entries are cached before being written out to disk. NO Journal entries are not cached before being written out to disk. YES Journal entries are cached before being written out to disk. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_JOB_NAME	FLDJOB	VARCHAR(3) Nullable	Indicates whether the job name will be stored when journal entries are deposited. NO The job name will not be stored when journal entries are deposited. YES The job name will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_USER_NAME	FLDUSR	VARCHAR(3) Nullable	Indicates whether the user name will be stored when journal entries are deposited. NO The user name will not be stored when journal entries are deposited. YES The user name will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_PROGRAM_NAME	FLDPGM	VARCHAR(3) Nullable	Indicates whether the program name will be stored when journal entries are deposited. NO The program name will not be stored when journal entries are deposited. YES The program name will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_PROGRAM_LIBRARY	FLDPGMLIB	VARCHAR(3) Nullable	Indicates whether the program library name will be stored when journal entries are deposited. NO The program library name will not be stored when journal entries are deposited. YES The program library name will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_SYSTEM_SEQUENCE_NUMBER	FLDSYSSEQ	VARCHAR(3) Nullable	Indicates whether the system sequence number will be stored when journal entries are deposited. NO The system sequence number will not be stored when journal entries are deposited. YES The system sequence number will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_REMOTE_ADDRESS	FLDRMTADR	VARCHAR(3) Nullable	Indicates whether the remote address will be stored when journal entries are deposited. NO The remote address will not be stored when journal entries are deposited. YES The remote address will be stored when journal entries are deposited. Contains the null value for a remote journal.

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FIXED_LENGTH_DATA_INCLUDES_THREAD_ID	FLDTHD	VARCHAR(3) Nullable	Indicates whether the thread identifier will be stored when journal entries are deposited. NO The thread identifier will not be stored when journal entries are deposited. YES The thread identifier will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_LOGICAL_UNIT_OF_WORK_ID	FLDLUW	VARCHAR(3) Nullable	Indicates whether the logical unit of work identifier will be stored when journal entries are deposited. NO The logical unit of work identifier will not be stored when journal entries are deposited. YES The logical unit of work identifier will be stored when journal entries are deposited. Contains the null value for a remote journal.
FIXED_LENGTH_DATA_INCLUDES_TRANSACTION_ID	FLDXID	VARCHAR(3) Nullable	Indicates whether the transaction identifier will be stored when journal entries are deposited. NO The transaction identifier will not be stored when journal entries are deposited. YES The transaction identifier will be stored when journal entries are deposited. Contains the null value for a remote journal.
JOURNALED_OBJECT_LIMIT	JRNOBJLMT	VARCHAR(10) Nullable	The number of objects that can be journaled to the journal. *MAX250K The maximum number of objects that can be journaled to the journal is 250,000. *MAX10M The maximum number of objects that can be journaled to the journal is 10,000,000. Contains the null value for a remote journal.
JOURNALED_OBJECTS	JRNALL	INTEGER Nullable	Total of all objects journaled to the journal. This count includes explicitly journaled objects such as files, file members, access paths, data areas, data queues, libraries, and integrated file system objects. This count also includes implicitly journaled objects such as journal receivers, commitment definitions, and objects journaled for system recovery purposes. Contains the null value for a remote journal.
JOURNALED_FILES	JRNFILE	INTEGER Nullable	The total number of files that are currently being journaled to this journal. Contains the null value for a remote journal.
JOURNALED_MEMBERS	JRNMBR	INTEGER Nullable	The total number of file members that are currently being journaled to this journal. Contains the null value for a remote journal.
JOURNALED_DATA_AREAS	JRNDAARA	INTEGER Nullable	The total number of data areas that are currently being journaled to this journal. Contains the null value for a remote journal.
JOURNALED_DATA_QUEUES	JRNDAQ	INTEGER Nullable	The total number of data queues that are currently being journaled to this journal. Contains the null value for a remote journal.
JOURNALED_IFS_OBJECTS	JRNIFS	INTEGER Nullable	The total number of integrated file system objects of type *DIR, *STMF, and *SYMLNK that are currently being journaled to this journal. Contains the null value for a remote journal.

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
JOURNALED_ACCESS_PATHS	JRNAP	INTEGER Nullable	The total number of access paths that are currently being journaled to this journal. Contains the null value for a remote journal.
JOURNALED_COMMITMENT_DEFINITIONS	JRCMTDFN	INTEGER Nullable	The total number of commitment definitions that are currently being implicitly journaled to this journal. Contains the null value for a remote journal.
JOURNALED_LIBRARIES	JRNLIB	INTEGER Nullable	The total number of libraries that are currently being journaled to this journal. Contains the null value for a remote journal.
JOURNAL_RECOVERY_COUNT	JRNRCYCNT	INTEGER Nullable	The approximate number of journaled changes that would need to be recovered during journal synchronization for this journal in the event of an abnormal IPL or vary on. Contains the null value for a local journal with the value *SYSDFT or for a remote journal.
REMOTE_JOURNAL_TYPE	RMTJRNTYPE	VARCHAR(10) Nullable	The type of remote journal. Values are *TYPE1 and *TYPE2. Contains the null value for a local journal.
JOURNAL_DELIVERY_MODE	DELIVMODE	VARCHAR(10) Nullable	The journal delivery mode that is being used to replicate journal entries to this journal. *ASYNC Journal entries are being delivered or replicated asynchronously. *ASYNCPEND Journal entries are to be delivered or replicated asynchronously, but the journal is currently in catch-up mode. *SYNC Journal entries are being delivered or replicated synchronously. *SYNCPEND Journal entries are to be delivered or replicated synchronously, but the journal is currently in catch-up mode. Contains the null value for a local journal or a remote journal whose JOURNAL_STATE field is not *ACTIVE or *CTLINACT.
LOCAL_JOURNAL_NAME	LCLJRNAME	VARCHAR(10) Nullable	The journal name of the local journal. The local journal is the journal that is the initiator of the original journal deposit that has been replicated downstream to this journal. Contains the null value for a local journal.
LOCAL_JOURNAL_LIBRARY	LCLJRNLIB	VARCHAR(10) Nullable	The library name of the local journal. Contains the null value for a local journal.
LOCAL_JOURNAL_SYSTEM	LCLJRNSYS	VARCHAR(8) Nullable	The name of the system for the local journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
LOCAL_JOURNAL_ASPGRP	LCLASPGRP	VARCHAR(10) Nullable	The name of the independent auxiliary storage pool (ASP) group of the local journal. *SYSBAS is used to indicate the system ASP and all basic user ASPs. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
SOURCE_JOURNAL_NAME	SRCJRNAME	VARCHAR(10) Nullable	The journal name of the source journal. The source journal is the journal that is directly upstream of this journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
SOURCE_JOURNAL_LIBRARY	SRCJRNLIB	VARCHAR(10) Nullable	The library name of the source journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
SOURCE_JOURNAL_SYSTEM	SRCJRNSYS	VARCHAR(8) Nullable	The name of the system for the source journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
SOURCE_JOURNAL_ASPPGRP	SRCASPPGRP	VARCHAR(10) Nullable	The name of the independent auxiliary storage pool (ASP) group of the source journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
LOCAL_RECEIVER_SYSTEM	LCLRCVSY	VARCHAR(8) Nullable	If this journal receiver is associated with a remote journal, the name of the system for the local journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
SOURCE_RECEIVER_SYSTEM	SRCRCVSY	VARCHAR(8) Nullable	If this journal receiver is associated with a remote journal, the name of the system for the source journal. Contains *UNKNOWN if journal is a remote journal and does not have an attached receiver. Contains the null value for a local journal.
ACTIVATION_TIME	ACTDT	TIMESTAMP Nullable	If the journal is a remote journal and it is currently active, the date and time the journal was activated. Contains the null value for a local journal or a remote journal whose JOURNAL_STATE field is not *ACTIVE or *CTLINACT.
ESTIMATED_TIME_BEHIND	ESTBEHIND	BIGINT Nullable	If the journal is an active remote journal and the delivery mode is asynchronous, this is the estimated amount of time, in milliseconds, between when the journal entries are written to disk on the source system and when they are received on the target system. Contains the null value for a local journal or a remote journal whose JOURNAL_STATE field is not *ACTIVE or *CTLINACT.
MAXIMUM_TIME_BEHIND	MAXBEHIND	BIGINT Nullable	The maximum value of ESTIMATED_TIME_BEHIND since the journal was activated. Contains the null value for a local journal or a remote journal whose JOURNAL_STATE field is not *ACTIVE or *CTLINACT.
MAXIMUM_BEHIND_TIMESTAMP	MAXBHNDTIM	TIMESTAMP Nullable	The date and time that the ESTIMATED_TIME_BEHIND occurred. Contains the null value for a local journal or a remote journal whose JOURNAL_STATE field is not *ACTIVE or *CTLINACT.

Table 143. JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
JOURNAL_ENTRY_FILTERING	FILTER	VARCHAR(3) Nullable	Indicates whether or not journal entry filtering is active for this journal. NO Journal entry filtering is not active for this journal. YES Journal entry filtering is active for this journal. Contains the null value for a local journal or a remote journal whose JOURNAL_STATE field is not *ACTIVE or *CTLINACT.

Examples

- List all journals that are falling behind sending entries to one or more remote journals:

```
SELECT JOURNAL_NAME, JOURNAL_LIBRARY,
       MAXIMUM_REMOTE_JOURNALS_ENTRIES_BEHIND,
       MAXIMUM_REMOTE_JOURNALS_TIME_BEHIND, MAXIMUM_REMOTE_JOURNALS_RETRANSMISSIONS
FROM QSYS2.JOURNAL_INFO
WHERE MAXIMUM_REMOTE_JOURNALS_ENTRIES_BEHIND > 0
ORDER BY MAXIMUM_REMOTE_JOURNALS_ENTRIES_BEHIND DESC
```

- Find any remote journals that are not currently active:

```
SELECT * FROM QSYS2.JOURNAL_INFO
WHERE JOURNAL_TYPE = '*REMOTE'
      AND JOURNAL_STATE <> '*ACTIVE'
ORDER BY JOURNAL_LIBRARY, JOURNAL_NAME,
```

- For security auditing reasons, find any journals that are not recording remote address info:

```
SELECT * FROM QSYS2.JOURNAL_INFO
WHERE REMOVE_FIXED_LENGTH_DETAIL = 'YES'
      OR FIXED_LENGTH_DATA_INCLUDES_REMOTE_ADDRESS = 'NO'
```

JOURNAL_RECEIVER_INFO view

The JOURNAL_RECEIVER_INFO view contains information about all journal receivers on the system.

The values returned for the columns in the view are closely related to the values returned by the Retrieve Journal Receiver Information (QjoRtvJrnReceiverInformation) API.

Authorization: The caller must have:

- *EXECUTE authority on the library containing the journal receiver, and
- *OBJOPR and some data authority other than *EXECUTE to the journal receiver.

If the journal receiver is associated with a journal, the caller requires:

- *EXECUTE authority on the library containing the journal, and
- *OBJOPR on the journal.

The following table describes the columns in the view. The system name is JRNRCV_INF. The schema is QSYS2.

Table 144. JOURNAL_RECEIVER_INFO view

Column Name	System Column Name	Data Type	Description
JOURNAL_RECEIVER_LIBRARY	JRNRCV_LIB	VARCHAR(10)	The name of the library that contains the journal receiver.
JOURNAL_RECEIVER	JRNRCV	VARCHAR(10)	The name of the journal receiver.

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
DESCRIPTIVE_TEXT	TEXT	VARCHAR(50) Nullable	The text description of the journal receiver.
JOURNAL_RECEIVER_ASP_NUMBER	RCV_ASPNUM	INTEGER	The number of the auxiliary storage pool to which storage for the journal receiver is allocated.
JOURNAL_LIBRARY	JRN_LIB	VARCHAR(10) Nullable	The name of the library that contains the journal. Contains the null value if the receiver has never been attached to a journal.
JOURNAL_NAME	JOURNAL	VARCHAR(10) Nullable	The name of the journal that the journal receiver is attached to or used to be attached to. Contains the null value if the receiver has never been attached to a journal.
THRESHOLD	THRESHOLD	INTEGER Nullable	An auxiliary disk storage space threshold value (in kilobytes) for the journal receiver. If the threshold value is exceeded during journaling and the journal has the MNGRCV(*USER) attribute, a message (CPF7099) is sent to the message queue that is specified on the Create Journal (CRTJRN) or the Change Journal (CHGJRN) command. If the journal has the MNGRCV(*SYSTEM) attribute, the system creates and attaches a new journal receiver, detaches the old journal receiver when the threshold is reached, and sends message CPF7020 to the journal message queue. Contains the null value for remote journal receivers.
SIZE	SIZE	INTEGER	The number of kilobytes of auxiliary disk storage used by this journal receiver.
STATUS	STATUS	VARCHAR(8)	The status of the journal receiver. ATTACHED The journal receiver is currently attached to the journal. ONLINE The journal receiver is online. The journal receiver has not been saved, and it has been detached from the journal. SAVED The journal receiver was saved after it was detached. The journal receiver storage was not freed when it was saved. FREED The journal receiver was saved after it was detached. The journal receiver storage was freed when it was saved. PARTIAL The journal receiver status is partial for one of the following reasons: <ul style="list-style-type: none"> • It was restored from a version that was saved while it was attached to the journal. Additional journal entries may have been written that were not restored. • It is associated with a remote journal and it does not contain all the journal entries that are in the associated journal receiver attached to the source journal. EMPTY The journal receiver has never been attached to a journal.
NUMBER_OF_JOURNAL_ENTRIES	ENTRIES	DECIMAL(20,0) Nullable	The number of journal entries that are contained in this journal receiver. Contains the null value if STATUS is EMPTY.

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FIRST_SEQUENCE_NUMBER	FIRST_SEQ	DECIMAL(21,0) Nullable	The journal sequence number of the first journal entry in this journal receiver. Contains the null value if STATUS is EMPTY.
LAST_SEQUENCE_NUMBER	LAST_SEQ	DECIMAL(21,0) Nullable	The journal sequence number of the last journal entry in this journal receiver. Contains the null value if STATUS is EMPTY.
MAXIMUM_ENTRY_SPECIFIC_DATA_LENGTH	MAX_ESD	DECIMAL(20,0) Nullable	The length in bytes of the longest entry-specific data among all journal entries in this journal receiver. Contains the null value if STATUS is EMPTY.
MAXIMUM_NULL_VALUE_INDICATORS	MAX_NVI	INTEGER Nullable	The maximum number of null value indicators among all journal entries in this journal receiver. Contains the null value if STATUS is EMPTY.
ATTACH_TIMESTAMP	ATTACHED	TIMESTAMP(0) Nullable	The date and time that this journal receiver was attached to the journal. For a journal receiver attached to a remote journal, this is the date and time that the journal receiver was attached on the local system. Contains the null value if STATUS is EMPTY.
DETACH_TIMESTAMP	DETACHED	TIMESTAMP(0) Nullable	The date and time that this journal receiver was detached from the journal. For a journal receiver attached to a *REMOTE journal, this is the date and time that the journal receiver was detached on the local system. Contains the null value if STATUS is EMPTY, ATTACHED, or PARTIAL.
SAVE_TIMESTAMP	SAVED	TIMESTAMP(0) Nullable	The date and time that the journal receiver was last saved. This value reflects when the receiver was saved from the system it exists on (either the source or target system time). Contains the null value if the journal receiver was never saved.
PREVIOUS_JOURNAL_RECEIVER_LIBRARY	PREV_RCVL	VARCHAR(10) Nullable	The name of the library of the previous journal receiver that is associated with the same journal. Contains the null value if STATUS is EMPTY, if there was no journal receiver attached to the journal prior to this journal receiver, or if it is currently associated with a remote journal.
PREVIOUS_JOURNAL_RECEIVER	PREV_RCV	VARCHAR(10) Nullable	The name of the previous journal receiver that is associated with the same journal. Contains the null value if STATUS is EMPTY, if there was no journal receiver attached to the journal prior to this journal receiver, or if it is currently associated with a remote journal.
NEXT_JOURNAL_RECEIVER_LIBRARY	NEXT_RCVL	VARCHAR(10) Nullable	The name of the library of the next journal receiver that is associated with the same journal. v
NEXT_JOURNAL_RECEIVER	NEXT_RCV	VARCHAR(10) Nullable	The name of the next journal receiver that is associated with the same journal. v

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
RECEIVER_MAXIMUM_SIZE	MAXOPT	pendin VARCHAR(8) Nullable	<p>The journal receiver sequence number and size option for this journal receiver. If this journal receiver is attached to a remote journal, the value is determined by the local journal.</p> <p>*NONE The journal receiver has a maximum journal receiver size of approximately 1.9 gigabytes and a maximum sequence number of 2,147,483,136.</p> <p>*MAXOPT1 The journal receiver has a maximum journal receiver size of approximately one terabyte (1,099,511,627,776 bytes) and a maximum sequence number of 9,999,999,999. Additionally, the maximum size of the journal entry that can be deposited is 15,761,440 bytes. This occurs if this receiver was attached when RCVSIZOPT(*MAXOPT1) was in effect for the journal.</p> <p>*MAXOPT2 The journal receiver has a maximum journal receiver size of approximately one terabyte (1,099,511,627,776 bytes) and a maximum sequence number of 9,999,999,999. This occurs if this receiver was attached when RCVSIZOPT(*MAXOPT2) was in effect for the journal. Additionally, the maximum size of the journal entry which can be deposited is 4,000,000,000 bytes.</p> <p>*MAXOPT3 The journal receiver has a maximum journal receiver size of approximately one terabyte (1,099,511,627,776 bytes) and a maximum sequence number of 18,446,744,073,709,551,600. This occurs if this receiver was attached when RCVSIZOPT(*MAXOPT3) was in effect for the journal. Additionally, the maximum size of the journal entry which can be deposited is 4,000,000,000 bytes.</p> <p>Contains the null value if STATUS is EMPTY.</p>
MINIMIZE_ESD_FOR_DATA_AREAS	MINDTAARA	VARCHAR(3) Nullable	<p>Whether the entry-specific data for data areas is minimized. If this journal receiver is attached to a remote journal, the value for is determined by the local journal.</p> <p>NO Journal entries for data areas have complete entry specific data.</p> <p>YES Journal entries for data areas may have minimized entry specific data.</p> <p>Contains the null value if STATUS is EMPTY.</p>

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MINIMIZE_ESD_FOR_FILES	MINFILE	VARCHAR(6) Nullable	<p>Whether the entry-specific data for files is minimized. If this journal receiver is attached to a remote journal, the value is determined by the local journal.</p> <p>FLDBDY Journal entries for files may have minimized entry specific data. The minimizing occurs on field boundaries. Therefore, the entry specific data will be viewable and may be used for auditing purposes.</p> <p>NO Journal entries for files have complete entry specific data.</p> <p>YES Journal entries for files may have minimized entry specific data. The minimizing does not occur on field boundaries. Therefore, the entry specific data may not be viewable and may not be used for auditing purposes.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_JOB_NAME	FLDJOB	VARCHAR(3) Nullable	<p>Indicates whether the job name is stored when journal entries are deposited.</p> <p>NO The job name is not stored when journal entries are deposited.</p> <p>YES The job name is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_USER_NAME	FLDUSR	VARCHAR(3) Nullable	<p>Indicates whether the user name is stored when journal entries were deposited.</p> <p>NO The user name is not stored when journal entries are deposited.</p> <p>YES The user name is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_PROGRAM_NAME	FLDPGM	VARCHAR(3) Nullable	<p>Indicates whether the program name is stored when journal entries were deposited.</p> <p>NO The program name is not stored when journal entries are deposited.</p> <p>YES The program name is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_PROGRAM_LIBRARY	FLDPGMLIB	VARCHAR(3) Nullable	<p>Indicates whether the program library name is stored when journal entries were deposited.</p> <p>NO The program library name is not stored when journal entries are deposited.</p> <p>YES The program library name is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_SYSTEM_SEQUENCE_NUMBER	FLDSYSSEQ	VARCHAR(3) Nullable	<p>Indicates whether the system sequence number is stored when journal entries were deposited.</p> <p>NO The system sequence number is not stored when journal entries are deposited.</p> <p>YES The system sequence number is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FIXED_LENGTH_DATA_INCLUDES_REMOTE_ADDRESS	FLDRMTADR	VARCHAR(3) Nullable	<p>Indicates whether the remote address is stored when journal entries were deposited.</p> <p>NO The remote address is not stored when journal entries are deposited.</p> <p>YES The remote address is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_THREAD_ID	FLDTHD	VARCHAR(3) Nullable	<p>Indicates whether the thread identifier is stored when journal entries were deposited.</p> <p>NO The thread identifier is not stored when journal entries are deposited.</p> <p>YES The thread identifier is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_LOGICAL_UNIT_OF_WORK_ID	FLDLUW	VARCHAR(3) Nullable	<p>Indicates whether the logical unit of work identifier is stored when journal entries were deposited.</p> <p>NO The logical unit of work identifier is not stored when journal entries are deposited.</p> <p>YES The logical unit of work identifier is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
FIXED_LENGTH_DATA_INCLUDES_TRANSACTION_ID	FLDXID	VARCHAR(3) Nullable	<p>Indicates whether the transaction identifier is stored when journal entries were deposited.</p> <p>NO The transaction identifier is not stored when journal entries are deposited.</p> <p>YES The transaction identifier is stored when journal entries are deposited.</p> <p>Contains the null value if STATUS is EMPTY.</p>
PENDING_TRANSACTIONS	PEND_TRANS	VARCHAR(3) Nullable	<p>Whether the journal receiver contains journal entries for commitment control transactions that have not yet been committed or rolled back.</p> <p>NO The journal receiver does not contain entries for pending commitment control transactions.</p> <p>YES The journal receiver contains entries for pending commitment control transactions.</p> <p>Contains the null value if STATUS is EMPTY or the journal receiver was attached to a remote journal.</p>
REMOTE_JOURNAL_TYPE	RMT_TYPE	VARCHAR(6) Nullable	<p>If this journal receiver was attached to a remote journal, this is the remote journal type for that journal when this journal receiver was attached. Values are *TYPE1 and *TYPE2.</p> <p>Contains the null value if the journal receiver was not attached to a remote journal.</p>
LOCAL_JOURNAL_SYSTEM	LCLJRNSYS	VARCHAR(8) Nullable	<p>The system name of the local journal. The local journal is the journal that is the initiator of the original journal deposit that has been replicated downstream to this journal.</p> <p>Contains the null value if the journal receiver was not attached to a remote journal.</p>
LOCAL_JOURNAL_LIBRARY	LCLJRNLIB	VARCHAR(10) Nullable	<p>The library name of the local journal.</p> <p>Contains the null value if the journal receiver was not attached to a remote journal.</p>

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOCAL_JOURNAL_NAME	LCLJRNAME	VARCHAR(10) Nullable	The name of the local journal. The local journal is the journal that is the initiator of the original journal deposit that has been replicated downstream to this journal. Contains the null value if the journal receiver was not attached to a remote journal.
LOCAL_JOURNAL_ASP_GROUP	LCLASPGRP	VARCHAR(10) Nullable	The independent auxiliary storage pool (IASP) group of the local journal. Contains the null value if the journal receiver was not attached to a remote journal.
LOCAL_JOURNAL_RECEIVER_LIBRARY	LCLRCVLIB	VARCHAR(10) Nullable	The library name of the journal receiver that is associated with the local journal. Contains the null value if the journal receiver was not attached to a remote journal.
SOURCE_JOURNAL_SYSTEM	SRCJRNSYS	VARCHAR(8) Nullable	The system name of the source journal. Contains the null value if the journal receiver was not attached to a remote journal.
SOURCE_JOURNAL_LIBRARY	SRCJRNLIB	VARCHAR(10) Nullable	The library name of the source journal. Contains the null value if the journal receiver was not attached to a remote journal.
SOURCE_JOURNAL_NAME	SRCJRNAME	VARCHAR(10) Nullable	The name of the source journal. The source journal is the journal that is directly upstream of this remote journal. Contains the null value if the journal receiver was not attached to a remote journal.
SOURCE_JOURNAL_ASP_GROUP	SRCASPGRP	VARCHAR(10) Nullable	The independent auxiliary storage pool (IASP) group of the source journal. Contains the null value if the journal receiver was not attached to a remote journal.
SOURCE_JOURNAL_RECEIVER_LIBRARY	SRCRCVLIB	VARCHAR(10) Nullable	The library name of the journal receiver that is associated with the source journal. Contains the null value if the journal receiver was not attached to a remote journal.
REDIRECTED_RECEIVER_LIBRARY	REDIR_LIB	VARCHAR(10) Nullable	The *TYPE1 receiver library redirection that was in effect when this journal receiver was attached. Contains the null value if the journal receiver was not attached to a remote journal or REMOTE_JOURNAL_TYPE is not *TYPE1.
FILTER_BY_OBJECT	FTR_OBJECT	VARCHAR(4) Nullable	Specifies whether journal entries sent to the remote journal will be filtered by object. *NO Journal entries sent to the remote journal will not be filtered by object. *YES Journal entries deposited for objects that indicated remote journal filtering at the time they were deposited will not be sent to the remote journal. Contains the null value if the journal receiver was not attached to a remote journal.
FILTER_IMAGES	FTR_IMAGE	VARCHAR(7) Nullable	Specifies whether before images will be sent to the remote journal. *NONE All journal entries will be sent to the remote journal, unless they are filtered by the Filter by object or Filter by program specifications. *BEFORE Before images will not be sent to the remote journal. Contains the null value if the journal receiver was not attached to a remote journal.

Table 144. JOURNAL_RECEIVER_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FILTER_PROGRAMS	FTR_PGM	INTEGER Nullable	The number of programs for which journal entries sent on behalf of these programs were filtered when sent to this journal receiver. These programs are listed in the FILTER_PROGRAM_ARRAY and FILTER_PROGRAM_JSON columns. Contains the null value if the journal receiver was not attached to a remote journal.
FILTER_PROGRAM_ARRAY	FTR_ARRAY	VARCHAR(219) Nullable	An array of up to 20 names of 22 characters apiece. The first 10 of each is a program name, followed by a blank, followed by 10 characters for the program library name. One comma separates the entries. The program library name can be *ALL. Journal entries sent on behalf of these programs will not be sent to the remote journal. Contains the null value if the journal receiver was not attached to a remote journal or FILTER_PROGRAMS is zero.
FILTER_PROGRAM_JSON	FTR_LISTJ	VARCHAR(498) CCSID 1208 Nullable	A list of programs and libraries. Journal entries sent on behalf of these programs were filtered when sent to this journal receiver. The program's library name may be *ALL. This list is returned as an array within a JSON object. Each entry in the JSON array contains two JSON objects: <ul style="list-style-type: none"> • An object with a name of "LIBRARY" and a value of the library name containing the program • An object with a name of "PROGRAM" and a value of the program name Contains the null value if the journal receiver was not attached to a remote journal or FILTER_PROGRAMS is zero.

Examples

- Return a list of journal receivers that have not been saved.

```
SELECT JOURNAL_RECEIVER_LIBRARY, JOURNAL_RECEIVER_NAME, STATUS
FROM QSYS2.JOURNAL_RECEIVER_INFO
WHERE STATUS = 'ONLINE';
```

- Return the length of time, in seconds, that journal receivers in RCVLIB were attached.

```
SELECT JOURNAL_RECEIVER_LIBRARY, JOURNAL_RECEIVER_NAME, TIMESTAMPDIFF(SECOND,
DETACH_TIMESTAMP, ATTACH_TIMESTAMP), SIZE
FROM QSYS2.JOURNAL_RECEIVER_INFO
WHERE DETACH_TIMESTAMP IS NOT NULL AND
JOURNAL_RECEIVER_LIBRARY = 'RCVLIB';
```

- Return a list of programs contained in the FILTER_PROGRAMS_JSON column related to journal receiver RCV1 in RCVLIB.

```
SELECT j.*, r.*
FROM QSYS2.JOURNAL_RECEIVER_INFO r,
JSON_TABLE(FILTER_PROGRAM_JSON, 'lax $.PROGRAM_LIST'
COLUMNS(LIBRARY VARCHAR(10),
PROGRAM VARCHAR(10))) j
WHERE JOURNAL_RECEIVER_LIBRARY = 'RCVLIB' AND JOURNAL_RECEIVER_NAME = 'RCV1';
```

JOURNALED_OBJECTS view

The JOURNALED_OBJECTS view returns information about journaled objects. Only information about external objects is returned. Internal objects such as commit blocks and access paths are not included.

The values returned for the columns in the view are closely related to the values returned by the Retrieve Journal Information (QjoRetrieveJournalInformation) API and the Work with Journal Attributes (WRKJRNA) CL command.

Authorization: The caller must have:

- *EXECUTE authority on the library containing the journal, and
- *OBJOPR and some data authority other than *EXECUTE to the journal

The following table describes the columns in the view. The system name is JRN_OBJC. The schema is QSYS2.

Table 145. JOURNALED_OBJECTS view

Column Name	System Column Name	Data Type	Description
JOURNAL_LIBRARY	JRNLIB	VARCHAR(10)	The name of the library that contains the journal.
JOURNAL_NAME	JRNNAME	VARCHAR(10)	The name of the journal.
IASP_NUMBER	IASPNUMBER	INTEGER	The number of the auxiliary storage pool to which storage for the journal is allocated.
OBJECT_TYPE	OBJ_TYPE	VARCHAR(7)	Type of object. *DIR Directory *DTAARA Data area *DTAQ Data queue *FILE Database file *JRNRCV Journal receiver *LIB Library *STMF Stream file *SYMLNK Symbolic link
OBJECT_LIBRARY	OBJ_LIB	VARCHAR(10) Nullable	The name of the library that contains the object. Contains the null value if OBJECT_TYPE is *DIR, *STMF, or *SYMLNK.
OBJECT_NAME	OBJ_NAME	VARCHAR(10) Nullable	The name of the object. Contains the null value if OBJECT_TYPE is *DIR, *STMF, or *SYMLNK.
FILE_TYPE	FILE_TYPE	VARCHAR(8) Nullable	The type of file that is journaled. LOGICAL Logical file PHYSICAL Physical file Contains the null value if OBJECT_TYPE is not *FILE.
PATH_NAME	PATH_NAME	DBCLOB(16M) CCSID 1200 Nullable	The path name of an integrated file system object. Contains the null value if OBJECT_TYPE is not *DIR, *STMF, or *SYMLNK.
FILE_IDENTIFIER	FILE_ID	BINARY(16) Nullable	The identifier associated with the integrated file system object. Contains the null value if OBJECT_TYPE is not *DIR, *STMF, or *SYMLNK.

Table 145. JOURNALED_OBJECTS view (continued)

Column Name	System Column Name	Data Type	Description
JOURNAL_IMAGES	IMAGES	VARCHAR(6) Nullable	<p>Specifies the kinds of images written to the journal for this object.</p> <p>*AFTER Only after images are written to the journal.</p> <p>*BOTH Both before and after images are written to the journal.</p> <p>Contains the null value if OBJECT_TYPE is *JRNRCV.</p>
OMIT_JOURNAL_ENTRY	OMIT_ENTRY	VARCHAR(10) Nullable	<p>Specifies the journal entries that are omitted.</p> <p>*NONE No entries are omitted.</p> <p>*OPNCLO Open and close entries are omitted. Open and close operations on the specified file members do not create open and close journal entries. Using this option prevents the use of TOJOB0 and TOJOB C entries on the Apply Journalized Changes (APYJRNCHG) and Remove Journalized Changes (RMVJRNCHG) commands, but it saves some storage space in the attached receivers.</p> <p>*OPNCLOSYN Open, close, and force entries are omitted. Open, close, and force operations on the specified objects do not generate open, close and force journal entries. Using this option prevents the use of TOJOB0 and TOJOB C entries on the Apply Journalized Changes (APYJRNCHG) command, but it saves some storage space in the journal receivers.</p> <p>Contains the null value if OBJECT_TYPE is *JRNRCV.</p>
INHERIT	INHERIT	VARCHAR(4) Nullable	<p>Specifies whether new objects created within this journalized directory or library, moved into this journalized directory or library, or restored into this journalized directory or library should inherit the journal state of the parent directory or library.</p> <p>*NO New objects will not inherit the journal state of the parent.</p> <p>*YES New objects will inherit the journal state of the parent.</p> <p>Contains the null value if OBJECT_TYPE is not *DIR or *LIB.</p>
REMOTE_JOURNAL_FILTER	RMT_FILTER	VARCHAR(4) Nullable	<p>Specifies whether the journal entries deposited for objects that inherit the journal state of the directory or library are eligible for remote journal filtering by object.</p> <p>*NO Journal entries deposited for objects will not be eligible for remote journal filtering by object.</p> <p>*YES Journal entries deposited for objects will be eligible for remote journal filtering by object. When using remote journal filtering by object, journal entries for the object will not be sent to the target system.</p> <p>Contains the null value if OBJECT_TYPE is *JRNRCV.</p>

Example

- Review all the objects journaled to APPLIB/APPJRN.

```
SELECT *
FROM QSYS2.JOURNALED_OBJECTS
WHERE JOURNAL_LIBRARY = 'APPLIB' AND JOURNAL_NAME = 'APPJRN'
ORDER BY OBJECT_TYPE, OBJECT_LIBRARY, OBJECT_NAME, PATH_NAME;
```

REMOTE_JOURNAL_INFO view

The REMOTE_JOURNAL_INFO view returns information about every remote journal defined for a local or remote journal on the IBM i where this view is referenced. The remote journals returned by this service exist on a different IBM i, as noted by the remote database (RDB) name.

This information is similar to what is returned by the Retrieve Journal Information (QjoRetrieveJournalInformation) API.

Authorization: The caller must have:

- *OBJOPR and some data authority other than *EXECUTE to the journal, and
- *EXECUTE authority to the library containing the journal.

The following table describes the columns in the view. The system name is RMT_JRNS. The schema is QSYS2.

Table 146. REMOTE_JOURNAL_INFO view

Column Name	System Column Name	Data Type	Description
SOURCE_JOURNAL_LIBRARY	SRC_JRNLIB	VARCHAR(10)	The library that contains SOURCE_JOURNAL.
SOURCE_JOURNAL	SRC_JRN	VARCHAR(10)	The name of the journal that has the remote journal defined.
REMOTE_DATABASE_NAME	RMT_RDB	VARCHAR(18)	The name of the remote database that is the target and contains the remote journal.
REMOTE_JOURNAL_LIBRARY	RMT_JRNLIB	VARCHAR(10)	The library that contains REMOTE_JOURNAL.
REMOTE_JOURNAL	RMT_JRN	VARCHAR(10)	The name of the remote journal that is the target for remote journaling.
REMOTE_JOURNAL_RECEIVER_LIBRARY	RMT_RCVLIB	VARCHAR(10)	The library name of the remote journal receiver that is the target for remote journaling. The special value of *SRCRCVLIB indicates the journal receivers are created on the target system in the same library as they exist on the source system.

Table 146. REMOTE_JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
REMOTE_JOURNAL_STATE	RMT_STATE	VARCHAR(8)	<p>The state of the remote journaling.</p> <p>ACTIVE The remote journal is ready to receive any journal entries from its source journal.</p> <p>CTLINACT The remote journal is in the process of a controlled inactivate. Therefore, the remote journal will be receiving those journal entries that were already queued for replication when the Change Remote Journal (CHGRMTJRN) command or the Change Journal State (QjoChangeJournalState) API requested to inactivate the remote journal. However, no entries deposited after that request will be replicated to the remote journal.</p> <p>FAILED The remote journal is not ready to receive any journal entries from its source journal due to a remote journal function failure, for example, a communications failure. You will need to inactivate the remote journal by using the Change Remote Journal (CHGRMTJRN) command or by calling the Change Journal State (QjoChangeJournalState) API.</p> <p>INACTIVE The remote journal is not ready to receive any journal entries from its source journal.</p> <p>PENDING The remote journal is transitioning from an INACTIVE state to an ACTIVE state.</p>
REMOTE_JOURNAL_TYPE	RMT_JTYPE	VARCHAR(6)	<p>The type of remote journal that was created. The type influences characteristics of the remote journal such as journal receiver restore options, redirection capabilities, and remote journal association support.</p> <p>*TYPE1 Type 1 remote journal</p> <p>*TYPE2 Type 2 remote journal</p>
DELIVERY_MODE	DELIVERY	VARCHAR(10) Nullable	<p>The remote journal delivery mode that is being used to replicate journal entries to the remote journal.</p> <p>*ASYNC Journal entries are being delivered or replicated asynchronously.</p> <p>*ASYNCPEND Journal entries are to be delivered or replicated asynchronously, but the remote journal is currently in catch-up mode.</p> <p>*SYNC Journal entries are being delivered or replicated synchronously.</p> <p>*SYNCPEND Journal entries are to be delivered or replicated synchronously, but the remote journal is currently in catch-up mode.</p> <p>Contains the null value if REMOTE_JOURNAL_STATE is not ACTIVE or CTLINACT.</p>

Table 146. REMOTE_JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
DATA_PORT_SERVICES_NODE_ID	NODE_ID	VARCHAR(8) Nullable	The node identifier being used by data port services to identify the target system in a cluster environment. If a node identifier and at least one internet address is retrieved, then data port services is being used as an alternate communication method to the target system. Contains the null value if the remote journal is not configured for data port services.
DATA_PORT_SERVICES_ADDRESS_1	IP_ADDR1	VARCHAR(45) Nullable	The first internet address being used by data port services to communicate to the target system. Contains the null value if the remote journal is not configured for data port services.
DATA_PORT_SERVICES_ADDRESS_2	IP_ADDR2	VARCHAR(45) Nullable	The second internet address being used by data port services to communicate to the target system. Contains the null value if the remote journal is not configured for data port services or there are less than two internet addresses being used.
DATA_PORT_SERVICES_ADDRESS_3	IP_ADDR3	VARCHAR(45) Nullable	The third internet address being used by data port services to communicate to the target system. Contains the null value if the remote journal is not configured for data port services or there are less than three internet addresses being used.
DATA_PORT_SERVICES_ADDRESS_4	IP_ADDR4	VARCHAR(45) Nullable	The fourth internet address being used by data port services to communicate to the target system. Contains the null value if the remote journal is not configured for data port services or there are less than four internet addresses being used.
VALIDITY_CHECKING	VALID_CHK	VARCHAR(9) Nullable	The validity checking status. When communications validity checking is turned on, the remote journal environment provides additional checking to verify that the data which is received by the target system matches the data that was sent from the source system. If the data does not match, the data will not be written to the target system, the remote journal environment will be inactivated, and messages indicating the communications failure will be issued to the journal message queue and QHST. *DISABLED Communications validity checking is turned off for this remote journal environment. *ENABLED Communications validity checking is turned on for this remote journal environment. Contains the null value if REMOTE_JOURNAL_STATE is not ACTIVE or CTLINACT.
SENDING_TASK_PRIORITY	PRIORITY	INTEGER Nullable	The priority of the sending task on the source system. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND.
SYNCHRONOUS_SENDING_TIME_OUT	TIME_OUT	INTEGER Nullable	The maximum amount of time, in seconds, to wait for a response from the remote system when a response is required in a synchronous remote journal environment. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *SYNC or *SYNCPEND.

Table 146. REMOTE_JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MAXIMUM_RESTART_ATTEMPTS	MAXRESTART	INTEGER Nullable	The number of times the operating system will attempt to reactivate the remote journal after a recoverable failure. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
RESTART_DELAY_TIME	RESTARTDLY	INTEGER Nullable	The number of seconds between attempts to reactivate the remote journal after a recoverable failure. Contains the null value if MAXIMUM_RESTART_ATTEMPTS is null or 0.
FILTER_BY_OBJECT	FTR_OBJECT	VARCHAR(4) Nullable	Specifies whether journal entries sent to the remote journal will be filtered by object. *NO Journal entries sent to the remote journal will not be filtered by object. *YES Journal entries deposited for objects that indicated remote journal filtering at the time they were deposited will not be sent to the remote journal. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
FILTER_IMAGES	FTR_IMAGE	VARCHAR(7) Nullable	Specifies whether before images will be sent to the remote journal. *NONE All journal entries will be sent to the remote journal, unless they are filtered by the Filter by object or Filter by program specifications. *BEFORE Before images will not be sent to the remote journal. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
FILTER_PROGRAMS	FTR_PGM	INTEGER Nullable	The number of programs returned in FILTER_PROGRAM_ARRAY. Journal entries sent on behalf of these programs will not be sent to the remote journal. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
FILTER_PROGRAM_ARRAY	FTR_ARRAY	VARCHAR(219) Nullable	An array of up to 20 names of 22 characters each. The first 10 of each is a program name, followed by a blank, followed by 10 characters for the program library name. One comma separates the entries. The program library name can be *ALL. Journal entries sent on behalf of these programs will not be sent to the remote journal. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or FILTER_PROGRAMS is zero.
FILTER_PROGRAM_JSON	FTR_LISTJ	VARCHAR(498) CCSID 1208 Nullable	A list of programs and libraries. Journal entries sent on behalf of these programs were filtered when sent to this journal receiver. The program's library name may be *ALL. This list is returned as an array within a JSON object. Each entry in the JSON array contains two JSON objects: <ul style="list-style-type: none"> • An object with a name of "LIBRARY" and a value of the library name containing the program • An object with a name of "PROGRAM" and a value of the program name Contains the null value if the journal receiver was not attached to a remote journal or FILTER_PROGRAMS is zero.

Table 146. REMOTE_JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
BUNDLES_SENT	BUNDLES	BIGINT Nullable	The number of bundles that have been sent to the target system since the source journal transitioned to an active state. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
MAXIMUM_BUNDLE_SIZE	BUNDLE_SIZ	BIGINT Nullable	The number of bytes in the largest bundle that has been sent to the target system since the source journal transitioned to an active state. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
MAXIMUM_BUNDLE_SIZE_TIME	BUNDLE_TIM	TIMESTAMP Nullable	The date and time that the maximum bundle size was sent to the target system. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or MAXIMUM_BUNDLE_SIZE is 0.
SUPER_BUNDLE_MODE_COUNT	SUPER_CNT	BIGINT Nullable	The number of times that the remote journal environment has automatically gone into super bundling mode. Super bundling mode helps the remote journal environment keep up with the local journal when the local journal has a high rate of journal entry deposits. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND.
ENTRIES_BEHIND	BEHIND	BIGINT Nullable	The number of entries that are waiting to be sent to the target system. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND.
MAXIMUM_ENTRIES_BEHIND	MBEHIND	BIGINT Nullable	The maximum number of entries that were waiting to be sent to the target system since the source journal transitioned to an active state. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND.
MAXIMUM_ENTRIES_BEHIND_TIME	MBEHINDT	TIMESTAMP Nullable	The date and time that MAXIMUM_ENTRIES_BEHIND occurred. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND or MAXIMUM_ENTRIES_BEHIND is 0.
TIME_BEHIND	TBEHIND	BIGINT Nullable	The value, in hundredths of seconds, that the source journal is behind in sending journal entries to the target system. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND.
MAXIMUM_TIME_BEHIND	MTBEHIND	BIGINT Nullable	The maximum value, in hundredths of seconds, that the source journal was behind in sending journal entries to the target system. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND.
MAXIMUM_TIME_BEHIND_TIME	MTBEHINDT	TIMESTAMP Nullable	The date and time that MAXIMUM_TIME_BEHIND occurred. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE or DELIVERY_MODE is not *ASYNC or *ASYNCPEND or MAXIMUM_TIME_BEHIND is 0.

Table 146. REMOTE_JOURNAL_INFO view (continued)

Column Name	System Column Name	Data Type	Description
RETRANSMISSIONS	RETRANS	BIGINT Nullable	The total number of times the local system retransmitted a segment because an acknowledgement was not received. This is a cumulative count of all segments resent since the remote journal was last activated. A value greater than zero may indicate a problem with the network. Contains the null value if this information is not available
LAST_CATCHUP_TIME	LAST_CATCH	TIMESTAMP Nullable	The date and time that the remote journal environment last transitioned to DELIVERY_MODE *ASYNCPEND or *SYNCPEND. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE.
LAST_ACTIVE_TIME	LAST_ACT	TIMESTAMP Nullable	The date and time that the remote journal environment last transitioned from DELIVERY_MODE *ASYNCPEND or *SYNCPEND to DELIVERY_MODE *ASYN or *SYN. Contains the null value if REMOTE_JOURNAL_STATE is INACTIVE
CONTROLLED_INACTIVATE_SEQUENCE_NUMBER	INACT_SEQ	DECIMAL(21,0) Nullable	The sequence number of the last journal entry that was queued for replication before the Change Remote Journal (CHGRMTJRN) command or the Change Journal State (QjoChangeJournalState) API was called to start a controlled inactivate of the remote journal. Contains the null value if REMOTE_JOURNAL_STATE is not CTLINACT.
CONTROLLED_INACTIVATE_JOURNAL_RECEIVER_LIBRARY	INACTJRCVL	VARCHAR(10) Nullable	The library of the journal receiver that contains the CONTROLLED_INACTIVATE_SEQUENCE_NUMBER. Contains the null value if REMOTE_JOURNAL_STATE is not CTLINACT.
CONTROLLED_INACTIVATE_JOURNAL_RECEIVER	INACTJRCV	VARCHAR(10) Nullable	The name of the journal receiver that contains the CONTROLLED_INACTIVATE_SEQUENCE_NUMBER. Contains the null value if REMOTE_JOURNAL_STATE is not CTLINACT.
SOURCE_JOURNAL_IASP_NUMBER	SRC_ASPNUM	INTEGER	The number of the auxiliary storage pool to which storage for the journal that has the remote journal defined is allocated.

Example

- Return information about remote journals defined for system RMTSYS1.

```
SELECT * FROM QSYS2.REMOTE_JOURNAL_INFO
WHERE REMOTE_DATABASE_NAME = 'RMTSYS1';
```

Librarian Services

These services provide object and library list information.

JOURNAL_INHERIT_RULES view

The JOURNAL_INHERIT_RULES view returns the journal inherit rules for libraries.

The rules define which objects created in a library, moved into a library, or restored into a library should inherit the journal state of the library and the inherited journal attributes.

Each rule defines object types, object names, and operations that the rule applies to. Multiple rules can be defined for the same set of objects. If multiple rules are defined for the same object and operation, the last rule defined for that object will be applied.

This information is similar to the journal inherit rules returned by the Retrieve Library Description (QLIRLIBD) API.

Authorization: The caller must have *READ authority to the library.

The following table describes the columns in the view. The system name is LIB_JRN. The schema is QSYS2.

Table 147. JOURNAL_INHERIT_RULES view

Column Name	System Column Name	Data Type	Description
LIBRARY_NAME	LIBRARY	VARCHAR(10)	The name of the library for this rule.
JOURNALED	JOURNALED	VARCHAR(3)	The current journaling status of the library. Rules are only in effect when the library is journaled. NO The library is not currently journaled. YES The library is currently journaled.
IASP_NUMBER	IASPNUMBER	INTEGER	The number of the auxiliary storage pool (ASP) from which the system allocates storage for the library.
IASP_NAME	IASP_NAME	VARCHAR(10)	The device description name of the independent auxiliary storage pool (IASP). The special value of *SYSBAS indicates SYSBASE, which includes the system ASP (ASP 1) and the basic user ASPs (ASPs 2-32).
ORDINAL_POSITION	ORDINAL	INTEGER	The order this rule was added for this library. The number starts at one for each library and increments by one for every additional rule for that library. If more than one rule is defined for any object, the one with the higher number will be used.
OBJECT_TYPE	TYPE	VARCHAR(7)	Specifies the object type of the objects that are identified by this rule. *ALL This rule applies to all object types that can be journaled. *DTAARA This rule applies to data areas. *DTAQ This rule applies to data queues. *FILE This rule applies to database physical files.

Table 147. JOURNAL_INHERIT_RULES view (continued)

Column Name	System Column Name	Data Type	Description
OPERATION	OPERATION	VARCHAR(10)	<p>The operations performed on the object type for which journaling is to be started if other criteria specified for the object type are also satisfied.</p> <p>*ALLOPR Journaling is started for all objects created in, moved into, or restored into the library. This is a combination of the values *CREATE, *MOVE, and *RESTORE.</p> <p>*CREATE Journaling is started for all objects created in the library.</p> <p>*MOVE Journaling is started for all objects moved into the library if they are not already journaled.</p> <p>*RESTORE If an object is restored over a currently existing object, the restored object will retain the same journal options and journal state of the object it was restored over. If an object was never journaled when it was saved, journaling is started for the object when it is restored into the library. If an object was journaled when it was saved, it will first attempt to start journaling to the journal it was journaled to when it was saved, with the same journal options it had when it was saved. If that journal does not exist, the object will start journaling to the same journal the library is journaled to, with the journal options defined by this rule.</p> <p>*RSTOVRJRN If an object is restored over a currently existing object, the restored object will retain the same journal options and journal state of the object it was restored over. Otherwise, journaling is started for all objects restored into the library, regardless of the journal options and journal state of the object when it was saved.</p>
RULE_ACTION	ACTION	VARCHAR(8)	<p>Indicates whether the objects that match the object type and operation in this rule will be included or omitted from the list of objects that inherit the journal options and journal state of the library.</p> <p>*INCLUDE All objects that match the object type and operation of this rule will inherit the journal options and journal state of the library.</p> <p>*OMIT All objects that match the object type and operation of this rule will not inherit any journaling attributes or state from the library. This overrides a *INCLUDE rule and therefore can be used to omit a subset of a previously defined *INCLUDE rule.</p>
NAME_FILTER	NAMEFILTER	VARCHAR(10) Nullable	<p>The object names of the objects that are identified by this rule.</p> <p>name This rule applies to all objects that match the other criteria and match the specified name.</p> <p>generic-name This rule applies to all objects that match the other criteria and match the specified generic name.</p> <p>A generic name is specified as a string that contains one or more characters followed by an asterisk (*). If a generic name is specified, then all objects that have names with the same prefix as the generic object name are selected.</p> <p>*ALL This rule applies to all objects that match the other criteria.</p>

Table 147. JOURNAL_INHERIT_RULES view (continued)

Column Name	System Column Name	Data Type	Description
JOURNAL_IMAGES	IMAGES	VARCHAR(7) Nullable	<p>The kinds of journal images that are written to the journal receiver for changes to objects that inherit the journal options and journal state from the library.</p> <p>*AFTER Only after images are journaled for an object for which journaling is started because it inherits the journaling attributes from the library.</p> <p>*BOTH Both before and after images are journaled for an object for which journaling is started because it inherits the journal option from the library. This value is only valid for data area (*DTAARA) and database file (*FILE) objects. If this value is specified and *ALL is specified for object type, the system will generate both before and after images for data areas and database files and the system will only generate after images for all other object types.</p> <p>*OBJDFT The default value for each object type will be used for this journal option when an object inherits the journaling attributes from the library. Database files (*FILE) will have both before and after images generated by the system. All other object types will have only after images generated by the system.</p> <p>Contains the null value if RULE_ACTION is *OMIT.</p>
OMIT_JOURNAL_ENTRY	OMIT_ENTRY	VARCHAR(7) Nullable	<p>The journal entries that are not to be written for changes to objects that inherit the journaling options and state of the library.</p> <p>*NONE No journal entries will be omitted for objects that inherit the journal options and journal state from the library.</p> <p>*OBJDFT The default value for each object type will be used for this journal option when an object inherits the journal options and journal state from the library. Open and close entries will be omitted for database files (*FILE). No other object types will omit journal entries.</p> <p>*OPNCLO Open and close entries are omitted for database file (*FILE) objects that inherit the journal options and journal state from the library.</p> <p>This prevents the use of TOJOB0 and TOJOB0C entries on the Apply Journalized Changes (APYJRNCHG) and Remove Journalized Changes (RMVJRNCHG) commands, but it saves some storage space in the journal receivers.</p> <p>This value is only valid for object type *FILE. If this value is specified and *ALL is specified for object type, database files will omit the entries. All other object types will not omit any journal entries.</p> <p>Contains the null value if RULE_ACTION is *OMIT.</p>

Table 147. JOURNAL_INHERIT_RULES view (continued)

Column Name	System Column Name	Data Type	Description
REMOTE_JOURNAL_FILTER	RMT_FILTER	VARCHAR(7) Nullable	<p>Specifies whether the journal entries written for the objects that inherit the journal state of the library should be eligible for remote journal filtering by object.</p> <p>*NO Journal entries deposited for the objects that inherit the journal state of the library will not be eligible for remote journal filtering by object.</p> <p>*OBJDFT The default value for each object type will be used for this journaling attribute when an object inherits the journal state of the library. For all object types, journal entries deposited for the objects that inherit the journal state of the library will not be eligible for remote journal filtering by object.</p> <p>*YES Journal entries deposited for the objects that inherit the journal state of the library will be eligible for remote journal filtering by object. When using remote journal filtering by object, most journal entries for the object will not be sent to the target system.</p>

Example

- Retrieve the journal inherit rules for library APPLIB.


```
SELECT * FROM QSYS2.JOURNAL_INHERIT_RULES
WHERE LIBRARY_NAME = 'APPLIB';
```

LIBRARY_INFO table function

The LIBRARY_INFO table function returns a result table that contains information about a specific library.

This information is similar to what is returned by the Retrieve Library Description (QLIRLIBD) API and the Retrieve Library Description (RTVLIBD) command

Authorization: The caller must have *READ authority to the library. The null value is returned for the OBJECT_AUDIT_CREATE column unless the caller has *ALLOBJ or *AUDIT special authority.



The schema is QSYS2.

library-name An expression that contains the name of a library.

ignore-errors A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned.

YES A warning is returned.

No row is returned when an error is encountered. This is the default.

detailed-info A character or graphic string expression that indicates which columns should be returned.

LIBRARY_SIZE All the columns including the LIBRARY_SIZE column are returned. For a library with a large number of objects, LIBRARY_SIZE can be a time consuming calculation. This is the default.

NO Null is returned for the LIBRARY_SIZE column.

The result of the function is a table containing one row with the format shown in the following table. All the columns are nullable.

Table 148. LIBRARY_INFO table function

Column Name	Data Type	Description
OBJECT_COUNT	INTEGER	The total number of external objects in the specified library. The count includes objects to which you may not be authorized.
LIBRARY_SIZE	BIGINT	The size of the library, in bytes, which includes the size of the objects in the library plus the size of the library object itself. This value will be a rounded value for a library larger than 1 000 000 000 bytes. Contains the null value if <i>detailed-info</i> is NO.
LIBRARY_SIZE_COMPLETE	VARCHAR(3)	Whether LIBRARY_SIZE includes all objects in the library. NO Some objects in the library were omitted because they were locked, or the user does not have any authority to the object. YES The size of all the objects in the library was used in determining the total library size. Contains the null value if <i>detailed-info</i> is NO.
LIBRARY_TYPE	CHAR(4)	The library type. PROD The library is a production library. Database files in production libraries cannot be opened for updating if a user, while in debug mode, requested that production libraries be protected. TEST The library is a test library. All objects in a test library can be updated during a test. See the Start Debug (STRDBG) command in the on-line help for more details.
TEXT_DESCRIPTION	VARCHAR(50)	The user-defined text that briefly describes the library and its function. Contains the null value if the library has no descriptive text.
IASP_NAME	VARCHAR(10)	The auxiliary storage pool (ASP) device name where the library is stored. If the library is in the system ASP or one of the basic user ASPs, contains *SYSBAS. Contains the null value if the name of the ASP device cannot be determined.
IASP_NUMBER	INTEGER	The number of the auxiliary storage pool (ASP) from which the system allocates storage for the library.

Table 148. LIBRARY_INFO table function (continued)

Column Name	Data Type	Description
CREATE_AUTHORITY	VARCHAR(10)	<p>The default public authority for an object created into the library. This is the authority given to a user who does not have specific authority to the object, who is not on an authorization list specified for the object, and whose user groups have no specific authority to the object. When you create an object into the library, the AUT parameter on the create command for the object determines the public authority for the object. If the AUT value on the create command for the object is *LIBCRTAUT, which is the default, the public authority for the object is set to the CRTAUT value for the library.</p> <p>*ALL The user can perform all authorized operations on an object created in this library.</p> <p>*CHANGE The user can read the object description and has read, add, update, and delete authority to an object created in this library.</p> <p>*EXCLUDE The user is prevented from accessing an object created in this library.</p> <p>*SYSVAL The default authority for an object created in this library is determined by the value specified by the QCRTAUT system value.</p> <p>*USE The user can read the object and its description but cannot change them for an object created in this library.</p> <p>Authorization list name The name of the authorization list that secures an object created in this library. The default public authority is taken from the authorization list, and the public authority for the object is specified as *AUTL.</p>
OBJECT_AUDIT_CREATE	VARCHAR(7)	<p>The auditing value for objects created in this library.</p> <p>*ALL All change or read access to the object is logged.</p> <p>*CHANGE All change access to the object by all users is logged.</p> <p>*NONE Use or change access to the object is not logged (no audit entry is sent to the security journal).</p> <p>*SYSVAL The value specified in the system value QCRTOBJAUD is used.</p> <p>*USRPRF The user profile of the user who accesses the object is used to determine if an audit record is sent for this access. The OBJAUD parameter of the Change User Auditing (CHGUSRAUD) command is used to turn auditing on for a specific user.</p> <p>Contains the null value if caller does not have either all object (*ALLOBJ) or audit (*AUDIT) special authority.</p>
JOURNALED	VARCHAR(3)	<p>Identifies the current journaling status of the library.</p> <p>NO The library is not currently journaled.</p> <p>YES The library is currently journaled.</p>
JOURNAL_LIBRARY	VARCHAR(10)	<p>The name of the library that contains the journal that receives the journaled changes to the library, if the library is currently journaled. If the library was previously journaled but is not currently journaled, contains the name of the library that contains the last journal to which the library was journaled.</p> <p>Contains the null value if journaling has never been started for this library.</p>
JOURNAL_NAME	VARCHAR(10)	<p>The name of the journal that receives the journaled changes to the library, if the library is currently journaled. If the library was previously journaled but is not currently journaled, contains the name of the last journal to which the library was journaled.</p> <p>Contains the null value if journaling has never been started for this library.</p>

Table 148. LIBRARY_INFO table function (continued)

Column Name	Data Type	Description
INHERIT_JOURNALING	VARCHAR(3)	Identifies whether new journal-eligible objects created into, moved into, or restored into this library should inherit journaling from the library according to the journal inherit rules. NO The new journal-eligible objects will not inherit journaling from the library. YES The new journal-eligible objects will inherit journaling from the library according to the journal inherit rules.
JOURNAL_INHERIT_RULES	INTEGER	The number of inherit rules defined for this library. The rules can be listed using the QSYS2.JOURNAL_INHERIT_RULES view.
JOURNAL_START_TIMESTAMP	TIMESTAMP	The timestamp journaling was last started for this library. Contains the null value if journaling has never been started for this library.
APPLY_STARTING_RECEIVER_LIBRARY	VARCHAR(10)	Specifies the library name of the oldest journal receiver needed to successfully use the Apply Journalized Changes (APYJRNCHG) command. Contains the null value if the object has not been journaled or it has not been saved and restored since journaling was started.
APPLY_STARTING_RECEIVER	VARCHAR(10)	Specifies the name of the oldest journal receiver needed to successfully use the Apply Journalized Changes (APYJRNCHG) command. Contains the null value if the object has not been journaled or it has not been saved and restored since journaling was started.
APPLY_STARTING_RECEIVER_ASP	VARCHAR(10)	The auxiliary storage pool (ASP) device name where APPLY_STARTING_RECEIVER is stored. If APPLY_STARTING_RECEIVER is in the system ASP or one of the basic user ASPs, contains *SYSBAS. Contains the null value if the name of the ASP device cannot be determined.

Example

- Retrieve information about library APPLIB.

```
SELECT * FROM TABLE(QSYS2.LIBRARY_INFO('APPLIB'));
```

LIBRARY_LIST_INFO view

The LIBRARY_LIST_INFO view contains information about the current job's library list.

The following table describes the columns in the view. The system name is LIBLIST. The schema is QSYS2.

Table 149. LIBRARY_LIST_INFO view

Column Name	System Column Name	Data Type	Description
ORDINAL_POSITION	COLNO	INTEGER	Position of this entry in the library list.
SCHEMA_NAME	NAME	VARCHAR(128) Nullable	Name of the schema.
SYSTEM_SCHEMA_NAME	SYS_NAME	VARCHAR(10)	System name of the schema.
TYPE	TYPE	VARCHAR(15)	The portion of the library list containing the selected library. Possible values are: USER The library is in the user portion of the library list. SYSTEM The library is in the system portion of the library list. PRODUCT The library is a product library in the library list. CURRENT The library is the current library entry in the library list.

Table 149. LIBRARY_LIST_INFO view (continued)

Column Name	System Column Name	Data Type	Description
IASP_NUMBER	IASP	SMALLINT Nullable	The number of the auxiliary storage pool where storage is allocated for the library.
TEXT_DESCRIPTION	TEXT	VARGRAPHIC(50) CCSID 1200 Nullable	The text description of the library. Contains the null value if there is no text description.

Example

- See the current library list for your job

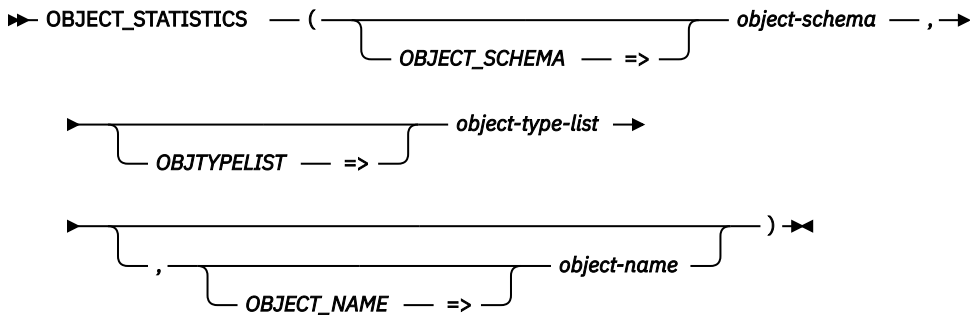
```
SELECT * FROM QSYS2.LIBRARY_LIST_INFO
```

OBJECT_STATISTICS table function

The OBJECT_STATISTICS table function returns information about objects in a library.

Authorization:

- For an object that is not a user profile:
 - If the caller has *EXECUTE authority to the library,
 - If the caller has *OBJOPR and *READ authority to an object, full details are returned.
 - Otherwise, partial information is returned along with an SQL warning of '01548'.
 - Otherwise, the object information is not returned.
- For a user profile object:
 - The caller must have at least one of the following:
 - Some authority to the user profile, or
 - Authorization to the QIBM_DB_SECADM function usage identifier.
 - Otherwise, the user profile object information is not returned.



The schema is QSYS2.

object-schema

A character or graphic string expression that identifies the name of a library. If the library's name is a delimited name, the delimited form of the name must be specified. It can be either a long or short library name.

The following special values are allowed for *object-schema*.

- *ALL** All libraries.
- *ALLAVL** All libraries in all available ASPs.
- *ALLUSR** All user libraries in *SYSBAS and the current thread's ASP group.

- *ALLUSRAVL** All user libraries in all available ASPs.
- *CURLIB** The job's current library.
- *LIBL** The library list.
- *USRLIBL** The job's current library and the user portion of the library list.

The following special value is allowed for *object-schema* when *object-type-list* is '*LIB' or 'LIB'.

- *ALLSIMPLE** The fastest approach to retrieving all user and system library names in *SYSBAS and the current thread's ASP group. Values are returned for the following columns: OBJNAME, OBJLONGNAME, OBJTYPE, OBJLIB, and OBJLONGSCHEMA. All other columns return NULL.

object-type-list A character or graphic string expression containing one or more system object types separated by either a blank or a comma. The object types can include or exclude the leading * character. The special value of '*ALL' or 'ALL' can be used to return all objects in the library *object-schema*.

object-name A character or graphic string expression that identifies the name of an object or a library. If the object's name is a delimited name, the delimited form of the name must be specified. It can be either a long or short object name. The name must be the valid system name for the object unless the object is a file or a library; for files and libraries the SQL name can be specified.

A generic name can be specified to find multiple system objects. If the last character in the name is an asterisk, the name is a generic name. For example, to return object names that start with 'EMP', specify an *object-name* value of 'EMP*'. When a generic name is specified, long SQL names are not included in the search.

If this parameter is specified, only objects with this name in *object-schema* corresponding to the object types in *object-type-list* are returned.

If this parameter is not specified, all objects in *object-schema* corresponding to the object types in *object-type-list* are returned.

If *object-schema* is *ALL, *ALLSIMPLE, *ALLAVL, *ALLUSR, or *ALLUSRAVL and *object-type* is *LIB or LIB, the *object-name* parameter is ignored.

The following special value is allowed for *object-name*.

- *ALLSIMPLE** The fastest approach to retrieving the system names for objects in a library. All objects in *object-schema* corresponding to the object types in *object-type-list* are returned. Values are returned for the following columns: OBJNAME, OBJTYPE, OBJLIB, and OBJLONGSCHEMA. All other columns return NULL.

The result of the function is a table containing a row for each object with the format shown in the following table. All the columns are null capable.

Table 150. OBJECT_STATISTICS table function

Column Name	Data Type	Description
OBJNAME	VARCHAR(10)	System name of the object.
OBJTYPE	VARCHAR(8)	System type of the object.
OBJOWNER	VARCHAR(10)	The user profile that owns the object. Contains the null value if no owner is available.
OBJDEFINER	VARCHAR(10)	The user profile that created the object. Contains the null value if the object definer is not known.
OBJCREATED	TIMESTAMP	Timestamp of when the object was created.
OBJSIZE	DECIMAL(15,0)	Size of the object, in bytes.

Table 150. OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
OBJTEXT	VARCHAR(50)	The description of the object. Contains the null value if the object has no text.
OBJLONGNAME	VARCHAR(128)	The SQL name for the object. For an external procedure or an external function, the name will be returned when a single procedure or function exists for that *PGM or *SRVPGM object. Contains the null value if an SQL name could not be returned.
LAST_USED_TIMESTAMP	TIMESTAMP	The date the object was used last. The time portion of the timestamp will always be 0. Contains the null value if the object has not been used.
LAST_USED_OBJECT	VARCHAR(4)	Indicates whether the LAST_USED_TIMESTAMP value is meaningful for this object. NO The object does not maintain the LAST_USED_TIMESTAMP. YES The object maintains the LAST_USED_TIMESTAMP.
DAYS_USED_COUNT	INTEGER	The number of days an object has been used on the system.
LAST_RESET_TIMESTAMP	TIMESTAMP	The date when the days used count was last reset to zero. The time portion of the timestamp will always be 0. Contains the null value if the days used count has not been reset.
IASP_NUMBER	SMALLINT	The auxiliary storage pool (ASP) where storage is allocated for the object.
OBJATTRIBUTE	VARCHAR(10)	The attribute for this object's type, if any. Contains an empty string if no attribute.
OBJLONGSCHEMA	VARCHAR(128)	The SQL schema name for this object.
TEXT	VARGRAPHIC(50) CCSID 1200	The description of the object for *LIB objects. Contains the null value if OBJTYPE is not *LIB.
SQL_OBJECT_TYPE	VARCHAR(9)	The SQL type of the object. Values are: <ul style="list-style-type: none"> • ALIAS • FUNCTION • INDEX • PACKAGE • PROCEDURE • ROUTINE • SEQUENCE • TABLE • TRIGGER • TYPE • VARIABLE • VIEW • XSR Contains the null value if the object is not an SQL object.
OBJLIB	VARCHAR(10)	System name of the schema.
CHANGE_TIMESTAMP	TIMESTAMP	The timestamp of the last time the object was changed.
USER_CHANGED	VARCHAR(3)	Indicates whether the object was modified by a user. Values are: NO The object was not modified by a user. YES The object was modified by a user. Contains the null value for certain objects that are installed as part of the operating system.
SOURCE_FILE	VARCHAR(10)	The name of the source file that was used to create the object. Contains the null value if a source file was not used.

Table 150. OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
SOURCE_LIBRARY	VARCHAR(10)	The name of the source file library that was used to create the object. Contains the null value if a source file was not used.
SOURCE_MEMBER	VARCHAR(10)	The name of the source file member that was used to create the object. Contains the null value if a source file was not used.
SOURCE_TIMESTAMP	TIMESTAMP	The last source update timestamp of the member in the source file at the time the object was created. Contains the null value if a source file was not used or if the source timestamp is not available.
CREATED_SYSTEM	VARCHAR(8)	The name of the system on which the object was created. Contains the null value if the system is not known.
CREATED_SYSTEM_VERSION	VARCHAR(9)	The version of the operating system when the object was created. The field has a VxRxMx format where: Vx The character V is followed by a version number. Rx The character R is followed by a release level. Mx The character M is followed by a modification level. Contains the null value if the version is not known.
LICENSED_PROGRAM	VARCHAR(7)	The name of the licensed program if the object is part of a licensed program. Contains the null value if the object is not a part of a licensed program.
LICENSED_PROGRAM_VERSION	VARCHAR(9)	The version number, release level, and modification level of the licensed program if the object is part of a licensed program. The field has a VxRxMx format where: Vx The character V is followed by a version number. Rx The character R is followed by a release level. Mx The character M is followed by a modification level. Contains the null value if the object is not a part of a licensed program.
COMPILER	VARCHAR(7)	The licensed program identifier of the compiler. Contains the null value if the object was not created with a compiler.
COMPILER_VERSION	VARCHAR(9)	The licensed program version number, release level, and modification level of the compiler. The field has a VxRxMx format where: Vx The character V is followed by a version number. Rx The character R is followed by a release level. Mx The character M is followed by a modification level. Contains the null value if the object was not created with a compiler.
OBJECT_CONTROL_LEVEL	CHAR(8)	The object control level for the object. Contains the null value if there is no object control level.
PTF_NUMBER	CHAR(7)	The Program Temporary Fix that resulted in the creation of this object. Contains the null value for a user created object.
APAR_ID	CHAR(6)	The authorized program analysis report (APAR) with this identification number associated with the last change. Will contain the value CHGDFT for a command that is changed using CHGCMDDFT. The Change Object Description (QLICOBJD) API can change this field to any value. Contains the null value is no value is available.
USER_DEFINED_ATTRIBUTE	VARCHAR(10)	Further defines an object type. This field is set by the user by using the QLICOBJD API. Contains the null value if the attribute has not been set.

Table 150. OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
ALLOW_CHANGE_BY_PROGRAM	VARCHAR(3)	<p>Identifies whether or not any changes other than the text or the days used count and reset date can be made to the object's description by the Change Object Description (QLICOBJD) API.</p> <p>NO The QLICOBJD API cannot be used to change fields in the object's description other than the text or the days used count and reset date.</p> <p>YES The QLICOBJD API can be used to change fields in the object's description.</p>
CHANGED_BY_PROGRAM	VARCHAR(3)	<p>Identifies whether the object has been modified by the Change Object Description (QLICOBJD) API.</p> <p>NO The object has not been modified by the QLICOBJD API.</p> <p>YES The object has been modified by the QLICOBJD API.</p>
COMPRESSED	VARCHAR(4)	<p>Indicates whether the object is compressed or decompressed. Values are:</p> <p>NO Permanently decompressed and compressible.</p> <p>YES Compressed.</p> <p>TEMP Temporarily decompressed.</p> <p>FREE Saved with storage freed; compression status cannot be determined.</p> <p>Contains the null value if the object is permanently decompressed and not compressible.</p>
PRIMARY_GROUP	VARCHAR(10)	<p>The name of the user profile that is the primary group for the object.</p> <p>Contains the null value if there is no primary group for the object.</p>
STORAGE_FREED	VARCHAR(3)	<p>The storage status of the object data.</p> <p>NO The storage for the object data has not been freed.</p> <p>YES The storage for the object data has been freed. See the SAVOBJ or SAVLIB command, STG parameter, for more details.</p>
ASSOCIATED_SPACE_SIZE	INTEGER	<p>The size, in bytes, of the primary associated space of the object.</p> <p>Contains the null value if the object has no primary associated space.</p>
OPTIMUM_SPACE_ALIGNMENT	VARCHAR(3)	<p>Identifies whether the primary associated space for the object has been optimally aligned. Optimum alignment may allow for better performance of applications that manipulate the object.</p> <p>NO The space associated with the object has not been optimally aligned.</p> <p>YES The space associated with the object has been optimally aligned.</p> <p>Contains the null value if the object has no associated space.</p>
OVERFLOW_STORAGE	VARCHAR(3)	<p>Indicates if the object has overflowed the auxiliary storage pool it resides in.</p> <p>NO The object has not overflowed the auxiliary storage pool.</p> <p>YES The object has overflowed the auxiliary storage pool.</p>
OBJECT_DOMAIN	VARCHAR(7)	<p>The domain that contains the object. Values are:</p> <p>*SYSTEM The object is in the system domain.</p> <p>*USER The object is in the user domain.</p>

Table 150. OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
OBJECT_AUDIT	VARCHAR(10)	<p>The type of auditing for an object. Values are:</p> <ul style="list-style-type: none"> *ALL Audit all access to this object by all users on the system. All access is defined as a read or change operation. *CHANGE Audit all change access to this object by all users on the system. *NONE No auditing occurs for this object when it is read or changed regardless of the user who is accessing the object. *USRPRF Audit this object only if the current user is being audited. The current user is tested to determine if auditing should be done for this object. The user profile can specify if only change access is audited or if both read and change accesses are audited for this object. <p>Contains the null value if you do not have either all object (*ALLOBJ) or audit (*AUDIT) special authority.</p>
OBJECT_SIGNED	VARCHAR(3)	<p>Indicates whether the object has a digital signature.</p> <ul style="list-style-type: none"> NO The object does not have a digital signature. YES The object has a digital signature.
SYSTEM_TRUSTED_SOURCE	VARCHAR(3)	<p>Indicates whether the object is signed by a source that is trusted by the system.</p> <ul style="list-style-type: none"> NO None of the object signatures came from a source that is trusted by the system. YES The object is signed by a source that is trusted by the system. If the object has multiple signatures, at least one of the signatures came from a source that is trusted by the system.
MULTIPLE_SIGNATURES	VARCHAR(3)	<p>Indicates whether the object has more than one digital signature.</p> <ul style="list-style-type: none"> NO The object has only one digital signature or does not have a digital signature. YES The object has more than one digital signature.
SAVE_TIMESTAMP	TIMESTAMP	<p>The timestamp the object was last saved.</p> <p>Contains the null value if the object has not been saved.</p>
RESTORE_TIMESTAMP	TIMESTAMP	<p>The timestamp the object was last restored.</p> <p>Contains the null value if the object has not been restored.</p>
SAVE_WHILE_ACTIVE_TIMESTAMP	TIMESTAMP	<p>The timestamp at which the object was saved while active.</p> <p>Contains the null value if the object has not been saved while active.</p>
SAVE_COMMAND	VARCHAR(10)	<p>The command used to save the object.</p> <p>Contains the null value if the object has not been saved.</p>
SAVE_DEVICE	VARCHAR(5)	<p>The type of the device to which the object was last saved. Valid values are:</p> <ul style="list-style-type: none"> *OPT The object was saved to optical. *SAVF The object was saved to a save file. *TAP The object was saved to tape. <p>Contains the null value if the object has not been saved.</p>
SAVE_FILE_NAME	VARCHAR(10)	<p>The save file used to save the object.</p> <p>Contains the null value if the object was not last saved to a save file.</p>
SAVE_FILE_LIBRARY	VARCHAR(10)	<p>The save file library used to save the object.</p> <p>Contains the null value if the object was not last saved to a save file.</p>
SAVE_VOLUME	VARCHAR(71)	<p>The tape or optical volumes used to save the object. A maximum of ten volumes is returned. The string contains one blank between volume identifiers. If more than ten volumes were used, an ellipsis (three periods) is returned to the right of the identifier of the tenth volume.</p> <p>Contains the null value if the object was not last saved to tape or optical.</p>

Table 150. OBJECT_STATISTICS table function (continued)

Column Name	Data Type	Description
SAVE_LABEL	VARCHAR(17)	The file label used when the object was saved to tape or optical. Contains the null value if the object was not last saved to tape or optical.
SAVE_SEQUENCE_NUMBER	DECIMAL(10,0)	The sequence number used to when the object was saved to tape. Contains the null value if the object was not last saved to tape.
LAST_SAVE_SIZE	DECIMAL(15,0)	The size of the object in bytes at the time of the last save. This value defines the amount of storage that is required if the object is restored. Contains the null value if the object has not been saved.
JOURNALED	VARCHAR(3)	Identifies the current journaling status of the object. Valid values are: NO The object is not currently journaled. YES The object is currently journaled.
JOURNAL_NAME	VARCHAR(10)	The name of journal that receives the journaled changes or the name of the last journal if the object is not currently journaled. Contains the null value if the object has not been journaled.
JOURNAL_LIBRARY	VARCHAR(10)	The name of journal library that receives the journaled changes or the name of the last journal library if the object is not currently journaled. Contains the null value if the object has not been journaled.
JOURNAL_IMAGES	VARCHAR(6)	Specifies the kinds of images that are generated for changes to the object. Valid values are: *AFTER Only after images are generated for changes to the object. *BOTH Both before and after images are generated for changes to the object. Contains the null value if the object has not been journaled.
OMIT_JOURNAL_ENTRY	VARCHAR(7)	Specifies the journal entries that are omitted. Valid values are: *NONE No journal entries are omitted. *OPNCLO Open and close entries are omitted. Open and close operations on the specified file members do not create open and close journal entries. Contains the null value if the object has not been journaled.
REMOTE_JOURNAL_FILTER	VARCHAR(3)	The remote journal filter value for the object. Valid values are: NO The journal entries deposited for the object will not be eligible for remote journal filtering. YES The journal entries deposited for the object will be eligible for remote journal filtering. Contains the null value if the object has not been journaled.
JOURNAL_START_TIMESTAMP	TIMESTAMP	The timestamp journaling was last started. Contains the null value if the object has not been journaled.
APPLY_STARTING_RECEIVER	VARCHAR(10)	Specifies the name of the oldest journal receiver needed to successfully use the Apply Journaled Changes (APYJRNCHG) or Remove Journaled Changes (RMVJRNCHG) command. Contains the null value if the object has not been journaled or it has not been saved and restored since journaling was started.
APPLY_STARTING_RECEIVER_LIBRARY	VARCHAR(10)	Specifies the library name of the oldest journal receiver needed to successfully use the Apply Journaled Changes (APYJRNCHG) or Remove Journaled Changes (RMVJRNCHG) command. Contains the null value if the object has not been journaled or it has not been saved and restored since journaling was started.

Example

- Find all journals in library MJATST.

```
SELECT * FROM TABLE (QSYS2.OBJECT_STATISTICS('MJATST ', 'JRN') ) AS X
```

or

```
SELECT * FROM TABLE (QSYS2.OBJECT_STATISTICS('MJATST ', '*JRN') ) AS X
```

- Find all journals and journal receivers in library MJATST.

```
SELECT * FROM TABLE (QSYS2.OBJECT_STATISTICS('MJATST ', 'JRN JRNRCV') ) AS X
```

or

```
SELECT * FROM TABLE (QSYS2.OBJECT_STATISTICS('MJATST ', '*JRN *JRNRCV') ) AS X
```

- Find all programs and service programs in library MYLIB. Use *ALLSIMPLE to return the list quickly, omitting the detail information.

```
SELECT * FROM TABLE (QSYS2.OBJECT_STATISTICS('MYLIB', 'PGM SRVPGM', '*ALLSIMPLE')) X
```

- Find any CL commands that have had their parameter defaults changed.

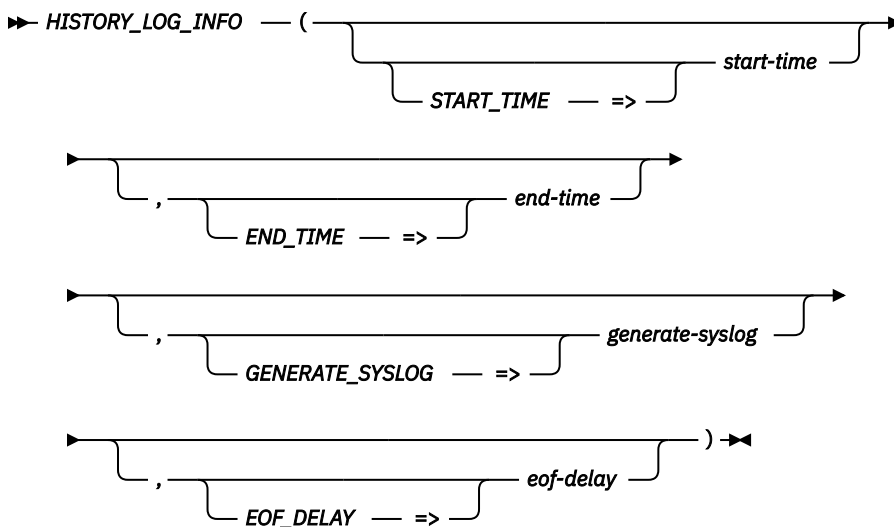
```
SELECT * FROM TABLE(QSYS2.OBJECT_STATISTICS('QSYS', '*CMD'))  
WHERE APAR_ID = 'CHGDFT';
```

Message Handling Services

These views and functions provide system message information.

HISTORY_LOG_INFO table function

The HISTORY_LOG_INFO table function returns one row for each message in the history log based on the timestamp range specified. It returns information similar to what is returned by the Display Log (DSPLOG) CL command and the Open List of History Log Messages (QMHOHST) API.



The schema is QSYS2.

Authorization: No authorization needed.

start-time A timestamp expression that indicates the starting timestamp to use when returning history log information.

If this parameter is omitted, the default of CURRENT DATE - 1 DAY is used.

end-time A timestamp expression that indicates the ending timestamp to use when returning history log information.

If this parameter is omitted, the default of '9999-12-30-00.00.00.000000' is used.

generate-syslog A character or graphic string expression that indicates whether to transform history log messages into syslog formatted detail. Values are:

NO No syslog information will be returned. The SYSLOG_EVENT, SYSLOG_FACILITY, SYSLOG_SEVERITY, and SYSLOG_PRIORITY columns will contain the null value.

RFC3164 Values will be returned for the SYSLOG_EVENT, SYSLOG_FACILITY, SYSLOG_SEVERITY, and SYSLOG_PRIORITY columns for each history log message. The SYSLOG_EVENT column will contain a syslog header that matches the RFC3164 format as described by the Internet Engineering Task Force (IETF) Request For Comments (RFC) 3164.

RFC5424 Values will be returned for the SYSLOG_EVENT, SYSLOG_FACILITY, SYSLOG_SEVERITY, and SYSLOG_PRIORITY columns for each history log message. The SYSLOG_EVENT column will contain a syslog header that matches the RFC5424 format as described by the Internet Engineering Task Force (IETF) Request For Comments (RFC) 5424.

If *generate-syslog* is not specified or is the null value, NO is used.

eof-delay An integer expression that specifies the number of seconds to sleep when all history log messages have been read. This delay allows the caller to establish a polling service that will continually return rows, sleeping for the specified interval whenever all messages have been processed.

A value of zero indicates no delay is used and a finite set of rows will be returned. A value greater than zero indicates that the table function will sleep, as needed, to wait for new history log messages and never end. If *eof-delay* is not specified or is the null value, zero is used.

If this parameter has a value greater than zero, the *generate-syslog* parameter must be RFC3164 or RFC5424, and the *end-time* parameter cannot be specified with a value other than its default value.

When using a non-zero *eof-delay* parameter, avoid using query clauses that depend on returning a finite number of rows. For example, using the FETCH FIRST n ROWS clause can cause the query to end when the requested number of rows has been satisfied. A query using the HISTORY_LOG_INFO function with a non-zero *eof-delay* parameter does not allow data to be copied (ALWCOPYDATA(*NO)). This means that a query requiring a copy of data, such as one using an ORDER BY clause or UNION DISTINCT, will issue an error and not be allowed. When using *eof-delay*, consider using a simple query to avoid blocking of rows. When rows are blocked for data transport efficiency, rows won't be returned until the block is full. Therefore, you should decide whether you favor data transport efficiency or moving events as soon as they occur.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 151. HISTORY_LOG_INFO table function

Column Name	Data Type	Description
ORDINAL_POSITION	INTEGER	A unique number for each row that indicates the time order of messages in the job log. The first (oldest) message returned from the history log will have a value of 1. Subsequent messages will have a value one greater than the previous message. Since these values are assigned when this catalog is queried, there will be no gaps in values.

Table 151. HISTORY_LOG_INFO table function (continued)

Column Name	Data Type	Description
MESSAGE_ID	VARCHAR(7)	The message ID for this message. Contains the null value if this is an impromptu message or MESSAGE_TYPE is REPLY.
MESSAGE_TYPE	VARCHAR(13)	Type of message. <ul style="list-style-type: none"> • COMPLETION • DIAGNOSTIC • ESCAPE • INFORMATIONAL • INQUIRY • NOTIFY • REPLY • REQUEST • SENDER
MESSAGE_SUBTYPE	VARCHAR(22)	Subtype of message. The values returned for REPLY messages: <ul style="list-style-type: none"> • FROM EXIT PROGRAM • FROM SYSTEM REPLY LIST • MESSAGE DEFAULT USED • NOT VALIDITY CHECKED • SYSTEM DEFAULT USED • VALIDITY CHECKED The value returned for some REQUEST messages: <ul style="list-style-type: none"> • WITH PROMPTING Contains the null value for other message types.
SEVERITY	SMALLINT	The severity assigned to the message.
MESSAGE_TIMESTAMP	TIMESTAMP	The timestamp when the message was sent.
FROM_USER	VARCHAR(10)	The current user of the job when the message was sent.
FROM_JOB	VARCHAR(28)	The qualified job name when the message was sent.
FROM_PROGRAM	VARCHAR(10)	The program that sent the message.
MESSAGE_LIBRARY	VARCHAR(10)	The name of the library containing the message file. Contains the null value if MESSAGE_ID is null.
MESSAGE_FILE	VARCHAR(10)	The message file containing the message. Contains the null value if MESSAGE_ID is null.
MESSAGE_TOKENS	VARCHAR(4096) FOR BIT DATA	The message token string. If the value is longer than 4096 characters, it will be truncated with no warning. Contains the null value if there are no message tokens.
MESSAGE_TEXT	VARGRAPHIC(1024) CCSID 1200	The first level text of the message including tokens, or the impromptu message text. Contains the null value if MESSAGE_ID is null or if the message file could not be accessed.
MESSAGE_SECOND_LEVEL_TEXT	VARGRAPHIC(4096) CCSID 1200	The second level text of the message including tokens. Contains the null value if MESSAGE_ID is null or if the message has no second level text or if the message file could not be accessed.
SYSLOG_EVENT	VARGRAPHIC(2048) CCSID 1200	The Common Event Format (CEF) syslog event for the message preceded by a header of the requested type. If a header-type of RFC3164 is requested, the maximum length is 1024 characters. If a header-type of RFC5424 is requested, the maximum length is 2048 characters. The string will be truncated with no warning if it exceeds the maximum length. The key names returned for history log information are listed in the Notes section. Contains the null value if NO was specified for the GENERATE_SYSLOG parameter.

Table 151. HISTORY_LOG_INFO table function (continued)

Column Name	Data Type	Description
SYSLOG_FACILITY	INTEGER	<p>The syslog facility assigned to the event.</p> <ul style="list-style-type: none"> 1 user-level messages 4 security/authorization messages <p>The facility assigned is defined in the Notes section.</p> <p>Contains the null value if NO was specified for the GENERATE_SYSLOG parameter.</p>
SYSLOG_SEVERITY	INTEGER	<p>The syslog severity assigned to the event.</p> <ul style="list-style-type: none"> 1 Alert: Action must be taken immediately 3 Error condition 4 Warning condition 5 Notice: A normal but significant condition 6 Informational message 7 Debug level message <p>The severity assigned is listed in the Notes section.</p> <p>Contains the null value if NO was specified for the GENERATE_SYSLOG parameter.</p>
SYSLOG_PRIORITY	INTEGER	<p>The syslog priority number assigned to the event.</p> <p>Contains the null value if NO was specified for the GENERATE_SYSLOG parameter.</p>

Notes

Syslog information: Syslog information is returned for all messages in the history log. Syslog information is also available for audit journal entries. See [DISPLAY_JOURNAL](#) table function for more details.

All history log messages return a SYSLOG_FACILITY value of 1 except as noted below. Messages are assigned a SYSLOG_SEVERITY value in the following way:

- Severity 1 Alert: Action must be taken immediately
 - MESSAGE_TYPE contains a value of INQUIRY, NOTIFY, or REPLY
- Severity 3 Error condition
 - MESSAGE_ID contains a value of CPF1164 with a job ending code value in the MESSAGE_TEXT column of 30 or higher
 - MESSAGE_TYPE contains a value of ESCAPE when the SEVERITY column contains a value of 50 or greater
- Severity 4 Warning condition
 - MESSAGE_ID contains a value of CPF1393. The SYSLOG_FACILITY column is set to 4.
 - MESSAGE_ID contains a value of CPF1164 with a job ending code value in the MESSAGE_TEXT column of 20
 - MESSAGE_TYPE contains a value of ESCAPE when the SEVERITY column contains a value of 30 or greater but less than 50
- Severity 5 Notice: A normal but significant condition
 - MESSAGE_ID contains a value of CPF1164 with a job ending code value in the MESSAGE_TEXT column of 10
 - MESSAGE_TYPE contains a value of INFORMATIONAL, COMPLETION, DIAGNOSTIC, or REQUEST when the SEVERITY column contains a value of 50 or greater
- Severity 6 Informational message

- MESSAGE_ID contains a value of CPF1164 with a job ending code value in the MESSAGE_TEXT column of 0
- MESSAGE_TYPE contains a value of ESCAPE when the SEVERITY column contains a value less than 30
- MESSAGE_TYPE contains a value of SENDER
- MESSAGE_TYPE contains a value of INFORMATIONAL, COMPLETION, DIAGNOSTIC, or REQUEST when the SEVERITY column contains a value less than 50
- Severity 7 Debug level message
 - MESSAGE_ID contains a value of CPF9897 or CPF9898 (regardless of severity or message type)

The Common Event Format key names that are generated within the SYSLOG_EVENT column are:

<i>Table 152. Common Event Format key names</i>	
Common Event Format key name	Description
msg	The message text (MESSAGE_TEXT column) from the history log message
reason	Text description of the history log message
sproc	The qualified job name (FROM_JOB column) from the history log message
suser	Current user name (FROM_USER column) from the history log message

Examples

- Return a list of history log messages for all of yesterday and today.

```
SELECT * FROM TABLE(QSYS2.HISTORY_LOG_INFO()) X
```

- Return a list of all history log messages for the last 24 hours.

```
SELECT * FROM TABLE(QSYS2.HISTORY_LOG_INFO(CURRENT_TIMESTAMP - 1 DAY)) X
```

- Return history log information since the last IPL, assuming that the last IPL timestamp is in a global variable named LAST_IPL_TIME.


```
SELECT * FROM TABLE(QSYS2.HISTORY_LOG_INFO(LAST_IPL_TIME, CURRENT_TIMESTAMP)) A
```

- Return syslog information formatted with an RFC3164 header for all history log messages from the start of today forward into the future. When all history log messages have been returned to the caller, the query will pause for 5 minutes (300 seconds) before checking again for messages.

```
SELECT syslog_facility, syslog_severity, syslog_event
FROM TABLE (QSYS2.HISTORY_LOG_INFO(START_TIME => CURRENT DATE,
GENERATE_SYSLOG => 'RFC3164',
EOF_DELAY => 300
) ) AS X;
```


JOBLOG_INFO table function

The JOBLOG_INFO table function returns one row for each message in a job log.



The schema is QSYS2.

job-name A character or graphic string expression that identifies the qualified name of a job. The special value of '*' indicates the current job.

ignore-errors A character or graphic string expression that identifies what to do when an error is encountered.

NO An error is returned. This is the default.

YES A warning is returned. No rows are returned when an error is encountered.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 153. JOBLOG_INFO table function

Column Name	Data Type	Description
ORDINAL_POSITION	INTEGER	A unique number for each row that indicates the time order of messages in the job log. The first (oldest) message in the job log will have a value of 1. Subsequent messages will have a value one greater than the previous message. Since these values are assigned when this catalog is queried, there will be no gaps in values. There is no visibility of messages that have been deleted from the job log.
MESSAGE_ID	VARCHAR(7)	The message ID for this message. Contains the null value if this is an impromptu message or a REQUEST message.
MESSAGE_TYPE	VARCHAR(13)	Type of message. Values are: <ul style="list-style-type: none"> • COMMAND • COMPLETION • DIAGNOSTIC • ESCAPE • INFORMATIONAL • INQUIRY • NOTIFY • REPLY • REQUEST • SCOPE • SENDER

Table 153. JOBLLOG_INFO table function (continued)

Column Name	Data Type	Description
MESSAGE_SUBTYPE	VARCHAR(22)	Subtype of message. Values for NOTIFY or ESCAPE messages are: <ul style="list-style-type: none"> EXCEPTION HANDLED EXCEPTION NOT HANDLED Values for REPLY messages are: <ul style="list-style-type: none"> FROM EXIT PROGRAM FROM SYSTEM REPLY LIST MESSAGE DEFAULT USED NOT VALIDITY CHECKED SYSTEM DEFAULT USED VALIDITY CHECKED Contains the null value for other message types.
SEVERITY	SMALLINT	The severity assigned to the message.
MESSAGE_TIMESTAMP	TIMESTAMP	The timestamp for when the message was issued.
FROM_LIBRARY	VARCHAR(10)	The library containing the program or service program that sent the message.
FROM_PROGRAM	VARCHAR(256)	The program or service program name that sent the message.
FROM_MODULE	VARCHAR(10)	The module that sent the message.
FROM_PROCEDURE	VARCHAR(4096)	The procedure that sent the message.
FROM_INSTRUCTION	VARCHAR(10)	The instruction that sent the message.
TO_LIBRARY	VARCHAR(10)	The library containing the program or service program that received the message
TO_PROGRAM	VARCHAR(10)	The program or service program name that received the message.
TO_MODULE	VARCHAR(10)	The module that received the message.
TO_PROCEDURE	VARCHAR(4096)	The procedure that received the message.
TO_INSTRUCTION	VARCHAR(10)	The instruction that received the message.
FROM_USER	VARCHAR(10)	The userid of the job when the message was sent.
MESSAGE_FILE	VARCHAR(10)	The message file containing the message.
MESSAGE_LIBRARY	VARCHAR(10)	The name of the library containing the message file.
MESSAGE_TOKEN_LENGTH	SMALLINT	The length of the MESSAGE_TOKENS string.
MESSAGE_TOKENS	VARCHAR(2048) FOR BIT DATA	The message token string. If the value is longer than 2048 characters, it will be truncated with no warning.
MESSAGE_TEXT	VARGRAPHIC(1024) CCSID 1200	The first level text of the message including tokens.
MESSAGE_SECOND_LEVEL_TEXT	VARGRAPHIC(4096) CCSID 1200	The second level text of the message including tokens.
MESSAGE_KEY	BINARY(4)	The key assigned to the message. The key is assigned by the command or API that sends the message.
QUALIFIED_JOB_NAME	VARCHAR(28)	The qualified name of the job for this job log.

Examples

- Return joblog information for job 347117/Quser/Qzdasoinit.

```
SELECT * FROM TABLE(QSYS2.JOBLLOG_INFO('347117/Quser/Qzdasoinit')) A
```

- Extract the last command entered by the user.

```
SELECT MESSAGE_TEXT FROM TABLE(QSYS2.JOBLLOG_INFO('817029/QUSER/QPADEV0004')) A
WHERE A.MESSAGE_TYPE = 'REQUEST'
ORDER BY ORDINAL_POSITION DESC
FETCH FIRST 1 ROW ONLY
```

MESSAGE_FILE_DATA view

The MESSAGE_FILE_DATA view returns one row for each message in a message file.

The information is similar to what is returned by the Display Message Description (DSPMSGD) CL command and the Retrieve Message (QMHRVTM) API.

Authorization: Rows will be returned for message files when the caller has *EXECUTE authority to the message file library and *USE authority to the message file.

The following table describes the columns in the view. The system name is MSGF_DATA. The schema is QSYS2.

Table 154. MESSAGE_FILE_DATA view

Column Name	System Column Name	Data Type	Description
MESSAGE_FILE_LIBRARY	MSGF_LIB	VARCHAR(10)	The library containing the message file.
MESSAGE_FILE	MSGF	VARCHAR(10)	The message file.
MESSAGE_ID	MSGID	CHAR(7)	The message identifier.
MESSAGE_TEXT	MSG_TEXT	VARGRAPHIC(132) CCSID 1200	The text of the message.
MESSAGE_SECOND_LEVEL_TEXT	SECLVL	VARGRAPHIC(300) 0) CCSID 1200 Nullable	The second-level message text of the message. Contains the null value if no second-level text is defined.
SEVERITY	SEVERITY	INTEGER	The severity of the message.
MESSAGE_DATA_COUNT	MSGDATACNT	INTEGER	The number of variables defined in MESSAGE_DATA.
MESSAGE_DATA	MSGDATA	VARCHAR(2078) Nullable	A string containing all the message data fields, also referred to as substitution variables, for the message. The format of the string is described following this table. Contains the null value if no message data fields are defined.
LOG_PROBLEM	LOGPRB	VARCHAR(4)	The log problem value for the message. *NO Problems are not logged. *YES Problems are logged.
CREATION_DATE	CRT_DATE	DATE	The date the message was created.
CREATION_LEVEL	CRT_LEVEL	INTEGER	The level number of the message. This is a value from 1 to 99.
MODIFICATION_DATE	MOD_DATE	DATE	The date the message was modified.
MODIFICATION_LEVEL	MOD_LEVEL	INTEGER	The modification level number of the message. This is a value from 1 to 99.
CCSID	CCSID	INTEGER	The CCSID that applies to the stored values of MESSAGE_TEXT and MESSAGE_SECOND_LEVEL_TEXT.
DEFAULT_PROGRAM_LIBRARY	DFT_PGMLIB	VARCHAR(10) Nullable	The library specified for the default program. Can contain the following special values: *CURLIB The current library is used. *LIBL The program is found using the library list. Contains the null value if no default program is defined.
DEFAULT_PROGRAM	DFT_PGM	VARCHAR(10) Nullable	The name of the program called to take default action if this message is sent as an escape message to a program or procedure that is not monitoring for it. Contains the null value if no default program is defined.

Table 154. MESSAGE_FILE_DATA view (continued)

Column Name	System Column Name	Data Type	Description
REPLY_TYPE	REPLY_TYPE	VARCHAR(6) Nullable	The type of valid values that can be made to an inquiry or notify message. *ALPHA Only an alphabetic string is valid. Blanks are not allowed. *CHAR Any character string is valid. If it is a quoted character string, the apostrophes are passed as part of the character string. *DEC Only a decimal number is a valid reply. *NAME Only a simple name is a valid reply. The name does not have to be an object name, but it must start with an alphabetic character; the remaining characters must be alphanumeric. Contains the null value if there is no reply type.
REPLY_LENGTH	REPLY_LEN	INTEGER Nullable	The maximum length of a reply to an inquiry or notify message. Contains the null value if there is no reply type.
REPLY_DECIMAL_POSITIONS	REPLY_DEC	INTEGER Nullable	The maximum number of decimal positions allowed in the message reply. Contains the null value if there is no reply type or if REPLY_TYPE is not *DEC.
DEFAULT_REPLY	DFT_REPLY	VARCHAR(132) Nullable	The default reply for the message. Contains the null value if no default reply is defined.
VALID_REPLY_VALUES_COUNT	REPLY_CNT	INTEGER	The number of entries returned in VALID_REPLY_VALUES.
VALID_REPLY_VALUES	REPLY_VALS	VARCHAR(659) Nullable	The list of valid reply values. Each value is a character string with a length of 32. One blank separates values. Contains the null value if no valid reply values are defined.
VALID_REPLY_LOWER_LIMIT	LOWERLIMIT	VARCHAR(32) Nullable	The lower value limit for a valid reply. Contains the null value if no range values for replies are defined.
VALID_REPLY_UPPER_LIMIT	UPPERLIMIT	VARCHAR(32) Nullable	The upper value limit for a valid reply. Contains the null value if no range values for replies are defined.
VALID_REPLY_RELATIONSHIP_OPERATOR	REL_OP	CHAR(3) Nullable	The relational operator for a relational test entry. *EQ Equal to *GE Greater than or equal to *GT Greater than *LE Less than or equal to *LT Less than *NE Not equal to Contains the null value if no relationship for valid replies is defined.
VALID_REPLY_RELATIONSHIP_VALUE	REL_VALUE	VARCHAR(32) Nullable	The value to be compared to the reply entered for a relational test entry. Contains the null value if no relationship for valid replies is defined.
SPECIAL_REPLY_VALUES_COUNT	SPECIALCNT	INTEGER	The number of entries returned in SPECIAL_REPLY_VALUES.
SPECIAL_REPLY_VALUES	SPECIALVAL	VARCHAR(1319) Nullable	The list of special reply values. Each pair of values consists of the <i>from-value</i> followed by a colon (:) followed by the <i>to-value</i> . One blank separates each pair of values. Contains the null value if no special reply values are defined.
DUMP_LIST_COUNT	DUMP_COUNT	INTEGER	The number of entries returned in DUMP_LIST.

Table 154. MESSAGE_FILE_DATA view (continued)

Column Name	System Column Name	Data Type	Description
DUMP_LIST	DUMP_LIST	VARCHAR(815) Nullable	<p>The list of data items to be dumped when the message is sent as an escape message to a program that is not monitoring for it. The list contains entries separated by a single blank.</p> <p>1-99 The number of the message data field that is to be dumped.</p> <p>*JOB The job information produced by the Display Job (DSPJOB) command is printed.</p> <p>*JOBDMPI The data areas of the job are dumped as specified by the Dump Job (DMPJOB) command.</p> <p>*JOBINT The internal machine data structures related to the machine process in which the job is running are dumped to the machine error log.</p> <p>Contains the null value if there is no dump list for this message.</p>
ALERT_OPTION	ALERTOPT	VARCHAR(9)	<p>Whether an alert is sent for the message.</p> <p>*DEFER An alert is sent after local problem analysis.</p> <p>*IMMED An alert is sent immediately when the message is sent to a message queue that has the allow alerts attribute set to *YES.</p> <p>*NO No alert is sent.</p> <p>*UNATTEND An alert is sent immediately when the system is running in unattended mode.</p>
ALERT_INDEX	ALERTINDEX	INTEGER Nullable	<p>The number of the message data field that is passed with the alert.</p> <p>Contains the null value if no alert is sent or no message data field is passed with the alert.</p>

The MESSAGE_DATA column contains all of the substitution variables formatted as follows. A single blank separates each variable attribute. There is one blank between each variable definition. The order of the attributes is:

1. Variable identifier, such as &1.
2. Data type of the variable.

- *BIN** A binary value formatted in the message as a signed decimal value.
- *CCHAR** A convertible character string.
- *CHAR** A character string formatted without enclosing apostrophes.
- *DEC** A packed decimal number that is formatted in the message as a signed decimal value with a decimal point.
- *DTS** An 8-byte field that contains a system date/time stamp and is formatted in the message as the date followed by one blank and then the time.
- *HEX** A string of bytes formatted as a hexadecimal value.
- *ITV** An 8-byte binary field that contains the time interval (in seconds) for wait time-out conditions.
- *QTDCHAR** A character string formatted with enclosing apostrophes.
- *SPP** A 16-byte space pointer to data in a space object.
- *SYP** A 16-byte system pointer to a system object.
- *UBIN** A binary value formatted in the message as an unsigned decimal value
- *UTC** An 8-byte field that contains a system date/time stamp in Coordinated Universal Time (UTC) and is formatted in the message as the date followed by one blank and then the

time. Before the output formatting the date/time stamp is adjusted from UTC using the time zone specified for the job.

***UTCD** An 8-byte field that contains a system date/time stamp in Coordinated Universal Time (UTC) and is formatted in the message as a date with no time. Before the output formatting the date/time stamp is adjusted from UTC using the time zone specified for the job.

***UTCT** An 8-byte field that contains a system date/time stamp in Coordinated Universal Time (UTC) and is formatted in the message as a time with no date. Before the output formatting the date/time stamp is adjusted from UTC using the time zone specified for the job.

3. Length, if applicable.

- *VARY or the length of the substitution variable.

4. Additional length information, if applicable.

- Fractional digits for a *DEC variable
- 2 or 4 when the length is *VARY

For example, when there are two substitution variables, one a varying character and one a four byte integer, the string might look like this.

```
&1 *CHAR *VARY 2 &2 *BIN 4
```

Example

Find any messages in the APPLIB/APPMSGs message file that contain the word VALUE in upper case, lower case, or mixed case in either the message text or the second level message text.


```
SELECT * FROM QSYS2.MESSAGE_FILE_DATA
WHERE MESSAGE_FILE_LIBRARY = 'APPLIB' AND MESSAGE_FILE = 'APPMSGs' AND
  (UPPER(MESSAGE_TEXT) LIKE '%VALUE%' OR
   UPPER(MESSAGE_SECOND_LEVEL_TEXT) LIKE '%VALUE%');
```

MESSAGE_QUEUE_INFO table function

The MESSAGE_QUEUE_INFO table function returns one row for each message in a specific message queue. It returns information similar to what is returned by the Display Messages (DSPMSG) CL command and the Receive Nonprogram Message (QMHRVCVM) and Open List of Messages (QGYOLMSG) APIs.

This table function does not change the contents of the message queue. The message is kept in the message queue without changing its new or old designation.

Authorization: The user must have *USE authority to the message queue and *EXECUTE authority to the message queue library.



The schema is QSYS2.

message-queue-library A character or graphic string expression that contains the name of the library containing *message-queue-name*. If omitted, the default is QSYS.

message-queue-name A character or graphic string expression that contains the name of a message queue. If omitted, the default is QSYSOPR.

message-filter A character or graphic string expression that indicates the type of messages to be returned. The string can contain one or more of the following values separated by blanks. Specifying 'COMPLETE INQUIRY SENDER' is the same as specifying ALL.

ALL All messages are returned. This is the default.

COMPLETE Only messages that do not require a reply are returned.

INQUIRY Only inquiry messages that require a reply are returned.

SENDER Only copies of the inquiry messages that were sent to other message queues and still require a reply are returned.

severity-filter An integer value indicating the minimum severity of messages to be returned. The value must be from 0 to 99. The default is 0, indicating all messages are to be returned.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 155. MESSAGE_QUEUE_INFO table function

Column Name	Data Type	Description
MESSAGE_ID	VARCHAR(7)	The message ID for this message. Contains the null value if this is an impromptu message or MESSAGE_TYPE is REPLY.
MESSAGE_TYPE	VARCHAR(13)	Type of message. Values are: <ul style="list-style-type: none"> • COMPLETION • DIAGNOSTIC • ESCAPE • INFORMATIONAL • INQUIRY • NOTIFY • REPLY • REQUEST • SENDER

Table 155. MESSAGE_QUEUE_INFO table function (continued)

Column Name	Data Type	Description
MESSAGE_SUBTYPE	VARCHAR(22)	<p>Subtype of message.</p> <p>The values returned for REPLY messages:</p> <ul style="list-style-type: none"> FROM EXIT PROGRAM FROM SYSTEM REPLY LIST MESSAGE DEFAULT USED NOT VALIDITY CHECKED SYSTEM DEFAULT USED VALIDITY CHECKED <p>The value returned for some REQUEST messages:</p> <ul style="list-style-type: none"> WITH PROMPTING <p>Contains the null value for other message types.</p>
MESSAGE_TEXT	VARGRAPHIC(1024) CCSID 1200	<p>The first level text of the message including tokens, or the impromptu message text.</p> <p>Contains the null value if MESSAGE_TYPE is REPLY or if the message file could not be accessed.</p>
SEVERITY	SMALLINT	The severity assigned to the message.
MESSAGE_TIMESTAMP	TIMESTAMP	The timestamp when the message was sent.
MESSAGE_KEY	BINARY(4)	<p>The key assigned to the message.</p> <p>The key is assigned by the command or API that sends the message. For details, see Message Types and Message Keys in the Qmhrvcvm API</p>
ASSOCIATED_MESSAGE_KEY	BINARY(4)	<p>For MESSAGE_TYPE of REPLY, contains the associated inquiry or notify message key.</p> <p>Contains the null value for other message types.</p>
FROM_USER	VARCHAR(10)	The current user of the thread when the message was sent.
FROM_JOB	VARCHAR(28)	The qualified job name of the job that sent the message.
FROM_PROGRAM	VARCHAR(10)	The program that sent the message.
MESSAGE_FILE_LIBRARY	VARCHAR(10)	<p>The name of the library containing the message file.</p> <p>Contains the null value if MESSAGE_ID is null or if the message file could not be accessed.</p>
MESSAGE_FILE_NAME	VARCHAR(10)	<p>The message file containing the message.</p> <p>Contains the null value if MESSAGE_ID is null.</p>
MESSAGE_TOKENS	VARCHAR(4096) FOR BIT DATA	<p>The message token string. If the value is longer than 4096 characters, it will be truncated with no warning.</p> <p>Contains the null value if there are no tokens.</p>
MESSAGE_SECOND_LEVEL_TEXT	VARGRAPHIC(4096) CCSID 1200	<p>The second level text of the message including tokens.</p> <p>Contains the null value if MESSAGE_ID is null or if the message has no second level text or if the message file could not be accessed.</p>

Example

- Return a list of inquiry messages in QSYSOPR that have not been replied to

```
SELECT MESSAGE_ID, MESSAGE_TEXT, MESSAGE_TIMESTAMP
FROM TABLE(QSYS2.MESSAGE_QUEUE_INFO(MESSAGE_FILTER => 'INQUIRY'));
```


MESSAGE_QUEUE_INFO view

The MESSAGE_QUEUE_INFO view returns one row for each message in a message queue. It returns information similar to what is returned by the Display Messages (DSPMSG) CL command and the Receive Nonprogram Message (QMHRVCV) API.

This view does not change the contents of the message queue. The message is kept in the message queue without changing its new or old designation. The view does not utilize the wait time parameter as described in the QMHRVCV API. A wait time of 0 is used.

Authorization: The user must have *USE authority to the message queue and *EXECUTE authority to the message queue library.

The following table describes the columns in the view. The system name is MSGQ_INFO. The schema is QSYS2.

Table 156. MESSAGE_QUEUE_INFO view

Column Name	System Column Name	Data Type	Description
MESSAGE_QUEUE_LIBRARY	MSGQ_LIB	VARCHAR(10)	The name of the library containing the message queue.
MESSAGE_QUEUE_NAME	MSGQ_NAME	VARCHAR(10)	The name of the message queue containing the message.
MESSAGE_ID	MSGID	VARCHAR(7) Nullable	The message ID for this message. Contains the null value if this is an impromptu message or MESSAGE_TYPE is REPLY.
MESSAGE_TYPE	MSG_TYPE	VARCHAR(13)	Type of message. Values are: <ul style="list-style-type: none"> • COMPLETION • DIAGNOSTIC • ESCAPE • INFORMATIONAL • INQUIRY • NOTIFY • REPLY • REQUEST • SENDER
MESSAGE_SUBTYPE	MSG_SUBTYP	VARCHAR(22) Nullable	Subtype of message. The values returned for REPLY messages: <ul style="list-style-type: none"> • FROM EXIT PROGRAM • FROM SYSTEM REPLY LIST • MESSAGE DEFAULT USED • NOT VALIDITY CHECKED • SYSTEM DEFAULT USED • VALIDITY CHECKED The value returned for some REQUEST messages: <ul style="list-style-type: none"> • WITH PROMPTING Contains the null value for other message types.
MESSAGE_TEXT	MSG_TEXT	VARGRAPHIC(1024) CCSID 1200 Nullable	The first level text of the message including tokens, or the impromptu message text. Contains the null value if the message file could not be accessed.
SEVERITY	SEVERITY	SMALLINT	The severity assigned to the message.
MESSAGE_TIMESTAMP	MSG_TIME	TIMESTAMP	The timestamp when the message was sent.
MESSAGE_KEY	MSG_KEY	BINARY(4)	The key assigned to the message. The key is assigned by the command or API that sends the message. For details, see Message Types and Message Keys in the Qmhrvcv API
ASSOCIATED_MESSAGE_KEY	ASSOC_KEY	BINARY(4) Nullable	For MESSAGE_TYPE of REPLY, contains the associated inquiry or notify message key. Contains the null value for other message types.

Table 156. MESSAGE_QUEUE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
FROM_USER	FROM_USER	VARCHAR(10)	The current user of the thread when the message was sent.
FROM_JOB	FROM_JOB	VARCHAR(28)	The qualified job name of the job that sent the message.
FROM_PROGRAM	FROM_PGM	VARCHAR(10)	The program that sent the message.
MESSAGE_FILE_LIBRARY	MSGF_LIB	VARCHAR(10) Nullable	The name of the library containing the message file. Contains the null value if MESSAGE_ID is null or if the message file could not be accessed.
MESSAGE_FILE_NAME	MSGF_NAME	VARCHAR(10) Nullable	The message file containing the message. Contains the null value if MESSAGE_ID is null.
MESSAGE_TOKENS	MSG_TOKENS	VARCHAR(4096) FOR BIT DATA Nullable	The message token string. If the value is longer than 4096 characters, it will be truncated with no warning. Contains the null value if there are no tokens.
MESSAGE_SECOND_LEVEL_TEXT	MSG_TEXT2	VARGRAPHIC(4096) CCSID 1200 Nullable	The second level text of the message including tokens. Contains the null value if MESSAGE_ID is null or if the message has no second level text or if the message file could not be accessed.

Example

- Examine all inquiry messages and their responses.

```
SELECT a.message_timestamp, a.message_text, a.from_user,
       b.message_timestamp, b.message_text, b.from_user
FROM qsys2.message_queue_info a INNER JOIN qsys2.message_queue_info b
   ON a.message_key = b.associated_message_key
WHERE a.message_type = 'INQUIRY' AND
      b.message_type = 'REPLY'
ORDER BY b.message_timestamp DESC;
```

REPLY_LIST_INFO view

The REPLY_LIST_INFO view contains information about the current job's reply list entries.

The following table describes the columns in the view. The system name is REPLYLIST. The schema is QSYS2.

Table 157. REPLY_LIST_INFO view

Column Name	System Column Name	Data Type	Description
SEQUENCE_NUMBER	SEQNO	SMALLINT	The number that specifies the search order of the entries in the reply list.
MESSAGE_ID	MSGID	VARCHAR(7)	The identifier of the inquiry message for which automatic system action is to be taken. A value of ANY indicates that this reply list entry matches any message identifier. Unless comparison data is specified for this reply list entry, all reply list entries with a sequence number greater than this one are ignored.

Table 157. REPLY_LIST_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MESSAGE_REPLY	REPLY	VARCHAR(32)	<p>When an inquiry message is received with a matching message identifier, this value defines whether an automatic reply to the message is given.</p> <p>DEFAULT The default reply to the inquiry message is sent.</p> <p>REQUIRED The inquiry message requires an explicit reply.</p> <p>character string The character string to be sent as the reply to the inquiry message.</p>
COMPARISON_DATA	COMPDATA	VARGRAPHIC(28) CCSID 1200 Nullable	<p>The character string that is compared with the message data of the inquiry message.</p> <p>Contains the null value if there is no comparison data.</p>
COMPARISON_DATA_OFFSET	OFFSET	SMALLINT Nullable	<p>The position in the message data of the inquiry message at which the comparison with the COMPARISON_DATA starts.</p> <p>Contains the null value if there is no comparison data.</p>
DUMP_JOB	DUMPJOB	VARCHAR(3)	<p>Specifies whether the job that sent the inquiry message is to be dumped.</p> <p>NO The job is not dumped.</p> <p>YES The job is dumped before control returns to the program that is sending the message.</p>

Example

- See the reply list entries for your job

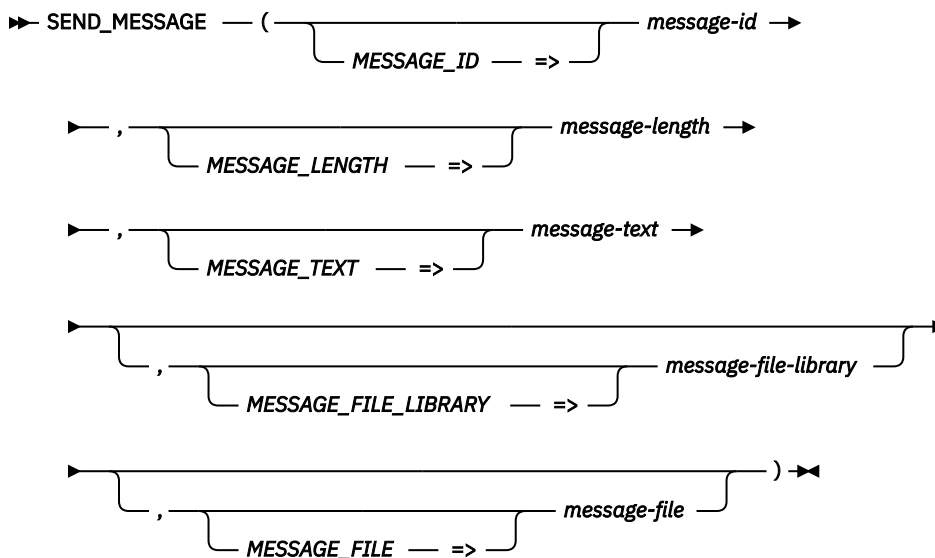
```
SELECT * FROM QSYS2.REPLY_LIST_INFO
```

SEND_MESSAGE procedure

The SEND_MESSAGE procedure sends an informational message to the QSYSOPR message queue.

Authorization: The caller must have:

- *OBJOPR and *ADD authorities to the QSYSOPR message queue,
- *EXECUTE authority to the library containing the message file, and
- *USE authority to the message file.



The schema is QSYS2.

- message-id** A character string expression that identifies the message ID to send to the QSYSOPR message queue.
- message-length** An integer value that defines the length of *message-text*.
- message-text** The message text for the message.
- message-file-library** The library that contains *message-file*. Can contain the special value *LIBL. The default is QSYS.
- message-file** The message file that contains *message-id*. The default is QSQLMSG.

Notes

This procedure is designed to send a predefined message to the QSYSOPR message queue. The message is expected to have one substitution variable defined as *CHAR *VARY 2.

When sending a message from a Query Supervisor exit program, SQL7064 in the QSYS/QSQLMSG message file is defined for this purpose. When sending a message in other situations, it is recommended that a different message should be used.

Example

- Send a message to the QSYSOPR message queue when a critical error is diagnosed.

```
CALL QSYS2.SEND_MESSAGE('APP1234', 31, 'Unexpected error in MYLIB/PGM2.' ,
                        'APPLIB', 'MYMSGF');
```

Performance Services

These services provide information related to system performance.

COLLECTION_SERVICES_INFO view

The COLLECTION_SERVICES_INFO view returns the configuration properties for Collection Services.

The values returned for the columns in the view are closely related to the values returned by the Configure Perf Collection (CFGPPRCOL) CL command and the Retrieve Collection Services Attributes (QypsRtvColSrvAttributes) API.

Authorization: The caller must have *USE authority on the QSYS/QYPSCOLL service program or be authorized through the QPMCCFCN authorization list.

The following table describes the columns in the view. The system name is CS_INFO. The schema is QSYS2.

Table 158. COLLECTION_SERVICES_INFO view

Column Name	System Column Name	Data Type	Description
ACTIVE_COLLECTION_LIBRARY	CURCOL_LIB	VARCHAR(10) Nullable	The name of the library where the currently active management collection object is stored. Contains the null value if there is no active collection.
ACTIVE_COLLECTION_NAME	CURCOL	VARCHAR(10) Nullable	The name of the current management collection object. Contains the null value if there is no active collection.
ACTIVE_COLLECTION_PROFILE	CURCOL_PRF	VARCHAR(10) Nullable	The name of the collection profile being used by the active collection. The collection profile defines which categories of data to collect. *CUSTOM A profile where both the categories to collect and category interval times may be customized. *ENHCPCPLN Enhanced Capacity Planning. Same as *STANDARDP plus the PEX Data - Processor Efficiency category. *MINIMUM The minimum set of categories required to support basic performance reporting functions. *STANDARD All categories that are typically used for performance reporting except for communications protocol related data. *STANDARDP Same as *STANDARD but communications protocol categories are included. Contains the null value if there is no active collection.
ACTIVE_COLLECTION_START_TIME	START_TIME	TIMESTAMP Nullable	The UTC timestamp when the current (active) collection was started. Contains the null value if there is no active collection.
COLLECTION_LIBRARY	LIB	VARCHAR(10)	The name of the library where performance data is stored.
COLLECTION_PROFILE	DFTCOLPRF	VARCHAR(10)	The name of the configured collection profile. The collection profile defines which categories of data to collect. *CUSTOM A profile where both the categories to collect and category interval times may be customized. *ENHCPCPLN Enhanced Capacity Planning. Same as *STANDARDP plus the PEX Data - Processor Efficiency category. *MINIMUM The minimum set of categories required to support basic performance reporting functions. *STANDARD All categories that are typically used for performance reporting except for communications protocol related data. *STANDARDP Same as *STANDARD but communications protocol categories are included.
DEFAULT_COLLECTION_INTERVAL	INTERVAL	INTEGER Nullable	The default interval, in seconds, used when collecting data for a category. Values are: 15, 30, 60, 300, 900, 1800, or 3600 seconds. Contains the null value if not collecting on an interval.
MGTCOL_RETENTION_PERIOD	RETPERIOD	INTEGER Nullable	The management collection (*MGTCOL) object retention period, in hours. *MGTCOL objects in the configured collection library that are older than the retention period are automatically deleted when Collection Services is started or cycled. Contains the null value for permanent retention.
CYCLE_TIME	CYCTIME	INTEGER	The number of minutes past midnight when the first cycle is to occur. The maximum value is 1439 minutes (which is one minute less than 24 hours).
CYCLE_INTERVAL	CYCITV	INTEGER	The number of hours between cycles. The cycle time can range from a minimum value of one hour to a maximum value of 24 hours.

Table 158. COLLECTION_SERVICES_INFO view (continued)

Column Name	System Column Name	Data Type	Description
CREATE_STANDARD_DB_FILES	CRTDBF	VARCHAR(3)	<p>Whether the standard database file collection is created by Collection Services while performance data is being collected.</p> <p>NO Collection Services will not create the standard database file collection.</p> <p>YES Collection Services will create the standard database file collection. A batch job named CRTPFRTA is submitted to process the data in the current management collection object as it is collected. The collection name (begins with a "Q") is the same as the name of the *MGTCOL the data was exported from.</p>
RETENTION_DAYS	STDDTARET	INTEGER Nullable	<p>This retention period, in days, used to determine how long standard database file collections should be retained on the system. When Collection Services is started or cycled, the Collection Services server job (QYSPFRCOL) will automatically delete standard database file collections in the configured collection library that are older than the current retention period.</p> <p>Contains the null value for permanent retention.</p>
CREATE_SUMMARY_DB_FILES	CRTPFRSUM	VARCHAR(3)	<p>Whether additional logical files are created as supported by the CRTPFRSUM command.</p> <p>NO Summary file data is not generated.</p> <p>YES Summary file data is generated.</p>
CREATE_SYSTEM_MONITOR_DB_FILES	CRTSYSMON	VARCHAR(3)	<p>Whether additional system monitor database files should be created by Collection Services while performance data is being collected.</p> <p>NO Collection Services will not create the system monitor database files for the standard database file collection.</p> <p>YES Collection Services will create the system monitor database files for the standard database file collection. The batch job named CRTPFRTA will populate the system monitor database files using the data in the current management collection object as it is collected.</p>
ENABLE_SYSTEM_MONITORING	ENBSYSMON	VARCHAR(3)	<p>Whether Collection Services is configured to collect and produce data for system monitors.</p> <p>NO System monitor support is not enabled.</p> <p>YES System monitor support is enabled. Collection Services will collect certain data categories more frequently than the configured default collection interval. A batch job named CRTPFRTA2 is submitted to produce a system monitor database file collection with the same name as the management collection object except it begins with an 'R'.</p>
SYSTEM_MONITOR_DB_FILE_RETENTION	SYSMONRET	INTEGER Nullable	<p>This retention period, in days, used to determine how long system monitor database file collections should be retained on the system. When Collection Services is started or cycled, the Collection Services server job (QYSPFRCOL) will automatically delete system monitor database file collections in the configured collection library that are older than the retention period.</p> <p>Contains the null value for permanent retention.</p>
CATEGORY_LIST_COUNT	CGY_COUNT	INTEGER	The number of entries in CATEGORY_LIST.
CATEGORY_LIST	SYSMONCGY	VARCHAR(2000) CCSID 1208 Nullable	<p>A list of collection categories that are included in the system monitor collection. Each category can have an independent collection interval. This list is returned as an array within a JSON object.</p> <p>Each entry in the JSON array contains two JSON objects:</p> <ul style="list-style-type: none"> An object with a name of "category" and a value of the category name An object with a name of "interval" and a value of the category's collection interval, in seconds <p>Contains the null value if CATEGORY_LIST_COUNT is 0.</p>
EXCLUDED_LINE_COUNT	EXC_COUNT	INTEGER	The number of entries in EXCLUDED_LINE_LIST.

Table 158. COLLECTION_SERVICES_INFO view (continued)

Column Name	System Column Name	Data Type	Description
EXCLUDED_LINE_LIST	EXC_LIST	VARCHAR(598) Nullable	A list of communication line names that are excluded from the calculation of communication line protocol metrics. Each entry is ten characters long with a comma and space separating entries. Contains the null value if EXCLUDED_LINE_COUNT is 0.
CREATE_HISTORICAL_DATA	CRTPFHST	VARCHAR(3)	Whether historical data will be created when Collection Services is cycled. NO Collection Services will not create historical data. YES Collection Services will create historical data by processing management collection objects that exist in the configured collection library. A batch job named QYSPFRHST is submitted by Collection Services at cycle time.
HISTORICAL_DATA_INTERVAL	HSTIVT	INTEGER	The time interval, in seconds, used to create historical data.
HISTORICAL_SUMMARY_DATA_RETENTION	HSTSUMRET	INTEGER	The historical summary data retention period, in months, that determines how long Collection Services historical summary data is to exist. Historical summary data is stored in files beginning with QAPMHMxxxx. Historical summary data older than the retention period is deleted.
CREATE_HISTORICAL_DETAILED_DATA	CRTHSTDTL	VARCHAR(3)	Whether historical detailed data will be created when Collection Services is cycled. Creating historical detailed data will allow detailed data for the top contributors of various metrics to be stored in the historical collection to be used as drill-down data. The number of top contributors is determined by the value of the historical detailed data filter. NO Collection Services will not create historical detailed data. YES Collection Services will create historical detailed data when processing management collection objects that exist in the configured collection library. A batch job named QYSPFRHST is submitted by Collection Services at cycle time.
HISTORICAL_DETAILED_DATA_RETENTION	HSTDTLRET	INTEGER	The historical detailed data retention period, in days, that determines how long Collection Services historical detailed data is to exist. Historical detailed data is stored in files beginning with QAPMHDxxxx. Historical detailed data older than the retention period is deleted.
HISTORICAL_DETAILED_DATA_FILTER	HSTFILTER	INTEGER Nullable	The number of top contributors for each metric to be stored in the historical data collection to be used as detailed drill-down data. Contains the null value if all detailed data is kept.

Example

- Return the Collection Services configuration properties.

```
SELECT * FROM QSYS2.COLLECTION_SERVICES_INFO;
```

PowerHA Services

These table functions and views provide information about PowerHA®.

[PowerHA SQL Services](#)


Product Services

These services provide information about licensed products.

LICENSE_EXPIRATION_CHECK procedure

The LICENSE_EXPIRATION_CHECK procedure sends a message to the QSYSOPR message queue for every license that corresponds to an installed product that has already expired or is set to expire within the specified number of days.

►► LICENSE_EXPIRATION_CHECK — (→



The schema is SYSTOOLS.

expiration-interval An integer value that indicates the number of days to use as the threshold for checking license information. If not specified, 30 will be used.

Authorization: None required.

Example

- Send an informational message to the system operator message queue, QSYS/QSYSOPR, for every installed product with a license that will expire in the next 10 days.

```
CALL SYSTOOLS.LICENSE_EXPIRATION_CHECK(10);
```

LICENSE_INFO view

The LICENSE_INFO view contains information about all products or features that contain license information.

The values returned for the columns in the view are similar to the values returned by the Work with License Information (WRKLICINF) CL command or [Retrieve License Information \(QLZARTV\) API](#). Refer to the API for more detailed information.

Authorization: None required.

The following table describes the columns in the view. The system name is LIC_INFO. The schema is QSYS2.

Table 159. LICENSE_INFO view

Column Name	System Column Name	Data Type	Description
PRODUCT_ID	LICPGM	VARCHAR(7)	The identifier of the product.
LICENSE_TERM	LIC_TERM	VARCHAR(6)	The license term indicates whether the authorized usage limit for a product exists until the next version, next release, or next modification level of the product. <ul style="list-style-type: none"> • Vx or vv for products licensed by version. • VxRy or vvrr for products licensed by release. • VxRyMz or vvrrmm for products licensed by modification.
RELEASE_LEVEL	RLS_LVL	VARCHAR(6)	The version, release, and modification level of the product in either VxRyMz or vvrrmm format.
FEATURE_ID	FEATURE	VARCHAR(4)	The feature number of the product.
INSTALLED	INSTALLED	VARCHAR(3)	Indicates whether this feature number of the product is installed. <p>NO The feature is not installed.</p> <p>YES The feature is installed.</p>
PROCESSOR_GROUP	PROC_GROUP	VARCHAR(3)	The processor group of this system.

Table 159. LICENSE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
PRODUCT_TEXT	LABEL	VARGRAPHIC(50) CCSID 1200 Nullable	The description of the product or feature. Contains null if there is no description text.
USAGE_LIMIT	USG_LIMIT	INTEGER Nullable	The usage limit of the product or feature that contains license information. Values are 0-999999 indicating the number of users allowed to access the product. Contains null if there is no usage limit.
USAGE_LIMIT_UPDATED	USG_UPDATE	TIMESTAMP(0) Nullable	The timestamp when the usage limit was last updated. Contains null if the usage limit has never been updated.
USAGE_TYPE	USAGE_TYPE	VARCHAR(11)	The usage type of the license. *CONCURRENT The usage type is concurrent, meaning the usage limit is for the number of uses held by unique jobs using the product or feature at the same time. *REGISTERED The usage type is registered, meaning the usage limit is for the number of uses held by license users registered to use the product or feature. *PROCESSOR The usage type is processor, meaning the usage limit is for the number of processors on system partitions where this product or feature is in use.
USAGE_COUNT	USG_COUNT	DECIMAL(8,2)	The current usage count for the product or feature. Valid values are 0 through 999999. If the product is using processor usage type, the usage count value is rounded up to the next whole number.
GLOBAL_COUNT	GLOB_COUNT	DECIMAL(8,2)	The number of jobs currently using this product or feature across all system partitions.
LICENSED_USER_COUNT	LIC_COUNT	INTEGER	The number of current license users.
THRESHOLD	THRESHOLD	DECIMAL(10,2) Nullable	The usage limit threshold for this product or feature. Contains null if there is no usage limit threshold.
PEAK_USAGE	PEAK_USAGE	DECIMAL(10,2)	For concurrent usage, the maximum number of uses held by license users of the product or feature at one time. For registered usage, the maximum number of uses that have been registered through license users for the product or feature. For processor usage, the maximum number of processors configured for this system partition while this product or feature was in use.
LAST_PEAK	LAST_PEAK	TIMESTAMP(0) Nullable	The timestamp when the peak usage of the product or feature last occurred since the peak usage was reset to zero. Contains null if the product has not been used since the peak usage was reset to zero.

Table 159. LICENSE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
COMPLIANCE_TYPE	COMP_TYPE	VARCHAR(10)	<p>The compliance type of the program determines the action taken when the usage limit of the product or feature is exceeded.</p> <p>*OPRACTION License requests are denied and failure messages are sent.</p> <p>*WARNING A warning message is sent.</p> <p>*KEYED Requests for licenses over the usage limit are allowed for the number of days in the product's grace period. Once the grace period ends, the license users holding uses over the usage limit will be released and no requests for uses over the limit will be granted until a new license key is received from the software provider. The expiration date is the date the license will expire. After the expiration date is reached, the default usage limit is in effect. When a request for a license is received after the usage limit has been reached, the system sends a warning message to the system operator message queue and to any additional message queues defined for the product.</p>
LOG_VIOLATION	LOG	VARCHAR(3)	<p>Specifies whether or not requests exceeding the usage limit are logged in the QUSRSYS/QLZALOG journal.</p> <p>NO The requests for a license when the usage count is greater than or equal to the usage limit will not be logged.</p> <p>YES The requests for a license when the usage count is greater than or equal to the usage limit will be logged.</p>
LICENSE_EXPIRATION	EXPIR_DATE	DATE Nullable	<p>The date the license will expire. After the expiration date is reached, the usage limit is reset to the default usage limit.</p> <p>Contains null if the license has no expiration date.</p>
GRACE_PERIOD	GRACE_PRD	INTEGER	<p>The number of days a user has to obtain a new license key after a product or feature exceeds its usage limit.</p>
GRACE_END	GRACE_END	DATE Nullable	<p>The date the grace period expires. When a request for license uses exceeds the usage limit for a product or feature, the date the grace period expires is determined by adding the number of days in the grace period to the current date.</p> <p>Contains null if there is no grace period or the grace period has expired.</p>
VENDOR_DATA	VENDOR	VARCHAR(8)	<p>Information the vendor defined when the key was added using the Add License Key Information (ADDLICKEY) command.</p>
MESSAGE_QUEUE1	MESSAGE_1	VARCHAR(10) Nullable	<p>The name of the first message queue to which messages will be sent.</p> <p>Each of these message queues, in addition to the system operator message queue, will be sent a messages if one of the following occurs:</p> <ul style="list-style-type: none"> • The usage count threshold is met. • A license request is made, and the usage count is equal to or greater than the usage limit. • The usage limit is changed. <p>The messages sent include:</p> <ul style="list-style-type: none"> • CPI9E10 - License usage limit changed for product &1. • CPI9E19 - Usage limit threshold exceeded. • CPI9E75 - Grace period will expire on &3. • CPI9E76 - Expiration date will be reached on &3. <p>Contains null if there is no first message queue.</p>
MESSAGE_QUEUE_LIBRARY1	LIBRARY_1	VARCHAR(10) Nullable	<p>The library containing the first message queue.</p> <p>Contains null if there is no first message queue.</p>

Table 159. LICENSE_INFO view (continued)

Column Name	System Column Name	Data Type	Description
MESSAGE_QUEUE2	MESSAGE_2	VARCHAR(10) Nullable	The name of the second message queue to which messages will be sent. Contains null if there is no second message queue.
MESSAGE_QUEUE_LIBRARY2	LIBRARY_2	VARCHAR(10) Nullable	The library containing the second message queue. Contains null if there is no second message queue.
MESSAGE_QUEUE3	MESSAGE_3	VARCHAR(10) Nullable	The name of the third message queue to which messages will be sent. Contains null if there is no third message queue.
MESSAGE_QUEUE_LIBRARY3	LIBRARY_3	VARCHAR(10) Nullable	The library containing the third message queue. Contains null if there is no third message queue.
MESSAGE_QUEUE4	MESSAGE_4	VARCHAR(10) Nullable	The name of the fourth message queue to which messages will be sent. Contains null if there is no fourth message queue.
MESSAGE_QUEUE_LIBRARY4	LIBRARY_4	VARCHAR(10) Nullable	The library containing the fourth message queue. Contains null if there is no fourth message queue.
MESSAGE_QUEUE5	MESSAGE_5	VARCHAR(10) Nullable	The name of the fifth message queue to which messages will be sent. Contains null if there is no fifth message queue.
MESSAGE_QUEUE_LIBRARY5	LIBRARY_5	VARCHAR(10) Nullable	The library containing the fifth message queue. Contains null if there is no fifth message queue.

Example

Return information about all licensed products and features that will expire within the next 2 weeks.

```
SELECT * FROM QSYS2.LICENSE_INFO
WHERE LICENSE_EXPIRATION <= CURRENT DATE + 14 DAYS;
```

SOFTWARE_PRODUCT_INFO view

The SOFTWARE_PRODUCT_INFO view returns information about software products.

The values returned for the columns in the view are closely related to the values returned by the Display Software Resources (DSPSFWRSC) command and the Retrieve Product Information (QSZRTVPR) API.

Authorization: None required.

The following table describes the columns in the view. The system name is SFW_PROD. The schema is QSYS2.

Table 160. SOFTWARE_PRODUCT_INFO view

Column Name	System Column Name	Data Type	Description
PRODUCT_ID	PRODUCT_ID	VARCHAR(7)	The product ID.
PRODUCT_OPTION	PROD_OPT	VARCHAR(5)	The product option of the product. Can contain the following special value: *BASE The base part of the product.
LOAD_ID	LOAD_ID	VARCHAR(4)	The load ID of the product load for which information was returned.
LOAD_TYPE	LOAD_TYPE	VARCHAR(8)	The type of load. CODE The load is a code load. LANGUAGE The load is a language load.

Table 160. SOFTWARE_PRODUCT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
RELEASE_LEVEL	RELEASE	VARCHAR(10)	The release level of the product selected.
INSTALLED	INSTALLED	VARCHAR(3)	<p>Whether the code load for this product option is installed. A load is installed if a product load (*PRDLOD) object is loaded on the system by the Restore Licensed Program (RSTLICPGM) command.</p> <p>NO The code load for this product option is not installed.</p> <p>YES The code load for this product option is installed.</p>
SYMBOLIC_LOAD_STATE	SYM_STATE	VARCHAR(9)	<p>The symbolic state of the load for which information was returned. This value, in conjunction with the LOAD_ERROR column, can be used to determine if the load is installed correctly.</p> <p>CREATED The product load object for this load exists.</p> <p>DAMAGED If this is for an option other than the base option or for a language load for the base option, the product load object has been damaged. If this is for the code load for the base option, one of the following happened:</p> <ul style="list-style-type: none"> • The product definition for this product ID and release level has been damaged. • The product load object has been damaged. <p>DEFINED The load is defined. The product load object for this load does not exist.</p> <p>INSTALLED The product load (*PRDLOD) object for this load was loaded onto the system by the RSTLICPGM command.</p> <p>LOADED Indicates one of the following:</p> <ul style="list-style-type: none"> • A restore licensed program function is in progress. • A delete licensed program function is in progress. <p>PACKAGED The product load object for this load has been packaged with the PKGPRDOPT command.</p>
LOAD_ERROR	LOAD_ERROR	VARCHAR(3)	<p>Whether there is a known error for this load.</p> <p>NO No error was found the last time that the state of this load was checked or updated.</p> <p>YES An error was found the last time that the state of this load was checked or updated.</p>

Table 160. SOFTWARE_PRODUCT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOAD_STATE	LOAD_STATE	CHAR(2)	<p>The state of the load for which information was returned.</p> <p>10 The load is defined.</p> <p>20 The product load object for this load exists.</p> <p>30 The product load object for this load has been packaged with the PKGPRDOPT command or the QSZPKGPO API.</p> <p>32 The product load object for this load has been packaged with the PKGPRDOPT command or the QSZPKGPO API. Either a development library or folder was renamed, but the product does not allow dynamic naming, or the product definition or product load for a packaged load was renamed or moved to another library.</p> <p>33 The product load object for this load has been packaged with the PKGPRDOPT command or the QSZPKGPO API. However, an object was found to be damaged the last time that the CHKPRDOPT command or SAVLICPGM command was used for this load.</p> <p>34 The product load object for this load has been packaged with the PKGPRDOPT command or the QSZPKGPO API. Either an attempt was made to delete the product load using the delete licensed program function and the function failed, or a packaged object was missing the last time that the CHKPRDOPT command or SAVLICPGM command was used for this load.</p> <p>35 A RSTLICPGM command is in progress. The product being replaced had been packaged, but not installed.</p> <p>38 A DLTLICPGM command is in progress. The product being deleted had been packaged, but not installed.</p> <p>3E A RSTLICPGM command did not complete successfully. A preoperation exit program failed. The product being replaced had been packaged, but not installed.</p> <p>3F A RSTLICPGM command failed. A preoperation exit program did not fail. The product being replaced had been packaged, but not installed.</p> <p>50 A RSTLICPGM command is in progress. The product being replaced had been installed.</p> <p>53 A DLTLICPGM command is in progress. The product being deleted had been installed.</p> <p>59 This product is an IBM-supplied product, and it is not compatible with the currently installed release level of the operating system.</p>

Table 160. SOFTWARE_PRODUCT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
LOAD_STATE (continued)			<p>60 The product load object for this load was loaded onto the system by the RSTLICPGM command.</p> <p>61 The product load object for this load was loaded onto the system by the RSTLICPGM command, but a postoperation exit program failed.</p> <p>62 An installed library or folder was renamed, but the product does not allow dynamic naming.</p> <p>63 The product load object for this load was installed by the RSTLICPGM command, but an object is damaged.</p> <p>64 The product load object for this load was installed by the RSTLICPGM command, but an object was found to be missing when CHKPRDOPT or SAVLICPGM was used, or an error occurred while DLTICPGM was being used.</p> <p>67 The CHKPRDOPT command was used for this product load, but the postoperation exit program failed or indicated that an error was found.</p> <p>6E A RSTLICPGM command did not complete successfully.</p> <p>6F A RSTLICPGM command failed.</p> <p>90 The product load was installed successfully. If an object was missing or was damaged, and the problem was corrected, using the CHKPRDOPT command sets the state back to 90.</p>
SUPPORTED	SUPPORTED	VARCHAR(3)	<p>Whether this load is currently supported. A load can be supported by using the Work with Supported Products (WRKSPTPRD) command in the System Manager for IBM i licensed program.</p> <p>NO This load is not currently supported.</p> <p>YES This load is currently supported.</p>
COMPATIBLE	COMPATIBLE	VARCHAR(3) Nullable	<p>Indicates whether this IBM product is compatible with the current release.</p> <p>NO The product is not compatible.</p> <p>YES The product is compatible.</p> <p>Contains the null value if this product does not have a compatibility value.</p>
PRODUCT_LIBRARY_COUNT	LIB_COUNT	INTEGER	The number of product libraries included in the PRODUCT_LIBRARIES column.
PRODUCT_LIBRARIES	PROD_LIB	VARCHAR(120) Nullable	<p>The list of product load libraries. Each library name entry is ten characters long with one blank separating entries.</p> <p>Contains the null value if PRODUCT_LIBRARY_COUNT is 0.</p>
PRODUCT_DIRECTORY_COUNT	DIR_COUNT	INTEGER	The number of product directories included in the PRODUCT_DIRECTORIES column.
PRODUCT_DIRECTORIES	PROD_DIR	DBCLOB(5M) CCSID 1200 Nullable	<p>The list of product directories. Directory entries are separated by the string '--' (a blank, two minus signs, and a blank).</p> <p>Contains the null value if PRODUCT_DIRECTORY_COUNT is 0.</p>

Table 160. SOFTWARE_PRODUCT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
TEXT_DESCRIPTION	TEXT	VARCHAR(132) Nullable	Text description for this product option. Contains the null value if there is no text description or the text description is not available.
PRIMARY_LANGUAGE_LOAD_ID	LANG_ID	CHAR(4) Nullable	For code loads, this field contains the primary language of the product option. Contains the null value for language loads and for code loads when no language is installed in the libraries for the code load.
RELEASE_DATE	REL_DATE	DATE Nullable	Indicates the value specified for the release date when the product definition for this product load was created. Contains the null value if no release date was specified.
MINIMUM_TARGET_RELEASE	MIN_TGTRLS	CHAR(6)	The minimum release of the operating system to which the Save Licensed Program (SAVLICPGM) command will allow the product to be saved.
MINIMUM_VRM_BASE	MIN_BASE	CHAR(6) Nullable	The minimum release level that is allowed for the *BASE option that will run with the current level of the option for the product. Can containing the following special value: *MATCH The release of the option matches that of the *BASE. Contains the null value for Licensed Internal Code.
REQUIREMENTS_MET	REQ_MET	INTEGER	The reason why the release requirements between the base and option may or may not be in error. 0 There is not enough information available to determine if the release requirements have been met. This value if always returned for a load type of *LANG. 1 The releases of the *BASE and option meet all requirements. 2 The release of the option is too old compared to the *BASE. 3 The release of the *BASE is too old compared to the option.
MIXED_RELEASES	MIXED_REL	VARCHAR(3)	Product allows mixed releases. NO The *BASE option and other options of this product cannot be at different release levels. YES The *BASE option and other options of this product can be at different release levels.
LEVEL_ID	LEVEL_ID	CHAR(3) Nullable	The level identifier of the product for which information was returned. The format is Lxx. Contains the null value for all products other than the operating system and Licensed Internal Code.
REGISTRATION_TYPE	REG_TYPE	CHAR(2)	The registration type associated with the product. The registration type and registration value together make up the registration ID for the product.
REGISTRATION_VALUE	REG_VALUE	VARCHAR(14)	The registration value associated with the product.

Example

- List any licensed programs that are in an error state.

```
SELECT *
FROM QSYS2.SOFTWARE_PRODUCT_INFO
WHERE LOAD_ERROR = 'YES';
```

PTF Services

These views provide PTF information.

ELECTRONIC_SERVICE_AGENT_INFO view

The ELECTRONIC_SERVICE_AGENT_INFO view returns detailed information about the Electronic Service Agent (ESA) connections.

The Electronic Service Agent must be configured and activated before this view can successfully return results. See [Electronic Service Agent](#) for information about ESA.

The values returned for the columns in the view are closely related to the detail generated by the VFYSRVAGT TYPE(*DETAIL) CL command.

When the ESA_CONNECTION column contains the null value, there is no connectivity. RESULT_BY_IP_ADDRESS and RESULT_BY_HOSTNAME will both return a value of FAILURE, and partial information will be returned in the other columns.

Authorization: The caller must have *ALLOBJ special authority.

The following table describes the columns in the view. The system name is ESA_INFO. The schema is QSYS2.

Table 161. ELECTRONIC_SERVICE_AGENT_INFO view

Column Name	System Column Name	Data Type	Description
ESA_STATUS	STATUS	VARCHAR(9)	The ESA status. ACTIVATED ESA has been activated on this system INACTIVE ESA has not been activated
ESA_CONNECTION	CONNECTION	VARCHAR(18) Nullable	The type of connection between the IBM i and IBM support servers. DIRECT LAN CONNECT A direct LAN connection is configured HTTP PROXY A connection through an HTTP/HTTPS proxy is configured Contains the null value if the type of connection is not available.
PROXY_HOST_OR_IP	PROXY_HOST	VARCHAR(256) Nullable	The proxy hostname or proxy IP address, depending on the configuration. Contains the null value if ESA_CONNECTION is DIRECT LAN CONNECT.
PROXY_PORT	PROXY_PORT	INTEGER Nullable	The proxy port. Contains the null value if ESA_CONNECTIVITY is DIRECT LAN CONNECT.
PROXY_ID	PROXY_ID	VARCHAR(16) Nullable	The proxy user ID. Contains the null value if the HTTP proxy is a non-authenticated HTTP proxy, if the proxy ID is not configured, or if ESA_CONNECTIVITY is DIRECT LAN CONNECT.

Table 161. ELECTRONIC_SERVICE_AGENT_INFO view (continued)

Column Name	System Column Name	Data Type	Description
RESULT_BY_IP_ADDRESS	RESULT_IP	VARCHAR(7)	<p>Connection status using IP address.</p> <p>SUCCESS Returned successful status</p> <p>FAILURE Returned failure status</p> <p>UNKNOWN No response</p> <p>If either RESULT_BY_IP_ADDRESS or RESULT_BY_HOSTNAME has a value of SUCCESS, ESA has a working connection.</p>
RESULT_BY_HOSTNAME	RESULT_HST	VARCHAR(7)	<p>Connection status using hostname.</p> <p>SUCCESS Returned successful status</p> <p>FAILURE Returned failure status</p> <p>UNKNOWN No response</p> <p>If either RESULT_BY_IP_ADDRESS or RESULT_BY_HOSTNAME has a value of SUCCESS, ESA has a working connection.</p>
SERVER_TYPE	SERV_TYPE	VARCHAR(25)	<p>The type of IBM support server.</p> <p>The following values are returned when using traditional path (Non-Edge) support. If using simplified path (Edge), the values are prefixed with EDGE. For example, EDGE BULK DATA.</p> <p>Appended at the end of the value is a blank followed by a number to make the value unique.</p> <p>BULK DATA Handles bulk data requests</p> <p>CONFIGURATION Handles service provider configuration requests</p> <p>FIX REPOSITORY Handles fix requests</p> <p>GATEWAY Distributes requests to the corresponding servers</p> <p>INVENTORY REPORT Handles inventory report requests</p> <p>PROBLEM REPORT Handles problem call home requests</p> <p>PROFILE Handles profile create/edit/delete requests</p> <p>STATUS REPORT Handles status requests</p> <p>UPDATE ORDER Handles PTF download requests</p>
SERVER_HOSTNAME	SERV_HOST	VARCHAR(30)	<p>The hostname of the server. When using Edge, SERVER_HOSTNAME will be ESUPPORT.IBM.COM.</p>
SERVER_IP_ADDRESS	SERV_IP	VARCHAR(45)	<p>The IP address of the server.</p>
SERVER_PORT	SERV_PORT	INTEGER	<p>The port of the server. The default port is 443.</p>

Example

- Return the connection information for the Electronic Service Agent.

```
SELECT * FROM QSYS2.ELECTRONIC_SERVICE_AGENT_INFO;
```

FIRMWARE_CURRENCY view

The FIRMWARE_CURRENCY view implements a live comparison of the firmware fix level installed on the partition against the level available through a feed from the Fix Level Recommendation Tool (FLRT).

When queried, the view uses Display Hardware Resources (DSPHDWRSC) and Display Firmware Status (DSPFMWSTS) commands and the JSON_TABLE and HTTPGETCLOB functions to consume a live JSON feed from the Fix Level Recommendation Tool. When querying this view, the job CCSID cannot be 65535 or the query will fail.

The result of the query shows the firmware fix level installed on the partition and the firmware latest fix levels made available by IBM. It also returns the general information shown by the DSPFMWSTS command.

Authorization: The DSPHDWRSC and DSPFMWSTS CL commands are used to obtain information. The user must be authorized to these commands or the query will fail.

The following table describes the columns in the view. The system name is FWCUR. The schema is SYSTOOLS.

Table 162. FIRMWARE_CURRENCY view

Column name	System column name	Data type	Description
FW_CURRENCY	FW_CRNCY	VARCHAR(28)	A description of the firmware status. Values returned are: INSTALLED LEVEL IS CURRENT The firmware fix level installed matches the most current available from IBM UPDATE AVAILABLE An update is available from IBM UPGRADE AVAILABLE An upgrade is available from IBM UPDATE AND UPGRADE AVAILABLE Both an update and an upgrade are available from IBM
FW_CURRENT_FIXPACK	FW_FIXPACK	VARCHAR(20)	The current fix level on the partition.
FW_RELEASE_DATE	FW_GA	DATE	The general availability date of the firmware on the partition.
FW_MACHINE_TYPE_MODEL	FW_MTM	VARCHAR(20)	The machine type model of the partition.
FW_RECOMMENDED_UPDATE	FW_RUPD	VARCHAR(20) Nullable	The update fix level of the firmware. Contains the null value if no update version is available.
FW_RECOMMENDED_UPGRADE	FW_RUPG	VARCHAR(20) Nullable	The upgrade fix level of the firmware. Contains the null value if no upgrade version is available.
FW_UPDATE_ACCESS_KEY_EXPIRATION	KEY_EXP	DATE Nullable	The expiration date of the Update access key. Server firmware fix packs with a later date are not activated until a valid Update access key expiration date is detected. Contains the null value if the update access key expiration date is not available.
FW_SERVICE_PARTITION	SVC_PART	VARCHAR(3)	Indicates whether the logical partition is operating as a service partition. NO The logical partition is not operating as a service partition. YES The logical partition is operating as a service partition.

Table 162. FIRMWARE_CURRENCY view (continued)

Column name	System column name	Data type	Description
FW_UPDATE_POLICY	UPD_POLICY	VARCHAR(5)	Indicates how changes are made to the server firmware. HMC The server firmware is managed by a Hardware Management Console (HMC). The operating system is not allowed to make changes to the server firmware. OPSYS The server firmware is managed by the operating system using Program Temporary Fixes (PTFs) for the specified server firmware product ID/release (FW_PRODUCT_ID and FW_PRODUCT_RELEASE).
FW_IPL_SOURCE	IPL_SRC	VARCHAR(9)	The copy of the server firmware that was used on the previous server IPL. PERMANENT The last server IPL used the permanent copy of the server firmware. TEMPORARY The last server IPL used the temporary copy of the server firmware.
FW_IPL_REQUIRED	IPL_REQD	VARCHAR(3)	Whether an IPL is required to activate PTFs for the server firmware product. NO PTFs are active. YES PTFs are applied but are not active.
FW_PRODUCT_ID	PRODUCT_ID	VARCHAR(7) Nullable	The IBM i product that matches the level of the server firmware on the system. Managing the server firmware level is performed by applying or removing PTFs for this product. Contains the null value if no product ID exists for the active server firmware.
FW_PRODUCT_RELEASE	RELEASE	VARCHAR(6) Nullable	The release corresponding to FW_PRODUCT_ID. Contains the null value if no product ID exists for the active server firmware.

Notes

- The view requires the system to have access to the internet with the ability to access the FLRT website. The FLRT website is:

```
https://www14.software.ibm.com/support/customer/care/flrt/liteTable?prodKey=fw&format=json
```

- If the FLRT website is relocated, this view can be updated by the user. Use the Insert Generated SQL feature in ACS to extract the source for the SYSTOOLS.FLRT_FW_INFO table function. Update the link and recreate the table function.
- If any of the CL commands or SQL functions used by the view encounter an error, an error will be returned to indicate the failure. The job log can be examined to determine the root cause of the problem.

Example

Examine the firmware fix level information for the partition.

```
SELECT * FROM SYSTOOLS.FIRMWARE_CURRENCY;
```

GROUP_PTF_CURRENCY view

The GROUP_PTF_CURRENCY is a view containing a query which implements a live comparison of the PTF Groups installed on the partition against the service levels listed on the IBM Preventive Service Planning website.

When queried, the view uses the XMLTable() and HTTPGETBLOB() table functions to consume a live XML feed from IBM Preventive Service Planning (PSP). If the partition cannot connect to the PSP website, the PTF_GROUP_CURRENCY column will contain PSP INFORMATION NOT AVAILABLE. When querying this view, the job CCSID cannot be 65535 or the query will fail.

The results of the query show which PTF Groups installed on the partition match the latest level made available by IBM and those which have a more recent version available.

The following table describes the columns in the view. The system name is GRPPTFCUR. The schema is SYSTOOLS.

Table 163. GROUP_PTF_CURRENCY view

Column name	System column name	Data type	Description
PTF_GROUP_CURRENCY	GRP_CRNCY	VARCHAR(46) Nullable	A description of the PTF group's status. Values returned are: INSTALLED LEVEL IS CURRENT Indicates that the PTF Group level installed matches the most current level available from IBM CURRENT AT THE NEXT IPL Indicates that the most current PTF Group level available from IBM is ready to be applied when the next IPL occurs. UPDATE AVAILABLE Indicates that a more recent PTF Group level is available from IBM PSP INFORMATION NOT AVAILABLE Indicates that the query is unable to connect to the external IBM PSP PTF Group level feed. Returns the null value if not applicable to this PTF group.
PTF_GROUP_ID	GRP_ID	CHAR(7) Nullable	The name of the PTF group.
PTF_GROUP_TITLE	GRP_TITLE	VARCHAR(1000) Nullable	The descriptive name of the PTF group.
PTF_GROUP_LEVEL_INSTALLED	GRP_LVL	INTEGER Nullable	The most recent level of this PTF Group installed on the partition.
PTF_GROUP_LEVEL_AVAILABLE	GRP_IBMLVL	INTEGER Nullable	The PTF Group level which is available from IBM PSP.
LAST_UPDATED_BY_IBM	GRP_UPDATE	DATE Nullable	The date that IBM made the latest PTF Group level available.
PTF_GROUP_RELEASE	GRP_RLS	VARCHAR(6) Nullable	The release level of the PTF Group. For example, 'R730' indicates IBM i 7.3 release level.
PTF_GROUP_STATUS_ON_SYSTEM	GRP_SYSSTS	VARCHAR(20) Nullable	This column will always contain the value 'INSTALLED'.
PTF_GROUP_LAST_UPDATED_BY_IBM	GRP_LSTUPD	CHAR(10) Nullable	The date that IBM made the latest PTF Group level available. This is the character form of the date formatted as MM/DD/YYYY.

Notes

- The PSP website is:

```
https://www.ibm.com/support/pages/sites/default/files/inline-files/xmldoc.xml
```

To determine the IP address for your geography, ping `www.ibm.com`.

- If the PSP website is relocated, this view can be updated by the user. Use the Insert Generated SQL feature in ACS to extract the source for the `SYSTOOLS.GROUP_PTF_CURRENCY` view. Update the link and recreate the view.
- The `PTF_GROUP_STATUS_ON_SYSTEM` column is included in this view to demonstrate that it would be possible to create your own version of this query or view which includes information about PTF Groups that are loaded, but not installed.

Example

Compare the PTF Group service level detail, ordering the results from furthest behind to current.

```
SELECT * FROM SYSTOOLS.GROUP_PTF_CURRENCY
ORDER BY PTF_GROUP_LEVEL_AVAILABLE - PTF_GROUP_LEVEL_INSTALLED DESC
```

Related reference

SYSTOOLS

`SYSTOOLS` is a set of Db2 for IBM i supplied examples and tools.

GROUP_PTF_DETAILS view

The `GROUP_PTF_DETAILS` is a view containing a query which implements a live comparison of the PTFs within PTF Groups installed on the partition against the service levels listed on the IBM Preventive Service Planning website.

When queried, the view uses the `XMLTable()` and `HTTPGETBLOB()` table functions to consume a live XML feed from IBM Preventive Service Planning (PSP). If the partition cannot connect to the PSP website, the query will fail with an SQL4302. When querying this view, the job CCSID cannot be 65535 or the query will fail.

The results of the query show which PTFs from all PTF Groups installed on the partition match the latest level made available by IBM and those which have a more recent version available.

The following table describes the columns in the view. The system name is `GRPPTFDTL`. The schema is `SYSTOOLS`.

Table 164. `GROUP_PTF_DETAILS` view

Column name	System column name	Data type	Description
<code>PTF_GROUP_DESCRIPTION</code>	<code>GRPDESC</code>	<code>VARCHAR(100)</code>	Description of the PTF group.
<code>PTF_GROUP_NAME</code>	<code>GRPNAME</code>	<code>CHAR(7)</code>	Name of the PTF group.
<code>PTF_STATUS</code>	<code>PTF_STATUS</code>	<code>VARCHAR(11)</code>	Status of the PTF. PTF APPLIED The PTF has been loaded and applied. PTF LOADED The PTF has been loaded but not applied. PTF MISSING The PTF does not exist on this partition.
<code>PTF_PRODUCT_ID</code>	<code>LICPGM</code>	<code>VARCHAR(7)</code>	The licensed program for this PTF.
<code>PTF_IDENTIFIER</code>	<code>PTFID</code>	<code>VARCHAR(7)</code>	The identifier of the PTF.
<code>APAR_NAME</code>	<code>APAR_NAME</code>	<code>VARCHAR(7)</code>	The APAR name associated with the PTF.
<code>INCLUDED_IN_GROUP</code>	<code>INCLUDED</code>	<code>DATE</code>	The date that this PTF was first made available in a group PTF.

Table 164. GROUP_PTF_DETAILS view (continued)

Column name	System column name	Data type	Description
PTF_CUM_PACKAGE	PTF_CUMPKG	VARCHAR(8)	The identifier of the cumulative PTF package containing this PTF.
PTF_PRODUCT_DESCRIPTION	PRODESC	VARCHAR(132) Nullable	Product description.
PTF_RELEASE_LEVEL	PTFRLS	VARCHAR(6) Nullable	The release level of the PTF.
PTF_PRODUCT_LOAD	PRODLOAD	VARCHAR(4) Nullable	The load ID of the product load for the PTF.
PTF_LOADED_STATUS	LOADSTAT	VARCHAR(19) Nullable	The current loaded status of the PTF. NOT LOADED The PTF has never been loaded. LOADED The PTF has been loaded. APPLIED The PTF has been applied. PERMANENTLY APPLIED The PTF has been applied permanently. PERMANENTLY REMOVED The PTF has been permanently removed. DAMAGED The PTF is damaged. An error occurred while applying the PTF. It needs to be reloaded and applied. SUPERCEDED The PTF is superseded. A PTF will have a status of superseded when one of the following situations occurs: <ul style="list-style-type: none"> • Another PTF with a more recent correction for the problem has been loaded on the system. The PTF ID that has been loaded can be found in the PTF_SUPERCEDED_BY_PTF column. • The PTF save file for another PTF with a more recent correction for the problem has been logged into *SERVICE on the system.
PTF_SAVE_FILE	SAVF	VARCHAR(3) Nullable	Indicates whether a save file exists for the PTF. NO The PTF has no save file. YES The PTF has a save file.
PTF_COVER_LETTER	COVER	VARCHAR(3) Nullable	Indicates whether a cover letter exists for the PTF. NO The PTF has no cover letter. YES The PTF has a cover letter.
PTF_ON_ORDER	ONORD	VARCHAR(3) Nullable	Indicates whether the PTF has been ordered. NO The PTF has not been ordered or has already been received. YES The PTF has been ordered.

Table 164. GROUP_PTF_DETAILS view (continued)

Column name	System column name	Data type	Description
PTF_IPL_ACTION	IPLACT	VARCHAR(19) Nullable	<p>The action to be taken on this PTF during the next unattended IPL.</p> <p>NONE No action occurs at the next IPL.</p> <p>TEMPORARILY APPLIED The PTF is temporarily applied at the next IPL.</p> <p>TEMPORARILY REMOVED The PTF is temporarily removed at the next IPL.</p> <p>PERMANENTLY APPLIED The PTF is permanently applied at the next IPL.</p> <p>PERMANENTLY REMOVED The PTF is permanently removed at the next IPL.</p>
PTF_ACTION_PENDING	ACTPEND	VARCHAR(3) Nullable	<p>Indicates whether a required action has yet to be performed to make this PTF active.</p> <p>NO No required actions are pending for this PTF.</p> <p>YES A required action needs to occur for this PTF to be active. Check the Activation Instructions section of the cover letter to determine what the action is. If the PTF_ACTION_REQUIRED column is set to IPL and the activation instructions have been performed, then the PTF is active. However, this column will not be updated until the next IPL.</p>
PTF_ACTION_REQUIRED	ACTREQ	VARCHAR(12) Nullable	<p>Indicates whether an action is required to make this PTF active when it is applied. See the cover letter to determine what action needs to be taken.</p> <p>NONE No activation instructions are needed for this PTF.</p> <p>EXIT PROGRAM This PTF was shipped with activation instructions in the cover letter. This value is returned for all PTFs that have an exit program to update the status of the PTF after the activation instructions have been performed.</p> <p>IPL This PTF was shipped with activation instructions in the cover letter. No exit program exists to verify the activation instructions were performed.</p>
PTF_IPL_REQUIRED	IPLREQ	VARCHAR(9) Nullable	<p>Indicates whether an IPL is required to apply this PTF.</p> <p>DELAYED The PTF is delayed. The PTF must be applied during an IPL.</p> <p>IMMEDIATE The PTF is immediate. No IPL is needed to apply the PTF.</p> <p>UNKNOWN The type of the PTF is not known.</p>
PTF_IS_RELEASED	RELEASED	VARCHAR(3) Nullable	<p>Indicates whether the PTF save file is available for distribution to another system. This is set to YES only when the System Manager for IBM i licensed program is on the system and the product is supported. The PTF_SAVE_FILE column must have a value of YES before using the value in this column.</p> <p>NO The PTF save file cannot be distributed.</p> <p>YES The PTF save file is released and can be distributed to another system.</p>
PTF_MINIMUM_LEVEL	MINLVL	VARCHAR(2) Nullable	<p>The indicator of the lowest level of the product to which this PTF can be applied. The level can be AA to 99.</p> <p>Contains the null value if the product does not have a level.</p>

Table 164. GROUP_PTF_DETAILS view (continued)

Column name	System column name	Data type	Description
PTF_MAXIMUM_LEVEL	MAXLVL	VARCHAR(2) Nullable	The indicator of the highest level of the product to which this PTF can be applied. The level can be AA to 99. Contains the null value if the product does not have a level.
PTF_STATUS_TIMESTAMP	STATTIME	TIMESTAMP Nullable	The date and time that the PTF status was last changed. Contains the null value when the status date and time is not available.
PTF_SUPERCEDED_BY_PTF	SUPERCEDE	VARCHAR(7) Nullable	The identifier of the PTF that has replaced this PTF. This field will be blank when the PTF is not superseded or when the superseding PTF has not been loaded on the system.
PTF_CREATION_TIMESTAMP	CRTTIME	TIMESTAMP Nullable	The date and time that the PTF was created. Contains the null value when the creation date and time cannot be determined.
PTF_INCLUDED_IN_GROUP_DATE	PTF_DATE	VARCHAR(10)	The date that this PTF was first made available in a group PTF. Contains the character form of a date formatted as MM/DD/YY.

Note

The PSP websites used by this service are found based upon the PTF groups that are currently installed on the partition. For each distinct PTF group, a unique PSP XML feed is accessed:

```
https://www.ibm.com/support/pages/sites/default/files/inline-files/<PTF-Group-Name>.xml
```

For example, the JAVA PTF group details can be accessed using:

```
https://www.ibm.com/support/pages/sites/default/files/inline-files/SF99572.xml
```

To determine the IP address for your geography, ping www.ibm.com.

Example

- Review the details for the PTFs which have not yet been applied for the PTF groups installed on this partition.

```
SELECT * FROM SYSTOOLS.GROUP_PTF_DETAILS
WHERE PTF_STATUS <> 'PTF APPLIED'
ORDER BY PTF_GROUP_NAME
```

Related reference

SYSTOOLS

SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

GROUP_PTF_INFO view

The GROUP_PTF_INFO view contains information about the group PTFs for the server.

The information returned is similar to the information available from [Work with PTF Groups \(WRKPTFGRP\)](#) CL command.

Authorization: The caller must have *USE authority to the Work with PTF Groups (WRKPTFGRP) command.

The following table describes the columns in the view. The system name is GRPPTFINFO. The schema is QSYS2.

Table 165. GROUP_PTF_INFO view

Column name	System column name	Data type	Description
COLLECTED_TIME	COLLE00001	TIMESTAMP	Date and time of when this row information was generated.
PTF_GROUP_NAME	PTF_G00001	VARCHAR(60) Nullable	Name of the PTF group.
PTF_GROUP_DESCRIPTION	PTF_G00002	VARCHAR(100) Nullable	Description of the PTF group.
PTF_GROUP_LEVEL	PTF_G00003	INTEGER Nullable	Level of the PTF group.
PTF_GROUP_TARGET_RELEASE	PTF_G00004	VARCHAR(6) Nullable	Release level for PTF group.
PTF_GROUP_STATUS	PTF_G00005	VARCHAR(20) Nullable	Status of the PTF group. UNKNOWN The PTF group status cannot be resolved because a related PTF group is either not found on the system or is in error. NOT APPLICABLE All PTFs in the PTF group and related PTF groups are for products that are not installed or supported on this system. SUPPORTED ONLY There are no PTFs in the PTF group or related PTF groups that are for installed products on this system. There is at least one PTF that is for a product, release, option, and load identifier that is supported on this system. NOT INSTALLED There is at least one PTF that is for an installed product on this system, and not all of the PTFs or their superseding PTFs are temporarily or permanently applied. INSTALLED All PTFs for products that are installed on this system are temporarily or permanently applied. If a PTF is superseded, a superseding PTF is either temporarily or permanently applied. ERROR The PTF group information is in error. Either delete the PTF group or replace the PTF group information that is currently on the system. APPLY AT NEXT IPL All PTFs for the installed products on the system are either set to be applied at the next IPL or are already temporarily or permanently applied. RELATED GROUP The PTF group does not have any PTFs for products installed or supported on the system. However, it is identified in another PTF group as a related PTF group. Deleting a PTF group in this status will cause the other PTF group to have a status of UNKNOWN. ON ORDER There is at least one PTF in the group that is on order and has not yet been installed on the system. It will be delivered on either physical or virtual media.

Example

Determine the level of the latest CUM PTF group installed on the system.

```
SELECT MAX(PTF_GROUP_LEVEL) AS CUM_LEVEL
FROM QSYS2.GROUP_PTF_INFO
WHERE PTF_GROUP_NAME IN ('SF99610', 'SF99710')
AND PTF_GROUP_STATUS = 'INSTALLED'
```

PTF_INFO view

The PTF_INFO view contains information about PTFs for the server.

The information returned is similar to [QpzListPTF API](#).

Authorization: The caller must have *USE authority to the Display Program Temporary Fix (DSPPTF) command.

The following table describes the columns in the view. The schema is QSYS2.

Table 166. PTF_INFO view

Column name	System column name	Data type	Description
PTF_PRODUCT_ID	LICPGM	VARCHAR(7) Nullable	Product identifier.
PTF_PRODUCT_OPTION	PRODOPT	VARCHAR(6) Nullable	Product option.
PTF_PRODUCT_RELEASE_LEVEL	PRODRLS	VARCHAR(6) Nullable	Product release level.
PTF_PRODUCT_DESCRIPTION	PRODESC	VARCHAR(132) Nullable	Product description.
PTF_IDENTIFIER	PTFID	VARCHAR(7) Nullable	The identifier of the PTF.
PTF_RELEASE_LEVEL	PTFRLS	VARCHAR(6) Nullable	The release level of the PTF.
PTF_PRODUCT_LOAD	PRODLOAD	VARCHAR(4) Nullable	The load ID of the product load for the PTF.

Table 166. PTF_INFO view (continued)

Column name	System column name	Data type	Description
PTF_LOADED_STATUS	LOADSTAT	VARCHAR(19)	The current loaded status of the PTF.
		Nullable	<p>NOT LOADED The PTF has never been loaded.</p> <p>LOADED The PTF has been loaded.</p> <p>APPLIED The PTF has been temporarily applied.</p> <p>PERMANENTLY APPLIED The PTF has been applied permanently.</p> <p>PERMANENTLY REMOVED The PTF has been permanently removed.</p> <p>DAMAGED The PTF is damaged. An error occurred while applying the PTF. It needs to be reloaded and applied.</p> <p>SUPERCEDED The PTF is superseded. A PTF will have a status of superseded when one of the following situations occurs:</p> <ul style="list-style-type: none"> • Another PTF with a more recent correction for the problem has been loaded on the system. The PTF ID that has been loaded can be found in the PTF_SUPERCEDED_BY_PTF column. • The PTF save file for another PTF with a more recent correction for the problem has been logged into *SERVICE on the system.
PTF_SAVE_FILE	SAVF	VARCHAR(3)	Indicates whether a save file exists for the PTF.
		Nullable	<p>NO The PTF has no save file.</p> <p>YES The PTF has a save file.</p>
PTF_COVER_LETTER	COVER	VARCHAR(3)	Indicates whether a cover letter exists for the PTF.
		Nullable	<p>NO The PTF has no cover letter.</p> <p>YES The PTF has a cover letter.</p>
PTF_ON_ORDER	ONORD	VARCHAR(3)	Indicates whether the PTF has been ordered.
		Nullable	<p>NO The PTF has not been ordered or has already been received.</p> <p>YES The PTF has been ordered.</p>
PTF_IPL_ACTION	IPLACT	VARCHAR(19)	The action to be taken on this PTF during the next unattended IPL.
		Nullable	<p>NONE No action occurs at the next IPL.</p> <p>TEMPORARILY APPLIED The PTF is temporarily applied at the next IPL.</p> <p>TEMPORARILY REMOVED The PTF is temporarily removed at the next IPL.</p> <p>PERMANENTLY APPLIED The PTF is permanently applied at the next IPL.</p> <p>PERMANENTLY REMOVED The PTF is permanently removed at the next IPL.</p>

Table 166. PTF_INFO view (continued)

Column name	System column name	Data type	Description
PTF_ACTION_PENDING	ACTPEND	VARCHAR(3) Nullable	<p>Indicates whether a required action has yet to be performed to make this PTF active.</p> <p>NO No required actions are pending for this PTF.</p> <p>YES A required action needs to occur for this PTF to be active. Check the Activation Instructions section of the cover letter to determine what the action is. If the PTF_ACTION_REQUIRED column is set to IPL and the activation instructions have been performed, then the PTF is active. However, this column will not be updated until the next IPL.</p>
PTF_ACTION_REQUIRED	ACTREQ	VARCHAR(12) Nullable	<p>Indicates whether an action is required to make this PTF active when it is applied. See the cover letter to determine what action needs to be taken.</p> <p>NONE No activation instructions are needed for this PTF.</p> <p>EXIT PROGRAM This PTF was shipped with activation instructions in the cover letter. This value is returned for all PTFs that have an exit program to update the status of the PTF after the activation instructions have been performed.</p> <p>IPL This PTF was shipped with activation instructions in the cover letter. No exit program exists to verify the activation instructions were performed.</p>
PTF_IPL_REQUIRED	IPLREQ	VARCHAR(9) Nullable	<p>Indicates whether an IPL is required to apply this PTF.</p> <p>DELAYED The PTF is delayed. The PTF must be applied during an IPL.</p> <p>IMMEDIATE The PTF is immediate. No IPL is needed to apply the PTF.</p> <p>UNKNOWN The type of the PTF is not known.</p>
PTF_IS_RELEASED	RELEASED	VARCHAR(3) Nullable	<p>Indicates whether the PTF save file is available for distribution to another system. This is set to YES only when the System Manager for IBM i licensed program is on the system and the product is supported. The PTF_SAVE_FILE column must have a value of YES before using the value in this column.</p> <p>NO The PTF save file cannot be distributed.</p> <p>YES The PTF save file is released and can be distributed to another system.</p>
PTF_MINIMUM_LEVEL	MINLVL	VARCHAR(2) Nullable	<p>The indicator of the lowest level of the product to which this PTF can be applied. The level can be AA to 99.</p> <p>Contains the null value if the product does not have a level.</p>
PTF_MAXIMUM_LEVEL	MAXLVL	VARCHAR(2) Nullable	<p>The indicator of the highest level of the product to which this PTF can be applied. The level can be AA to 99.</p> <p>Contains the null value if the product does not have a level.</p>
PTF_STATUS_TIMESTAMP	STATTIME	TIMESTAMP Nullable	<p>The date and time that the PTF status was last changed.</p> <p>Contains the null value when the status date and time is not available.</p>
PTF_SUPERCEDED_BY_PTF	SUPERCEDE	VARCHAR(7) Nullable	<p>The identifier of the PTF that has replaced this PTF.</p> <p>This field will be blank when the PTF is not superseded or when the superseding PTF has not been loaded on the system.</p>

Table 166. PTF_INFO view (continued)

Column name	System column name	Data type	Description
PTF_CREATION_TIMESTAMP	CRTTIME	TIMESTAMP Nullable	The date and time that the PTF was created. Contains the null value when the creation date and time cannot be determined.
PTF_TECHNOLOGY_REFRESH_PTF	TRPTF	VARCHAR(3) Nullable	Indicates whether this is a technology refresh PTF. NO This is not a technology refresh PTF. YES This is a technology refresh PTF.
PTF_TEMPORARY_APPLY_TIMESTAMP	TMPTIME	TIMESTAMP Nullable	The date and time that the PTF was temporarily applied. Contains the null value if the PTF has not been temporarily applied.

Examples

- Find which PTFs will be impacted by the next IPL.

```
SELECT PTF_IDENTIFIER, PTF_IPL_ACTION, A.*
FROM QSYS2.PTF_INFO A
WHERE PTF_IPL_ACTION <> 'NONE'
```

- Find which PTFs are loaded but not applied.

```
SELECT PTF_IDENTIFIER, PTF_PRODUCT_DESCRIPTION, A.*
FROM QSYS2.PTF_INFO A
WHERE PTF_LOADED_STATUS = 'LOADED'
ORDER BY PTF_PRODUCT_ID
```

Security Services

These views, procedures, and functions provide security information.

AUTHORITY_COLLECTION view

The AUTHORITY_COLLECTION view returns information about the authority check for an object.

See [AUTHORITY_COLLECTION view](#) for the description of the view.

AUTHORIZATION_LIST_INFO view

The AUTHORIZATION_LIST_INFO view returns a list of all objects secured by an authorization list.

The information returned is similar to the information available through the Display Authorization List Objects (DSPAUTLOBJ) CL command and the List Objects Secured by Authorization List (QSYLATLO) API.

Authorization: Detail is returned when one of the following is true:

- The caller has *READ authority to the authorization list.
- The caller is authorized to the QIBM_DB_SECADM function usage identifier.
- The caller has *ALLOBJ special authority.

The following table describes the columns in the view. The system name is AUTHL_INFO. The schema is QSYS2.

Table 167. AUTHORIZATION_LIST_INFO view

Column name	System column name	Data type	Description
AUTHORIZATION_LIST	AUTH_LIST	VARCHAR(10)	The authorization list for this object.

Table 167. AUTHORIZATION_LIST_INFO view (continued)

Column name	System column name	Data type	Description
SYSTEM_OBJECT_SCHEMA	SYS_DNAME	VARCHAR(10) Nullable	The library that contains the object. Returns the null value if the object is not in the QSYS or QDLS file system.
SYSTEM_OBJECT_NAME	SYS_ONAME	VARCHAR(10) Nullable	The object that is secured by the authorization list. Returns the null value if the object is not in the QSYS or QDLS file system.
SYSTEM_OBJECT_TYPE	SYS_OTYPE	VARCHAR(8)	The system object type of the secured object.
OBJECT_ATTRIBUTE	OBJATTR	VARCHAR(5) Nullable	The attribute for the secured object's type. Returns the null value if the object has no attribute or if it is not in the QSYS or QDLS file system.
OBJECT_SCHEMA	OSHEMA	VARCHAR(128) Nullable	The SQL schema name for this object. Returns the null value if the object is not in the QSYS or QDLS file system.
OBJECT_NAME	ONAME	VARCHAR(128) Nullable	The SQL name of the object. For an external procedure or an external function, the name will be returned when a single procedure or function exists for that *PGM or *SRVPGM object. Contains the null value if an SQL name could not be returned.
OBJECT_TYPE	OTYPE	VARCHAR(9) Nullable	The SQL object type. The following values can be returned. ALIAS The object is an SQL alias. FUNCTION The object is an SQL function. INDEX The object is an SQL index. PACKAGE The object is an SQL package. PROCEDURE The object is an SQL procedure. ROUTINE The object is used in SQL by one or more external functions and/or external procedures. SEQUENCE The object is an SQL sequence. TABLE The object is an SQL table. TRIGGER The object is an SQL trigger. TYPE The object is an SQL type. VARIABLE The object is an SQL global variable. VIEW The object is an SQL view. XSR The object is an XML schema repository object. Returns the null value if the object is not an SQL object.
OBJECT_OWNER	OWNER	VARCHAR(10)	The owner of the object.
PRIMARY_GROUP	GROUP	VARCHAR(10) Nullable	The user who is the primary group for the object. Returns the null value if there is no primary group for the object.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The descriptive text for the secured object. Returns the null value if the object is not in the QSYS or QDLS file system.
ASPGRP	ASPGRP	VARCHAR(10)	The name of the ASP device containing the object. A value of *SYSBAS indicates the system ASP and all basic user ASPs.

Table 167. AUTHORIZATION_LIST_INFO view (continued)

Column name	System column name	Data type	Description
AUTHORITY HOLDER	AUT HOLDER	VARCHAR(3)	Indicates whether the object is an authority holder. NO The object is not an authority holder. YES The object is an authority holder.
PATH_NAME	PATH_NAME	DBCLOB(16M) CCSID 1200 Nullable	The path name for the object that is secured by the authorization list. Returns the null value if the object is in the QSYS or QDLS file system.
DLO_NAME	DLO_NAME	VARCHAR(12) Nullable	The document library object (DLO) name for the object. Returns the null value if OBJECT_TYPE is not *DOC (document) or *FLR (folder).
FOLDER_PATH	FOLDER	VARCHAR(63) Nullable	The name of the folder that contains the DLO object. Returns the null value if the object is not in a folder.

Example

Return information about all the object secured by authorization list APP1.

```
SELECT * FROM QSYS2.AUTHORIZATION_LIST_INFO WHERE AUTHORIZATION_LIST = 'APP1';
```

AUTHORIZATION_LIST_USER_INFO view

The AUTHORIZATION_LIST_USER_INFO view returns a list of all authorization lists and their authorities.

The information returned is similar to the information available through the Display Authorization List (DSPAUTL) CL command.

Authorization: None required.

The following table describes the columns in the view. The system name is AUTL_USERS. The schema is QSYS2.

Table 168. AUTHORIZATION_LIST_USER_INFO view

Column name	System column name	Data type	Description
AUTHORIZATION_LIST	AUTL	VARCHAR(10)	The name of the authorization list.
AUTHORIZATION_NAME	USER_NAME	VARCHAR(10)	User profile name. Can contain the following special value. *PUBLIC This row contains the public authority for the object.

Table 168. AUTHORIZATION_LIST_USER_INFO view (continued)

Column name	System column name	Data type	Description
OBJECT_AUTHORITY	OBJ_AUTH	VARCHAR(12)	<p>The authority that the user has to the object. Contains one of the following values:</p> <p>*ALL Allows all operations on the object except those that are limited to the owner or controlled by authorization list management authority.</p> <p>*CHANGE Allows all operations on the object except those that are limited to the owner or controlled by object existence authority, object alter authority, object reference authority, and object management authority.</p> <p>*EXCLUDE All operations on the object are prohibited.</p> <p>*USE Allows access to the object attributes and use of the object. The user cannot change the object.</p> <p>USER DEFINED The specific object authorities and data authorities do not match any of the predefined object authority levels.</p>
AUTHORIZATION_LIST_MANAGEMENT	AUTL_MGMT	VARCHAR(3)	<p>The authorization list management authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OWNER	OWNER	VARCHAR(10)	The owner of the authorization list.
OBJECT_OPERATIONAL	OBJOPER	VARCHAR(3)	<p>The object operational authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_MANAGEMENT	OBJMGT	VARCHAR(3)	<p>The object management authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_EXISTENCE	OBJEXIST	VARCHAR(3)	<p>The object existence authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_ALTER	OBJALTER	VARCHAR(3)	<p>The object alter authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
OBJECT_REFERENCE	OBJREF	VARCHAR(3)	<p>The object reference authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>
DATA_READ	DATA_READ	VARCHAR(3)	<p>The data read authority for AUTHORIZATION_NAME.</p> <p>NO The user does not have this authority.</p> <p>YES The user has this authority.</p>

Table 168. AUTHORIZATION_LIST_USER_INFO view (continued)

Column name	System column name	Data type	Description
DATA_ADD	DATA_ADD	VARCHAR(3)	The data add authority for AUTHORIZATION_NAME. NO The user does not have this authority. YES The user has this authority.
DATA_UPDATE	DATA_UPD	VARCHAR(3)	The data update authority for AUTHORIZATION_NAME. NO The user does not have this authority. YES The user has this authority.
DATA_DELETE	DATA_DEL	VARCHAR(3)	The data delete authority for AUTHORIZATION_NAME. NO The user does not have this authority. YES The user has this authority.
DATA_EXECUTE	DATA_EXEC	VARCHAR(3)	The data execute authority for AUTHORIZATION_NAME. NO The user does not have this authority. YES The user has this authority.
TEXT_DESCRIPTION	TEXT	VARCHAR(50) Nullable	The descriptive text for the authorization list. Contains null if the authorization list has no text description.

Example

List the public security settings for all authorization lists.

```
SELECT *
FROM QSYS2.AUTHORIZATION_LIST_USER_INFO
WHERE AUTHORIZATION_NAME = '*PUBLIC';
```


CERTIFICATE_INFO table function

The CERTIFICATE_INFO table function returns a result table that contains information about server or Certificate Authority (CA) certificates.

This information is similar to what is returned by the Retrieve Certificate Information (QYCURTVCI, QycuRetrieveCertificateInfo) API.

Authorization: The caller must provide the password for the certificate store. In addition, the caller must have *ALLOBJ and *SECADM special authorities.

►► CERTIFICATE_INFO — (→



The schema is QSYS2.

certificate-store-password An expression that contains the password for the specified certificate store. It is recommended that the password be passed as a variable, not as a string visible as clear text.

The following value can be used instead of specifying the password.

***NOPWD** The certificate store password will be retrieved from the stashed password file.

certificate-store A character or graphic string expression that indicates the certificate store from which a list of certificates is to be retrieved. The value can either be a fully qualified Integrated File System (IFS) directory path and file name of the certificate store, starting with a leading forward slash (/), or one of the following special values. If the file name does not identify a certificate store, no rows are returned. If *certificate-store* is not specified, *SYSTEM is the default.

***OBJECTSIGNING** The *OBJECTSIGNING certificate store.

***SIGNATUREVERIFICATION** The *SIGNATUREVERIFICATION certificate store.

***SYSTEM** The *SYSTEM certificate store.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 169. CERTIFICATE_INFO table function

Column Name	Data Type	Description
CERTIFICATE_LABEL	VARCHAR(256)	The label for the certificate.
SERIAL_NUMBER	VARCHAR(64)	Serial number.
VALIDITY_START	TIMESTAMP(0)	The beginning date of the validity period.
VALIDITY_END	TIMESTAMP(0)	The ending date of the validity period.
TRUSTED	VARCHAR(3)	Indicates if the certificate is trusted. NO The certificate is not trusted. YES The certificate is trusted.
KEY_SIZE	INTEGER	The size of the key, in bytes.
PRIVATE_KEY	VARCHAR(3)	Indicates if the certificate has a private key. NO The certificate does not have a private key. YES The certificate has a private key.
PRIVATE_KEY_LABEL	VARCHAR(64)	The private key label. Contains the null value if PRIVATE_KEY is NO or if PRIVATE_KEY_STORAGE_LOCATION is SOFTWARE.
PRIVATE_KEY_STORAGE_LOCATION	VARCHAR(19)	Where the key is stored. HARDWARE The key is stored in hardware. HARDWARE ENCRYPTION The key is stored in hardware encryption. SOFTWARE The key is stored is software. Contains the null value if PRIVATE_KEY is NO.
DIGITAL_SIGNATURE	VARCHAR(3)	The certificate has the digital signature extension. NO The certificate does not have the digital signature extension. YES The certificate has the digital signature extension.
NONREPUDIATION	VARCHAR(3)	Indicates if the certificate has the nonrepudiation extension. NO The certificate does not have the nonrepudiation extension. YES The certificate has the nonrepudiation extension.

Table 169. CERTIFICATE_INFO table function (continued)

Column Name	Data Type	Description
KEY_ENCIPHERMENT	VARCHAR(3)	Indicates if the certificate has the key encipherment extension. NO The certificate does not have the key encipherment extension. YES The certificate has the key encipherment extension.
DATA_ENCIPHERMENT	VARCHAR(3)	Indicates if the certificate has the data encipherment extension. NO The certificate does not have the data encipherment extension. YES The certificate has the data encipherment extension.
KEY_AGREEMENT	VARCHAR(3)	Indicates if the certificate has the key agreement extension NO The certificate does not have the key agreement extension. YES The certificate has the key agreement extension.
KEY_CERTIFICATE_SIGNATURE	VARCHAR(3)	Indicates if the certificate has the key certificate signature extension. NO The certificate does not have the key certificate signature extension. YES The certificate has the key certificate signature extension.
CRL_SIGNATURE	VARCHAR(3)	Indicates if the certificate has the Certificate Revocation List (CRL) signature extension. NO The certificate does not have the CRL signature extension. YES The certificate has the CRL signature extension.
CRL_LOCATION	VARCHAR(50)	The CRL location. Contains the null value if no value is available.
ENCIPHER_ONLY	VARCHAR(3)	Indicates if the certificate has the encipher only extension. NO The certificate does not have the encipher only extension. YES The certificate has the encipher only extension.
DECIPHER_ONLY	VARCHAR(3)	Indicates if the certificate has the decipher only extension. NO The certificate does not have the decipher only extension. YES The certificate has the decipher only extension.
LDAP_SERVER_NAME	VARCHAR(900)	The LDAP server name. Contains the null value if no value is available.
IP_ADDRESS_COUNT	INTEGER	The number of addresses in the IP_ADDRESSES column. Currently, only one IP address is returned. Contains the null value if no IP addresses are available.
IP_ADDRESSES	VARCHAR(45)	The IP address. Contains the null value if no value is available.
DOMAIN_NAME_COUNT	INTEGER	The number of domain names in the DOMAIN_NAMES column. Currently, only one domain name is returned. Contains the null value if no domain names are available.
DOMAIN_NAMES	VARCHAR(256)	The domain name. Contains the null value if no value is available.
EMAIL_ADDRESS	VARCHAR(256)	The email address. Contains the null value if no value is available.
CRYPTOGRAPHIC_DEVICE_COUNT	INTEGER	The number of cryptographic devices.
CRYPTOGRAPHIC_DEVICES	VARCHAR(109)	A list of cryptographic device descriptions. Each entry is ten characters long. A single blank separates entries. Contains the null value if CRYPTOGRAPHIC_DEVICE_COUNT is 0.

Table 169. CERTIFICATE_INFO table function (continued)

Column Name	Data Type	Description
SUBJECT_COMMON_NAME	VARCHAR(256)	The subject's common name. The SUBJECT set of columns define information about the end-entity that is being described for the certificate. Contains the null value if no value is available.
SUBJECT_ORGANIZATIONAL_UNIT	VARCHAR(256)	The subject's organizational unit. Contains the null value if no value is available.
SUBJECT_ORGANIZATION	VARCHAR(256)	The subject's organization. Contains the null value if no value is available.
SUBJECT_LOCALITY	VARCHAR(128)	The subject's locality. Contains the null value if no value is available.
SUBJECT_STATE_PROVINCE	VARCHAR(128)	The subject's state or province. Contains the null value if no value is available.
SUBJECT_POSTAL_CODE	VARCHAR(16)	The subject's postal code. Contains the null value if no value is available.
SUBJECT_COUNTRY_REGION	VARCHAR(3)	The subject's country or region. Contains the null value if no value is available.
ISSUER_COMMON_NAME	VARCHAR(256)	The issuer's common name. The ISSUER set of columns define information about the Certificate Authority that signed the end-entity certificate. Contains the null value if no value is available.
ISSUER_ORGANIZATIONAL_UNIT	VARCHAR(256)	The issuer's organizational unit. Contains the null value if no value is available.
ISSUER_ORGANIZATION	VARCHAR(256)	The issuer's organization. Contains the null value if no value is available.
ISSUER_LOCALITY	VARCHAR(128)	The issuer's locality. Contains the null value if no value is available.
ISSUER_STATE_PROVINCE	VARCHAR(128)	The issuer's state or province. Contains the null value if no value is available.
ISSUER_POSTAL_CODE	VARCHAR(16)	The issuer's postal code. Contains the null value if no value is available.
ISSUER_COUNTRY_REGION	VARCHAR(3)	The issuer's country or region. Contains the null value if no value is available.

Example

Retrieve all the certificates for the *SYSTEM certificate store that will be expired within the next month. Use a password that has been set in a global variable.

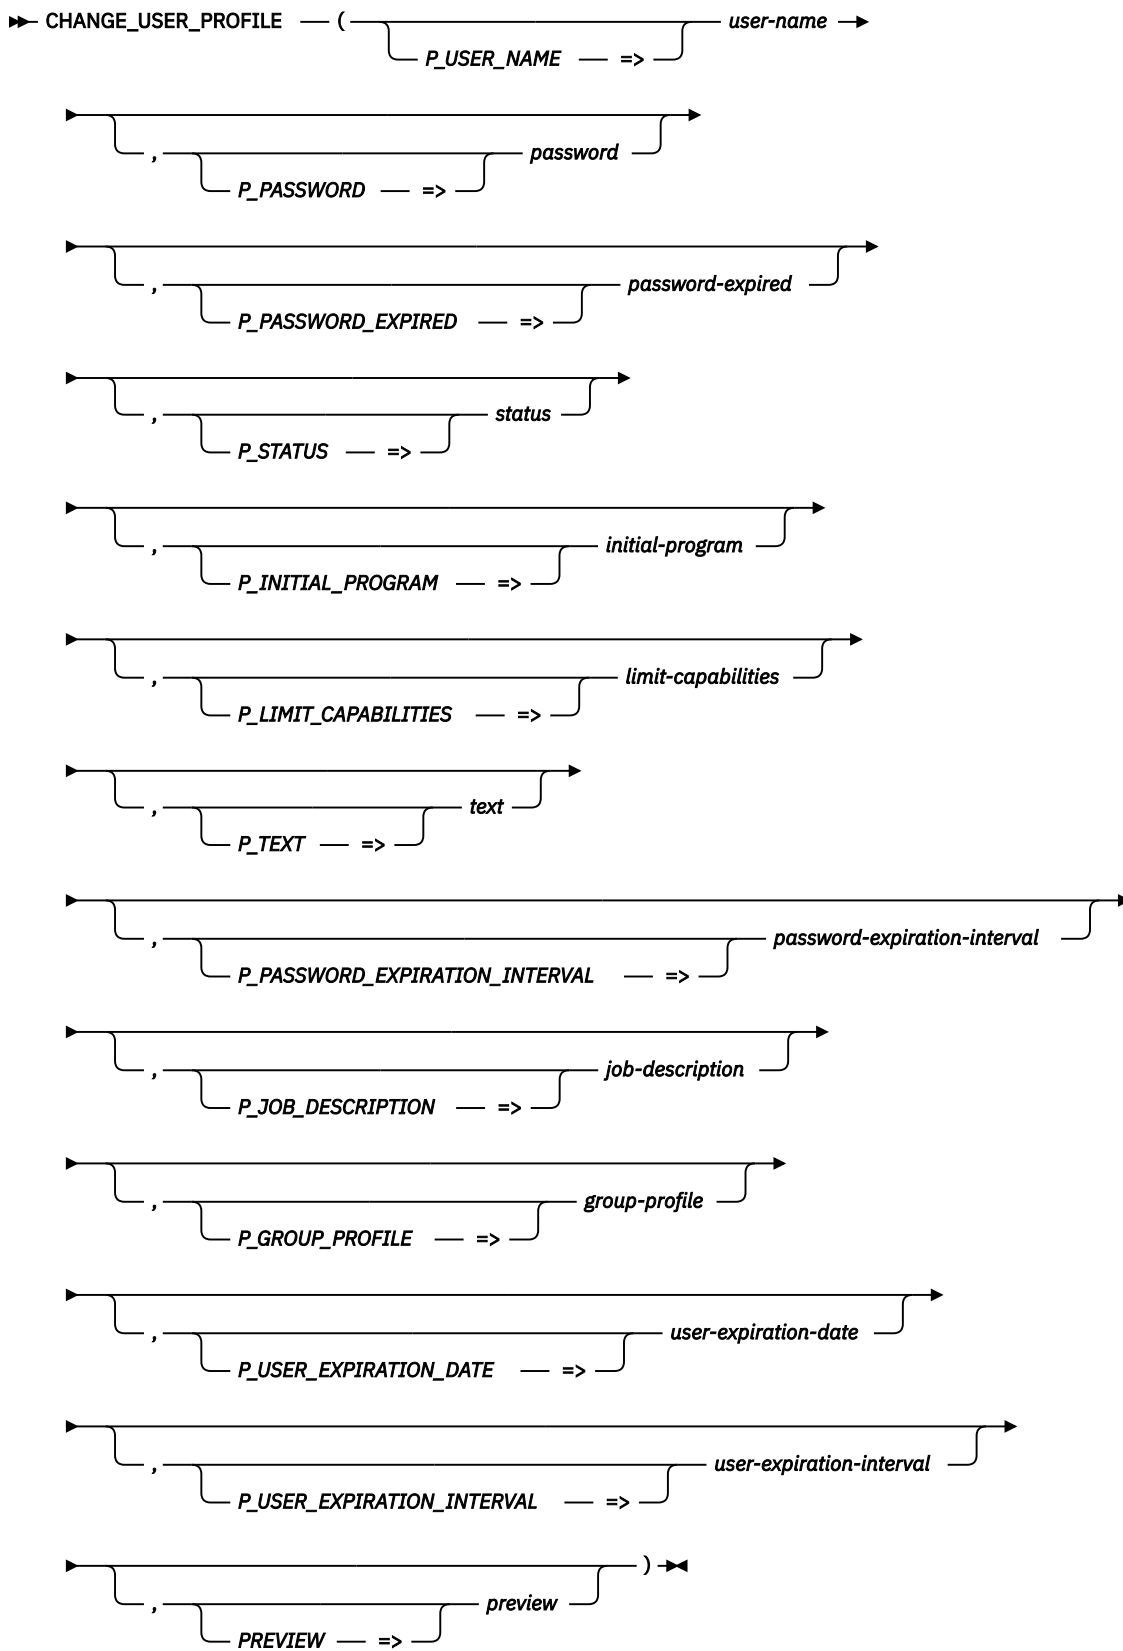
- ```
CREATE VARIABLE MYLIB.SYSTEM_CERT_PW VARCHAR(30);
SET MYLIB.SYSTEM_CERT_PW = 'cert_pwd';
SELECT * FROM TABLE(QSYS2.CERTIFICATE_INFO(CERTIFICATE_STORE_PASSWORD=> MYLIB.SYSTEM_CERT_PW))
WHERE VALIDITY_END < CURRENT DATE + 1 MONTH;
```

## CHANGE\_USER\_PROFILE table function

The CHANGE\_USER\_PROFILE table function changes a subset of user profile attributes.

For a detailed description of the parameters and their values, refer to the [CHGUSRPRF CL command](#).

**Authorization:** This table function calls the Change User Profile (CHGUSRPRF) CL command. Any authority requirements for the CL command apply to the use of this function.



The schema is SYSTOOLS.

**user-name** A character string containing the name of the user profile whose values are to be changed.

- password** A character string containing a new password value for the user profile. This is the PASSWORD parameter. The default is \*SAME.
- password-expired** A character string that specifies whether the password for this user is set to expired. This is the PWDEXP parameter. The default is \*SAME.
- status** A character string that specifies the status of the user profile. This is the STATUS parameter. The default is \*SAME.
- initial-program** A character string that specifies the initial program to call. This is the INLPGM parameter. The default is \*SAME.
- limit-capabilities** A character string that specifies the capabilities for a user. This is the LMTCPB parameter. The default is \*SAME.
- text** A character string that specifies the descriptive text for the user profile. This is the TEXT parameter. The default is \*SAME.
- password-expiration-interval** A character string that specifies the password expiration interval, in days. This is the PWDEXPITV parameter. The default is \*SAME.
- job-description** A character string that specifies the job description associated with this user profile. This is the JOBD parameter. The default is \*SAME.
- group-profile** A character string that specifies the group profile associated with this user profile. This is the GRPPRF parameter. The default is \*SAME.
- user-expiration-date** A character string that specifies the date when the user profile expires and is automatically disabled. This is the USREXPDATE parameter. The default is \*SAME.
- user-expiration-interval** An integer value that specifies the expiration interval, in days, before the user profile is automatically disabled. This is the USREXPITV parameter. The default is NULL, meaning the value will not be changed.
- preview** A character string that indicates whether the table function should execute the CHGUSRPRF command that has been constructed based upon the input parameters or whether only a preview of the potential action should be shown.
- NO** The table function should change the user profiles
- YES** The table function should return a preview of the changes. This is the default.

The result of the function is a table containing a single row with the details about the user profile change. The columns of the result table are described in the following table. The result columns are nullable.

Table 170. CHANGE\_USER\_PROFILE table function

| Column Name        | Data Type     | Description                                                                                                                                                                           |
|--------------------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USER_NAME          | VARCHAR(10)   | The user profile being changed or previewed.                                                                                                                                          |
| CHANGE_ATTEMPTED   | VARCHAR(3)    | Indicates whether the change was attempted.<br><b>NO</b> The change was not attempted<br><b>YES</b> The change was attempted                                                          |
| CHANGE_SUCCESSFUL  | VARCHAR(3)    | Indicates whether the change was successful.<br><b>NO</b> The change was not successful<br><b>YES</b> The change was successful<br>Contains the null value if CHANGE_ATTEMPTED is NO. |
| CHGUSRPRF_COMMAND  | VARCHAR(1000) | The CHGUSRPRF command string.                                                                                                                                                         |
| FAILURE_MESSAGE_ID | CHAR(7)       | The message ID.<br>Contains the null value if CHANGE_ATTEMPTED is NO or CHANGE_SUCCESSFUL is YES.                                                                                     |

Table 170. CHANGE\_USER\_PROFILE table function (continued)

| Column Name          | Data Type        | Description                                                                    |
|----------------------|------------------|--------------------------------------------------------------------------------|
| FAILURE_MESSAGE_TEXT | VARGRAPHIC(1024) | The text of the message.                                                       |
|                      | CCSID 1200       | Contains the null value if CHANGE_ATTEMPTED is NO or CHANGE_SUCCESSFUL is YES. |

## Note

This function is provided in the SYSTOOLS schema as a helper function to manage user profiles. Similar to other Db2 for i provided tools within SYSTOOLS, the SQL source can be extracted and used as a model for building similar helper functions, or to create a customized version within a user-specified schema.

## Example

Disable all enabled user profiles that have no password set.

Preview the list of profiles that will be affected.

```
SELECT * FROM QSYS2.USER_INFO,
 TABLE(SYSTOOLS.CHANGE_USER_PROFILE(
 P_USER_NAME => AUTHORIZATION_NAME,
 P_STATUS => '*DISABLED',
 PREVIEW => 'YES'))
WHERE STATUS = '*ENABLED' AND
 NO_PASSWORD_INDICATOR = 'YES';
```

Build and execute the CHGUSRPRF commands..

```
SELECT * FROM QSYS2.USER_INFO,
 TABLE(SYSTOOLS.CHANGE_USER_PROFILE(
 P_USER_NAME => AUTHORIZATION_NAME,
 P_STATUS => '*DISABLED',
 PREVIEW => 'NO'))
WHERE STATUS = '*ENABLED' AND
 NO_PASSWORD_INDICATOR = 'YES';
```

## DRDA\_AUTHENTICATION\_ENTRY\_INFO view

The DRDA\_AUTHENTICATION\_ENTRY\_INFO view returns user server authentication entry information.

A server authentication entry defines a userid and password to send on a connect request over TCP/IP. A server authentication list is associated with every user profile on the system. The Add Server Authentication Entry (ADDSVRAUTE) command is used to add entries.

When a DRDA connection over TCP/IP is attempted without specifying a userid and password, and password authentication is required, the Db2 for i client checks the server authentication list for the user profile under which the client job is running. If it finds a match between the RDB name on the CONNECT statement and the server name in an authentication entry, or the server name is the special value QDDMDRDASERVER, the associated userid (and password if one exists) is used for the connection.

A server authentication entry can also be used to specify a userid and password to be used for a DDM connection over TCP/IP. When a DDM connection is attempted over TCP/IP, and password authentication is required, the Db2 for i client checks the server authentication list for the user profile under which the client job is running. If it finds a match between the RDB name specified in the DDM file and the server name in an authentication entry, or the server name is the special value QDDMDRDASERVER, the associated userid (and password if one exists) is used for the connection. If no RDB name is specified in the DDM file and the server name is either of the special values QDDMDRDASERVER or QDDMSERVER, the associated userid (and password if one exists) is used for the connection.

**Authorization:** Only \*USRPRF objects that the caller has \*OBJOPR and \*READ authority to will be returned.

The following table describes the columns in the view. The system name is DRDA\_AUTHE. The schema is QSYS2.

Table 171. DRDA\_AUTHENTICATION\_ENTRY\_INFO view

| Column Name               | System Column Name | Data Type                      | Description                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------|--------------------|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUTHORIZATION_NAME        | USER_NAME          | VARCHAR(10)                    | The user profile on the client system.                                                                                                                                                                                                                                                                                                                                           |
| SERVER_NAME               | SRVR_NAME          | VARGRAPHIC(200)<br>CCSID 1200  | The target system for the authentication entry.<br><br>This is the name of the RDB or QDDMDRDASERVER that is used for connections made on behalf of RDB DDM files or DRDA connections. For a non-RDB DDM file that does not use the RDB directory, the value will be QDDMDRDASERVER or QDDMSERVER. See <a href="#">Client security in a TCP/IP network</a> for more information. |
| SERVER_AUTHORIZATION_NAME | SRVR_USER          | VARGRAPHIC(1000)<br>CCSID 1200 | The user profile on the target system.                                                                                                                                                                                                                                                                                                                                           |
| PASSWORD_STORED           | PW_STORED          | VARCHAR(3)                     | Indicates whether a password is stored for the authentication entry.<br><br><b>YES</b> A password is stored for the authentication entry.<br><b>NO</b> A password is not stored for the authentication entry.                                                                                                                                                                    |

## Example

For an auditor, generate a list of user profiles that have authentication entries on the system:

```
SELECT DISTINCT(AUTHORIZATION_NAME)
FROM QSYS2.DRDA_AUTHENTICATION_ENTRY_INFO
```

## FUNCTION\_INFO view

The FUNCTION\_INFO view contains details about function usage identifiers.

**Authorization:** None required.

The following table describes the columns in the view. The system name is FCN\_INFO. The schema is QSYS2.

Table 172. FUNCTION\_INFO view

| Column Name                | System Column Name | Data Type                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|--------------------|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FUNCTION_ID                | FCNID              | VARCHAR(30)<br>Nullable                    | The function ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| FUNCTION_CATEGORY          | FCNCAT             | VARCHAR(10)<br>Nullable                    | Indicates whether the function is a client or host function.<br><br><b>1 - CLIENT</b> The function is a locally managed client function within IBM i Navigator.<br><b>2 - CLIENT</b> The function is a locally managed client function, not within IBM i Navigator.<br><b>3 - HOST</b> The function is a host function.<br><b>4 - CLIENT</b> The function is a centrally managed client function within IBM i Navigator.<br><b>5 - CLIENT</b> The function is a centrally managed client function, not within IBM i Navigator. |
| FUNCTION_TYPE              | FCNTYP             | VARCHAR(13)<br>Nullable                    | The type of function.<br><br><b>PRODUCT</b> The function is a function product.<br><b>GROUP</b> The function is a function group.<br><b>ADMINISTRABLE</b> The function is an administrable function.                                                                                                                                                                                                                                                                                                                           |
| FUNCTION_NAME_MESSAGE_TEXT | FCNMSGTXT          | VARGRAPHIC(330)<br>CCSID(1200)<br>Nullable | The first-level text for the function-name message ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| FUNCTION_NAME              | FCNNAM             | VARGRAPHIC(330)<br>CCSID(1200)<br>Nullable | The text for the function name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |



Table 172. FUNCTION\_INFO view (continued)

| Column Name                       | System Column Name | Data Type                                      | Description                                                                                                                                                                                                                                                         |
|-----------------------------------|--------------------|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FUNCTION_DESCRIPTION_MESSAGE_TEXT | FCNDESCTXT         | VARGRAPHIC(330)<br>CCSID(1200)<br><br>Nullable | The first-level text for the function-description message ID.                                                                                                                                                                                                       |
| FUNCTION_DESCRIPTION              | FCNDESC            | VARGRAPHIC(330)<br>CCSID(1200)<br><br>Nullable | The text for the function description.                                                                                                                                                                                                                              |
| FUNCTION_PRODUCT_ID               | FCNPRDID           | VARCHAR(30)<br><br>Nullable                    | The ID of the product that the function is registered for.                                                                                                                                                                                                          |
| FUNCTION_GROUP_ID                 | FCNGRPID           | VARCHAR(30)<br><br>Nullable                    | The ID of the function group that the function is grouped with. If the function is not grouped with a function group, this field is set to *NONE.                                                                                                                   |
| DEFAULT_USAGE                     | DFTUG              | VARCHAR(7)<br><br>Nullable                     | The default usage for the function.<br><br><b>DENIED</b> The default usage does not allow usage of the function.<br><b>ALLOWED</b> The default usage allows usage of the function.                                                                                  |
| ALLOBJ_INDICATOR                  | ALLOBJ             | VARCHAR(8)<br><br>Nullable                     | Indicates whether a user with *ALLOBJ special authority can use the function.<br><br><b>NOT USED</b> The user, its groups, or default must allow usage of the function.<br><b>USED</b> A user with *ALLOBJ special authority is always allowed to use the function. |
| USAGE_INFORMATION_INDICATOR       | USGINFO            | VARCHAR(3)<br><br>Nullable                     | Indicates whether there is usage information defined for the function.<br><br><b>NO</b> There is no usage information defined for the function.<br><b>YES</b> There is usage information defined for the function.                                                  |

## Example

Determine what function usage IDs exist and their default configuration.

```
SELECT * FROM QSYS2.FUNCTION_INFO ORDER BY FUNCTION_ID
```

## FUNCTION\_USAGE view

The FUNCTION\_USAGE view contains function usage configuration details.

**Authorization:** Only callers with \*SECADM user special authority can examine the function usage configuration details returned with this view.

The following table describes the columns in the view. The system name is FCN\_USAGE. The schema is QSYS2.

Table 173. FUNCTION\_USAGE view

| Column Name | System Column Name | Data Type   | Description                                                                                                                                                 |
|-------------|--------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FUNCTION_ID | FCNID              | VARCHAR(30) | The ID of the function.                                                                                                                                     |
| USER_NAME   | USER_NAME          | VARCHAR(10) | The name of the user profile that has a usage setting for this function                                                                                     |
| USAGE       | USAGE              | VARCHAR(7)  | Usage setting.<br><br><b>ALLOWED</b> The user profile is allowed to use the function.<br><b>DENIED</b> The user profile is not allowed to use the function. |
| USER_TYPE   | USER_TYPE          | VARCHAR(5)  | Type of user profile.<br><br><b>USER</b> The user profile is a user.<br><b>GROUP</b> The user profile is a group.                                           |

## Example

Determine what function usage has been granted or revoked.

```
SELECT * FROM QSYS2.FUNCTION_USAGE ORDER BY FUNCTION_ID, USER_NAME
```

## GROUP\_PROFILE\_ENTRIES view

The GROUP\_PROFILE\_ENTRIES view contains one row for each user profile that is part of a group profile.

Both group profile (GRPPRF) and supplemental group profile (SUPGRPPRF) information is considered for each user profile.

**Authorization:** Only \*USRPRF objects that the caller has \*READ authority to will be returned.

The following table describes the columns in the view. The system name is GROUPLIST. The schema is QSYS2.

Table 174. GROUP\_PROFILE\_ENTRIES view

| Column Name        | System Column Name | Data Type               | Description                    |
|--------------------|--------------------|-------------------------|--------------------------------|
| GROUP_PROFILE_NAME | GROUPNAME          | VARCHAR(128)            | Group profile name             |
| USER_PROFILE_NAME  | USERNAME           | VARCHAR(128)            | User profile name              |
| USER_TEXT          | USER_TEXT          | VARCHAR(50)<br>Nullable | User profile text description. |

## OBJECT\_OWNERSHIP view

The OBJECT\_OWNERSHIP view returns ownership information for all objects.

The values returned for the columns in the view are closely related to the values returned by the WRKOBJOWN CL command and the List Objects User Is Authorized to, Owns, or Is Primary Group of (QSYLOBJA) API.

**Authorization:** The caller must have read authority to the user profile that owns the object.

The following table describes the columns in the view. The system name is OBJ\_OWN. The schema is QSYS2.

Table 175. OBJECT\_OWNERSHIP view

| Column Name        | System Column Name | Data Type                             | Description                                                                                                                   |
|--------------------|--------------------|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| AUTHORIZATION_NAME | USER_NAME          | VARCHAR(10)                           | User profile that owns the object.                                                                                            |
| OBJECT_TYPE        | OBJ_TYPE           | VARCHAR(7)                            | The type of object.                                                                                                           |
| OBJECT_LIBRARY     | LIBNAME            | VARCHAR(10)<br>Nullable               | The name of the library containing the object.<br>Contains the null value if OBJECT_NAME is null.                             |
| OBJECT_NAME        | NAME               | VARCHAR(10)<br>Nullable               | The name of the object.<br>Contains the null value if OBJECT_TYPE is *BLKSF, *CHRSF, *DIR, *FIFO, *STMF, or *SYMLNK.          |
| PATH_NAME          | PATH_NAME          | DBCLOB(16M)<br>CCSID 1200<br>Nullable | The path name of the object.<br>Contains the null value if OBJECT_TYPE is not *BLKSF, *CHRSF, *DIR, *FIFO, *STMF, or *SYMLNK. |
| OBJECT_ATTRIBUTE   | ATTRIBUTE          | VARCHAR(10)<br>Nullable               | The object's attribute.<br>Contains the null value if there is no attribute for the object.                                   |
| TEXT_DESCRIPTION   | TEXT               | VARCHAR(50)<br>Nullable               | The descriptive text for this object.<br>Contains the null value if the object has no text description.                       |

Table 175. OBJECT\_OWNERSHIP view (continued)

| Column Name                   | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------|--------------------|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IASP_NAME                     | IASP_NAME          | VARCHAR(10)            | The auxiliary storage pool (ASP) device name where the object is stored. If the object is in the system ASP or one of the basic user ASPs, contains *SYSBAS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| AUTHORITY HOLDER              | AUT HOLDER         | VARCHAR(3)             | Whether the object is an authority holder.<br><br><b>NO</b> The object is not an authority holder.<br><b>YES</b> The object is an authority holder.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| AUTHORIZATION_LIST_MANAGEMENT | AUTL_MGMT          | VARCHAR(3)<br>Nullable | Whether AUTHORIZATION_NAME has authorization list management authority to the object.<br><br><b>NO</b> The user does not have authorization list management authority.<br><b>YES</b> The user has authorization list management authority.<br><br>Contains the null value if OBJECT_TYPE is not *AUTL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| OBJECT_AUTHORITY              | OBJ_AUTH           | VARCHAR(12)            | The authority that AUTHORIZATION_NAME has to the object. Contains one of the following special values:<br><br><b>*ALL</b> Allows all operations on the object except those that are limited to the owner or controlled by authorization list management authority.<br><br><b>*CHANGE</b> Allows all operations on the object except those that are limited to the owner or controlled by object existence authority, object alter authority, object reference authority, and object management authority.<br><br><b>*EXCLUDE</b> All operations on the object are prohibited.<br><br><b>*USE</b> Allows access to the object attributes and use of the object. The user cannot change the object.<br><br><b>USER DEFINED</b> The specific object authorities and data authorities do not match any of the predefined object authority levels. |
| OBJECT_OPERATIONAL            | OBJOPER            | VARCHAR(3)             | Indicates the object operational authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| OBJECT_MANAGEMENT             | OBJMGT             | VARCHAR(3)             | The object management authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| OBJECT_EXISTENCE              | OBJEXIST           | VARCHAR(3)             | The object existence authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 175. OBJECT\_OWNERSHIP view (continued)

| Column Name      | System Column Name | Data Type  | Description                                                                                                                                               |
|------------------|--------------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBJECT_ALTER     | OBJALTER           | VARCHAR(3) | The object alter authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.     |
| OBJECT_REFERENCE | OBJREF             | VARCHAR(3) | The object reference authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority. |
| DATA_READ        | DATA_READ          | VARCHAR(3) | The data read authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.        |
| DATA_ADD         | DATA_ADD           | VARCHAR(3) | The data add authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.         |
| DATA_UPDATE      | DATA_UPD           | VARCHAR(3) | The data update authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.      |
| DATA_DELETE      | DATA_DEL           | VARCHAR(3) | The data delete authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.      |
| DATA_EXECUTE     | DATA_EXEC          | VARCHAR(3) | The data execute authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.     |

## Examples

- Return a list of all objects owned by user FRANKDBA.

```
SELECT * FROM QSYS2.OBJECT_OWNERSHIP
WHERE AUTHORIZATION_NAME = 'FRANKDBA';
```

- Return a list of only objects in the IFS that are owned by user FRANKDBA.

```
SELECT * FROM QSYS2.OBJECT_OWNERSHIP
WHERE AUTHORIZATION_NAME = 'FRANKDBA'
AND PATH_NAME IS NOT NULL;
```

## OBJECT\_PRIVILEGES table function

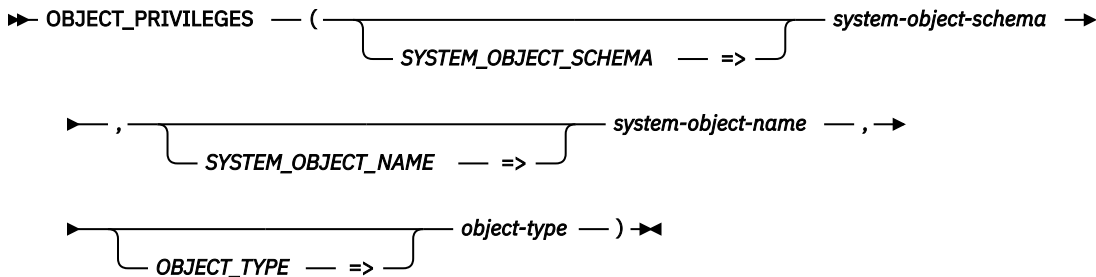
The OBJECT\_PRIVILEGES table function returns a row for every user authorized to the specified object, along with their associated object and data authorities.

The information returned is similar to the information available through the Display Object Authority (DSPOBJAUT) CL command.

**Authorization:** All authorized users are returned for an object when at least one of the following is true:

- The caller has \*OBJMGT authority.
- The caller is the owner of the object.
- The object is an authorization list.
- The caller is authorized to the QIBM\_DB\_SECADM function usage identifier.

Otherwise, only authorizations for the caller are returned.



The schema is QSYS2.

**system-object-schema** A character or graphic string expression that identifies the library that contains *system-object-name*.

**system-object-name** A character or graphic string expression that identifies the object.

**object-type** A character or graphic string expression that specifies the system object type of *system-object-name*.

The result of the function is a table containing rows with the format shown in the following table. All columns are nullable.

Table 176. `OBJECT_PRIVILEGES` table function

| Column name        | Data type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUTHORIZATION_USER | VARCHAR(10) | User profile name. Can contain the following special value.<br><br><b>*PUBLIC</b> This row contains the public authority for the object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| OBJECT_AUTHORITY   | VARCHAR(12) | The authority that the user has to the object. Contains one of the following special values:<br><br><b>*ALL</b> Allows all operations on the object except those that are limited to the owner or controlled by authorization list management authority.<br><br><b>*AUTL</b> The public authority specified in the authorization list used by this object is used.<br><br><b>*CHANGE</b> Allows all operations on the object except those that are limited to the owner or controlled by object existence authority, object alter authority, object reference authority, and object management authority.<br><br><b>*EXCLUDE</b> All operations on the object are prohibited.<br><br><b>*USE</b> Allows access to the object attributes and use of the object. The user cannot change the object.<br><br><b>USER DEFINED</b> The specific object authorities and data authorities do not match any of the predefined object authority levels. |
| AUTHORIZATION_LIST | VARCHAR(10) | The name of the authorization list if the object is secured by an authorization list.<br><br>Contains null if the object is not secured by an authorization list.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 176. OBJECT\_PRIVILEGES table function (continued)

| Column name                   | Data type   | Description                                                                                                                                                                                                             |
|-------------------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRIMARY_GROUP                 | VARCHAR(10) | The user who is the primary group for the object.<br>Returns the null value if there is no primary group for the object.                                                                                                |
| OWNER                         | VARCHAR(10) | The user profile that owns the object.                                                                                                                                                                                  |
| AUTHORIZATION_LIST_MANAGEMENT | VARCHAR(3)  | The authorization list management authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                  |
| OBJECT_OWNER                  | VARCHAR(3)  | Indicates whether the AUTHORIZATION_USER for this row is also the OWNER of the object.<br><br><b>NO</b> AUTHORIZATION_USER is not the owner of the object.<br><b>YES</b> AUTHORIZATION_USER is the owner of the object. |
| OBJECT_OPERATIONAL            | VARCHAR(3)  | Indicates the object operational authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                   |
| OBJECT_MANAGEMENT             | VARCHAR(3)  | The object management authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                              |
| OBJECT_EXISTENCE              | VARCHAR(3)  | The object existence authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                               |
| OBJECT_ALTER                  | VARCHAR(3)  | The object alter authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                   |
| OBJECT_REFERENCE              | VARCHAR(3)  | The object reference authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                               |
| DATA_READ                     | VARCHAR(3)  | The data read authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                      |
| DATA_ADD                      | VARCHAR(3)  | The data add authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                       |
| DATA_UPDATE                   | VARCHAR(3)  | The data update authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                    |
| DATA_DELETE                   | VARCHAR(3)  | The data delete authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                    |

Table 176. OBJECT\_PRIVILEGES table function (continued)

| Column name  | Data type  | Description                                                                                                                                           |
|--------------|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| DATA_EXECUTE | VARCHAR(3) | The data execute authority for AUTHORIZATION_USER.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority. |

## Example

Return authority information for the file APPLIB/EMPLOYEE.

```
SELECT *
FROM TABLE(QSYS2.OBJECT_PRIVILEGES('APPLIB', 'EMPLOYEE', '*FILE'));
```

## OBJECT\_PRIVILEGES view

The OBJECT\_PRIVILEGES view returns a row for every user authorized to an object, along with their associated object and data authorities.

The information returned is similar to the information available through the Display Object Authority (DSPOBJAUT) CL command.

**Authorization:** All authorized users are returned for an object when at least one of the following is true:

- The caller has \*OBJMGT authority.
- The caller is the owner of the object.
- The object is an authorization list.
- The caller is authorized to the QIBM\_DB\_SECADM function usage identifier.

Otherwise, only authorizations for the caller are returned.

The following table describes the columns in the view. The system name is OBJ\_PRIV. The schema is QSYS2.

Table 177. OBJECT\_PRIVILEGES view

| Column name          | System column name | Data type                    | Description                                                                                                                                                                                                                                                      |
|----------------------|--------------------|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBJECT_SCHEMA        | OSHEMA             | VARCHAR(128)                 | The SQL schema name for this object.                                                                                                                                                                                                                             |
| OBJECT_NAME          | NAME               | VARCHAR(128)<br><br>Nullable | The SQL name of the object.<br><br>For an external procedure or an external function, the name will be returned when a single procedure or function exists for that *PGM or *SRVPGM object.<br><br>Contains the null value if an SQL name could not be returned. |
| SYSTEM_OBJECT_SCHEMA | SYS_DNAME          | VARCHAR(10)                  | The library that contains the object.                                                                                                                                                                                                                            |
| SYSTEM_OBJECT_NAME   | SYS_ONAME          | VARCHAR(10)                  | The system object name.                                                                                                                                                                                                                                          |
| OBJECT_TYPE          | OBJTYPE            | VARCHAR(8)                   | The system object type.                                                                                                                                                                                                                                          |

Table 177. OBJECT\_PRIVILEGES view (continued)

| Column name        | System column name | Data type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SQL_OBJECT_TYPE    | SQLTYPE            | VARCHAR(9)<br>Nullable  | <p>The SQL object type. The following values can be returned.</p> <p><b>ALIAS</b> The object is an SQL alias.</p> <p><b>FUNCTION</b> The object is an SQL function.</p> <p><b>INDEX</b> The object is an SQL index.</p> <p><b>PACKAGE</b> The object is an SQL package.</p> <p><b>PROCEDURE</b> The object is an SQL procedure.</p> <p><b>ROUTINE</b> The object is used in SQL by one or more external functions and/or external procedures.</p> <p><b>SEQUENCE</b> The object is an SQL sequence.</p> <p><b>TABLE</b> The object is an SQL table.</p> <p><b>TRIGGER</b> The object is an SQL trigger.</p> <p><b>TYPE</b> The object is an SQL type.</p> <p><b>VARIABLE</b> The object is an SQL global variable.</p> <p><b>VIEW</b> The object is an SQL view.</p> <p><b>XSR</b> The object is an XML schema repository object.</p> <p>Contains the null value if the object is not an SQL object or if the user does not have sufficient authority to the object.</p> |
| AUTHORIZATION_NAME | USER_NAME          | VARCHAR(10)             | <p>User profile name. Can contain the following special value.</p> <p><b>*PUBLIC</b> This row contains the public authority for the object.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| OBJECT_AUTHORITY   | OBJ_AUTH           | VARCHAR(12)             | <p>The authority that the user has to the object. Contains one of the following special values:</p> <p><b>*ALL</b> Allows all operations on the object except those that are limited to the owner or controlled by authorization list management authority.</p> <p><b>*AUTL</b> The public authority specified in the authorization list used by this object is used.</p> <p><b>*CHANGE</b> Allows all operations on the object except those that are limited to the owner or controlled by object existence authority, object alter authority, object reference authority, and object management authority.</p> <p><b>*EXCLUDE</b> All operations on the object are prohibited.</p> <p><b>*USE</b> Allows access to the object attributes and use of the object. The user cannot change the object.</p> <p><b>USER DEFINED</b> The specific object authorities and data authorities do not match any of the predefined object authority levels.</p>                     |
| OWNER              | OWNER              | VARCHAR(10)             | The user profile that owns the object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| AUTHORIZATION_LIST | AUTL               | VARCHAR(10)<br>Nullable | <p>The name of the authorization list if the object is secured by an authorization list.</p> <p>Contains null if the object is not secured by an authorization list.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |



Table 177. OBJECT\_PRIVILEGES view (continued)

| Column name                   | System column name | Data type               | Description                                                                                                                                                                                                             |
|-------------------------------|--------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRIMARY_GROUP                 | GROUP              | VARCHAR(10)<br>Nullable | The user who is the primary group for the object.<br>Contains the null value if there is no primary group for the object.                                                                                               |
| AUTHORIZATION_LIST_MANAGEMENT | AUTL_MGMT          | VARCHAR(3)              | The authorization list management authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                  |
| OBJECT_OWNER                  | OBJ_OWNER          | VARCHAR(3)              | Indicates whether the AUTHORIZATION_NAME for this row is also the OWNER of the object.<br><br><b>NO</b> AUTHORIZATION_NAME is not the owner of the object.<br><b>YES</b> AUTHORIZATION_NAME is the owner of the object. |
| OBJECT_OPERATIONAL            | OBJOPER            | VARCHAR(3)              | Indicates the object operational authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                   |
| OBJECT_MANAGEMENT             | OBJMGT             | VARCHAR(3)              | The object management authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                              |
| OBJECT_EXISTENCE              | OBJEXIST           | VARCHAR(3)              | The object existence authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                               |
| OBJECT_ALTER                  | OBJALTER           | VARCHAR(3)              | The object alter authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                   |
| OBJECT_REFERENCE              | OBJREF             | VARCHAR(3)              | The object reference authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                               |
| DATA_READ                     | DATA_READ          | VARCHAR(3)              | The data read authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                      |
| DATA_ADD                      | DATA_ADD           | VARCHAR(3)              | The data add authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                       |
| DATA_UPDATE                   | DATA_UPD           | VARCHAR(3)              | The data update authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                                                                    |

Table 177. OBJECT\_PRIVILEGES view (continued)

| Column name      | System column name | Data type               | Description                                                                                                                                                             |
|------------------|--------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DATA_DELETE      | DATA_DEL           | VARCHAR(3)              | The data delete authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                    |
| DATA_EXECUTE     | DATA_EXEC          | VARCHAR(3)              | The data execute authority for AUTHORIZATION_NAME.<br><br><b>NO</b> The user does not have this authority.<br><b>YES</b> The user has this authority.                   |
| TEXT_DESCRIPTION | TEXT               | VARCHAR(50)<br>Nullable | The descriptive text for this object.<br>Contains the null value if the object has no text description or if the user does not have sufficient authority to the object. |

## Example

Find user profiles that are publicly accessible.

```
SELECT *
FROM QSYS2.OBJECT_PRIVILEGES
WHERE SYSTEM_OBJECT_SCHEMA = 'QSYS' AND
 OBJECT_TYPE = '*USRPRF' AND
 AUTHORIZATION_NAME = '*PUBLIC' AND
 OBJECT_AUTHORITY <> '*EXCLUDE';
```

## SECURITY\_INFO view

The SECURITY\_INFO view returns one row containing information about the IBM i security configuration.

The values returned for the columns in the view are closely related to the values returned by the Display Security Attributes (DSPSECA) and Display Security Auditing (DSPSECAUD) CL commands and by the Retrieve Security Attributes (QSYRTVSA) API.

**Authorization:** The caller must have \*AUDIT special authority to see the system values for QAUDCTL, QAUDLVL, QAUDLVL2, and QCRTOBJAUD.

To see the AUDIT\_JOURNAL\_RECEIVER\_LIBRARY and AUDIT\_JOURNAL\_RECEIVER values, the caller must have:

- \*OBJOPR and some data authority other than \*EXECUTE to journal QSYS/QAUDJRN.

The following table describes the columns in the view. The system name is SEC\_INFO. The schema is QSYS2.

Table 178. SECURITY\_INFO view

| Column Name            | System Column Name | Data Type           | Description                                                                                                                                                                                                                                                                           |
|------------------------|--------------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SECURITY_LEVEL         | SECLVL             | INTEGER             | The security level that is currently being used by the system.<br><br><b>20</b> Password security only<br><b>30</b> Password and object security<br><b>40</b> Password, object, and operating system integrity<br><b>50</b> Password, object, and enhanced operating system integrity |
| PENDING_SECURITY_LEVEL | PENDESECLVL        | INTEGER<br>Nullable | The security level that the system will use after the next IPL.<br>Contains the null value if the security level will not change after the next IPL.                                                                                                                                  |

Table 178. SECURITY\_INFO view (continued)

| Column Name                  | System Column Name | Data Type           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------|--------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PASSWORD_LEVEL               | PWDLVL             | INTEGER             | The password level that is currently being used by the system.<br><br><b>0</b> User profile passwords with a length of 1-10 characters are supported.<br><b>1</b> User profile passwords with a length of 1-10 characters are supported. IBM i NetServer passwords for Windows 95/98/ME clients will be removed from the system.<br><b>2</b> User profile passwords with a length of 1-128 characters are supported.<br><b>3</b> User profile passwords with a length of 1-128 characters are supported. IBM i NetServer passwords for Windows 95/98/ME clients will be removed from the system. |
| PENDING_PASSWORD_LEVEL       | PENDPWDLVL         | INTEGER<br>Nullable | The password level that the system will use after the next IPL.<br><br>Contains the null value if the password level will not change after the next IPL.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| AUDIT_JOURNAL_EXISTS         | QAUDJRN            | VARCHAR(3)          | Whether the security journal QAUDJRN exists.<br><br><b>NO</b> The security journal QAUDJRN does not exist.<br><b>YES</b> The security journal QAUDJRN exists.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| PASSWORD_CHANGE_BLOCK        | QPWDCHGBLK         | VARCHAR(5)          | The current setting for the block password change (QPWDCHGBLK) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PASSWORD_EXPIRATION_INTERVAL | QPWDEXPITV         | VARCHAR(6)          | The current setting for the password expiration interval (QPWDEXPITV) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| PASSWORD_EXPIRATION_WARNING  | QPWDEXPWRN         | INTEGER             | The current setting for the password expiration warning (QPWDEXPWRN) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| PASSWORD_LIMIT_DIGITS        | QPWDLMTAJC         | INTEGER             | The current setting for the limit adjacent digits in password (QPWDLMTAJC) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| PASSWORD_LIMIT_CHARACTERS    | QPWDLMTCHR         | VARCHAR(10)         | The current setting for the limit characters in password (QPWDLMTCHR) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| PASSWORD_LIMIT_REPEAT        | QPWDLMTREP         | INTEGER             | The current setting for the limit repeating characters in password (QPWDLMTREP) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| PASSWORD_LIMIT_POSITIONS     | QPWDPOSDIF         | INTEGER             | The current setting for the limit password character positions (QPWDPOSDIF) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| PASSWORD_REQUIRE_DIGIT       | QPWDRQDDGT         | INTEGER             | The current setting for the require digit in password (QPWDRQDDGT) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| PASSWORD_MAXIMUM_LENGTH      | QPWDMAXLEN         | INTEGER             | The current setting for the maximum password length (QPWDMAXLEN) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| PASSWORD_MINIMUM_LENGTH      | QPWDMINLEN         | INTEGER             | The current setting for the minimum password length (QPWDMINLEN) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| PASSWORD_DUPLICATION         | QPWDRQDDIF         | INTEGER             | The current setting for the duplicate password control (QPWDRQDDIF) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| PASSWORD_RULES               | QPWDRULES          | VARCHAR(750)        | The current setting for the password rules (QPWDRULES) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| PASSWORD_VALIDATION_PROGRAM  | QPWDVLDPGM         | VARCHAR(20)         | The current setting for the password validation program (QPWDVLDPGM) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| CREATE_PUBLIC_AUTHORITY      | QCRTAUT            | VARCHAR(8)          | The current setting for the create default public authority (QCRTAUT) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| CREATE_OBJECT_AUDITING       | QCRTOBJAUD         | VARCHAR(7)          | The current setting for the create object auditing (QCRTOBJAUD) system value.<br><br>Returns the value *NOTAVL if caller does not have *AUDIT special authority.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| MAXIMUM_SIGNON_ATTEMPTS      | QMAXSIGN           | VARCHAR(6)          | The current setting for the maximum sign-on attempts allowed (QMAXSIGN) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| MAXIMUM_SIGNON_ACTION        | QMAXSGNACN         | INTEGER             | The current setting for the action to take for failed sign-on attempts (QMAXSGNACN) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 178. SECURITY\_INFO view (continued)

| Column Name                      | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VERIFY_OBJECT_RESTORE            | QVYOBJRST          | INTEGER                 | The current setting for the verify object on restore (QVYOBJRST) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| ALLOW_OBJECT_RESTORE             | QALWOBJRST         | VARCHAR(150)            | The current setting for the allow object restore (QALWOBJRST) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| USE_ADOPTED_AUTHORITY            | QUSEADPAUT         | VARCHAR(10)             | The current setting for the use adopted authority (QUSEADPAUT) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| ALLOW_USER_DOMAIN                | QALWUSRDMN         | VARCHAR(500)            | The current setting for the allow user domain objects in libraries (QALWUSRDMN) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| LIMIT_SECOFR_ACCESS              | QLMTSECOFR         | INTEGER                 | The current setting for the limit security officer device access (QLMTSECOFR) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| INACTIVE_JOB_TIMEOUT             | QINACTITV          | VARCHAR(5)              | The current setting for the inactive job time-out (QINACTITV) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| INACTIVE_JOB_MESSAGE_QUEUE       | QINACTMSGQ         | VARCHAR(20)             | The current setting for the inactive job message queue (QINACTMSGQ) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| DISCONNECTED_JOB_INTERVAL        | QDSCJOBITV         | VARCHAR(5)              | The current setting for the time interval before disconnected jobs end (QDSCJOBITV) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| AUTOCONFIGURE_DEVICES            | QAUTOCFG           | INTEGER                 | The current setting for the autoconfigure devices (QAUTOCFG) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| AUTOCONFIGURE_REMOTE_CONTROLLERS | QAUTORMT           | INTEGER                 | The current setting for the autoconfigure of remote controllers (QAUTORMT) system value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| AUDITING_CONTROL                 | QAUDCTL            | VARCHAR(50)             | The current setting for the auditing control (QAUDCTL) system value.<br>Returns the value *NOTAVL if caller does not have *AUDIT special authority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| AUDITING_LEVEL                   | QAUDLVL            | VARCHAR(160)            | The current setting for the auditing level (QAUDLVL) system value.<br>Returns the value *NOTAVL if caller does not have *AUDIT special authority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| AUDITING_LEVEL_EXTENSION         | QAUDLVL2           | VARCHAR(990)            | The current setting for the auditing level extension (QAUDLVL2) system value.<br>Returns the value *NOTAVL if caller does not have *AUDIT special authority.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| AUDIT_JOURNAL_RECEIVER_LIBRARY   | JRNRCV_LIB         | VARCHAR(10)<br>Nullable | The name of the library that contains the journal receiver attached to the security journal.<br>Contains the null value if AUDIT_JOURNAL_EXISTS is NO or if caller is not authorized.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| AUDIT_JOURNAL_RECEIVER           | JRNRCV             | VARCHAR(10)<br>Nullable | The name of the journal receiver attached to the security journal.<br>Contains the null value if AUDIT_JOURNAL_EXISTS is NO or if caller is not authorized.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| ALLOW_DIGITAL_CERTIFICATE_ADD    | DCM_ADD            | VARCHAR(3)              | Whether digital certificates can be added to a certificate store using the Add Verifier (QYDOADDV, QydoAddVerifier) API, and whether the password for a certificate store can be reset using Digital Certificate Manager (DCM).<br><br><b>NO</b> Digital certificates cannot be added to a certificate store using the QYDOADDV API, and certificate store passwords cannot be reset using DCM.<br><br><b>YES</b> Digital certificates can be added to a certificate store using the QYDOADDV API, and certificate store passwords can be reset using DCM.<br><br>The Change SST Security Attributes (CHGSSTSECA) command can be used to change this attribute. |

Table 178. SECURITY\_INFO view (continued)

| Column Name                         | System Column Name | Data Type  | Description                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------|--------------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ALLOW_SECURITY_SYSVAL_CHANGE        | SYSVAL_CHG         | VARCHAR(3) | Whether the security related system values can be changed.<br><br><b>NO</b> The security related system values cannot be changed.<br><b>YES</b> The security related system values can be changed.<br><br>The Change SST Security Attributes (CHGSSTSECA) command can be used to change this attribute.                                                                                                                               |
| ALLOW_SERVICE_TOOLS_PASSWORD_CHANGE | SSTPWD_CHG         | VARCHAR(3) | Whether a service tools user ID with a default password that is expired can change its own password.<br><br><b>NO</b> A service tools user ID with a default password that is expired cannot change its own password.<br><b>YES</b> A service tools user ID with a default password that is expired can change its own password.<br><br>The Change SST Security Attributes (CHGSSTSECA) command can be used to change this attribute. |
| NEXT_USER_ID                        | NEXT_UID           | BIGINT     | The value that will be used the next time a user ID number (UID) is generated for a user profile.                                                                                                                                                                                                                                                                                                                                     |
| NEXT_GROUP_ID                       | NEXT_GID           | BIGINT     | The value that will be used the next time a group ID number (GID) is generated for a user profile.                                                                                                                                                                                                                                                                                                                                    |

### Example

- Return the security and password levels for the system.

```
SELECT SECURITY_LEVEL, PASSWORD_LEVEL FROM QSYS2.SECURITY_INFO;
```

## SET\_COLUMN\_ATTRIBUTE procedure

The SET\_COLUMN\_ATTRIBUTE procedure sets the SECURE attribute for a column so variable values used for the column cannot be seen in the database monitor or plan cache.

**Authorization:** The caller must have:

- \*OBJOPR and \*OBJALTER authority to the table, and
- \*EXECUTE authority to the library containing the table.

►► SET\_COLUMN\_ATTRIBUTE — ( →

► — *schema-name* — , — *table-name* — , — *column-name* — , — *attribute* — ) ►►

The schema is SYSPROC.

**schema-name** A character string expression containing the system name of a schema.

**table-name** A character string expression containing the system name of a table.

**column-name** A character string expression containing the system name of a column.

**attribute** A character string expression containing the attribute to set for the column.

Valid values are:

**SECURE NO** This column does not contain data that needs to be secured in a database monitor or plan cache

**SECURE YES** This column contains data that needs to be secured in a database monitor or plan cache.

All variable values for any query that references this column will not be visible in a database monitor or plan cache unless the security officer has started the database monitor or the security officer is accessing the plan cache. All host variable values will appear as \*SECURE when examined from the monitor and plan cache unless the user is the QSECOFR user.

The secure setting for a column is shown in the SECURE column of the QSYS2/SYSCOLUMNS2 catalog.

## Example

Set the credit card column in the ORDERS table so it is secure.

```
CALL SYSPROC.SET_COLUMN_ATTRIBUTE('LIB1', 'ORDERS', 'CCNBR', 'SECURE YES');
```

## SQL\_CHECK\_AUTHORITY scalar function

The SQL\_CHECK\_AUTHORITY scalar function returns an indication of whether the user is authorized to query the specified \*FILE object.

►► SQL\_CHECK\_AUTHORITY ( ( — *library-name* — , — *file-name* — ) ) ►►

The schema is QSYS2.

***library-name*** Library name containing the file.  
***file-name*** File name for which authority will be examined.

The result of the function is a SMALLINT.

The returned value is:

- 0 If the user does not have authority to query the file, the object is not a \*FILE object, or the object does not exist.
- 1 If the user is authorized to query the file.

## SQL\_CHECK\_FUNCTION\_USAGE scalar function

The SQL\_CHECK\_FUNCTION\_USAGE scalar function returns an indication of whether the effective user of the thread is authorized to the specified function usage identifier, either directly, as a member of a group profile, or through adopted authority.

The validation is performed using the Check User Function Usage (QSYCKUFU, QsyCheckUserFunctionUsage) API. The usage setting for the user, the user's group, the default usage value, and the allow \*ALLOBJ indicator for the function usage identifier are taken into account.

**Authorization:** None required.

►► SQL\_CHECK\_FUNCTION\_USAGE ( ( — *function-usage* — ) ) ►►

FUNCTION\_USAGE — =>

The schema is QSYS2.

***function-usage*** A character string containing one function usage identifier.

The result of the function is INTEGER.

The function returns the following values:

- 0 The user is not authorized to the function usage identifier.

- 1 The user is authorized to the function usage identifier.

### Example

- Check whether the current user is authorized to the QIBM\_DB\_SQLADM function usage identifier.


```
VALUES SQL_CHECK_FUNCTION_USAGE('QIBM_DB_SQLADM');
```

### SQL\_CHECK\_SPECIAL\_AUTHORITY scalar function

The SQL\_CHECK\_SPECIAL\_AUTHORITY scalar function returns an indication of whether the effective user of the thread has the specified special authority, either directly, as a member of a group profile, or through adopted authority.

If the user does not have the special authority, no AF-K audit journal entry is produced.

**Authorization:** None required.



The schema is QSYS2.

**special-authority** A character string containing one system special authority value. It must include the leading \* character.

The result of the function is INTEGER.

The function returns the following values:

- 0 The user does not have the special authority.
- 1 The user has the special authority.

### Example

- Check whether the current user has \*JOBCTL special authority.

```
VALUES SQL_CHECK_SPECIAL_AUTHORITY('*JOBCTL');
```

### USER\_INFO view

The USER\_INFO view contains information about user profiles.

The USER\_INFO\_BASIC view returns a subset of this information and is faster.

The values returned for the columns in the view are closely related to the values returned by [Retrieve User Information \(QSYRUSRI\) API](#). Refer to the API for more detailed information.

**Authorization:** Only \*USRPRF objects that the caller has \*OBJOPR and \*READ authority to will be returned. To see a non-null value for the USER\_DEFAULT\_PASSWORD column, the caller must have \*ALLOBJ and \*SECADM authority.

The following table describes the columns in the view. The schema is QSYS2.

Table 179. USER\_INFO view

| Column Name        | System Column Name | Data Type               | Description        |
|--------------------|--------------------|-------------------------|--------------------|
| AUTHORIZATION_NAME | USER_NAME          | VARCHAR(10)<br>Nullable | User profile name. |

Table 179. USER\_INFO view (continued)

| Column Name                  | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                |
|------------------------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PREVIOUS_SIGNON              | PRVSIGNON          | TIMESTAMP<br>Nullable   | The date and time the user last signed on.<br>Contains null if the profile has never been used to sign on.                                                                                                                                                                                                                 |
| SIGN_ON_ATTEMPTS_NOT_VALID   | SIGNONINV          | INTEGER<br>Nullable     | The number of sign-on attempts that were not valid since the last successful sign-on.                                                                                                                                                                                                                                      |
| STATUS                       | STATUS             | VARCHAR(10)<br>Nullable | The status of the user profile. Contains one of the following values:<br><br><b>*DISABLED</b> The user profile is disabled; therefore, the user cannot sign on.<br><br><b>*ENABLED</b> The user profile is enabled; therefore, the user is able to sign on.                                                                |
| NETSERVER_DISABLED           | NETSERVER          | VARCHAR(3)              | Whether this user profile is disabled for IBM i NetServer use.<br><br><b>NO</b> The user profile is not disabled for IBM i NetServer use.<br><br><b>YES</b> The user profile is disabled for IBM i NetServer use.                                                                                                          |
| PASSWORD_CHANGE_DATE         | PWDCHGDAT          | TIMESTAMP<br>Nullable   | The date the user's password was last changed.                                                                                                                                                                                                                                                                             |
| NO_PASSWORD_INDICATOR        | NOPWD              | VARCHAR(3)<br>Nullable  | Indicates whether *NONE is specified for the password in the user profile.<br><br><b>NO</b> The password in the user profile is not *NONE.<br><br><b>YES</b> The password in the user profile is *NONE.                                                                                                                    |
| PASSWORD_LEVEL_0_1           | PWD_0_1            | VARCHAR(3)<br>Nullable  | Indicates whether the user profile has a password that can be used for a system at QPWDLVL 0 or 1.<br><br><b>NO</b> The password cannot be used.<br><br><b>YES</b> The password can be used.<br><br>Contains the null value if NO_PASSWORD_INDICATOR has a value of YES.                                                   |
| PASSWORD_LEVEL_2_3           | PWD_2_3            | VARCHAR(3)<br>Nullable  | Indicates whether the user profile has a password that can be used for a system at QPWDLVL 2 or 3.<br><br><b>NO</b> The password cannot be used.<br><br><b>YES</b> The password can be used.<br><br>Contains the null value if NO_PASSWORD_INDICATOR has a value of YES.                                                   |
| PASSWORD_EXPIRATION_INTERVAL | PWDEXPITV          | SMALLINT<br>Nullable    | The number of days (from 1 through 366) the user's password can remain active before it must be changed. Can also be one of the following values:<br><br><b>0</b> The system value QPWDEXPITV is used to determine the user's password expiration interval.<br><br><b>-1</b> The user's password does not expire (*NOMAX). |
| DATE_PASSWORD_EXPIRES        | PWDEXPDAT          | TIMESTAMP<br>Nullable   | The date the user's password expires.<br>Contains null if the password will not expire.                                                                                                                                                                                                                                    |



Table 179. USER\_INFO view (continued)

| Column Name                 | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------|--------------------|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DAYS_UNTIL_PASSWORD_EXPIRES | PWDDAYSEXP         | INTEGER<br>Nullable      | The number of days until the password will expire. A value of 0 indicates the password is expired.<br><br>Contains null if the password will not expire within the number of days specified by the password expiration warning (QPWDEXPWARN) system value.                                                                                                      |
| SET_PASSWORD_TO_EXPIRE      | PWDEXP             | VARCHAR(3)<br>Nullable   | Indicates whether the user's password is set to expire, requiring the user to change the password when signing on. Contains one of the following values:<br><br><b>NO</b> The user's password is not set to expire.<br><b>YES</b> The user's password is set to expire.                                                                                         |
| USER_CLASS_NAME             | USRCLS             | VARCHAR(10)<br>Nullable  | The user's class name. Contains one of the following values:<br><br><b>*PGMR</b> The user has a class of programmer.<br><b>*SECADM</b> The user has a class of security administrator.<br><b>*SECOFR</b> The user has a class of security officer.<br><b>*SYSOPR</b> The user has a class of system operator.<br><b>*USER</b> The user has a class of end user. |
| SPECIAL_AUTHORITIES         | SPCAUT             | VARCHAR(88)<br>Nullable  | A list of the special authorities the user has. Up to 8 authorities are returned. Each entry is padded with blanks to fill 11 characters.<br><br>Contains null if the user has no special authorities.                                                                                                                                                          |
| GROUP_PROFILE_NAME          | GRPPRF             | VARCHAR(10)<br>Nullable  | The name of the group profile. Contains the value *NONE if the user does not have a group profile.                                                                                                                                                                                                                                                              |
| SUPPLEMENTAL_GROUP_COUNT    | SUPGRPCNT          | SMALLINT                 | The number of supplemental groups in the SUPPLEMENTAL_GROUP_LIST column.                                                                                                                                                                                                                                                                                        |
| SUPPLEMENTAL_GROUP_LIST     | SUPGRPLIST         | VARCHAR(150)<br>Nullable | A list of supplemental groups for the user profile. Up to 15 supplemental groups are returned. Each entry except for the last one is padded with blanks to fill 10 characters.<br><br>Contains null if the user has no supplemental groups.                                                                                                                     |
| OWNER                       | OWNER              | VARCHAR(10)<br>Nullable  | This field indicates who is to own objects created by this user. Contains one of the following values:<br><br><b>*GRPPRF</b> The user's group profile owns any objects the user creates.<br><b>*USRPRF</b> The user owns any objects the user creates.                                                                                                          |

Table 179. USER\_INFO view (continued)

| Column Name                  | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GROUP_AUTHORITY              | GRPAUT             | VARCHAR(10)<br>Nullable | The authority the user's group profile has to objects the user creates. Contains one of the following values:<br><br><b>*ALL</b> The group profile has all authority to the objects the user creates.<br><br><b>*CHANGE</b> The group profile has change authority to the objects the user creates.<br><br><b>*EXCLUDE</b> The group profile has exclude authority to the objects the user creates.<br><br><b>*NONE</b> The group profile has no authority to the objects the user creates. If the user does not have a group profile, this value is returned.<br><br><b>*USE</b> The group profile has use authority to the objects the user creates. |
| ASSISTANCE_LEVEL             | ASTLVL             | VARCHAR(10)<br>Nullable | The user interface that the user will use. Contains one of the following values:<br><br><b>*ADVANCED</b> The expert system user interface.<br><br><b>*BASIC</b> The Operational Assistant user interface.<br><br><b>*INTERMED</b> The system user interface.<br><br><b>*SYSVAL</b> The system value QASTLVL determines which user interface the user is using.                                                                                                                                                                                                                                                                                         |
| CURRENT_LIBRARY_NAME         | CURLIB             | VARCHAR(10)<br>Nullable | The name of the user's current library. Contains *CRTDFT if the user does not have a current library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| INITIAL_MENU_NAME            | INLMNU             | VARCHAR(10)<br>Nullable | The initial menu for the user. Can contain the special value *SIGNOFF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| INITIAL_MENU_LIBRARY_NAME    | INLMNULIB          | VARCHAR(10)<br>Nullable | The name of the library that the initial menu is in. Can contain the special value *LIBL. Contains null if the menu name is *SIGNOFF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| INITIAL_PROGRAM_NAME         | INITPGM            | VARCHAR(10)<br>Nullable | The initial program for the user. Contains *NONE if there is no initial program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| INITIAL_PROGRAM_LIBRARY_NAME | INITPGMLIB         | VARCHAR(10)<br>Nullable | The name of the library that the initial program is in. Can contain the special value *LIBL. Contains null if the initial program name is *NONE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| LIMIT_CAPABILITIES           | LMTCPB             | VARCHAR(10)<br>Nullable | Indicates whether the user has limited capabilities. Contains one of the following values:<br><br><b>*NO</b> The user is not limited.<br><br><b>*PARTIAL</b> The user cannot change his initial program or current library.<br><br><b>*YES</b> The user cannot change his initial menu, initial program, or current library. The user cannot run commands from the command line.                                                                                                                                                                                                                                                                       |
| TEXT_DESCRIPTION             | TEXT               | VARCHAR(50)<br>Nullable | The descriptive text for the user profile.<br><br>Contains null if the user profile has no text description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 179. USER\_INFO view (continued)

| Column Name                  | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DISPLAY_SIGNON_INFORMATION   | DSPSGNINF          | VARCHAR(10)<br>Nullable | Indicates whether the sign-on information display is shown when the user signs on. Contains one of the following special values:<br><br><b>*NO</b> The sign-on information display is not shown when the user signs on.<br><br><b>*SYSVAL</b> The system value QDSPSGNINF determines if the sign-on information display is shown when the user signs on.<br><br><b>*YES</b> The sign-on information display is shown when the user signs on.                                                                    |
| LIMIT_DEVICE_SESSIONS        | LMTDEVSSN          | VARCHAR(10)<br>Nullable | Specifies if the number of device sessions allowed for a user is limited. Can contain one of the following special values:<br><br><b>*NO</b> The user is not limited to a specific number of device sessions.<br><br><b>*SYSVAL</b> The system value QLMTDEVSSN determines if the user is limited to a specific number of device sessions.<br><br><b>*YES</b> The user is limited to one device session.                                                                                                        |
| KEYBOARD_BUFFERING           | KBDBUF             | VARCHAR(10)<br>Nullable | The keyboard buffering value that is used when a job is initialized for this user. Contains one of the following special values:<br><br><b>*NO</b> The type-ahead and attention-key buffering options are not on.<br><br><b>*SYSVAL</b> The system value QKBDBUF determines the keyboard buffering value for this user.<br><br><b>*TYPEAHEAD</b> The type-ahead option is on, but the attention-key buffering option is not.<br><br><b>*YES</b> The type-ahead and attention-key buffering options are both on. |
| MAXIMUM_ALLOWED_STORAGE      | MAXSTGLRG          | BIGINT<br>Nullable      | The maximum amount of auxiliary storage (in kilobytes) that can be assigned to store permanent objects owned by the user. Contains null if the user has no maximum storage.                                                                                                                                                                                                                                                                                                                                     |
| STORAGE_USED                 | STGUSED            | BIGINT<br>Nullable      | The amount of auxiliary storage (in kilobytes) occupied by this user's owned objects on *SYSBAS. The QSYS2.USER_STORAGE catalog should be used to determine the storage consumed on all ASPs.                                                                                                                                                                                                                                                                                                                   |
| HIGHEST_SCHEDULING_PRIORITY  | PTYLMT             | CHAR(1)<br>Nullable     | The highest scheduling priority the user is allowed to have for each job submitted to the system.                                                                                                                                                                                                                                                                                                                                                                                                               |
| JOB_DESCRIPTION_NAME         | JOBID              | VARCHAR(10)<br>Nullable | The name of the job description used for jobs that start through subsystem work station entries.                                                                                                                                                                                                                                                                                                                                                                                                                |
| JOB_DESCRIPTION_LIBRARY_NAME | JOBDLIB            | VARCHAR(10)<br>Nullable | Job description library name. Can contain the special value *LIBL.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| ACCOUNTING_CODE              | ACGCDE             | VARCHAR(15)<br>Nullable | The accounting code that is associated with this user.<br><br>Contains null if there is no accounting code.                                                                                                                                                                                                                                                                                                                                                                                                     |
| MESSAGE_QUEUE_NAME           | MSGQ               | VARCHAR(10)<br>Nullable | The name of the message queue that is used by this user.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 179. USER\_INFO view (continued)

| Column Name                         | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MESSAGE_QUEUE_LIBRARY_NAME          | MSGQLIB            | VARCHAR(10)<br>Nullable | The name of the library the message queue is in. Can contain the special value *LIBL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MESSAGE_QUEUE_DELIVERY_METHOD       | DLVRY              | VARCHAR(10)<br>Nullable | How the messages are delivered to the message queue used by the user. Contains one of the following values:<br><br><b>*BREAK</b> The job to which the message queue is assigned is interrupted when a message arrives on the message queue.<br><br><b>*DFT</b> Messages requiring replies are answered with their default reply.<br><br><b>*HOLD</b> The messages are held in the message queue until they are requested by the user or program.<br><br><b>*NOTIFY</b> The job to which the message queue is assigned is notified when a message arrives on the message queue. |
| MESSAGE_QUEUE_SEVERITY              | SEV                | SMALLINT<br>Nullable    | The lowest severity that a message can have and still be delivered to a user in break or notify mode                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| OUTPUT_QUEUE_NAME                   | OUTQ               | VARCHAR(10)<br>Nullable | The output queue used by this user. Can contain one of the following special values:<br><br><b>*DEV</b> An output queue with the same name as the device specified in the printer device parameter is used by the user.<br><br><b>*WRKSTN</b> The output queue assigned to the user's work station is used.                                                                                                                                                                                                                                                                    |
| OUTPUT_QUEUE_LIBRARY_NAME           | OUTQLIB            | VARCHAR(10)<br>Nullable | The name of the library where the output queue is located.<br><br>Contains null if the output queue name is *DEV or *WRKSTN.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| PRINT_DEVICE                        | PRTDEV             | VARCHAR(10)<br>Nullable | The printer used to print for this user. Can contain one of the following special values:<br><br><b>*SYSVAL</b> The default system printer specified in the system value QPRTDEV is used.<br><br><b>*WRKSTN</b> The printer assigned to the user's work station is used.                                                                                                                                                                                                                                                                                                       |
| SPECIAL_ENVIRONMENT                 | SPCENV             | VARCHAR(10)<br>Nullable | The special environment the user operates in after signing on. Contains one of the following special values:<br><br><b>*NONE</b> The user operates in the IBM i environment.<br><br><b>*SYSVAL</b> The system value QSPCENV is used to determine the user's special environment.<br><br><b>*S36</b> The user operates in the System/36 environment.                                                                                                                                                                                                                            |
| ATTENTION_KEY_HANDLING_PROGRAM_NAME | ATNPGM             | VARCHAR(10)<br>Nullable | The attention key handling program for this user. Can contain one of the following special values:<br><br><b>*NONE</b> No Attention-key-handling program is used.<br><br><b>*SYSVAL</b> The system value QATNPGM determines the user's Attention-key-handling program.                                                                                                                                                                                                                                                                                                         |

Table 179. USER\_INFO view (continued)

| Column Name                                     | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------|--------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ATTENTION_KEY_HANDLING_<br>PROGRAM_LIBRARY_NAME | ATNPGMLIB          | VARCHAR(10)<br>Nullable | The name of the library where the program is located. Can contain the special value *LIBL.<br><br>Contains null if the attention key handling program is *NONE or *SYSVAL.                                                                                                                                                                                                                                                                                                                            |
| LANGUAGE_ID                                     | LANGID             | VARCHAR(10)<br>Nullable | The language ID used by the system for this user. Can contain the following special value:<br><br><b>*SYSVAL</b> The system value QLANGID is used to determine the user's language ID.                                                                                                                                                                                                                                                                                                                |
| COUNTRY_OR_REGION_ID                            | CNTRYID            | VARCHAR(10)<br>Nullable | Country or region ID. Can contain the following special value:<br><br><b>*SYSVAL</b> The system value QCNTRYID is used to determine the user's country or region ID.                                                                                                                                                                                                                                                                                                                                  |
| CHARACTER_CODE_SET_ID                           | CCSID              | VARCHAR(6)<br>Nullable  | The CCSID for the user. Can contain the following special value:<br><br><b>QCCSID</b> The system value QCCSID is used to determine the user's character code set ID.                                                                                                                                                                                                                                                                                                                                  |
| USER_OPTIONS                                    | USROPT             | VARCHAR(77)<br>Nullable | A list of the options for users to customize their environment. Up to 7 options are returned. Each entry is padded with blanks to fill 11 characters.<br><br>Contains null if there are no user options.                                                                                                                                                                                                                                                                                              |
| SORT_SEQUENCE_TABLE_NAME                        | SRTSEQ             | VARCHAR(10)<br>Nullable | The name of the sort sequence table used for string comparisons. Can contain one of the following special values:<br><br><b>*HEX</b> The hexadecimal values of the characters are used to determine the sort sequence.<br><br><b>*LANGIDSHR</b> A shared-weight sort table associated with the language specified.<br><br><b>*LANGIDUNQ</b> A unique-weight sort table associated with the language specified.<br><br><b>*SYSVAL</b> The system value QSRTSEQ.                                        |
| SORT_SEQUENCE_TABLE_LIBRARY_NAME                | SRTSEQLIB          | VARCHAR(10)<br>Nullable | The name of the library that is used to locate the sort sequence table.<br><br>Contains null if the sort sequence table is a special value.                                                                                                                                                                                                                                                                                                                                                           |
| OBJECT_AUDITING_VALUE                           | OBJAUD             | VARCHAR(10)<br>Nullable | The object auditing value for this user. Contains one of the following values:<br><br><b>*ALL</b> Object read and change operations are audited for the current user if the object's auditing value is *USRPRF.<br><br><b>*CHANGE</b> Object changes are audited for the current user if the object's auditing value is *USRPRF.<br><br><b>*NONE</b> No additional object auditing is done for the current user.<br><br><b>*NOTAVL</b> The user is not allowed to retrieve the object auditing value. |

Table 179. USER\_INFO view (continued)

| Column Name                   | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USER_ACTION_AUDIT_LEVEL       | AUDLVL             | VARCHAR(363)<br>Nullable | The action audit values for this user. Up to 31 options are returned. Each entry is padded with blanks to fill 11 characters.<br><br>Contains null if there are no action values or if the caller is not authorized to retrieve the action audit level.                                                                                                                                                                                                                                                                                      |
| GROUP_AUTHORITY_TYPE          | GRPAUTYP           | VARCHAR(10)<br>Nullable  | The type of authority the user's group profile has to objects the user creates. Contains one of the following special values:<br><br><b>*PGP</b> The group profile will be the primary group for objects the user creates.<br><br><b>*PRIVATE</b> The group profile has a private authority to the objects the user creates. If the user does not have a group profile, this value is returned.                                                                                                                                              |
| USER_ID_NUMBER                | UID                | BIGINT<br>Nullable       | The user ID number for the user profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| GROUP_ID_NUMBER               | GID                | BIGINT<br>Nullable       | The group ID number for the user profile. The value 0 is returned if the user has no group ID number.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| LOCALE_JOB_ATTRIBUTES         | SETJOBATR          | VARCHAR(88)<br>Nullable  | A list of the job attributes that are taken from the user's locale path. This column contains a list of up to 8 items. Each entry is padded with blanks to fill 11 characters.                                                                                                                                                                                                                                                                                                                                                               |
| GROUP_MEMBER_INDICATOR        | GRPMBR             | VARCHAR(3)<br>Nullable   | Whether this user is a group that has members. Contains one of the following values:<br><br><b>NO</b> The user is not a group, or is a group but does not have any members.<br><br><b>YES</b> The user is a group that has members.                                                                                                                                                                                                                                                                                                          |
| DIGITAL_CERTIFICATE_INDICATOR | DCIND              | VARCHAR(3)<br>Nullable   | Whether there are digital certificates associated with this user. Contains one of the following values:<br><br><b>NO</b> There are no digital certificates associated with this user.<br><br><b>YES</b> There is at least one digital certificate associated with this user.                                                                                                                                                                                                                                                                 |
| CHARACTER_IDENTIFIER_CONTROL  | CHRIDCTL           | VARCHAR(10)<br>Nullable  | The character identifier control for the user. Can contain the following special values:<br><br><b>*DEVD</b> The *DEVD special value performs the same function as on the CHRID command parameter for display files, printer files, and panel groups.<br><br><b>*JOBCCSID</b> The *JOBCCSID special value performs the same function as on the CHRID command parameter for display files, printer files, and panel groups.<br><br><b>*SYSVAL</b> The value QCHRIDCTL system value will be used to determine the CHRID control for this user. |
| LOCAL_PASSWORD_MANAGEMENT     | LCLPDMGT           | VARCHAR(3)<br>Nullable   | Indicates if password is managed locally. Contains one of the following values:<br><br><b>NO</b> The password is not managed locally.<br><br><b>YES</b> The password is managed locally.                                                                                                                                                                                                                                                                                                                                                     |

Table 179. USER\_INFO view (continued)

| Column Name               | System Column Name | Data Type                                      | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------|--------------------|------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BLOCK_PASSWORD_CHANGE     | PWDCHGBLK          | VARCHAR(10)<br>Nullable                        | Specifies the time period, in hours, during which a password is blocked from being changed following the prior successful password change operation. Can contain one of the following special values:<br><br><b>*NONE</b> The password can be changed at any time.<br><br><b>*SYSVAL</b> The system value QPWDCHGBLK is used to determine the password change limit.                               |
| USER_ENTITLEMENT_REQUIRED | ENTITLERQD         | VARCHAR(3)<br>Nullable                         | Whether a user entitlement is required for this user profile. Contains one of the following values:<br><br><b>NO</b> A user entitlement is not required for this user profile.<br><br><b>YES</b> A user entitlement is required for this user profile.                                                                                                                                             |
| USER_EXPIRATION_INTERVAL  | USREXPITV          | SMALLINT<br>Nullable                           | The number of days (from 1 through 366) before the user profile is automatically disabled. The value 0 is returned if no expiration interval is defined.                                                                                                                                                                                                                                           |
| USER_EXPIRATION_DATE      | USREXPDATE         | TIMESTAMP<br>Nullable                          | The date when the user profile expires and is automatically disabled or deleted.<br><br>Contains null if the user profile will not expire.                                                                                                                                                                                                                                                         |
| USER_EXPIRATION_ACTION    | ACTION             | VARCHAR(8)<br>Nullable                         | The action that will occur when the user profile has expired. Contains one of the following values:<br><br><b>*DELETE</b> The user profile will be deleted. If the user profile cannot be deleted, it will be disabled.<br><br><b>*DISABLE</b> The user profile will be disabled.<br><br><b>*NONE</b> The user profile will not expire.                                                            |
| HOME_DIRECTORY            | HOMEDIR            | VARGRAPHIC(1024)<br>CCSID 1200<br><br>Nullable | The home directory for this user profile.                                                                                                                                                                                                                                                                                                                                                          |
| LOCALE_PATH_NAME          | LOCALE             | VARGRAPHIC(1024)<br>CCSID 1200<br><br>Nullable | The locale path name that is assigned to the user profile when a job is started. Can contain one of the following special values:<br><br><b>*C</b> The C locale path name is assigned.<br><br><b>*NONE</b> No locale path name is assigned.<br><br><b>*POSIX</b> The POSIX locale path name is assigned.<br><br><b>*SYSVAL</b> The QLOCALE system value is used to determine the locale path name. |
| USER_DEFAULT_PASSWORD     | DFTPWD             | VARCHAR(3)<br>Nullable                         | The password is the default password.<br><br><b>NO</b> The password is not the default password.<br><br><b>YES</b> The password appears to be the default password since it matches the user profile name.<br><br>Contains null if not authorized to view this information.                                                                                                                        |
| USER_OWNER                | USER_OWNER         | VARCHAR(10)<br>Nullable                        | The user profile that owns this user profile.                                                                                                                                                                                                                                                                                                                                                      |

Table 179. USER\_INFO view (continued)

| Column Name                            | System Column Name | Data Type                               | Description                                                                                                                                                                                                                                                                             |
|----------------------------------------|--------------------|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USER_CREATOR                           | CREATOR            | VARCHAR(10)<br>Nullable                 | The user profile that created this user profile.                                                                                                                                                                                                                                        |
| SIZE                                   | SIZE               | DECIMAL(15,0)<br>Nullable               | Size of the user profile, in bytes.                                                                                                                                                                                                                                                     |
| CREATION_TIMESTAMP                     | TIMESTAMP          | TIMESTAMP<br>Nullable                   | Timestamp of when the user profile was created.                                                                                                                                                                                                                                         |
| LAST_USED_TIMESTAMP                    | LASTUSED           | TIMESTAMP<br>Nullable                   | The date the user profile was used last. The time portion of the timestamp will always be 0.                                                                                                                                                                                            |
| DAYS_USED_COUNT                        | DAYSUSED           | INTEGER<br>Nullable                     | The number of days the user profile has been used on the system.                                                                                                                                                                                                                        |
| LAST_RESET_TIMESTAMP                   | LASTRESET          | TIMESTAMP<br>Nullable                   | The date when the days used count was last reset to zero. The time portion of the timestamp will always be 0.                                                                                                                                                                           |
| AUTHORITY_COLLECTION_ACTIVE            | AUTCOLACT          | VARCHAR(3)                              | Whether authority collection is active for this user.<br><br><b>NO</b> Authority collection is not active for this user.<br><br><b>YES</b> Authority collection is active for this user.                                                                                                |
| AUTHORITY_COLLECTION_REPOSITORY_EXISTS | AUTCOLREP          | VARCHAR(3)                              | Whether an authority collection repository exists for this user.<br><br><b>NO</b> An authority collection repository does not exist for this user.<br><br><b>YES</b> An authority collection repository exists for this user.                                                           |
| PASE_SHELL_PATH                        | SHELL_PATH         | VARCHAR(1024)<br>CCSID 1208<br>Nullable | Path to the user's PASE shell. If AUTHORIZATION_NAME is QSYS, this column contains the default shell path used for all user profiles that have not had a value explicitly set.<br><br>Returns the null value if a value has not been set using the QSYS2.SET_PASE_SHELL_INFO procedure. |

## Example

Determine which users are having trouble signing on.

```
SELECT * FROM QSYS2.USER_INFO
WHERE SIGN_ON_ATTEMPTS_NOT_VALID > 0
```

## USER\_INFO\_BASIC view

The USER\_INFO\_BASIC view contains information about user profiles.

This view returns a subset of columns that are returned by the USER\_INFO view. It does not return: USER\_OWNER, USER\_CREATOR, SIZE, CREATION\_TIMESTAMP, LAST\_USED\_TIMESTAMP, DAYS\_USED\_COUNT, or LAST\_RESET\_TIMESTAMP. USER\_INFO\_BASIC typically performs much better than USER\_INFO.

The values returned for the columns in the view are closely related to the values returned by [Retrieve User Information \(QSYRUSRI\) API](#). Refer to the API for more detailed information.

**Authorization:** Only \*USRPRF objects that the caller has \*OBJOPR and \*READ authority to will be returned. To see a non-null value for the USER\_DEFAULT\_PASSWORD column, the caller must have \*ALLOBJ and \*SECADM authority.



The following table describes the columns in the view. The system name is USER\_INFOB. The schema is QSYS2.

Table 180. USER\_INFO\_BASIC view

| Column Name                | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                              |
|----------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUTHORIZATION_NAME         | USER_NAME          | VARCHAR(10)<br>Nullable | User profile name.                                                                                                                                                                                                                                                       |
| PREVIOUS_SIGNON            | PRVSIGNON          | TIMESTAMP<br>Nullable   | The date and time the user last signed on.<br>Contains null if the profile has never been used to sign on.                                                                                                                                                               |
| SIGN_ON_ATTEMPTS_NOT_VALID | SIGNONINV          | INTEGER<br>Nullable     | The number of sign-on attempts that were not valid since the last successful sign-on.                                                                                                                                                                                    |
| STATUS                     | STATUS             | VARCHAR(10)<br>Nullable | The status of the user profile. Contains one of the following values:<br><br><b>*DISABLED</b> The user profile is disabled; therefore, the user cannot sign on.<br><br><b>*ENABLED</b> The user profile is enabled; therefore, the user is able to sign on.              |
| NETSERVER_DISABLED         | NETSERVER          | VARCHAR(3)              | Whether this user profile is disabled for IBM i NetServer use.<br><br><b>NO</b> The user profile is not disabled for IBM i NetServer use.<br><br><b>YES</b> The user profile is disabled for IBM i NetServer use.                                                        |
| PASSWORD_CHANGE_DATE       | PWDCHGDAT          | TIMESTAMP<br>Nullable   | The date the user's password was last changed.                                                                                                                                                                                                                           |
| NO_PASSWORD_INDICATOR      | NOPWD              | VARCHAR(3)<br>Nullable  | Indicates whether *NONE is specified for the password in the user profile.<br><br><b>NO</b> The password in the user profile is not *NONE.<br><br><b>YES</b> The password in the user profile is *NONE.                                                                  |
| PASSWORD_LEVEL_0_1         | PWD_0_1            | VARCHAR(3)<br>Nullable  | Indicates whether the user profile has a password that can be used for a system at QPWDLVL 0 or 1.<br><br><b>NO</b> The password cannot be used.<br><br><b>YES</b> The password can be used.<br><br>Contains the null value if NO_PASSWORD_INDICATOR has a value of YES. |
| PASSWORD_LEVEL_2_3         | PWD_2_3            | VARCHAR(3)<br>Nullable  | Indicates whether the user profile has a password that can be used for a system at QPWDLVL 2 or 3.<br><br><b>NO</b> The password cannot be used.<br><br><b>YES</b> The password can be used.<br><br>Contains the null value if NO_PASSWORD_INDICATOR has a value of YES. |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                  | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|--------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PASSWORD_EXPIRATION_INTERVAL | PWDEXPITV          | SMALLINT<br>Nullable     | The number of days (from 1 through 366) the user's password can remain active before it must be changed. Can also be one of the following values:<br><br><b>0</b> The system value QPWDEXPITV is used to determine the user's password expiration interval.<br><br><b>-1</b> The user's password does not expire (*NOMAX).                                                      |
| DATE_PASSWORD_EXPIRES        | PWDEXPDAT          | TIMESTAMP<br>Nullable    | The date the user's password expires.<br>Contains null if the password will not expire.                                                                                                                                                                                                                                                                                         |
| DAYS_UNTIL_PASSWORD_EXPIRES  | PWDDAYSEXP         | INTEGER<br>Nullable      | The number of days until the password will expire. A value of 0 indicates the password is expired.<br>Contains null if the password will not expire within the number of days specified by the password expiration warning (QPWDEXPWRN) system value.                                                                                                                           |
| SET_PASSWORD_TO_EXPIRE       | PWDEXP             | VARCHAR(3)<br>Nullable   | Indicates whether the user's password is set to expire, requiring the user to change the password when signing on. Contains one of the following values:<br><br><b>NO</b> The user's password is not set to expire.<br><b>YES</b> The user's password is set to expire.                                                                                                         |
| USER_CLASS_NAME              | USRCLS             | VARCHAR(10)<br>Nullable  | The user's class name. Contains one of the following values:<br><br><b>*PGMR</b> The user has a class of programmer.<br><br><b>*SECADM</b> The user has a class of security administrator.<br><br><b>*SECOFR</b> The user has a class of security officer.<br><br><b>*SYSOPR</b> The user has a class of system operator.<br><br><b>*USER</b> The user has a class of end user. |
| SPECIAL_AUTHORITIES          | SPCAUT             | VARCHAR(88)<br>Nullable  | A list of the special authorities the user has. Up to 8 authorities are returned. Each entry is padded with blanks to fill 11 characters.<br>Contains null if the user has no special authorities.                                                                                                                                                                              |
| GROUP_PROFILE_NAME           | GRPPRF             | VARCHAR(10)<br>Nullable  | The name of the group profile. Contains the value *NONE if the user does not have a group profile.                                                                                                                                                                                                                                                                              |
| SUPPLEMENTAL_GROUP_COUNT     | SUPGRPCNT          | SMALLINT                 | The number of supplemental groups in the SUPPLEMENTAL_GROUP_LIST column.                                                                                                                                                                                                                                                                                                        |
| SUPPLEMENTAL_GROUP_LIST      | SUPGRPLIST         | VARCHAR(150)<br>Nullable | A list of supplemental groups for the user profile. Up to 15 supplemental groups are returned. Each entry except for the last one is padded with blanks to fill 10 characters.<br>Contains null if the user has no supplemental groups.                                                                                                                                         |
| OWNER                        | OWNER              | VARCHAR(10)<br>Nullable  | This field indicates who is to own objects created by this user. Contains one of the following values:<br><br><b>*GRPPRF</b> The user's group profile owns any objects the user creates.<br><br><b>*USRPRF</b> The user owns any objects the user creates.                                                                                                                      |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                  | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GROUP_AUTHORITY              | GRPAUT             | VARCHAR(10)<br>Nullable | The authority the user's group profile has to objects the user creates. Contains one of the following values:<br><br><b>*ALL</b> The group profile has all authority to the objects the user creates.<br><br><b>*CHANGE</b> The group profile has change authority to the objects the user creates.<br><br><b>*EXCLUDE</b> The group profile has exclude authority to the objects the user creates.<br><br><b>*NONE</b> The group profile has no authority to the objects the user creates. If the user does not have a group profile, this value is returned.<br><br><b>*USE</b> The group profile has use authority to the objects the user creates. |
| ASSISTANCE_LEVEL             | ASTLVL             | VARCHAR(10)<br>Nullable | The user interface that the user will use. Contains one of the following values:<br><br><b>*ADVANCED</b> The expert system user interface.<br><br><b>*BASIC</b> The Operational Assistant user interface.<br><br><b>*INTERMED</b> The system user interface.<br><br><b>*SYSVAL</b> The system value QASTLVL determines which user interface the user is using.                                                                                                                                                                                                                                                                                         |
| CURRENT_LIBRARY_NAME         | CURLIB             | VARCHAR(10)<br>Nullable | The name of the user's current library. Contains *CRTDFT if the user does not have a current library.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| INITIAL_MENU_NAME            | INLMNU             | VARCHAR(10)<br>Nullable | The initial menu for the user. Can contain the special value *SIGNOFF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| INITIAL_MENU_LIBRARY_NAME    | INLMNULIB          | VARCHAR(10)<br>Nullable | The name of the library that the initial menu is in. Can contain the special value *LIBL. Contains null if the menu name is *SIGNOFF.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| INITIAL_PROGRAM_NAME         | INITPGM            | VARCHAR(10)<br>Nullable | The initial program for the user. Contains *NONE if there is no initial program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| INITIAL_PROGRAM_LIBRARY_NAME | INITPGMLIB         | VARCHAR(10)<br>Nullable | The name of the library that the initial program is in. Can contain the special value *LIBL. Contains null if the initial program name is *NONE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| LIMIT_CAPABILITIES           | LMTCPB             | VARCHAR(10)<br>Nullable | Indicates whether the user has limited capabilities. Contains one of the following values:<br><br><b>*NO</b> The user is not limited.<br><br><b>*PARTIAL</b> The user cannot change his initial program or current library.<br><br><b>*YES</b> The user cannot change his initial menu, initial program, or current library. The user cannot run commands from the command line.                                                                                                                                                                                                                                                                       |
| TEXT_DESCRIPTION             | TEXT               | VARCHAR(50)<br>Nullable | The descriptive text for the user profile.<br><br>Contains null if the user profile has no text description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                  | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DISPLAY_SIGNON_INFORMATION   | DSPSGNINF          | VARCHAR(10)<br>Nullable | Indicates whether the sign-on information display is shown when the user signs on. Contains one of the following special values:<br><br><b>*NO</b> The sign-on information display is not shown when the user signs on.<br><br><b>*SYSVAL</b> The system value QDSPSGNINF determines if the sign-on information display is shown when the user signs on.<br><br><b>*YES</b> The sign-on information display is shown when the user signs on.                                                                    |
| LIMIT_DEVICE_SESSIONS        | LMTDEVSSN          | VARCHAR(10)<br>Nullable | Specifies if the number of device sessions allowed for a user is limited. Can contain one of the following special values:<br><br><b>*NO</b> The user is not limited to a specific number of device sessions.<br><br><b>*SYSVAL</b> The system value QLMTDEVSSN determines if the user is limited to a specific number of device sessions.<br><br><b>*YES</b> The user is limited to one device session.                                                                                                        |
| KEYBOARD_BUFFERING           | KBDBUF             | VARCHAR(10)<br>Nullable | The keyboard buffering value that is used when a job is initialized for this user. Contains one of the following special values:<br><br><b>*NO</b> The type-ahead and attention-key buffering options are not on.<br><br><b>*SYSVAL</b> The system value QKBDBUF determines the keyboard buffering value for this user.<br><br><b>*TYPEAHEAD</b> The type-ahead option is on, but the attention-key buffering option is not.<br><br><b>*YES</b> The type-ahead and attention-key buffering options are both on. |
| MAXIMUM_ALLOWED_STORAGE      | MAXSTGLRG          | BIGINT<br>Nullable      | The maximum amount of auxiliary storage (in kilobytes) that can be assigned to store permanent objects owned by the user. Contains null if the user has no maximum storage.                                                                                                                                                                                                                                                                                                                                     |
| STORAGE_USED                 | STGUSED            | BIGINT<br>Nullable      | The amount of auxiliary storage (in kilobytes) occupied by this user's owned objects on *SYSBAS. The QSYS2.USER_STORAGE catalog should be used to determine the storage consumed on all ASPs.                                                                                                                                                                                                                                                                                                                   |
| HIGHEST_SCHEDULING_PRIORITY  | PTYLMT             | CHAR(1)<br>Nullable     | The highest scheduling priority the user is allowed to have for each job submitted to the system.                                                                                                                                                                                                                                                                                                                                                                                                               |
| JOB_DESCRIPTION_NAME         | JOBID              | VARCHAR(10)<br>Nullable | The name of the job description used for jobs that start through subsystem work station entries.                                                                                                                                                                                                                                                                                                                                                                                                                |
| JOB_DESCRIPTION_LIBRARY_NAME | JOBDLIB            | VARCHAR(10)<br>Nullable | Job description library name. Can contain the special value *LIBL.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| ACCOUNTING_CODE              | ACGCDE             | VARCHAR(15)<br>Nullable | The accounting code that is associated with this user.<br><br>Contains null if there is no accounting code.                                                                                                                                                                                                                                                                                                                                                                                                     |
| MESSAGE_QUEUE_NAME           | MSGQ               | VARCHAR(10)<br>Nullable | The name of the message queue that is used by this user.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                         | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MESSAGE_QUEUE_LIBRARY_NAME          | MSGQLIB            | VARCHAR(10)<br>Nullable | The name of the library the message queue is in. Can contain the special value *LIBL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| MESSAGE_QUEUE_DELIVERY_METHOD       | DLVRY              | VARCHAR(10)<br>Nullable | How the messages are delivered to the message queue used by the user. Contains one of the following values:<br><br><b>*BREAK</b> The job to which the message queue is assigned is interrupted when a message arrives on the message queue.<br><br><b>*DFT</b> Messages requiring replies are answered with their default reply.<br><br><b>*HOLD</b> The messages are held in the message queue until they are requested by the user or program.<br><br><b>*NOTIFY</b> The job to which the message queue is assigned is notified when a message arrives on the message queue. |
| MESSAGE_QUEUE_SEVERITY              | SEV                | SMALLINT<br>Nullable    | The lowest severity that a message can have and still be delivered to a user in break or notify mode                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| OUTPUT_QUEUE_NAME                   | OUTQ               | VARCHAR(10)<br>Nullable | The output queue used by this user. Can contain one of the following special values:<br><br><b>*DEV</b> An output queue with the same name as the device specified in the printer device parameter is used by the user.<br><br><b>*WRKSTN</b> The output queue assigned to the user's work station is used.                                                                                                                                                                                                                                                                    |
| OUTPUT_QUEUE_LIBRARY_NAME           | OUTQLIB            | VARCHAR(10)<br>Nullable | The name of the library where the output queue is located.<br><br>Contains null if the output queue name is *DEV or *WRKSTN.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| PRINT_DEVICE                        | PRTDEV             | VARCHAR(10)<br>Nullable | The printer used to print for this user. Can contain one of the following special values:<br><br><b>*SYSVAL</b> The default system printer specified in the system value QPRTDEV is used.<br><br><b>*WRKSTN</b> The printer assigned to the user's work station is used.                                                                                                                                                                                                                                                                                                       |
| SPECIAL_ENVIRONMENT                 | SPCENV             | VARCHAR(10)<br>Nullable | The special environment the user operates in after signing on. Contains one of the following special values:<br><br><b>*NONE</b> The user operates in the IBM i environment.<br><br><b>*SYSVAL</b> The system value QSPCENV is used to determine the user's special environment.<br><br><b>*S36</b> The user operates in the System/36 environment.                                                                                                                                                                                                                            |
| ATTENTION_KEY_HANDLING_PROGRAM_NAME | ATNPGM             | VARCHAR(10)<br>Nullable | The attention key handling program for this user. Can contain one of the following special values:<br><br><b>*NONE</b> No Attention-key-handling program is used.<br><br><b>*SYSVAL</b> The system value QATNPGM determines the user's Attention-key-handling program.                                                                                                                                                                                                                                                                                                         |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                                 | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------------------|--------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ATTENTION_KEY_HANDLING_PROGRAM_LIBRARY_NAME | ATNPGMLIB          | VARCHAR(10)<br>Nullable | The name of the library where the program is located. Can contain the special value *LIBL.<br><br>Contains null if the attention key handling program is *NONE or *SYSVAL.                                                                                                                                                                                                                                                                                                                            |
| LANGUAGE_ID                                 | LANGID             | VARCHAR(10)<br>Nullable | The language ID used by the system for this user. Can contain the following special value:<br><br><b>*SYSVAL</b> The system value QLANGID is used to determine the user's language ID.                                                                                                                                                                                                                                                                                                                |
| COUNTRY_OR_REGION_ID                        | CNTRYID            | VARCHAR(10)<br>Nullable | Country or region ID. Can contain the following special value:<br><br><b>*SYSVAL</b> The system value QCNTRYID is used to determine the user's country or region ID.                                                                                                                                                                                                                                                                                                                                  |
| CHARACTER_CODE_SET_ID                       | CCSID              | VARCHAR(6)<br>Nullable  | The CCSID for the user. Can contain the following special value:<br><br><b>QCCSID</b> The system value QCCSID is used to determine the user's character code set ID.                                                                                                                                                                                                                                                                                                                                  |
| USER_OPTIONS                                | USROPT             | VARCHAR(77)<br>Nullable | A list of the options for users to customize their environment. Up to 7 options are returned. Each entry is padded with blanks to fill 11 characters.<br><br>Contains null if there are no user options.                                                                                                                                                                                                                                                                                              |
| SORT_SEQUENCE_TABLE_NAME                    | SRTSEQ             | VARCHAR(10)<br>Nullable | The name of the sort sequence table used for string comparisons. Can contain one of the following special values:<br><br><b>*HEX</b> The hexadecimal values of the characters are used to determine the sort sequence.<br><br><b>*LANGIDSHR</b> A shared-weight sort table associated with the language specified.<br><br><b>*LANGIDUNQ</b> A unique-weight sort table associated with the language specified.<br><br><b>*SYSVAL</b> The system value QSRTSEQ.                                        |
| SORT_SEQUENCE_TABLE_LIBRARY_NAME            | SRTSEQLIB          | VARCHAR(10)<br>Nullable | The name of the library that is used to locate the sort sequence table.<br><br>Contains null if the sort sequence table is a special value.                                                                                                                                                                                                                                                                                                                                                           |
| OBJECT_AUDITING_VALUE                       | OBJAUD             | VARCHAR(10)<br>Nullable | The object auditing value for this user. Contains one of the following values:<br><br><b>*ALL</b> Object read and change operations are audited for the current user if the object's auditing value is *USRPRF.<br><br><b>*CHANGE</b> Object changes are audited for the current user if the object's auditing value is *USRPRF.<br><br><b>*NONE</b> No additional object auditing is done for the current user.<br><br><b>*NOTAVL</b> The user is not allowed to retrieve the object auditing value. |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                   | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USER_ACTION_AUDIT_LEVEL       | AUDLVL             | VARCHAR(363)<br>Nullable | The action audit values for this user. Up to 31 options are returned. Each entry is padded with blanks to fill 11 characters.<br><br>Contains null if there are no action values or if the caller is not authorized to retrieve the action audit level.                                                                                                                                                                                                                                                                                      |
| GROUP_AUTHORITY_TYPE          | GRPAUTYP           | VARCHAR(10)<br>Nullable  | The type of authority the user's group profile has to objects the user creates. Contains one of the following special values:<br><br><b>*PGP</b> The group profile will be the primary group for objects the user creates.<br><br><b>*PRIVATE</b> The group profile has a private authority to the objects the user creates. If the user does not have a group profile, this value is returned.                                                                                                                                              |
| USER_ID_NUMBER                | UID                | BIGINT<br>Nullable       | The user ID number for the user profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| GROUP_ID_NUMBER               | GID                | BIGINT<br>Nullable       | The group ID number for the user profile. The value 0 is returned if the user has no group ID number.                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| LOCALE_JOB_ATTRIBUTES         | SETJOBATR          | VARCHAR(88)<br>Nullable  | A list of the job attributes that are taken from the user's locale path. This column contains a list of up to 8 items. Each entry is padded with blanks to fill 11 characters.                                                                                                                                                                                                                                                                                                                                                               |
| GROUP_MEMBER_INDICATOR        | GRPMBR             | VARCHAR(3)<br>Nullable   | Whether this user is a group that has members. Contains one of the following values:<br><br><b>NO</b> The user is not a group, or is a group but does not have any members.<br><br><b>YES</b> The user is a group that has members.                                                                                                                                                                                                                                                                                                          |
| DIGITAL_CERTIFICATE_INDICATOR | DCIND              | VARCHAR(3)<br>Nullable   | Whether there are digital certificates associated with this user. Contains one of the following values:<br><br><b>NO</b> There are no digital certificates associated with this user.<br><br><b>YES</b> There is at least one digital certificate associated with this user.                                                                                                                                                                                                                                                                 |
| CHARACTER_IDENTIFIER_CONTROL  | CHRIDCTL           | VARCHAR(10)<br>Nullable  | The character identifier control for the user. Can contain the following special values:<br><br><b>*DEVD</b> The *DEVD special value performs the same function as on the CHRID command parameter for display files, printer files, and panel groups.<br><br><b>*JOBCCSID</b> The *JOBCCSID special value performs the same function as on the CHRID command parameter for display files, printer files, and panel groups.<br><br><b>*SYSVAL</b> The value QCHRIDCTL system value will be used to determine the CHRID control for this user. |
| LOCAL_PASSWORD_MANAGEMENT     | LCLPDMGT           | VARCHAR(3)<br>Nullable   | Indicates if password is managed locally. Contains one of the following values:<br><br><b>NO</b> The password is not managed locally.<br><br><b>YES</b> The password is managed locally.                                                                                                                                                                                                                                                                                                                                                     |

Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                 | System Column Name | Data Type                                      | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------|--------------------|------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BLOCK_PASSWORD_CHANGE       | PWDCHGBLK          | VARCHAR(10)<br>Nullable                        | Specifies the time period, in hours, during which a password is blocked from being changed following the prior successful password change operation. Can contain one of the following special values:<br><br><b>*NONE</b> The password can be changed at any time.<br><br><b>*SYSVAL</b> The system value QPWDCHGBLK is used to determine the password change limit.                               |
| USER_ENTITLEMENT_REQUIRED   | ENTITLERQD         | VARCHAR(3)<br>Nullable                         | Whether a user entitlement is required for this user profile. Contains one of the following values:<br><br><b>NO</b> A user entitlement is not required for this user profile.<br><br><b>YES</b> A user entitlement is required for this user profile.                                                                                                                                             |
| USER_EXPIRATION_INTERVAL    | USREXPITV          | SMALLINT<br>Nullable                           | The number of days (from 1 through 366) before the user profile is automatically disabled. The value 0 is returned if no expiration interval is defined.                                                                                                                                                                                                                                           |
| USER_EXPIRATION_DATE        | USREXPDATE         | TIMESTAMP<br>Nullable                          | The date when the user profile expires and is automatically disabled or deleted.<br><br>Contains null if the user profile will not expire.                                                                                                                                                                                                                                                         |
| USER_EXPIRATION_ACTION      | ACTION             | VARCHAR(8)<br>Nullable                         | The action that will occur when the user profile has expired. Contains one of the following values:<br><br><b>*DELETE</b> The user profile will be deleted. If the user profile cannot be deleted, it will be disabled.<br><br><b>*DISABLE</b> The user profile will be disabled.<br><br><b>*NONE</b> The user profile will not expire.                                                            |
| HOME_DIRECTORY              | HOMEDIR            | VARGRAPHIC(1024)<br>CCSID 1200<br><br>Nullable | The home directory for this user profile.                                                                                                                                                                                                                                                                                                                                                          |
| LOCALE_PATH_NAME            | LOCALE             | VARGRAPHIC(1024)<br>CCSID 1200<br><br>Nullable | The locale path name that is assigned to the user profile when a job is started. Can contain one of the following special values:<br><br><b>*C</b> The C locale path name is assigned.<br><br><b>*NONE</b> No locale path name is assigned.<br><br><b>*POSIX</b> The POSIX locale path name is assigned.<br><br><b>*SYSVAL</b> The QLOCALE system value is used to determine the locale path name. |
| USER_DEFAULT_PASSWORD       | DFTPWD             | VARCHAR(3)<br>Nullable                         | The password is the default password.<br><br><b>NO</b> The password is not the default password.<br><br><b>YES</b> The password appears to be the default password since it matches the user profile name.<br><br>Contains null if not authorized to view this information.                                                                                                                        |
| AUTHORITY_COLLECTION_ACTIVE | AUTCOLACT          | VARCHAR(3)                                     | Whether authority collection is active for this user.<br><br><b>NO</b> Authority collection is not active for this user.<br><br><b>YES</b> Authority collection is active for this user.                                                                                                                                                                                                           |



Table 180. USER\_INFO\_BASIC view (continued)

| Column Name                            | System Column Name | Data Type                               | Description                                                                                                                                                                                                                                                                                    |
|----------------------------------------|--------------------|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUTHORITY_COLLECTION_REPOSITORY_EXISTS | AUTCOLREP          | VARCHAR(3)                              | <p>Whether an authority collection repository exists for this user.</p> <p><b>NO</b> An authority collection repository does not exist for this user.</p> <p><b>YES</b> An authority collection repository exists for this user.</p>                                                           |
| PASE_SHELL_PATH                        | SHELL_PATH         | VARCHAR(1024)<br>CCSID 1208<br>Nullable | <p>Path to the user's PASE shell. If AUTHORIZATION_NAME is QSYS, this column contains the default shell path used for all user profiles that have not had a value explicitly set.</p> <p>Returns the null value if a value has not been set using the QSYS2.SET_PASE_SHELL_INFO procedure.</p> |

### Example

Determine which users have \*ALLOBJ special authority.

```
SELECT * FROM QSYS2.USER_INFO_BASIC
WHERE SPECIAL_AUTHORITIES LIKE '%*ALLOBJ%';
```

## Spool Services

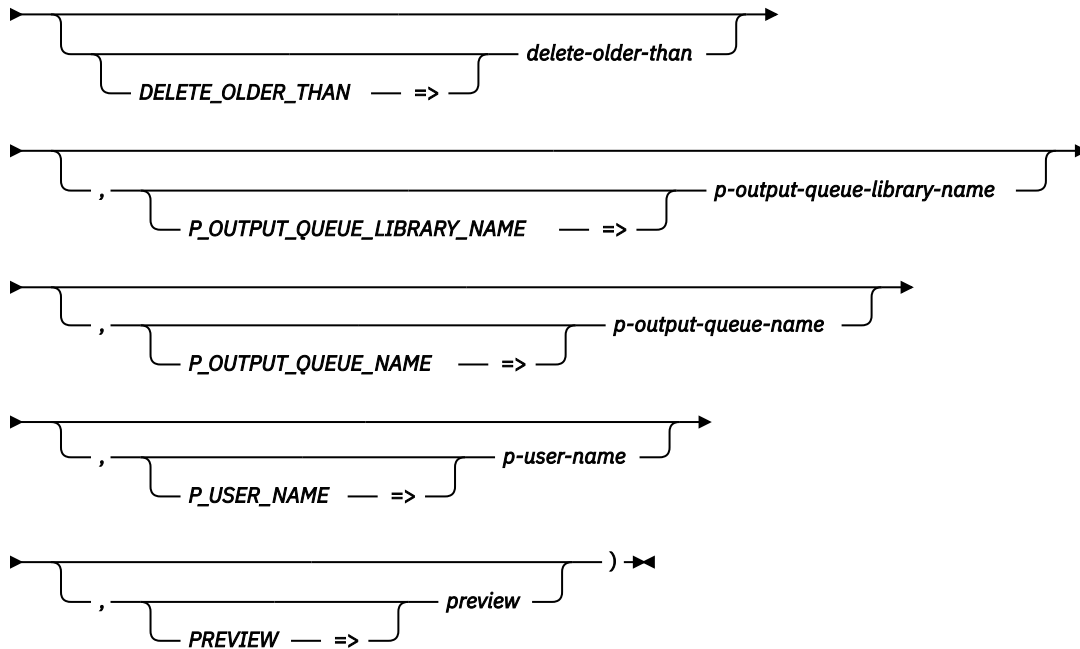
These views and functions provide information about spooled files.

### DELETE\_OLD\_SPOOLED\_FILES procedure

The DELETE\_OLD\_SPOOLED\_FILES procedure deletes spooled files according to filtering criteria. It can optionally return a preview of the files that meet the filtering criteria without performing the delete.

**Authorization:** The user must have the authorizations required to use the QSYS2.OUTPUT\_QUEUE\_ENTRIES\_BASIC view and the DLTSPF CL command. Without sufficient authority, the procedure appears to run successfully but no spooled files are processed.

➤ DELETE\_OLD\_SPOOLED\_FILES — ( →



The schema is QSYS2.

- delete-older-than*** A timestamp value that defines the starting point for deleting spooled files. Any spooled file older than this timestamp is eligible for deletion. The default is CURRENT TIMESTAMP - 3 MONTHS.
- p-output-queue-library-name*** A character or graphic string that specifies the name of a library. Any spooled file in any output queue in this library is eligible for deletion. The default is \*ALL.
- p-output-queue-name*** A character or graphic string that specifies an output queue name. Any spooled file in this output queue is eligible for deletion. The default is \*ALL.
- p-user-name*** A character or graphic string that specifies the name of a user whose spooled files are to be deleted. Any spooled file with this user name is eligible for deletion. The default is \*ALL.
- preview*** A character or graphic string that indicates whether the identified spooled files should be deleted or returned as a result set.
  - NO** The spooled files will be deleted. This is the default.
  - YES** A result set list of spooled files will be returned. No files will be deleted.

## Note

This procedure is provided in the SYSTOOLS schema as an example of how to delete spooled files and how to return a result set using an SQL procedure. Creating customized versions of this procedure to better suit a specific need is encouraged. Use the Insert Generated SQL feature in IBM i Access Client Solutions (ACS) to extract the source for this procedure. Then modify it and create a new procedure in a user-specified schema.

## Example

- List all the spooled files in PRT01 that are older than 30 days.

```
CALL SYSTOOLS.DELETE_OLD_SPOOLED_FILES(DELETE_OLDER_THAN => CURRENT DATE - 30 DAYS,
 P_OUTPUT_QUEUE_NAME => 'PRT01',
 PREVIEW => 'YES');
```

- Delete all the spooled files in PRT01 that are older than 30 days.

```
CALL SYSTOOLS.DELETE_OLD_SPOOLED_FILES(DELETE_OLDER_THAN => CURRENT DATE - 30 DAYS,
 P_OUTPUT_QUEUE_NAME => 'PRT01',
 PREVIEW => 'NO');
```

## GENERATE\_PDF scalar function

The GENERATE\_PDF scalar function generates a PDF file in the Integrated File System containing the content of a spooled file.

This function requires the following product: 5770TS1 - Option 1 - Transform Services - AFP to PDF Transform

**Authorization:** This scalar function uses the Copy Spooled File (CPYSPLF) CL command. Any authority requirements for the CL command apply to the use of this function.

```
➤ GENERATE_PDF ((job-name →
 { JOB_NAME — => }
 , (spooled-file-name →
 { SPOOLED_FILE_NAME — => }
 , (spooled-file-number →
 { SPOOLED_FILE_NUMBER — => }
 , (path-name —) ➤
 { PATH_NAME — => }
```

The schema is SYSTOOLS.

***job-name*** A character string containing a qualified job name. Can contain the following special value:

- \* Use the name of the current job.

***spooled-file-name*** A character string containing the name of the spooled file.

***spooled-file-number*** The number of the spooled file. Can contain the following special value:

- \***LAST** The spooled file with the highest number with a name matching *spooled-file-name* is selected.

***path-name*** A character string containing the name of the path where the result PDF file is to be written.

The result of the function is an integer. If the command is successful, the function returns a value of 1. If the command returns an error, the function returns a value of -1.

## Note

This function is provided in the SYSTOOLS schema as an example of how spooled file data can be converted to PDF format by embedding the CPYSPLF CL command in an SQL scalar function. Similar to other Db2 for i provided tools within SYSTOOLS, the SQL source can be extracted and used as a model for building similar helper functions, or to create a customized version within a user-specified schema.

## Example

Convert the specified spooled file to PDF format and save it as /usr/listing1.

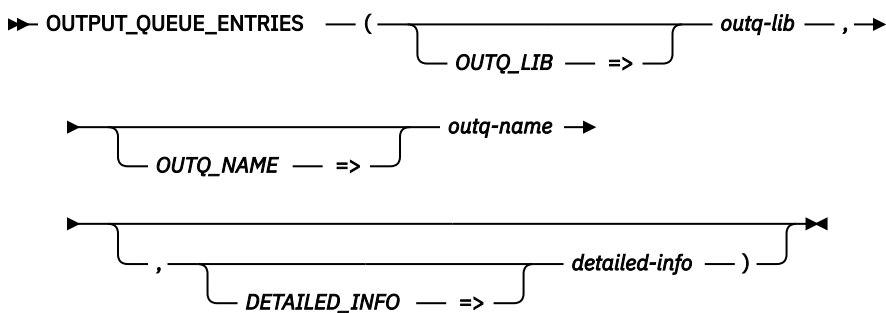
```
VALUES SYSTOOLS.GENERATE_PDF(
 JOB_NAME => '908049/QUSER/QZDASOINIT',
 SPOOLED_FILE_NAME => 'PGMA',
 SPOOLED_FILE_NUMBER => 2,
 PATH_NAME => '/usr/listing1');
```

## OUTPUT\_QUEUE\_ENTRIES table function

The OUTPUT\_QUEUE\_ENTRIES table function returns one row for each spooled file in an output queue.

**Authorization:** The caller must have:

- Execute authority to the output queue library and
  - Read authority to the output queue object, or
  - \*JOBCTL special authority and the output queue has OPRCTL(\*YES), or
  - \*SPLCTL special authority



The schema is QSYS2.

**outq-lib** A character or graphic string expression that identifies the name of the library containing *outq-name*. If this parameter is blank, the default of \*LIBL is used.

**outq-name** A character or graphic string expression that identifies the name of an output queue.

**detailed-info** A character or graphic string expression that indicates the type of information to be returned.

**YES** All the information available for the output queue is returned.

**NO** Only the general information is returned for the output queue. This is the information in the columns prior to the ACCOUNTING\_CODE column. This is the default.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 181. OUTPUT\_QUEUE\_ENTRIES table function

| Column Name       | Data Type   | Description                                                                                                                              |
|-------------------|-------------|------------------------------------------------------------------------------------------------------------------------------------------|
| CREATE_TIMESTAMP  | TIMESTAMP   | The timestamp when the file was created.                                                                                                 |
| SPOOLED_FILE_NAME | VARCHAR(10) | The file name that was specified by the user program when the file was created, or the name of the device file used to create this file. |
| USER_NAME         | VARCHAR(10) | The name of the user profile that produced the file.                                                                                     |
| USER_DATA         | VARCHAR(10) | The user-specified data that describes this file. Contains null if there is no user-specified data.                                      |

Table 181. OUTPUT\_QUEUE\_ENTRIES table function (continued)

| Column Name                                                                                                                                       | Data Type                                                   | Description                                                                                                                                                                   |                                                                                                                                                  |
|---------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| STATUS                                                                                                                                            | VARCHAR(15)                                                 | Status of the spooled file.                                                                                                                                                   |                                                                                                                                                  |
|                                                                                                                                                   |                                                             | <b>CLOSED</b>                                                                                                                                                                 | The file has been completely processed by a program but SCHEDULE(*JOBEND) was specified and the job that produced the file has not yet finished. |
|                                                                                                                                                   |                                                             | <b>DEFERRED</b>                                                                                                                                                               | Printing of the file has been deferred.                                                                                                          |
|                                                                                                                                                   |                                                             | <b>DELETED</b>                                                                                                                                                                | The file has been deleted.                                                                                                                       |
|                                                                                                                                                   |                                                             | <b>HELD</b>                                                                                                                                                                   | The file has been held.                                                                                                                          |
|                                                                                                                                                   |                                                             | <b>MESSAGE WAITING</b>                                                                                                                                                        | This file has a message which needs a reply or an action to be taken.                                                                            |
|                                                                                                                                                   |                                                             | <b>OPEN</b>                                                                                                                                                                   | The file has not been completely processed and is not ready to be selected by a writer.                                                          |
|                                                                                                                                                   |                                                             | <b>PENDING</b>                                                                                                                                                                | The file is pending to be printed.                                                                                                               |
|                                                                                                                                                   |                                                             | <b>PRINTING</b>                                                                                                                                                               | The file has been completely sent to the printer but print complete status has not been sent back.                                               |
|                                                                                                                                                   |                                                             | <b>READY</b>                                                                                                                                                                  | The file is available to be written.                                                                                                             |
|                                                                                                                                                   |                                                             | <b>SAVED</b>                                                                                                                                                                  | The file has been printed and then saved. This file remains saved until it is released.                                                          |
| <b>SENDING</b>                                                                                                                                    | The file is being sent or has been sent to a remote system. |                                                                                                                                                                               |                                                                                                                                                  |
| <b>WRITING</b>                                                                                                                                    | This file is currently being produced by the writer.        |                                                                                                                                                                               |                                                                                                                                                  |
| SIZE                                                                                                                                              | INTEGER                                                     | The size of the spooled file, in kilobytes.                                                                                                                                   |                                                                                                                                                  |
| TOTAL_PAGES                                                                                                                                       | INTEGER                                                     | The total number of pages in the file.                                                                                                                                        |                                                                                                                                                  |
| COPIES                                                                                                                                            | SMALLINT                                                    | The number of copies remaining to print.                                                                                                                                      |                                                                                                                                                  |
| FORM_TYPE                                                                                                                                         | VARCHAR(10)                                                 | The type of form that should be loaded in the printer to print this file.                                                                                                     |                                                                                                                                                  |
| JOB_NAME                                                                                                                                          | VARCHAR(28)                                                 | The qualified job name that produced the file.                                                                                                                                |                                                                                                                                                  |
| DEVICE_TYPE                                                                                                                                       | VARCHAR(10)                                                 | The type of data stream used to represent the file.                                                                                                                           |                                                                                                                                                  |
|                                                                                                                                                   |                                                             | <b>*AFPDS</b>                                                                                                                                                                 | Advanced Function Presentation data stream                                                                                                       |
|                                                                                                                                                   |                                                             | <b>*AFPDSLINE</b>                                                                                                                                                             | AFPDS data mixed with 1403 line data                                                                                                             |
|                                                                                                                                                   |                                                             | <b>*IPDS</b>                                                                                                                                                                  | Intelligent printer data stream                                                                                                                  |
|                                                                                                                                                   |                                                             | <b>*LINE</b>                                                                                                                                                                  | 1403 line data                                                                                                                                   |
|                                                                                                                                                   |                                                             | <b>*SCS</b>                                                                                                                                                                   | Systems Network Architecture (SNA) character stream                                                                                              |
| <b>*USERASCII</b>                                                                                                                                 | ASCII data                                                  |                                                                                                                                                                               |                                                                                                                                                  |
| OUTPUT_PRIORITY                                                                                                                                   | SMALLINT                                                    | The priority of the spooled file.                                                                                                                                             |                                                                                                                                                  |
| FILE_NUMBER                                                                                                                                       | INTEGER                                                     | The spooled file number of the specified file.                                                                                                                                |                                                                                                                                                  |
| SYSTEM                                                                                                                                            | VARCHAR(8)                                                  | The name of the system where the job that created the spooled file ran.                                                                                                       |                                                                                                                                                  |
| <b>Values for the following columns are returned when the DETAILED_INFO parameter is YES. Otherwise, the columns will contain the null value.</b> |                                                             |                                                                                                                                                                               |                                                                                                                                                  |
| ACCOUNTING_CODE                                                                                                                                   | VARCHAR(15)                                                 | An identifier assigned by the system to record the resources used to write this file.                                                                                         |                                                                                                                                                  |
| EXPIRATION_DATE                                                                                                                                   | DATE                                                        | The date the file will be eligible for removal from the system by the Delete Expired Spooled Files (DLTEXPSPLF) command. Contains the null value if the file will not expire. |                                                                                                                                                  |
| SAVE_AFTER_WRITE                                                                                                                                  | VARCHAR(4)                                                  | Indicates whether this file is to be saved after it is written.                                                                                                               |                                                                                                                                                  |
|                                                                                                                                                   |                                                             | <b>*NO</b>                                                                                                                                                                    | The file is deleted after it has been written.                                                                                                   |
|                                                                                                                                                   |                                                             | <b>*YES</b>                                                                                                                                                                   | The file is set to save status after it has been written.                                                                                        |
| PAGE_LENGTH                                                                                                                                       | INTEGER                                                     | The page length, in lines per page, used by the spooled file.                                                                                                                 |                                                                                                                                                  |
| LINES_PER_INCH                                                                                                                                    | DECIMAL(5,1)                                                | The number of lines per vertical inch defined in the printer file.                                                                                                            |                                                                                                                                                  |
| PAGE_WIDTH                                                                                                                                        | INTEGER                                                     | The page width, in characters per printed line, used by the spooled file.                                                                                                     |                                                                                                                                                  |
| CHARACTERS_PER_INCH                                                                                                                               | DECIMAL(5,1)                                                | The number of characters per horizontal inch, defined in the printer file.                                                                                                    |                                                                                                                                                  |

Table 181. OUTPUT\_QUEUE\_ENTRIES table function (continued)

| Column Name                      | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRINT_FIDELITY                   | VARCHAR(8)  | The kind of error handling that is performed when printing.<br><br><b>*ABSOLUTE</b> The file is printed only if it can be printed exactly as specified in the data stream.<br><br><b>*CONTENT</b> The printing overrides errors in the data stream and continues printing with the printers best quality based on the content fidelity.                                                                                                                                                                                                                                           |
| PAGE_ROTATION                    | VARCHAR(5)  | The degree of rotation of the text on the page, with respect to the way the form is loaded into the printer.<br><br><b>*AUTO</b> Computer output reduction is done automatically if the output is too large to fit on the form, regardless of the print quality.<br><br><b>*DEVD</b> The operating system sends a device default rotation value to the printer. Page rotation is dependent on the printer's specifications.<br><br><b>*COR</b> Output created for a form 13.2 inches wide by 11.0 inches long is adjusted to print on a form 11.0 inches wide by 8.5 inches long. |
| PRINT_BOTH_SIDES                 | VARCHAR(7)  | How the information prints.<br><br><b>*FORMDF</b> The file uses a user-specified form definition. This value is used only for *LINE, *AFPDS, and *AFPDSLIN printer device type files.<br><br><b>*NO</b> The printing on the page is on one side only.<br><br><b>*YES</b> The printing is on both sides of the page with the top of each page the same for both sides.<br><br><b>*TUMBLE</b> The printing is on both sides with the top of one printed page at the opposite end from the top of the other printed page.                                                            |
| FILE_AVAILABLE                   | VARCHAR(8)  | The time when this file becomes available to an output device for processing.<br><br><b>*IMMED</b> The file is available as soon as the file is opened.<br><br><b>*FILEEND</b> The file is available as soon as the file is closed.<br><br><b>*JOBEND</b> The file is available when the job that owns the file is completed.                                                                                                                                                                                                                                                     |
| STARTING_PAGE                    | VARCHAR(10) | The page at which printing is to start for the file. Can contain the following special value:<br><br><b>*ENDPAGE</b> Printing starts with the last page.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| ENDING_PAGE                      | VARCHAR(10) | The page at which printing is to end for the file. Can contain the following special value:<br><br><b>*END</b> Printing ends with the last page.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| DEVICE_FILE_LIBRARY              | VARCHAR(10) | The name of the library that contains the device file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| DEVICE_FILE_NAME                 | VARCHAR(10) | The name of the device file used to create the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| PROGRAM_THAT_OPENED_FILE_LIBRARY | VARCHAR(10) | The name of the library that contains the program that opened the file. Contains null when the program is not known.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| PROGRAM_THAT_OPENED_FILE_NAME    | VARCHAR(10) | The name of the program that opened the spooled file. Contains null when the program is not known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| FORM_DEFINITION_LIBRARY          | VARCHAR(10) | The name of the library that contains the form definition. Contains null if FORM_DEFINITION_NAME is a special value or if no form definition is specified for this spooled file.                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 181. OUTPUT\_QUEUE\_ENTRIES table function (continued)

| Column Name             | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FORM_DEFINITION_NAME    | VARCHAR(10)  | <p>The name of the form definition to use for this print request. Can contain one of the following special values:</p> <ul style="list-style-type: none"> <li><b>*DEV D</b> The form definition in the printer device description will be used.</li> <li><b>*IN LINE</b> The form definition defined in the spooled file data stream will be used.</li> <li><b>*IN LINED</b> The form definition defined in the spooled file data stream will be used. If a form definition is not found, the form definition in the printer device description will be used.</li> <li><b>F1DFLT</b> The form definition defined in the spooled file data stream will be used.</li> </ul> <p>Contains null when no form definition is specified for this spooled file.</p> |
| PAGE_DEFINITION_LIBRARY | VARCHAR(10)  | <p>The name of the library containing the page definition. Contains the null value for *LINE or *AFPDSLINE printer device type files.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| PAGE_DEFINITION_NAME    | VARCHAR(10)  | <p>The name of the page definition to use for the file. Contains the null value for *LINE or *AFPDSLINE printer device type files.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| FRONT_OVERLAY_LIBRARY   | VARCHAR(10)  | <p>The name of the library containing the front overlay. Can contain one of these special values:</p> <ul style="list-style-type: none"> <li><b>*CURLIB</b> The current library is searched the front overlay.</li> <li><b>*LIBL</b> The library list is used to locate the front overlay.</li> </ul> <p>Contains null when FRONT_OVERLAY_NAME is *NONE.</p>                                                                                                                                                                                                                                                                                                                                                                                               |
| FRONT_OVERLAY_NAME      | VARCHAR(10)  | <p>The name of the front overlay. Can contain the following special value:</p> <ul style="list-style-type: none"> <li><b>*NONE</b> The file does not use the front overlay.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| BACK_OVERLAY_LIBRARY    | VARCHAR(10)  | <p>The name of the library containing the back overlay. Contains null when BACK_OVERLAY_NAME is a special value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| BACK_OVERLAY_NAME       | VARCHAR(10)  | <p>The name of the back overlay. Can contain the following special values:</p> <ul style="list-style-type: none"> <li><b>*FRONTOVL</b> The back overlay is the same as the front overlay.</li> <li><b>*NONE</b> The file does not use the back overlay.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| CHARACTER_SET_LIBRARY   | VARCHAR(10)  | <p>The name of the library containing the font character set object. Can contain one of these special values:</p> <ul style="list-style-type: none"> <li><b>*CURLIB</b> The current library is searched for the font character set object.</li> <li><b>*LIBL</b> The library list is used to locate the font character set object.</li> </ul> <p>Contains null when CHARACTER_SET_NAME is *FONT.</p>                                                                                                                                                                                                                                                                                                                                                       |
| CHARACTER_SET_NAME      | VARCHAR(10)  | <p>The name of the font character set object used to print this file. Can contain the following special value:</p> <ul style="list-style-type: none"> <li><b>*FONT</b> The information specified on the font parameter is used instead of the character set and code page.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CODE_PAGE_LIBRARY       | VARCHAR(10)  | <p>The name of the library containing the code page used to print this spooled file. Can contain one of these special values:</p> <ul style="list-style-type: none"> <li><b>*CURLIB</b> The current library is searched for the code page name.</li> <li><b>*LIBL</b> The library list is used to locate the code page name.</li> </ul> <p>Contains null when no code page is specified for this spooled file.</p>                                                                                                                                                                                                                                                                                                                                         |
| CODE_PAGE_NAME          | VARCHAR(10)  | <p>The name of the code page used to print this spooled file. Contains null when no code page is specified for this spooled file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CHARACTER_SET_POINTSIZE | DECIMAL(5,1) | <p>The point size in which this file's characters should be printed. Contains null if the character set does not have a point size.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 181. OUTPUT\_QUEUE\_ENTRIES table function (continued)

| Column Name               | Data Type    | Description                                                                                                                                                                                                                                                                                                                       |
|---------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CODED_FONT_LIBRARY        | VARCHAR(10)  | The name of the library containing the coded font used to print this spooled file. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched for the coded font.<br><b>*LIBL</b> The library list is used to locate the coded font.<br><br>Contains null when CODED_FONT_NAME is *FNTCHRSET. |
| CODED_FONT_NAME           | VARCHAR(10)  | The name of the coded font used to print this spooled file. Can contain the following special value:<br><br><b>*FNTCHRSET</b> The values used are the values specified on the character set name and library name and code page name and library name fields.                                                                     |
| CODED_FONT_POINTSIZE      | DECIMAL(5,1) | The point size in which this file's characters should be printed. Contains null if the coded font does not have a point size.                                                                                                                                                                                                     |
| MULTIBYTE_DATA            | VARCHAR(10)  | Whether the file can contain double-byte character set (DBCS) data, Unicode data, or both. Values are *YES and *NO.                                                                                                                                                                                                               |
| DBCS_CODED_FONT_LIBRARY   | VARCHAR(10)  | The name of the library containing the DBCS-coded font. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched for the DBCS-coded font.<br><b>*LIBL</b> The library list is used to locate the DBCS-coded font.<br><br>Contains null when DBCS_CODED_FONT_NAME is *SYSVAL.                |
| DBCS_CODED_FONT_NAME      | VARCHAR(10)  | The name of the DBCS-coded font used to print DBCS-coded data on printers configured as AFP(*YES). Can contain the following special value:<br><br><b>*SYSVAL</b> The DBCS-coded font specified in the system value is used.                                                                                                      |
| DBCS_CODED_FONT_POINTSIZE | DECIMAL(5,1) | The point size in which this file's DCBS characters should be printed. Contains null if the DBCS-coded font does not have a point size.                                                                                                                                                                                           |

## Example

Find the 100 largest spool files in the QEZJOBLOG output queue. Since no detailed information is needed, specify NO to avoid the additional processing.

```
SELECT *
FROM TABLE(QSYS2.OUTPUT_QUEUE_ENTRIES('*LIBL', 'QEZJOBLOG', 'NO')) A
ORDER BY SIZE DESC
FETCH FIRST 100 ROWS ONLY
```

## OUTPUT\_QUEUE\_ENTRIES view

The OUTPUT\_QUEUE\_ENTRIES view returns one row for each spooled file in every output queue. This view uses the QSYS2.OUTPUT\_QUEUE\_ENTRIES table function with DETAILED\_INFO => 'YES'.

**Authorization:** The caller must have:

- Execute authority to the output queue library and
  - Read authority to the output queue object, or
  - \*JOBCTL special authority and the output queue has OPRCTL(\*YES), or
  - \*SPLCTL special authority

To achieve the best performance when querying the OUTPUT\_QUEUE\_ENTRIES view, the use of a WHERE clause is recommended if you are interested in examining specific output queue libraries or output queues. OUTPUT\_QUEUE\_ENTRIES\_BASIC typically performs much better than OUTPUT\_QUEUE\_ENTRIES. OUTPUT\_QUEUE\_ENTRIES should only be used when OUTPUT\_QUEUE\_ENTRIES\_BASIC does not include the columns needed by the query.



The following table describes the columns in the view. The system name is OUTQ\_INFO. The schema is QSYS2.

Table 182. OUTPUT\_QUEUE\_ENTRIES view

| Column Name               | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OUTPUT_QUEUE_NAME         | OUTQ               | VARCHAR(10)             | Name of the output queue containing the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| OUTPUT_QUEUE_LIBRARY_NAME | OUTQLIB            | VARCHAR(10)             | The name of the library that contains the output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| CREATE_TIMESTAMP          | CREATED            | TIMESTAMP               | The timestamp when the file was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| SPOOLED_FILE_NAME         | SPOOLNAME          | VARCHAR(10)             | The file name that was specified by the user program when the file was created, or the name of the device file used to create this file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| USER_NAME                 | USER_NAME          | VARCHAR(10)             | The name of the user profile that produced the file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| USER_DATA                 | USER_DATA          | VARCHAR(10)<br>Nullable | The user-specified data that describes this file. Contains null if there is no user-specified data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| STATUS                    | STATUS             | VARCHAR(15)             | Status of the spooled file.<br><br><b>CLOSED</b> The file has been completely processed by a program but SCHEDULE(*JOBEND) was specified and the job that produced the file has not yet finished.<br><br><b>DEFERRED</b> Printing of the file has been deferred.<br><br><b>DELETED</b> The file has been deleted.<br><br><b>HELD</b> The file has been held.<br><br><b>MESSAGE WAITING</b> This file has a message which needs a reply or an action to be taken.<br><br><b>OPEN</b> The file has not been completely processed and is not ready to be selected by a writer.<br><br><b>PENDING</b> The file is pending to be printed.<br><br><b>PRINTING</b> The file has been completely sent to the printer but print complete status has not been sent back.<br><br><b>READY</b> The file is available to be written.<br><br><b>SAVED</b> The file has been printed and then saved. This file remains saved until it is released.<br><br><b>SENDING</b> The file is being sent or has been sent to a remote system.<br><br><b>WRITING</b> This file is currently being produced by the writer. |
| SIZE                      | SIZE               | INTEGER                 | The size of the spooled file, in kilobytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| TOTAL_PAGES               | PAGES              | INTEGER                 | The total number of pages in the file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| COPIES                    | COPIES             | SMALLINT                | The number of copies remaining to print.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| FORM_TYPE                 | FORM_TYPE          | VARCHAR(10)             | The type of form that should be loaded in the printer to print this file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| JOB_NAME                  | JOB_NAME           | VARCHAR(28)             | The qualified job name that produced the file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| DEVICE_TYPE               | DEVTYPE            | VARCHAR(10)             | The type of data stream used to represent the file.<br><br><b>*AFPDS</b> Advanced Function Presentation data stream<br><b>*AFPDSLNE</b> AFPDS data mixed with 1403 line data<br><b>*IPDS</b> Intelligent printer data stream<br><b>*LINE</b> 1403 line data<br><b>*SCS</b> Systems Network Architecture (SNA) character stream<br><b>*USERASCII</b> ASCII data                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| OUTPUT_PRIORITY           | OUTPTY             | SMALLINT                | The priority of the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 182. OUTPUT\_QUEUE\_ENTRIES view (continued)

| Column Name         | System Column Name | Data Type        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|--------------------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FILE_NUMBER         | FILENUM            | INTEGER          | The spooled file number of the specified file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| SYSTEM              | SYSTEM             | VARCHAR(8)       | The name of the system where the job that created the spooled file ran.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ACCOUNTING_CODE     | ACGCDE             | VARCHAR(15)      | An identifier assigned by the system to record the resources used to write this file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| EXPIRATION_DATE     | EXPDATE            | DATE<br>Nullable | The date the file will be eligible for removal from the system by the Delete Expired Spooled Files (DLTEXPSPLF) command. Contains the null value if the file will not expire.                                                                                                                                                                                                                                                                                                                                                                                             |
| SAVE_AFTER_WRITE    | SAVEAFTER          | VARCHAR(4)       | Indicates whether this file is to be saved after it is written.<br><br><b>*NO</b> The file is deleted after it has been written.<br><b>*YES</b> The file is set to save status after it has been written.                                                                                                                                                                                                                                                                                                                                                                 |
| PAGE_LENGTH         | PAGELEN            | INTEGER          | The page length, in lines per page, used by the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| LINES_PER_INCH      | LPI                | DECIMAL(5,1)     | The number of lines per vertical inch defined in the printer file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| PAGE_WIDTH          | WIDTH              | INTEGER          | The page width, in characters per printed line, used by the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| CHARACTERS_PER_INCH | CPI                | DECIMAL(5,1)     | The number of characters per horizontal inch, defined in the printer file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| PRINT_FIDELITY      | FIDELITY           | VARCHAR(8)       | The kind of error handling that is performed when printing.<br><br><b>*ABSOLUTE</b> The file is printed only if it can be printed exactly as specified in the data stream.<br><b>*CONTENT</b> The printing overrides errors in the data stream and continues printing with the printers best quality based on the content fidelity.                                                                                                                                                                                                                                       |
| PAGE_ROTATION       | ROTATION           | VARCHAR(5)       | The degree of rotation of the text on the page, with respect to the way the form is loaded into the printer.<br><br><b>*AUTO</b> Computer output reduction is done automatically if the output is too large to fit on the form, regardless of the print quality.<br><b>*DEVD</b> The operating system sends a device default rotation value to the printer. Page rotation is dependent on the printer's specifications.<br><b>*COR</b> Output created for a form 13.2 inches wide by 11.0 inches long is adjusted to print on a form 11.0 inches wide by 8.5 inches long. |
| PRINT_BOTH_SIDES    | BOTHSIDES          | VARCHAR(7)       | How the information prints.<br><br><b>*FORMDF</b> The file uses a user-specified form definition. This value is used only for *LINE, *AFPDS, and *AFPDSLIN printer device type files.<br><b>*NO</b> The printing on the page is on one side only.<br><b>*YES</b> The printing is on both sides of the page with the top of each page the same for both sides.<br><b>*TUMBLE</b> The printing is on both sides with the top of one printed page at the opposite end from the top of the other printed page.                                                                |
| FILE_AVAILABLE      | FILEAVAIL          | VARCHAR(8)       | The time when this file becomes available to an output device for processing.<br><br><b>*IMMED</b> The file is available as soon as the file is opened.<br><b>*FILEEND</b> The file is available as soon as the file is closed.<br><b>*JOBEND</b> The file is available when the job that owns the file is completed.                                                                                                                                                                                                                                                     |

Table 182. OUTPUT\_QUEUE\_ENTRIES view (continued)

| Column Name                      | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STARTING_PAGE                    | STARTPAGE          | VARCHAR(10)             | The page at which printing is to start for the file. Can contain the following special value:<br><br><b>*ENDPAGE</b> Printing starts with the last page.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| ENDING_PAGE                      | ENDPAGE            | VARCHAR(10)             | The page at which printing is to end for the file. Can contain the following special value:<br><br><b>*END</b> Printing ends with the last page.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| DEVICE_FILE_LIBRARY              | DEVLIB             | VARCHAR(10)             | The name of the library that contains the device file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| DEVICE_FILE_NAME                 | DEVFILE            | VARCHAR(10)             | The name of the device file used to create the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PROGRAM_THAT_OPENED_FILE_LIBRARY | LIBOPEN            | VARCHAR(10)<br>Nullable | The name of the library that contains the program that opened the file. Contains null when the program is not known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| PROGRAM_THAT_OPENED_FILE_NAME    | PGMOPEN            | VARCHAR(10)<br>Nullable | The name of the program that opened the spooled file. Contains null when the program is not known.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| FORM_DEFINITION_LIBRARY          | FORMLIB            | VARCHAR(10)<br>Nullable | The name of the library that contains the form definition. Contains null if FORM_DEFINITION_NAME is a special value or if no form definition is specified for this spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| FORM_DEFINITION_NAME             | FORMNAME           | VARCHAR(10)<br>Nullable | The name of the form definition to use for this print request. Can contain one of the following special values:<br><br><b>*DEV D</b> The form definition in the printer device description will be used.<br><br><b>*INLINE</b> The form definition defined in the spooled file data stream will be used.<br><br><b>*INLINED</b> The form definition defined in the spooled file data stream will be used. If a form definition is not found, the form definition in the printer device description will be used.<br><br><b>F1DFLT</b> The form definition defined in the spooled file data stream will be used.<br><br>Contains null when no form definition is specified for this spooled file. |
| PAGE_DEFINITION_LIBRARY          | PAGELIB            | VARCHAR(10)<br>Nullable | The name of the library containing the page definition. Contains the null value for *LINE or *AFPDLINE printer device type files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| PAGE_DEFINITION_NAME             | PAGENAME           | VARCHAR(10)<br>Nullable | The name of the page definition to use for the file. Contains the null value for *LINE or *AFPDLINE printer device type files.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| FRONT_OVERLAY_LIBRARY            | FRONTLIB           | VARCHAR(10)<br>Nullable | The name of the library containing the front overlay. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched the front overlay.<br><br><b>*LIBL</b> The library list is used to locate the front overlay.<br><br>Contains null when FRONT_OVERLAY_NAME is *NONE.                                                                                                                                                                                                                                                                                                                                                                                         |
| FRONT_OVERLAY_NAME               | FRONTNAME          | VARCHAR(10)             | The name of the front overlay. Can contain the following special value:<br><br><b>*NONE</b> The file does not use the front overlay.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| BACK_OVERLAY_LIBRARY             | BACKLIB            | VARCHAR(10)<br>Nullable | The name of the library containing the back overlay. Contains null when BACK_OVERLAY_NAME is a special value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| BACK_OVERLAY_NAME                | BACKNAME           | VARCHAR(10)             | The name of the back overlay. Can contain the following special values:<br><br><b>*FRONTOVL</b> The back overlay is the same as the front overlay.<br><br><b>*NONE</b> The file does not use the front overlay.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 182. OUTPUT\_QUEUE\_ENTRIES view (continued)

| Column Name             | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHARACTER_SET_LIBRARY   | CHRSETLIB          | VARCHAR(10)<br>Nullable  | The name of the library containing the font character set object. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched for the font character set object.<br><br><b>*LIBL</b> The library list is used to locate the font character set object.<br><br>Contains null when CHARACTER_SET_NAME is *FONT.               |
| CHARACTER_SET_NAME      | CHRSETNAME         | VARCHAR(10)              | The name of the font character set object used to print this file. Can contain the following special value:<br><br><b>*FONT</b> The information specified on the font parameter is used instead of the character set and code page.                                                                                                                            |
| CODE_PAGE_LIBRARY       | CODELIB            | VARCHAR(10)<br>Nullable  | The name of the library containing the code page used to print this spooled file. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched for the code page name.<br><br><b>*LIBL</b> The library list is used to locate the code page name.<br><br>Contains null when no code page is specified for this spooled file. |
| CODE_PAGE_NAME          | CODENAME           | VARCHAR(10)<br>Nullable  | The name of the code page used to print this spooled file. Contains null when no code page is specified for this spooled file.                                                                                                                                                                                                                                 |
| CHARACTER_SET_POINTSIZE | CHARSIZE           | DECIMAL(5,1)<br>Nullable | The point size in which this file's characters (defined by CHARACTER_SET) should be printed. Contains null if the character set does not have a point size.                                                                                                                                                                                                    |
| CODED_FONT_LIBRARY      | FONTLIB            | VARCHAR(10)<br>Nullable  | The name of the library containing the coded font used to print this spooled file. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched for the coded font.<br><br><b>*LIBL</b> The library list is used to locate the coded font.<br><br>Contains null when CODED_FONT_NAME is *FNTCHRSET.                          |
| CODED_FONT_NAME         | FONTNAME           | VARCHAR(10)              | The name of the coded font used to print this spooled file. Can contain the following special value:<br><br><b>*FNTCHRSET</b> The values used are the values specified on the character set name and library name and code page name and library name fields.                                                                                                  |
| CODED_FONT_POINTSIZE    | FONTSIZE           | DECIMAL(5,1)<br>Nullable | The point size in which this file's characters (defined by CODED_FONT) should be printed. Contains null if the coded font does not have a point size.                                                                                                                                                                                                          |
| MULTIBYTE_DATA          | MULTIBYTE          | VARCHAR(10)              | Whether the file can contain double-byte character set (DBCS) data, Unicode data, or both. Values are *YES and *NO.                                                                                                                                                                                                                                            |
| DBCS_CODED_FONT_LIBRARY | DBCSSLIB           | VARCHAR(10)<br>Nullable  | The name of the library containing the DBCS-coded font. Can contain one of these special values:<br><br><b>*CURLIB</b> The current library is searched for the DBCS-coded font.<br><br><b>*LIBL</b> The library list is used to locate the DBCS-coded font.<br><br>Contains null when DBCS_CODED_FONT_NAME is *SYSVAL.                                         |
| DBCS_CODED_FONT_NAME    | DBCSNAME           | VARCHAR(10)              | The name of the DBCS-coded font used to print DBCS-coded data on printers configured as AFP(*YES). Can contain the following special value:<br><br><b>*SYSVAL</b> The DBCS-coded font specified in the system value is used.                                                                                                                                   |

Table 182. OUTPUT\_QUEUE\_ENTRIES view (continued)

| Column Name               | System Column Name | Data Type                | Description                                                                                                                                                          |
|---------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DBCS_CODED_FONT_POINTSIZE | DBCSSIZE           | DECIMAL(5,1)<br>Nullable | The point size in which this file's DCBS characters (defined by DBCS_CODED_FONT) should be printed. Contains null if the DBCS-coded font does not have a point size. |

## Example

For the output queue with the largest number of files, determine how many kilobytes of data would be deleted by running the Delete Expired Spooled Files (DLTEXPSPLF) CL command.

```
WITH MOSTFILES(OUTQ_LIB, OUTQ_NAME) AS (
 SELECT OUTPUT_QUEUE_LIBRARY_NAME, OUTPUT_QUEUE_NAME
 FROM QSYS2.OUTPUT_QUEUE_INFO
 ORDER BY NUMBER_OF_FILES DESC
 LIMIT 1
)
SELECT SUM(SIZE) AS KB_TO_CLEAR
FROM MOSTFILES, QSYS2.OUTPUT_QUEUE_ENTRIES
WHERE OUTPUT_QUEUE_LIBRARY_NAME = OUTQ_LIB AND
 OUTPUT_QUEUE_NAME = OUTQ_NAME AND
 EXPIRATION_DATE IS NOT NULL
 AND
 EXPIRATION_DATE < CURRENT DATE;
```

## OUTPUT\_QUEUE\_ENTRIES\_BASIC view

The OUTPUT\_QUEUE\_ENTRIES\_BASIC view returns one row for each spooled file in every output queue. This view uses the QSYS2.OUTPUT\_QUEUE\_ENTRIES table function with DETAILED\_INFO => 'NO'.

**Authorization:** The caller must have:

- Execute authority to the output queue library and
  - Read authority to the output queue object, or
  - \*JOBCTL special authority and the output queue has OPRCTL(\*YES), or
  - \*SPLCTL special authority

To achieve the best performance when querying the OUTPUT\_QUEUE\_ENTRIES\_BASIC view, the use of a WHERE clause is recommended if you are interested in examining specific output queue libraries or output queues. OUTPUT\_QUEUE\_ENTRIES\_BASIC typically performs much better than OUTPUT\_QUEUE\_ENTRIES. OUTPUT\_QUEUE\_ENTRIES should only be used when OUTPUT\_QUEUE\_ENTRIES\_BASIC does not include the columns needed by the query.

The following table describes the columns in the view. The system name is OUTQ\_INFOB. The schema is QSYS2.

Table 183. OUTPUT\_QUEUE\_ENTRIES\_BASIC view

| Column Name               | System Column Name | Data Type               | Description                                                                                                                              |
|---------------------------|--------------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| OUTPUT_QUEUE_NAME         | OUTQ               | VARCHAR(10)             | Name of the output queue containing the spooled file.                                                                                    |
| OUTPUT_QUEUE_LIBRARY_NAME | OUTQLIB            | VARCHAR(10)             | The name of the library that contains the output queue.                                                                                  |
| CREATE_TIMESTAMP          | CREATED            | TIMESTAMP               | The timestamp when the file was created.                                                                                                 |
| SPOOLED_FILE_NAME         | SPOOLNAME          | VARCHAR(10)             | The file name that was specified by the user program when the file was created, or the name of the device file used to create this file. |
| USER_NAME                 | USER_NAME          | VARCHAR(10)             | The name of the user profile that produced the file.                                                                                     |
| USER_DATA                 | USER_DATA          | VARCHAR(10)<br>Nullable | The user-specified data that describes this file. Contains null if there is no user-specified data.                                      |

Table 183. OUTPUT\_QUEUE\_ENTRIES\_BASIC view (continued)

| Column Name     | System Column Name | Data Type   | Description                                                               |                                                                                                                                                  |
|-----------------|--------------------|-------------|---------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| STATUS          | STATUS             | VARCHAR(15) | Status of the spooled file.                                               |                                                                                                                                                  |
|                 |                    |             | <b>CLOSED</b>                                                             | The file has been completely processed by a program but SCHEDULE(*JOBEND) was specified and the job that produced the file has not yet finished. |
|                 |                    |             | <b>DEFERRED</b>                                                           | Printing of the file has been deferred.                                                                                                          |
|                 |                    |             | <b>DELETED</b>                                                            | The file has been deleted.                                                                                                                       |
|                 |                    |             | <b>HELD</b>                                                               | The file has been held.                                                                                                                          |
|                 |                    |             | <b>MESSAGE WAITING</b>                                                    | This file has a message which needs a reply or an action to be taken.                                                                            |
|                 |                    |             | <b>OPEN</b>                                                               | The file has not been completely processed and is not ready to be selected by a writer.                                                          |
|                 |                    |             | <b>PENDING</b>                                                            | The file is pending to be printed.                                                                                                               |
|                 |                    |             | <b>PRINTING</b>                                                           | The file has been completely sent to the printer but print complete status has not been sent back.                                               |
|                 |                    |             | <b>READY</b>                                                              | The file is available to be written.                                                                                                             |
|                 |                    |             | <b>SAVED</b>                                                              | The file has been printed and then saved. This file remains saved until it is released.                                                          |
|                 |                    |             | <b>SENDING</b>                                                            | The file is being sent or has been sent to a remote system.                                                                                      |
|                 |                    |             | <b>WRITING</b>                                                            | This file is currently being produced by the writer.                                                                                             |
| SIZE            | SIZE               | INTEGER     | The size of the spooled file, in kilobytes.                               |                                                                                                                                                  |
| TOTAL_PAGES     | PAGES              | INTEGER     | The total number of pages in the file.                                    |                                                                                                                                                  |
| COPIES          | COPIES             | SMALLINT    | The number of copies remaining to print.                                  |                                                                                                                                                  |
| FORM_TYPE       | FORM_TYPE          | VARCHAR(10) | The type of form that should be loaded in the printer to print this file. |                                                                                                                                                  |
| JOB_NAME        | JOB_NAME           | VARCHAR(28) | The qualified job name that produced the file.                            |                                                                                                                                                  |
| DEVICE_TYPE     | DEVTYPE            | VARCHAR(10) | The type of data stream used to represent the file.                       |                                                                                                                                                  |
|                 |                    |             | <b>*AFPDS</b>                                                             | Advanced Function Presentation data stream                                                                                                       |
|                 |                    |             | <b>*AFPDSLIN</b>                                                          | AFPDS data mixed with 1403 line data                                                                                                             |
|                 |                    |             | <b>*IPDS</b>                                                              | Intelligent printer data stream                                                                                                                  |
|                 |                    |             | <b>*LINE</b>                                                              | 1403 line data                                                                                                                                   |
|                 |                    |             | <b>*SCS</b>                                                               | Systems Network Architecture (SNA) character stream                                                                                              |
|                 |                    |             | <b>*USERASCII</b>                                                         | ASCII data                                                                                                                                       |
| OUTPUT_PRIORITY | OUTPTY             | SMALLINT    | The priority of the spooled file.                                         |                                                                                                                                                  |
| FILE_NUMBER     | FILENUM            | INTEGER     | The spooled file number of the specified file.                            |                                                                                                                                                  |
| SYSTEM          | SYSTEM             | VARCHAR(8)  | The name of the system where the job that created the spooled file ran.   |                                                                                                                                                  |

## Examples

- Find the 100 largest spool files in the QEZJOBLOG output queue.

```
SELECT * FROM QSYS2.OUTPUT_QUEUE_ENTRIES_BASIC
WHERE OUTPUT_QUEUE_NAME = 'QEZJOBLOG'
ORDER BY SIZE DESC
FETCH FIRST 100 ROWS ONLY
```

- Find the top 10 consumers of SPOOL storage.

```

SELECT USER_NAME, SUM(SIZE) AS TOTAL_SPOOL_SPACE
FROM QSYS2.OUTPUT_QUEUE_ENTRIES_BASIC
WHERE USER_NAME NOT LIKE 'Q%'
GROUP BY USER_NAME
ORDER BY TOTAL_SPOOL_SPACE DESC LIMIT 10;

```

## OUTPUT\_QUEUE\_INFO view

The OUTPUT\_QUEUE\_INFO view returns one row for each output queue.

The values returned for the columns in the view are similar to the values returned by the [Retrieve Output Queue Information \(QSPROUTQ\) API](#). Refer to the API for more detailed information.

**Authorization:** Rows will be returned for output queues when the caller has:

- Execute authority to the output queue library and
  - Read authority to the output queue object, or
  - \*JOBCTL special authority and the output queue has OPRCTL(\*YES), or
  - \*SPLCTL special authority

The following table describes the columns in the view. The system name is OUTQ\_DTL. The schema is QSYS2.

Table 184. OUTPUT\_QUEUE\_INFO view

| Column Name               | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OUTPUT_QUEUE_NAME         | OUTQ               | VARCHAR(10)             | Name of the output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| OUTPUT_QUEUE_LIBRARY_NAME | OUTQLIB            | VARCHAR(10)             | The name of the library that contains the output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| NUMBER_OF_FILES           | FILES              | INTEGER                 | The total number of spooled files currently on this output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| NUMBER_OF_WRITERS         | WRITERS            | INTEGER                 | The number of printer writers that have been started to this output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| WRITERS_TO_AUTOSTART      | AUTOSTART          | INTEGER                 | The number of remote printer writers to autostart to this output queue at system IPL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| PRINTER_DEVICE_NAME       | DEV_NAME           | VARCHAR(10)<br>Nullable | The name of the printer device. If more than one writer is started, this is the printer device name of the first writer.<br>Contains the null value if WRITER_TYPE is not PRINTER.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| ORDER_OF_FILES            | FILE_ORDER         | VARCHAR(7)              | The order of the spooled files on the output queue.<br><br><b>*FIFO</b> The queue is first-in first-out for each file. That is, on the queue, new spooled files are placed behind all other spooled files that have the same priority.<br><br><b>*JOBNBR</b> The queue entries for the spooled files are sorted in priority sequence using the job number (the date and time that the job entered the system) of the job that created the spooled file.                                                                                                                                                                                                                                                                                                                        |
| DISPLAY_ANY_FILE          | ANYFILE            | VARCHAR(6)              | Whether users who have authority to read this output queue can display the output data of any output file on this queue, or only the data in their own files.<br><br><b>*NO</b> Users authorized to the queue can only display, copy, or send their own spooled files, unless one of the following applies: <ul style="list-style-type: none"> <li>• they have *SPLCTL special authority, or</li> <li>• they have *JOBCTL special authority and OPERATOR_CONTROLLED is *YES.</li> </ul> <b>*OWNER</b> Only the owner of a file or a user with *SPLCTL authority can display, copy, send, or move their own spooled files to another output queue.<br><br><b>*YES</b> Any user having authority to read the queue can display, copy, or send the data of any file on the queue. |

Table 184. OUTPUT\_QUEUE\_INFO view (continued)

| Column Name                 | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------|--------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_SEPARATORS              | JOB_SEP            | VARCHAR(4)              | The number of job separators (0-9) to be placed at the beginning of the output for each job having spooled file entries on this output queue. Can also contain the following special value:<br><br><b>*MSG</b> No job separators are used; instead a message is sent to the writer's message queue at the end of each job indicating that the output can be removed.                                                                                                                  |
| MAXIMUM_PAGES               | MAX_PAGES          | INTEGER<br>Nullable     | Only spooled files with this number of pages or less will print between MAXIMUM_PAGES_STARTING_TIME and MAXIMUM_PAGES_ENDING_TIME. If more than one set of maximum spooled file size values is defined for this output queue, only information for the first set is returned.<br><br>Contains the null value if no maximum spooled file size is defined.                                                                                                                              |
| MAXIMUM_PAGES_STARTING_TIME | MAX_START          | TIME<br>Nullable        | The starting time, in local job time, that spooled files exceeding MAXIMUM_PAGES will be restricted from printing. If a spooled file exceeds the page limit it will be in deferred status until ENDING_TIME.<br><br>Contains the null value if no maximum spooled file size is defined.                                                                                                                                                                                               |
| MAXIMUM_PAGES_ENDING_TIME   | MAX_END            | TIME<br>Nullable        | The ending time, in local job time, when spooled files exceeding MAXIMUM_PAGES will be allowed to print.<br><br>Contains the null value if no maximum spooled file size is defined.                                                                                                                                                                                                                                                                                                   |
| OPERATOR_CONTROLLED         | OPR_CTRL           | VARCHAR(4)              | Whether users with job control authority are allowed to manage or control the files on this queue. Users have job control authority if SPCAUT(*JOBCTL) is specified in their user profile.<br><br><b>*NO</b> This queue and its entries cannot be controlled or changed by users with job control authority unless they also have some other special authority.<br><br><b>*YES</b> Users with job control authority can control the queue and make changes to the files on the queue. |
| AUTHORITY_TO_CHECK          | ALL_AUTH           | VARCHAR(7)              | Indicates what type of authorities to the output queue allow the user to control all the files on the queue.<br><br><b>*DTAAUT</b> Any user with *READ, *ADD, and *DELETE authority to the output queue can control all output files on the queue.<br><br><b>*OWNER</b> Only the owner of the output queue can control all the output files on the queue.                                                                                                                             |
| DATA_QUEUE_LIBRARY          | DTAQ_LIB           | VARCHAR(10)<br>Nullable | The name of the library containing the data queue.<br><br>Contains the null value if no data queue is associated with this output queue.                                                                                                                                                                                                                                                                                                                                              |
| DATA_QUEUE_NAME             | DTAQ_NAME          | VARCHAR(10)<br>Nullable | The name of the data queue associated with this output queue.<br><br>Contains the null value if no data queue is associated with this output queue.                                                                                                                                                                                                                                                                                                                                   |
| OUTPUT_QUEUE_STATUS         | STATUS             | VARCHAR(8)              | The status of the output queue.<br><br><b>HELD</b> The queue is held.<br><br><b>RELEASED</b> The queue is released.                                                                                                                                                                                                                                                                                                                                                                   |
| WRITER_JOB_NAME             | WRITER_JOB         | VARCHAR(28)<br>Nullable | The qualified job name of the writer job. If more than one writer is started, this is the name of the first writer.<br><br>Contains the null value if a writer job is not started for this queue.                                                                                                                                                                                                                                                                                     |



Table 184. OUTPUT\_QUEUE\_INFO view (continued)

| Column Name                | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|--------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WRITER_JOB_STATUS          | WRITER_STS         | VARCHAR(4)<br>Nullable  | The status of the writer job. If more than one writer is started, this is the status of the first writer.<br><br><b>END</b> The writer job has ended.<br><b>HLD</b> The writer job is held.<br><b>JOBQ</b> The writer job is on the job queue.<br><b>MSGW</b> The writer job is waiting for a message response.<br><b>STR</b> The writer job is started for the output queue.<br><br>Contains the null value if a writer job is not started for this queue. |
| WRITER_TYPE                | WRITER_TYP         | VARCHAR(7)<br>Nullable  | The type of writer started for this output queue.<br><br><b>PRINTER</b> Printer writer.<br><b>REMOTE</b> Remote writer.<br><br>Contains the null value if a writer job is not started for this queue.                                                                                                                                                                                                                                                       |
| SPOOLED_FILE_ASP_ATTRIBUTE | ASP_ATTR           | VARCHAR(8)              | The auxiliary storage pool (ASP) where the spooled files are to reside.<br><br><b>*OUTQASP</b> The spooled files reside in the auxiliary storage pool in which the output queue resides.<br><b>*SYSTEM</b> The spooled files reside in the system auxiliary storage pool.                                                                                                                                                                                   |
| SPOOLED_FILE_ASP_NUMBER    | ASPNUM             | INTEGER                 | The number of the auxiliary storage pool (ASP) where the spooled files reside.                                                                                                                                                                                                                                                                                                                                                                              |
| SPOOLED_FILE_ASPGRP        | ASPGRP             | VARCHAR(10)<br>Nullable | The name of the auxiliary storage pool (ASP) device where the spooled files reside. Can also contain the following special value:<br><br><b>*SYSBAS</b> The spooled files resides in the system ASP (ASP 1) or one of the defined basic user ASPs (ASPs 2-32).<br><br>Contains the null value if the name is not available.                                                                                                                                 |
| TEXT_DESCRIPTION           | TEXT               | VARCHAR(50)<br>Nullable | The text description of the output queue.<br><br>Contains the null value if the output queue has no description.                                                                                                                                                                                                                                                                                                                                            |
| MESSAGE_QUEUE_LIBRARY      | MSGQ_LIB           | VARCHAR(10)<br>Nullable | The name of the library containing the message queue. Can contain the following special value:<br><br><b>*LIBL</b> The library list is searched to find the message queue.<br><br>Contains the null value if the output queue is not a remote output queue or if WRITER_TYPE is PRINTER.                                                                                                                                                                    |
| MESSAGE_QUEUE_NAME         | MSGQ_NAME          | VARCHAR(10)<br>Nullable | The name of the message queue to which messages, created by the remote writer started to this output queue, are sent.<br><br>Contains the null value if the output queue is not a remote output queue or if WRITER_TYPE is PRINTER.                                                                                                                                                                                                                         |
| HOST_PRINT_TRANSFORM       | TRANSFORM          | VARCHAR(4)<br>Nullable  | Whether to use the host print transform function to transform a spooled file.<br><br><b>*NO</b> Do not transform data streams using host print transform.<br><b>*YES</b> Transform data streams using host print transform.<br><br>Contains the null value if NETWORK_CONNECTION_TYPE is *SNA and USER_DRIVER_PROGRAM_NAME is null.                                                                                                                         |
| IMAGE_CONFIGURATION_NAME   | IMAGE_NAME         | VARCHAR(10)<br>Nullable | The name of the image configuration.<br><br>Contains the null value if no image configuration is used when transforming the spooled file before sending.                                                                                                                                                                                                                                                                                                    |

Table 184. OUTPUT\_QUEUE\_INFO view (continued)

| Column Name                            | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------|--------------------|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MANUFACTURER_TYPE_AND_MODEL            | TYPE_MODEL         | VARCHAR(17)<br>Nullable  | The manufacturer, type, and model for a printer using the host print transform function.<br><br>See <a href="#">Printer Model Settings for Host Print Transform (HPT)</a> in the IBM Support Portal for the list of supported values.<br><br>Contains the null value when NETWORK_CONNECTION_TYPE is *SNA, or when NETWORK_CONNECTION_TYPE is *IP and HOST_PRINT_TRANSFORM is *NO.                                                                                                                                                                                                                                                                        |
| WORKSTATION_CUSTOMIZING_OBJECT_LIBRARY | CUSTOM_LIB         | VARCHAR(10)<br>Nullable  | The library name for the workstation customizing object.<br><br>Contains the null value if there is no workstation customizing object or if NETWORK_CONNECTION_TYPE is *SNA and USER_DRIVER_PROGRAM_NAME is null.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| WORKSTATION_CUSTOMIZING_OBJECT_NAME    | CUSTOM_NAM         | VARCHAR(10)<br>Nullable  | The name of an object that consists of a table of attributes used to customize a given ASCII device.<br><br>Contains the null value if there is no workstation customizing object or if NETWORK_CONNECTION_TYPE is *SNA and USER_DRIVER_PROGRAM_NAME is null.                                                                                                                                                                                                                                                                                                                                                                                             |
| NETWORK_CONNECTION_TYPE                | NET_TYPE           | VARCHAR(7)<br>Nullable   | The type of network connection to the remote system.<br><br><p><b>*IP</b> The TCP/IP network is used as the connectivity to the remote system.</p> <p><b>*SNA</b> The SNADS network is used as the connectivity to the remote system.</p> <p><b>*USRDFN</b> A user-defined connectivity is used as the connectivity to the remote system.</p> <p>Contains the null value if the output queue is not a remote output queue.</p>                                                                                                                                                                                                                            |
| DESTINATION_TYPE                       | DEST_TYPE          | VARCHAR(8)<br>Nullable   | The type of destination system that spooled files on this output queue are being sent to.<br><br><p><b>*NDS</b> The destination is Novell NetWare 3 or 4, and the connection type is *USRDFN.</p> <p><b>*OS400</b> The destination system is an IBM i.</p> <p><b>*OTHER</b> The destination system does not match any of the other special values. This is commonly used when the destination is a printer.</p> <p><b>*PSF2</b> The destination system is a PC using Print Services Facility/2.</p> <p><b>*S390</b> This destination system is a System/390® system.</p> <p>Contains the null value if the output queue is not a remote output queue.</p> |
| REMOTE_SYSTEM_NAME                     | REMOTE_NAM         | VARCHAR(255)<br>Nullable | The name of the remote system.<br><br>Contains the null value if the output queue is not a remote output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| REMOTE_PRINTER_QUEUE                   | REMOTE_PRT         | VARCHAR(255)<br>Nullable | The name of the remote printer. Can also contain one of these special values:<br><br><p><b>*SYSTEM</b> The default system printer on the remote system will determine the printer queue.</p> <p><b>*USER</b> The user profile that creates the spooled file will determine the user ID on the remote system that it is sent to.</p> <p>Contains the null value if the output queue is not a remote output queue.</p>                                                                                                                                                                                                                                      |

Table 184. OUTPUT\_QUEUE\_INFO view (continued)

| Column Name                    | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DESTINATION_OPTIONS            | DEST_OPT           | VARCHAR(128)<br>Nullable | <p>Destination-dependent options that are specific to a particular implementation of an LPR Print Server. Can also contain the special values:</p> <p><b>*NOWAIT</b> The remote writer will not wait for confirmation that the destination system has finished processing the spooled file.</p> <p><b>*USRDFNTXT</b> Use the value for the user-defined text of the user profile when the spooled file was created.</p> <p>Contains the null value if the output queue is not a remote output queue.</p> |
| USER_DRIVER_PROGRAM_LIBRARY    | UDP_LIB            | VARCHAR(10)<br>Nullable  | <p>The name of the library that contains the user driver program. Can also be one of these special values:</p> <p><b>*CURLIB</b> The current library for the job is used to locate the user driver program.</p> <p><b>*LIBL</b> The library list used to locate the user driver program.</p> <p>Contains the null value if no user driver program is specified.</p>                                                                                                                                      |
| USER_DRIVER_PROGRAM_NAME       | UDP_NAME           | VARCHAR(10)<br>Nullable  | <p>The name of the user-specified driver program that is used to process the spooled files on the output queue.</p> <p>Contains the null value if no user driver program is specified.</p>                                                                                                                                                                                                                                                                                                               |
| USER_DEFINED_OBJECT_LIBRARY    | UDO_LIB            | VARCHAR(10)<br>Nullable  | <p>The name of the library that contains the user-defined object. Can also be one of these special values:</p> <p><b>*CURLIB</b> The current library for the job is used to locate the user-defined object.</p> <p><b>*LIBL</b> The library list used to locate the user-defined object.</p> <p>Contains the null value if no user-defined object is specified.</p>                                                                                                                                      |
| USER_DEFINED_OBJECT_NAME       | UDO_NAME           | VARCHAR(10)<br>Nullable  | <p>The name of the user-defined object that is used by user applications or user-specified programs that process spooled files.</p> <p>Contains the null value if no user-defined object is specified.</p>                                                                                                                                                                                                                                                                                               |
| USER_DEFINED_OBJECT_TYPE       | UDO_TYPE           | VARCHAR(7)<br>Nullable   | <p>The type of the user-defined object.</p> <p><b>*DTAARA</b> Data area.</p> <p><b>*DTAQ</b> Data queue.</p> <p><b>*FILE</b> File.</p> <p><b>*PSFCFG</b> PSF configuration object.</p> <p><b>*USRIDX</b> User index.</p> <p><b>*USRQ</b> User queue.</p> <p><b>*USRSPC</b> User space.</p> <p>Contains the null value if no user-defined object is specified.</p>                                                                                                                                        |
| DATA_TRANSFORM_PROGRAM_LIBRARY | DTP_LIB            | VARCHAR(10)<br>Nullable  | <p>The name of the library that contains the data transform program. Can also be one of these special values:</p> <p><b>*CURLIB</b> The current library for the job is used to locate the data transform program.</p> <p><b>*LIBL</b> The library list used to locate the data transform program.</p> <p>Contains the null value if no data transform program is specified.</p>                                                                                                                          |
| DATA_TRANSFORM_PROGRAM_NAME    | DTP_NAME           | VARCHAR(10)<br>Nullable  | <p>The name of the user-specified data transform program that is used by the driver program.</p> <p>Contains the null value if no data transform program is specified.</p>                                                                                                                                                                                                                                                                                                                               |

Table 184. OUTPUT\_QUEUE\_INFO view (continued)

| Column Name            | System Column Name | Data Type                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------|--------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| USER_DEFINED_OPTION_1  | UDEF_OPT1          | VARCHAR(10)<br>Nullable     | The first user-defined option.<br>Contains the null value if there are no user-defined options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| USER_DEFINED_OPTION_2  | UDEF_OPT2          | VARCHAR(10)<br>Nullable     | The second user-defined option.<br>Contains the null value if there are not at least two user-defined options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| USER_DEFINED_OPTION_3  | UDEF_OPT3          | VARCHAR(10)<br>Nullable     | The third user-defined option.<br>Contains the null value if there are not at least three user-defined options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| USER_DEFINED_OPTION_4  | UDEF_OPT4          | VARCHAR(10)<br>Nullable     | The fourth user-defined option.<br>Contains the null value if there are not at least four user-defined options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| USER_DEFINED_DATA      | UDEF_DATA          | VARBINARY(5000)<br>Nullable | Data defined by the user to be used by user applications or user-specified programs that process spooled files.<br>Contains the null value if there is no user-defined data.                                                                                                                                                                                                                                                                                                                                                                                                                              |
| LDAP_PUBLISHING_STATUS | PUBLISHED          | VARCHAR(3)                  | Whether the output queue is published in the network directory.<br><br><b>NO</b> Output queue is not published.<br><b>YES</b> Output queue is published.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| FORMS_CONTROL_BUFFER   | FORMS_BUF          | VARCHAR(8)<br>Nullable      | The forms control buffer (FCB) for files sent to a VM/MVS host system. Contains either the name of the FCB or one of the following special values:<br><br><b>*PRTF</b> The first 8 characters of the printer file used to pool the file determines the name of the FCB.<br><br><b>*USRDTA</b> The first 8 characters of the user data (USRDATA) spooled file attribute determines the name of the FCB. If the user data is blank, no FCB is used.<br><br>Contains the null value if no FCB is used when sending spooled files or if NETWORK_CONNECTION_TYPE is not *SNA or DESTINATION_TYPE is not *S390. |
| VM_MVS_CLASS           | VM_CLASS           | CHAR(1)<br>Nullable         | The VM/MVS SYSOUT class for distributions sent to a VM host system or to a MVS host system. Values are A-Z, 0-9 to indicate the distribution class.<br><br>Contains the null value if not defined for this output queue.                                                                                                                                                                                                                                                                                                                                                                                  |

## Example

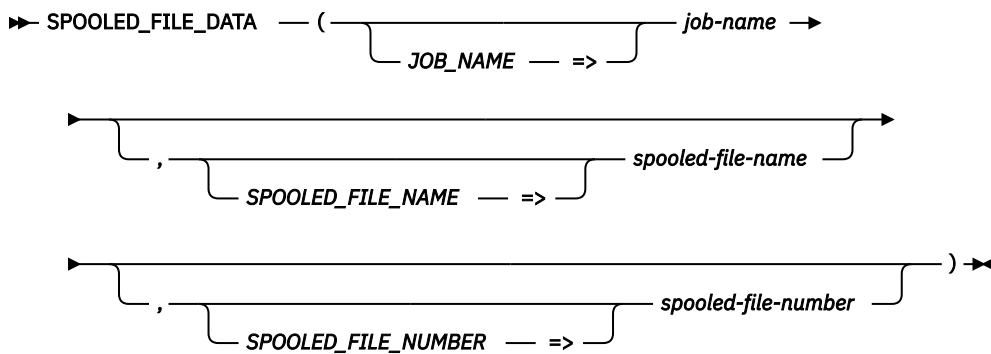
```
SELECT * FROM QSYS2.OUTPUT_QUEUE_INFO
```

## SPOOLED\_FILE\_DATA table function

The SPOOLED\_FILE\_DATA table function returns the content of a spooled file.

If the spooled file contains double byte data, the job CCSID must be a mixed CCSID.

**Authorization:** This table function uses the CPYSPLF CL command. Any authority requirements for the CL command apply to the use of this function.



The schema is SYSTOOLS.

**job-name** A character string containing a qualified job name. Can contain the following special value:

- \* Use the name of the current job.

**spooled-file-name** A character string containing the name of the spooled file. If this parameter is omitted, QPJOBLOG is used.

**spooled-file-number** The number of the spooled file. If this parameter is omitted, the spooled file with the highest number matching *spooled-file-name* is used.

The result of the function is a table containing a row for each record in the specified spooled file. The columns of the result table are described in the following table. The result columns are nullable.

Table 185. SPOOLED\_FILE\_DATA table function

| Column Name      | Data Type    | Description                                        |
|------------------|--------------|----------------------------------------------------|
| ORDINAL_POSITION | INTEGER      | Relative position of this row in the spooled file. |
| SPOOLED_DATA     | VARCHAR(200) | The data for this row in the spooled file.         |

## Note

This function is provided in the SYSTOOLS schema as an example of how spooled file data can be returned by embedding the CPYSPLF CL command in an SQL table function. Creating customized versions of this table function to better suit a specific need is encouraged. Use the Insert Generated SQL feature in ACS to extract the source for this function. Then modify it and create a new procedure in a user-specified schema.

## Example

Return the most recent QSYSPRT file for a specific job:

```
SELECT * FROM TABLE(SYSTOOLS.SPOOLED_FILE_DATA(
 JOB_NAME => '193846/SLROMANO/QPADEV0009',
 SPOOLED_FILE_NAME => 'QSYSPRT'))
ORDER BY ORDINAL_POSITION;
```

## SPOOLED\_FILE\_INFO table function

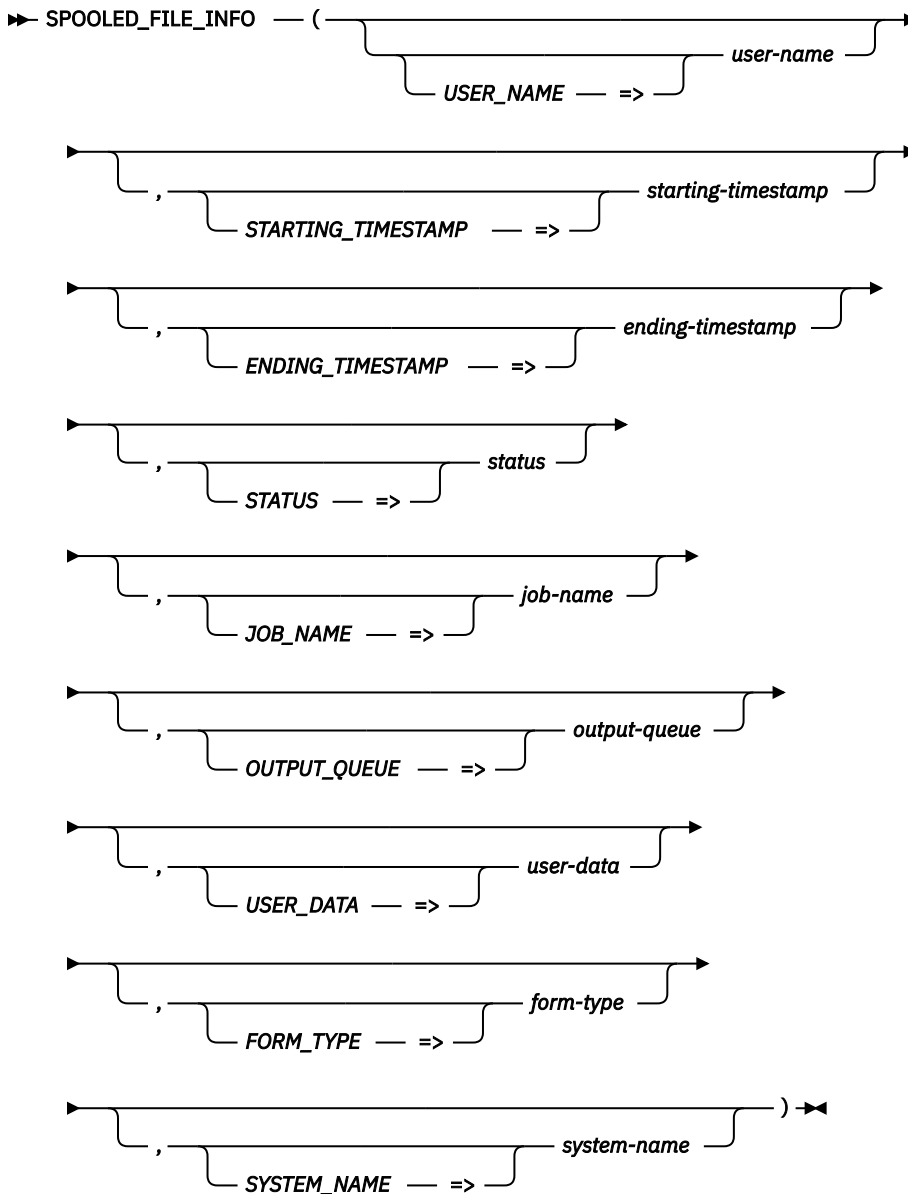
The SPOOLED\_FILE\_INFO table function returns a list of spooled files on the system.

This information is similar to what is returned by the Work with Spooled Files (WRKSPLF) CL command and the Open List of Spooled Files (QGYOLSPL) API.

**Authorization:** The caller must have:

- Execute authority to the output queue library containing a spooled file and

- \*USE authority to the output queue object



The schema is QSYS2.

Filtering on multiple values of `user-name`, `output-queue`, or `user-data` is slower.

Additionally, filtering on `status` is slower.

**user-name** A character string containing user profile names to filter on. The list can contain up to 20 values separated by blanks.

Can be one of the following special values:

**\*ALL** The list is not filtered by user profile.

**\*CURRENT** Spooled files owned by the current user profile are included in the list. This is the default.

**starting-timestamp** A timestamp value for the earliest spooled file to return. If the parameter is omitted or the null value, all spooled files with a create timestamp less than or equal to `ending-timestamp` are returned.

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ending-timestamp</b> | A timestamp value for the latest spooled file to return. If the parameter is omitted or the null value, all spooled files with a create timestamp greater than or equal to <i>starting-timestamp</i> are returned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>status</b>           | A character string containing the spooled file status to filter on. The list can contain one or more status values separated by blanks. <ul style="list-style-type: none"> <li><b>*ALL</b> Spooled files are included in the list regardless of the current status. This is the default.</li> <li><b>*CLOSED</b> The file has been completely processed by a program, but SCHEDULE(*JOBEND) was specified. The job that produced the spooled file has not finished.</li> <li><b>*DEFERRED</b> This spooled file has been deferred from printing.</li> <li><b>*FINISHED</b> This spooled file is no longer in the system. These spooled files are included in the list of spooled files only if the qualified job name is specified.</li> <li><b>*HELD</b> The file has been held.</li> <li><b>*MESSAGE</b> This file has a message that needs a reply or needs an action to be taken.</li> <li><b>*OPEN</b> The file has not been completely processed and is not ready to be selected by a writer.</li> <li><b>*PENDING</b> This file is pending (waiting) to be printed.</li> <li><b>*PRINTER</b> The file has been completely sent to the printer, but the print complete status has not been sent back.</li> <li><b>*READY</b> The file is available to be written to an output device by a writer.</li> <li><b>*SAVED</b> The file has been written and then saved. This file remains saved until it is released.</li> <li><b>*SENDING</b> This spooled file is being sent or has been sent to a remote system.</li> <li><b>*WRITING</b> This file is currently being produced by the writer on an output device.</li> </ul> |
| <b>job-name</b>         | A character string containing qualified job name.<br>Can be one of the following special values: <ul style="list-style-type: none"> <li><b>*</b> The current job.</li> <li><b>*ALL</b> All jobs matching the other criteria are returned. This is the default.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>output-queue</b>     | A character string containing the name of the output queue to filter on. An output queue name must be in the format <i>library-name/outq-name</i> . The list can contain up to 20 values separated by blanks.<br>*LIBL or *CURLIB can be specified for the library name.<br>Instead of a list of qualified output queue names, the following special value can be used: <ul style="list-style-type: none"> <li><b>*ALL</b> Spooled files are included in the list regardless of the output queue. This is the default.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>user-data</b>        | A character string containing the value of the user-specified data or file name for spooled files to filter on.<br>Can be the following special value: <ul style="list-style-type: none"> <li><b>*ALL</b> Spooled files are included in the list regardless of the value for user-specified data. This is the default.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>form-type</b>        | A character string containing the form type value to filter on.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Can be one of the following special values:

- \*ALL** Spooled files are included in the list regardless of the value for form type. This is the default.
- \*STD** Only files that specify the standard form type are included in the list.

**system-name**

A character string containing the name of the system where the job that created the spooled file ran.

Can be one of the following special values:

- \*ALL** The list is not filtered based on job system name. This is the default.
- \*CURRENT** Only spooled files created on the current system are to be returned.

The result of the function is a table containing rows with the format shown in the following table. All columns are nullable.

Table 186. SPOOLED\_FILE\_INFO table function

| Column Name         | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SPOOLED_FILE_NAME   | VARCHAR(10)  | The name of the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| SPOOLED_FILE_NUMBER | INTEGER      | The number of the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| STATUS              | VARCHAR(15)  | The status of the file.<br><br><b>CLOSED</b> The file has been processed completely by a program, but SCHEDULE(*JOBEND) was specified. The job that produced the spooled file has not finished.<br><br><b>DEFERRED</b> Printing of the file has been deferred.<br><br><b>DELETED</b> This spooled file is no longer in the system. These spooled files are included in the list of spooled files only if the qualified job name is specified.<br><br><b>HELD</b> The file has been held.<br><br><b>MESSAGE WAITING</b> This file has a message which needs a reply or an action to be taken.<br><br><b>OPEN</b> The file has not been processed completely and is not ready to be selected by a writer.<br><br><b>PENDING</b> This file is pending (waiting) to be printed.<br><br><b>PRINTING</b> The file has been completely sent to the printer, but the print complete status has not been sent back.<br><br><b>READY</b> The file is available to be written.<br><br><b>SAVED</b> The file has been written and then saved. This file remains saved until it is released.<br><br><b>SENDING</b> This spooled file is being sent or has been sent to a remote system.<br><br><b>WRITING</b> This file currently is being produced by the writer on an output device. |
| OUTPUT_PRIORITY     | INTEGER      | The priority of the spooled file. The priority ranges from 1 (highest) to 9 (lowest).<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| CREATION_TIMESTAMP  | TIMESTAMP(0) | The timestamp, based on local job time, when the file was opened.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| USER_DATA           | VARCHAR(10)  | The user-specified data that describes the file.<br><br>Contains the null value if there is no user-specified data or STATUS is DELETED.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



Table 186. SPOOLED\_FILE\_INFO table function (continued)

| Column Name                    | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SIZE                           | BIGINT      | The spooled file size in bytes. The size of the spooled file is the data stream size plus the spooled file's attributes, plus the "overhead" storage used to store the spooled file's data stream.<br><br>Contains the null value if the size is not available or STATUS is DELETED.                                                                                                                                                                            |
| TOTAL_PAGES                    | INTEGER     | The total number of pages or number of records for the spooled file.<br><br>Contains the null value if the total number of pages is not available or STATUS is DELETED.                                                                                                                                                                                                                                                                                         |
| COPIES                         | INTEGER     | The number of copies remaining to print.<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                                                                                                                                                                                                 |
| QUALIFIED_JOB_NAME             | VARCHAR(28) | The qualified job name of the job that owns the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                   |
| JOB_NAME                       | VARCHAR(10) | The name of the job that owns the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| JOB_USER                       | VARCHAR(10) | The name of the user that owns the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                                |
| JOB_NUMBER                     | VARCHAR(6)  | The number of the job that owns the spooled file.                                                                                                                                                                                                                                                                                                                                                                                                               |
| FILE_AVAILABLE                 | VARCHAR(8)  | The schedule of the spooled file.<br><br><b>*IMMED</b> The spooled file is schedule immediate. A spooling writer can process the spooled file immediately.<br><br><b>*FILEEND</b> The spooled file is schedule file end. A spooling writer cannot process the spooled file until it has been closed.<br><br><b>*JOBEND</b> The spooled file is schedule job end. A spooling writer cannot process the spooled file until the job of the spooled file has ended. |
| FORM_TYPE                      | VARCHAR(10) | Spooled file form type. The type of form to load in the printer to print this file.<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                                                                                                                                                      |
| OUTPUT_QUEUE_LIBRARY           | VARCHAR(10) | The library where the output queue is located.<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                                                                                                                                                                                           |
| OUTPUT_QUEUE                   | VARCHAR(10) | The name of the output queue in which the spooled file is located.<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                                                                                                                                                                       |
| ASP_NUMBER                     | INTEGER     | The auxiliary storage pool in which the spooled file resides.<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                                                                                                                                                                            |
| SYSTEM                         | VARCHAR(8)  | The name of the system where the job that created the spooled file ran.                                                                                                                                                                                                                                                                                                                                                                                         |
| INTERNET_PRINT_PROTOCOL_JOB_ID | INTEGER     | The IPP job identifier assigned by the system based on the output queue to which the file was added or moved. This value ranges from 1 to 2147483647 and is not guaranteed to be unique for a given output queue.<br><br>Contains the null value when STATUS is DELETED.                                                                                                                                                                                        |

## Example

- List all the spooled files for the current job.

```
SELECT * FROM TABLE(QSYS2.SPOOLED_FILE_INFO(JOB_NAME => '*'));
```

## Storage Services

These views provide information about storage and storage devices.

## ASP\_INFO view

The ASP\_INFO view returns information about auxiliary storage pools (ASPs).

The values returned for the columns in the view are similar to the values returned by the Work with Configuration Status (WRKCFGSTS) CL command and the [Open List of ASPs \(QYASPOL\) API](#).

**Authorization:** None required.

The following table describes the columns in the view. The system name is ASP\_INFO. The schema is QSYS2.

Table 187. ASP\_INFO view

| Column Name             | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DEVICE_DESCRIPTION_NAME | DEVD_NAME          | VARCHAR(10)<br>Nullable | The name of the device description that most recently brought the independent ASP (IASP) to varyon/active state.<br><br>Contains the null value if the ASP is not an IASP.                                                                                                                                                                                                                                               |
| ASP_NUMBER              | ASP_NUM            | INTEGER                 | A unique identifier for an ASP. Possible values are 1 through 255.<br><br><b>1</b> The system ASP<br><b>2-32</b> User ASPs<br><b>33-255</b> IASPs                                                                                                                                                                                                                                                                        |
| ASP_STATE               | ASP_STATE          | VARCHAR(10)             | The device configuration status of an ASP.<br><br><b>ACTIVE</b> The status of the ASP is active.<br><b>AVAILABLE</b> The status of the ASP is available.<br><b>FAILED</b> The status of the ASP is failed.<br><b>NONE</b> There is no status. This value is used for the system ASP and any basic user ASPs.<br><b>VARIED OFF</b> The status of the ASP is varyoff.<br><b>VARIED ON</b> The status of the ASP is varyon. |
| ASP_TYPE                | ASP_TYPE           | VARCHAR(9)<br>Nullable  | The use that is assigned to the ASP.<br><br><b>PRIMARY</b> The ASP is a primary ASP.<br><b>SECONDARY</b> The ASP is a secondary ASP.<br><b>SYSTEM</b> The ASP is the system ASP.<br><b>UDFS</b> The ASP is a user-defined file system ASP.<br><b>USER</b> The ASP is a user ASP.<br><br>Contains the null value for an IASP when the type cannot be determined.                                                          |
| RDB_NAME                | RDB_NAME           | VARCHAR(18)<br>Nullable | The name that is assigned to the database that this ASP defines.<br><br>Contains the null value if ASP_TYPE is UDFS or USER.                                                                                                                                                                                                                                                                                             |
| NUMBER_OF_DISK_UNITS    | DISK_UNITS         | INTEGER                 | The total number of disk units in the ASP. If mirroring is active for disk units within the ASP, the mirrored pair of units is counted as one.                                                                                                                                                                                                                                                                           |
| DISK_UNITS_PRESENT      | PRESENT            | VARCHAR(4)              | Indicates whether disk units in the ASP were found.<br><br><b>ALL</b> All disk units were found.<br><b>NONE</b> No disk units were found.<br><b>SOME</b> The disk unit that is used to provide the identity of the ASP was found but some other disk units were not found.                                                                                                                                               |

Table 187. ASP\_INFO view (continued)

| Column Name                    | System Column Name | Data Type           | Description                                                                                                                                                                                                                                                                |
|--------------------------------|--------------------|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOTAL_CAPACITY                 | TOTCAP             | BIGINT<br>Nullable  | The total number of used and unused megabytes in the ASP. A special value of -2 is returned if the size of this field is exceeded.<br><br>Contains the null value if the capacity cannot be determined.                                                                    |
| TOTAL_CAPACITY_AVAILABLE       | TOTCAPA            | BIGINT<br>Nullable  | The total number of unused megabytes in the ASP. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if the capacity cannot be determined.                                                                                |
| PROTECTED_CAPACITY             | PROTCAP            | BIGINT<br>Nullable  | The total number of used and unused megabytes in the ASP that are protected by mirroring or device parity. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if the capacity cannot be determined.                      |
| PROTECTED_CAPACITY_AVAILABLE   | PROTCAPA           | BIGINT<br>Nullable  | The number of unused megabytes in the ASP that are protected by mirroring or device parity. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if the capacity cannot be determined.                                     |
| UNPROTECTED_CAPACITY           | UNPROTCAP          | BIGINT<br>Nullable  | The total number of used and unused megabytes in the ASP that are not protected by mirroring or device parity. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if the capacity cannot be determined.                  |
| UNPROTECTED_CAPACITY_AVAILABLE | UNPROTCAPA         | BIGINT<br>Nullable  | The number of unused megabytes in the ASP that are not protected by mirroring or device parity. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if the capacity cannot be determined.                                 |
| SYSTEM_STORAGE                 | SYS_STG            | INTEGER<br>Nullable | The amount of storage in megabytes currently allocated in the ASP for operating system use.<br><br>Contains the null value if this is not the system ASP.                                                                                                                  |
| OVERFLOW_STORAGE               | OVER_STG           | BIGINT<br>Nullable  | The number of megabytes of storage that has overflowed from the user ASP into the system ASP. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if this is an IASP.                                                     |
| STORAGE_THRESHOLD_PERCENTAGE   | THRESHOLD          | INTEGER             | When the storage in the ASP reaches this percentage, a warning message is sent to the QSYSOPR message queue. When this percentage is reached for the system ASP (ASP 1), message CPF0907 is sent. When this percentage is reached for a user ASP, message CPI0953 is sent. |

Table 187. ASP\_INFO view (continued)

| Column Name              | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------|--------------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OVERFLOW_RECOVERY_RESULT | OVER_RES           | VARCHAR(7)<br>Nullable | <p>An indicator of the result of the ASP overflow recovery operation, which is performed during IPL at the user's request. When this operation is requested, an attempt is made to recover the user ASP from an overflow condition by moving overflowed auxiliary storage from the system ASP back to the user ASP during the storage management recovery step of an IPL.</p> <p><b>CANCEL</b> ASP overflow recovery was canceled prior to completion.</p> <p><b>FAIL</b> ASP overflow recovery failed due to insufficient space in the user ASP.</p> <p><b>SUCCESS</b> All overflowed storage was successfully moved.</p> <p>Contains the null value if this is an IASP.</p> |
| ERROR_LOG_SPACE          | ERR_SPACE          | INTEGER<br>Nullable    | <p>The number of megabytes of auxiliary storage allocated to the error log.</p> <p>Contains the null value if this is not the system ASP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| MACHINE_LOG_SPACE        | LOG_SPACE          | INTEGER<br>Nullable    | <p>The number of megabytes of auxiliary storage allocated to the machine log.</p> <p>Contains the null value if this is not the system ASP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| MACHINE_TRACE_SPACE      | TRC_SPACE          | INTEGER<br>Nullable    | <p>The number of megabytes of auxiliary storage allocated to the machine trace.</p> <p>Contains the null value if this is not the system ASP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MAIN_STORAGE_DUMP_SPACE  | MSD_SPACE          | INTEGER<br>Nullable    | <p>The number of megabytes of auxiliary storage allocated to the main storage dump space.</p> <p>Contains the null value if this is not the system ASP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| MICROCODE_SPACE          | MIC_SPACE          | INTEGER<br>Nullable    | <p>The number of megabytes of auxiliary storage allocated to the microcode and space used by the microcode.</p> <p>Contains the null value if this is an IASP that is varied off.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| END_IMMEDIATE            | END_IMMED          | VARCHAR(3)<br>Nullable | <p>This column only applies to the system ASP (ASP 1).</p> <p><b>NO</b> If a request for space in the system ASP cannot be satisfied because there is not enough storage, the system will be allowed to continue running.</p> <p><b>YES</b> If a request for space in the system ASP cannot be satisfied because there is not enough storage, the system will be ended immediately.</p> <p>Contains the null value if this is not the system ASP.</p>                                                                                                                                                                                                                         |

Table 187. ASP\_INFO view (continued)

| Column Name                    | System Column Name | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|--------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| COMPRESSION_RECOVERY_POLICY    | COMP_RECOV         | VARCHAR(18) | <p>The compression recovery policy for the ASP. If the ASP has compressed drives as part of its configuration, this value controls how overflow situations are handled for this ASP. The following policies allow the user to control what is done when the ASP appears full.</p> <p><b>OVERFLOW DELAY</b> When the ASP capacity is about to be exceeded, the operating system posts system reference code (SRC) A6xx 0277 in the system control panel and waits for a limited time for space to become available. If space becomes available before the limited time ends, the SRC is removed from the system control panel and normal operations resume. If space does not become available before the limited time ends, data overflows into the system ASP.</p> <p><b>OVERFLOW IMMEDIATE</b> When the ASP capacity is about to be exceeded, the data immediately overflows into the system ASP.</p> <p><b>WAIT</b> When the ASP capacity is about to be exceeded, the operating system posts SRC A6xx 0277 in the system control panel and waits indefinitely for space to become available. The user must take action before normal operation resumes. Possible actions include deleting objects from the ASP or changing the compression recovery policy to a value that allows the ASP to overflow.</p> |
| COMPRESSED_DISK_UNITS          | COMPRESSED         | VARCHAR(4)  | <p>Whether there are compressed disk units in the ASP.</p> <p><b>ALL</b> All disk units in this ASP are compressed.</p> <p><b>NONE</b> No compressed disk units in this ASP.</p> <p><b>SOME</b> Compressed and uncompressed disk units in this ASP.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| CHANGES_WRITTEN_TO_DISK        | WRITTEN            | VARCHAR(3)  | <p>An indicator of whether all changes made the previous time the IASP was online were written to disk. Varyoff processing attempts to write changed IASP storage but, in some failures, it may not be successful.</p> <p><b>NO</b> Not all changes were written to disk.</p> <p><b>YES</b> All changes were written to disk.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| MULTIPLE_CONNECTION_DISK_UNITS | MULT_CONN          | VARCHAR(3)  | <p>A disk unit may have multiple resource names. Each resource name represents a unique connection to the disk unit. All active connections are used to communicate with the disk unit. This attribute indicates whether the disk unit has more than one connection.</p> <p><b>NO</b> The disk unit has only one connection.</p> <p><b>YES</b> The disk unit has more than one connection.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 187. ASP\_INFO view (continued)

| Column Name    | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|--------------------|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BALANCE_STATUS | BALANCE            | VARCHAR(8)<br>Nullable | <p>The current status of the balance function for this ASP.</p> <p><b>COMPLETE</b> The ASP balance function has completed running. The ASP is completely balanced.</p> <p><b>ENDED</b> The ASP balance function has run, but was ended before the ASP was completely balanced. The Start ASP Balance (STRASPBAL) command can be used to restart the balance function.</p> <p><b>ENDING</b> The ASP balance function is currently in the process of ending. Either the time limit has run out or the End ASP Balance (ENDASPBAL) command was issued for this ASP.</p> <p><b>NONE</b> No balance activity has occurred for this ASP.</p> <p><b>RUNNING</b> The ASP balance function is currently running for this ASP.</p> <p>Contains the null value if ASP_STATE is not ACTIVE or AVAILABLE for IASP.</p> |

Table 187. ASP\_INFO view (continued)

| Column Name  | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|--------------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BALANCE_TYPE | BAL_TYPE           | VARCHAR(21)<br>Nullable | <p>The type of balance activity that is currently running or was done last.</p> <p><b>CAPACITY BALANCING</b> Capacity balancing. Capacity balancing redistributes data so that the percentage of disk space used is the same on all disk units within the ASP.</p> <p><b>CLEAR COLLECTION DATA</b> Clear collection data. Clear collection data removes the trace data created by running the Trace ASP Balance (TRCASPBAL) command.</p> <p><b>HSM BALANCING</b> Hierarchical Storage Management (HSM) balancing. HSM balancing can be run only on an ASP that contains a mixture of high-performance and low-performance disk units. An example of low-performance disk units is compressed disk units. The HSM balance function moves high-use data to high-performance units and moves low-use data to low-performance units. The high-use and low-use data is identified by running the Trace ASP Balance (TRCASPBAL) command.</p> <p><b>MOVE DATA</b> Move data. Move data is used to reduce the down time associated with removing a disk unit. The Check ASP Balance (CHKASPBAL) command can be used to determine which units are currently marked to no longer receive new allocations and to have their existing allocations moved to other disk units.</p> <p><b>MP BALANCING</b> Media Preference (MP) balancing. MP balancing can be run only on an ASP that contains a mixture of Solid State Disk (SSD) units and Hard Disk Drive (HDD) units. The goal of the MP balance function is to have high-use data on SSD units and low-use data on HDD units. The high-use and low-use data is identified by running the Trace ASP Balance (TRCASPBAL) command.</p> <p><b>NONE</b> No ASP balance activity was requested for the ASP.</p> <p><b>USAGE BALANCING</b> Usage balancing. Usage balancing redistributes data so that the percentage of disk activity is the same on all disk units within the ASP. High-use and low-use data is identified by running the Trace ASP Balance (TRCASPBAL) command. Usage balancing moves data among the disk units, guided by the trace results, in an attempt to equalize the utilizations.</p> <p>Contains the null value if ASP_STATE is not ACTIVE or AVAILABLE for IASP.</p> |

Table 187. ASP\_INFO view (continued)

| Column Name               | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------|--------------------|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BALANCE_DATA_MOVED        | BAL_MOVED          | BIGINT<br>Nullable       | The number of megabytes that have been moved by the balance function. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if BALANCE_STATUS is not RUNNING.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| BALANCE_DATA_REMAINING    | BAL_REMAIN         | BIGINT<br>Nullable       | The number of megabytes that remain to be moved by the balance function before the move is considered complete. A special value of -2 is returned if the value was too big to return.<br><br>Contains the null value if BALANCE_STATUS is not RUNNING.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| BALANCE_TIMESTAMP         | BAL_TIME           | TIMESTAMP(0)<br>Nullable | The timestamp of the last status change for the balance function.<br><br>Contains the null value when BALANCE_TYPE is NONE or the null value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| TRACE_STATUS              | TRC_STATUS         | VARCHAR(10)<br>Nullable  | The current status of the trace function. The trace gathers statistics about the data on the disk units within the ASP. This data is used by the balance functions.<br><br><p><b>CLEARING</b> The trace data for this ASP is being cleared.</p> <p><b>COMPLETE 1</b> The trace function has completed running. The statistics for the ASP have been gathered and are ready for the balance function to start.</p> <p><b>COMPLETE 2</b> The trace function has completed and the statistics for the ASP have been gathered. The ASP is ready for further collection or for the balance function to start.</p> <p><b>ENDING</b> The trace function is currently in the process of ending. Either the time limit has run out or the trace was stopped through use of the Trace ASP Balance (TRCASPBAL) command.</p> <p><b>NONE</b> There is no current trace data for this ASP.</p> <p><b>RUNNING</b> The trace function is currently running for this ASP.</p> <p>Contains the null value if ASP_STATE is not ACTIVE or AVAILABLE for IASP.</p> |
| TRACE_DURATION            | TRC_DUR            | INTEGER<br>Nullable      | The number of minutes that the trace function has run collecting data for this ASP. The trace can be run multiple times for an ASP.<br><br>Contains the null value when TRACE_STATUS is NONE or the null value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| TRACE_TIMESTAMP           | TRC_TIME           | TIMESTAMP(0)<br>Nullable | The timestamp of the last status change for the trace function.<br><br>Contains the null value when TRACE_STATUS is NONE or the null value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| RESOURCE_NAME             | RESOURCE           | VARCHAR(10)<br>Nullable  | The resource name that identifies the ASP by which a collection of disks is known.<br><br>Contains the null value for the system ASP, any user ASPs, and for an IASP where the name cannot be determined.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PRIMARY_ASP_RESOURCE_NAME | PRIMARY            | VARCHAR(10)<br>Nullable  | The resource name of the primary ASP for a secondary ASP.<br><br>Contains the null value if ASP_TYPE is not SECONDARY.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |



## Example

- Show ASP information for the partition.

```
SELECT * FROM QSYS2.ASP_INFO;
```

## ASP\_JOB\_INFO view

The ASP\_JOB\_INFO view returns information about active jobs that are using an independent auxiliary storage pool (IASP).

The information is similar to what is returned by the Work with ASP Jobs (WRKASPJOB) CL command.

**Authorization:** None required to see information for jobs where the caller's user profile is the same as the job user identity of the job for which the information is being returned. Otherwise, the caller must have \*JOBCTL special authority, or QIBM\_DB\_SQLADM or QIBM\_DB\_SYSMON function usage authority.

The following table describes the columns in the view. The system name is ASPJ\_INFO. The schema is QSYS2.

Table 188. ASP\_JOB\_INFO view

| Column Name          | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IASP_NAME            | IASP_NAME          | VARCHAR(10)             | The name of the independent ASP (IASP) device description.                                                                                                                                                                                                                                                                                                                                                                        |
| IASP_NUMBER          | IASPNUM            | INTEGER                 | The number associated with the ASP device.                                                                                                                                                                                                                                                                                                                                                                                        |
| JOB_NAME             | JOB_NAME           | VARCHAR(28)             | The qualified job name.                                                                                                                                                                                                                                                                                                                                                                                                           |
| JOB_STATUS           | JOB_STATUS         | VARCHAR(4)<br>Nullable  | The status of the initial thread of the job.<br><br>For the list of values see <a href="#">Work Management API Attribute Descriptions in Application Programming Interfaces</a> and search on "Active job status".                                                                                                                                                                                                                |
| JOB_TYPE             | JOB_TYPE           | VARCHAR(3)<br>Nullable  | Type of active job.<br><br><b>ASJ</b> Autostart<br><b>BCH</b> Batch<br><b>BCI</b> Batch Immediate<br><b>EVK</b> Started by a procedure start request<br><b>INT</b> Interactive<br><b>M36</b> Advanced 36 server job<br><b>MRT</b> Multiple requester terminal<br><b>PDJ</b> Print driver job<br><b>PJ</b> Prestart job<br><b>RDR</b> Spool reader<br><b>SBS</b> Subsystem monitor<br><b>SYS</b> System<br><b>WTR</b> Spool writer |
| AUTHORIZATION_NAME   | USER_NAME          | VARCHAR(10)<br>Nullable | The user profile under which the initial thread is running at this time.                                                                                                                                                                                                                                                                                                                                                          |
| SUBSYSTEM_NAME       | SUB_NAME           | VARCHAR(10)<br>Nullable | Name of subsystem where job is running.                                                                                                                                                                                                                                                                                                                                                                                           |
| SQL_STATEMENT_STATUS | SQL_STATUS         | VARCHAR(8)<br>Nullable  | The status of SQL within this job.<br><br><b>ACTIVE</b> An SQL statement is currently running<br><b>COMPLETE</b> At least one SQL statement has run and has completed<br><br>Contains the null value if no SQL statement has been run.                                                                                                                                                                                            |

Table 188. ASP\_JOB\_INFO view (continued)

| Column Name                   | System Column Name | Data Type                  | Description                                                                                                                                                                                                                            |
|-------------------------------|--------------------|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SQL_STATEMENT_TEXT            | SQL_STMT           | VARCHAR(10000)<br>Nullable | Statement text of the last SQL statement to run or the SQL statement that is currently running. The statement text will be truncated if it is longer than the column.<br><br>Contains the null value if no SQL statement has been run. |
| SQL_STATEMENT_START_TIMESTAMP | SQL_TIME           | TIMESTAMP<br>Nullable      | The timestamp of the execution start for an active SQL statement.<br><br>Contains the null value if there is no active SQL statement.                                                                                                  |
| ASP_TYPE                      | ASP_TYPE           | VARCHAR(9)<br>Nullable     | The use that is assigned to the ASP.<br><br><b>PRIMARY</b> The ASP is a primary ASP.<br><b>SECONDARY</b> The ASP is a secondary ASP.<br><b>UDFS</b> The ASP is a user-defined file system ASP.                                         |
| RDB_NAME                      | RDB_NAME           | VARCHAR(18)<br>Nullable    | The name that is assigned to the database that this ASP defines.<br><br>Contains the null value if ASP_TYPE is not PRIMARY or SECONDARY.                                                                                               |

## Example

List all the jobs that are active for IASP33.

```
SELECT JOB_NAME, JOB_STATUS, JOB_TYPE, AUTHORIZATION_NAME
FROM QSYS2.ASP_JOB_INFO
WHERE IASP_NAME = 'IASP33';
```

## ASP\_VARY\_INFO view

The ASP\_VARY\_INFO view returns one row for each step associated with a vary on or vary off operation for all independent ASP devices.

The values returned for the columns in the view are similar to the values returned by the Display ASP Status (DSPASPSTS) CL command.

**Authorization:** The privileges held by the authorization ID of the statement must have \*USE authority to the independent ASP device description. If the user does not have \*USE authority to all independent ASP device descriptions, a warning is returned to indicate that partial data is returned.

The following table describes the columns in the view. The system name is VARY\_INFO. The schema is QSYS2.

Table 189. ASP\_VARY\_INFO view

| Column Name      | System Column Name | Data Type   | Description                                                                                                                                                    |
|------------------|--------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IASP_NAME        | IASP_NAME          | VARCHAR(10) | The name of the ASP device description.                                                                                                                        |
| OPERATION_NUMBER | OP_NUMBER          | INTEGER     | A value for an instance of a vary on or vary off operation, where the highest number is the most recent operation. The most recent 64 operations are returned. |
| OPERATION_TYPE   | OP_TYPE            | VARCHAR(8)  | The type of vary operation.<br><br><b>VARY OFF</b><br><b>VARY ON</b>                                                                                           |

Table 189. ASP\_VARY\_INFO view (continued)

| Column Name     | System Column Name | Data Type                    | Description                                                                                                                                                                                                                                                                                |
|-----------------|--------------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OPERATION_STATE | OP_STATE           | VARCHAR(8)                   | The state of the entire operation.<br><br><b>ACTIVE</b> The operation is active. STEP_STATE shows the status of the steps that are part of the operation.<br><br><b>COMPLETE</b> The operation completed successfully.<br><br><b>FAILED</b> The operation failed to complete successfully. |
| STEP            | STEP               | VARGRAPHIC(50)<br>CCSID 1200 | The description of the operation step.                                                                                                                                                                                                                                                     |
| STEP_STATE      | STEP_STATE         | VARCHAR(8)                   | The state of the operation step.<br><br><b>ACTIVE</b> The step is active.<br><br><b>COMPLETE</b> The step completed successfully.<br><br><b>FAILED</b> The step failed to complete successfully.                                                                                           |
| START_TIMESTAMP | START              | TIMESTAMP                    | The timestamp for the start of this operation step.                                                                                                                                                                                                                                        |
| END_TIMESTAMP   | END                | TIMESTAMP<br>Nullable        | The timestamp for the end of this operation step.<br>Contains the null value if the operation step has not completed or may never complete.                                                                                                                                                |
| DURATION        | DURATION           | DECIMAL(12,6)<br>Nullable    | The time duration, in seconds, of this operation step.<br>Contains the null value if the operation step has not completed or may never complete.                                                                                                                                           |
| JOB_NAME        | JOB_NAME           | VARCHAR(28)<br>Nullable      | The qualified job name that initiated this vary operation.<br>Contains the null value if the job name is not available.                                                                                                                                                                    |
| IASP_NUMBER     | IASPNUM            | INTEGER                      | The number associated with the ASP device.                                                                                                                                                                                                                                                 |

## Example

- Return the steps from available vary on operations, listed from most expensive to least expensive.

```
SELECT * FROM QSYS2.ASP_VARY_INFO
WHERE OPERATION_TYPE = 'VARY ON'
ORDER BY IASP_NAME, DURATION DESC;
```

- Create a table to retain vary on historical data. Populate it with the current available values.

```
CREATE TABLE VARY_HISTORY AS
(SELECT * FROM QSYS2.ASP_VARY_INFO) WITH DATA;
```

- Update the table that contains vary on historical data with any new rows.

```
MERGE INTO VARY_HISTORY H
USING QSYS2.ASP_VARY_INFO N
ON H.OPERATION_NUMBER = N.OPERATION_NUMBER
WHEN NOT MATCHED THEN
INSERT VALUES (N.IASP_NAME, N.OPERATION_NUMBER, N.OPERATION_TYPE,
N.OPERATION_STATE, N.STEP, N.STEP_STATE,
N.START_TIMESTAMP, N.END_TIMESTAMP, N.DURATION,
N.JOB_NAME, N.IASP_NUMBER);
```

## MEDIA\_LIBRARY\_INFO view

The MEDIA\_LIBRARY\_INFO view returns information that can also be seen through the Work with Media Library Status (WRKMLBSTS) command interface.

The following table describes the columns in the view. The system name is MEDIA\_INFO. The schema is QSYS2.

Table 190. MEDIA\_LIBRARY\_INFO view

| Column Name     | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------|--------------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DEVICE_NAME     | DEVICE             | VARCHAR(10)             | The name of the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| DEVICE_STATUS   | DEVICE_STS         | VARCHAR(20)             | The status of the device. The most common values are:<br><br><b>VARIED ON</b> The media library device is varied on.<br><b>VARIED OFF</b> The media library device is varied off.<br><b>ACTIVE</b> The resource is currently in use by a job under this media library.<br><br>See <a href="#">List Configuration Descriptions API</a> for a complete list of status values.                                                                                                                                                                                                                                                                                                        |
| DEVICE_TYPE     | DEVICE_TYP         | VARCHAR(10)             | The type of device. Contains the special value *RSRCNAME if the device type is determined by the resource in the RESOURCE_NAME column.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| DEVICE_MODEL    | DEVICE_MDL         | VARCHAR(10)             | The model number of the device. Contains the special value *RSRCNAME if the device model is determined by the resource in the RESOURCE_NAME column.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| RESOURCE_NAME   | RESOURCE           | VARCHAR(10)<br>Nullable | The name of the resource.<br><br>Contains the null value if the DEVICE_STATUS column has a value of VARIED_OFF, or if the tape library does not have any associated tape resources.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| RESOURCE_STATUS | RSRC_STS           | VARCHAR(11)<br>Nullable | The status of the resource.<br><br><b>OPERATIONAL</b> The resource is working and the system can address the tape drive resource.<br><br><b>ACTIVE</b> The resource is currently in use by a job under this media library.<br><br><b>UNAVAILABLE</b> The resource is currently not available because it may be in use by another object, another client, or DST.<br><br><b>FAILED</b> The resource is not operational and the system can no longer communicate with that resource. A hardware problem may have occurred.<br><br>Contains the null value if the DEVICE_STATUS column has a value of VARIED_OFF, or if the tape library does not have any associated tape resources. |

Table 190. MEDIA\_LIBRARY\_INFO view (continued)

| Column Name                   | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|--------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RESOURCE_ALLOCATION_STATUS    | ALLOCATION         | VARCHAR(11)<br>Nullable | <p>Current allocation status for the resource.</p> <p><b>ALLOCATED</b> For a tape media library device the resource is exclusively assigned to this system and cannot be accessed by another system. For an optical media library device the drive is available for use by this media library.</p> <p><b>UNPROTECTED</b> A tape resource is not exclusively assigned to this system. This resource can be assigned to this system when no other system has already assigned the resource.</p> <p><b>DEALLOCATED</b> For a tape media library the resource is not assigned to this system and is not available to respond to requests. For an optical media library the device is not available for use by this media library.</p> <p><b>STAND-ALONE</b> A tape resource is not available. The tape resource is reserved by a varied on stand-alone tape device description for non-library mode use.</p> <p><b>*UNKNOWN</b> An optical media library is varied off or failed. The current allocation for a resource cannot be determined.</p> <p>Contains the null value if the DEVICE_STATUS column has a value of VARIED_OFF, or if the tape library does not have any associated tape resources.</p> |
| RESOURCE_ALLOCATION_PRIORITY  | ALLOC_PRTY         | VARCHAR(4)              | <p>The priority of a job when requesting a resource. 1 is highest priority, 99 is lowest. Can contain the following special value:</p> <p><b>*JOB</b> The priority of the job is used as the resource allocation priority.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| INITIAL_MOUNT_WAIT_TIME       | INIT_WAIT          | VARCHAR(6)              | <p>The maximum amount of time a request will wait for allocation of a tape resource for the initial mount. Contains either a numeric string representing the number of minutes or one of the following special values:</p> <p><b>*JOB</b> The allocation wait time is determined by the default wait time attribute of the job requesting the allocation, rounded up to the nearest minute.</p> <p><b>*IMMED</b> The request will not wait for a tape resource to become available.</p> <p><b>*NOMAX</b> The request will wait until a tape resource is available.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| END_OF_VOLUME_MOUNT_WAIT_TIME | END_WAIT           | VARCHAR(6)              | <p>The maximum amount of time a request will wait for allocation of a tape resource for the end of volume mount. Contains either a numeric string representing the number of minutes or one of the following special values:</p> <p><b>*JOB</b> The allocation wait time is determined by the default wait time attribute of the job requesting the allocation, rounded up to the nearest minute.</p> <p><b>*IMMED</b> The request will not wait for a tape resource to become available.</p> <p><b>*NOMAX</b> The request will wait until a tape resource is available.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| DEVICE_DESCRIPTION            | DEVICE_DES         | VARCHAR(50)             | The text description of the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Example


Return information about all media library devices.

```
SELECT * FROM QSYS2.MEDIA_LIBRARY_INFO
```

## SYSDISKSTAT table function

The SYSDISKSTAT table function contains information about disks. It provides an option to reset the baseline for collecting statistical information.

**Authorization:** None required.



The schema is QSYS2.

### **reset-statistics**

A character or graphic string expression that contains a value of YES or NO.

If this parameter has a value of YES, statistics are reset such that the time of this query execution is used as the new baseline. The columns that contain this statistical data have names that are prefixed with ELAPSED\_. Future invocations of SYSDISKSTAT within this connection will return statistical detail relative to the new baseline. If this parameter has a value of NO, statistics are not reset for the invocation. If this parameter is not specified, the default is NO.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 191. SYSDISKSTAT table function

| Column Name             | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASP_NUMBER              | SMALLINT     | Specifies the independent auxiliary storage pool (IASP) number.                                                                                                                                                                                                                                                                                                                     |
| DISK_TYPE               | VARCHAR(4)   | Disk type number of the disk.                                                                                                                                                                                                                                                                                                                                                       |
| DISK_MODEL              | VARCHAR(4)   | Model number of the disk.                                                                                                                                                                                                                                                                                                                                                           |
| UNIT_NUMBER             | SMALLINT     | Unit number of the disk.                                                                                                                                                                                                                                                                                                                                                            |
| SERIAL_NUMBER           | VARCHAR(15)  | The serial number of the disk unit.                                                                                                                                                                                                                                                                                                                                                 |
| RESOURCE_NAME           | VARCHAR(10)  | The unique system-assigned name of the disk unit.                                                                                                                                                                                                                                                                                                                                   |
| RESOURCE_STATUS         | VARCHAR(7)   | The status of the resource.<br><br><b>ACTIVE</b> RESOURCE_NAME is active.<br><b>PASSIVE</b> RESOURCE_NAME is not active.<br>Contains the null value if the path status is not known.                                                                                                                                                                                                |
| MULTIPLE_PATH_UNIT      | VARCHAR(3)   | A disk unit may have multiple resource names. Each resource name represents a unique connection to the disk unit. All active connections are used to communicate with the disk unit. This attribute indicates whether the disk unit has more than one connection.<br><br><b>NO</b> The disk unit has only one connection.<br><b>YES</b> The disk unit has more than one connection. |
| UNIT_TYPE               | SMALLINT     | Indicates the type of disk unit:<br><br><b>0</b> Not solid state disk<br><b>1</b> Solid state disk (SSD)                                                                                                                                                                                                                                                                            |
| UNIT_STORAGE_CAPACITY   | BIGINT       | Unit storage capacity has the same value as the unit media capacity for configured disk units. This value is 0 for non-configured units.                                                                                                                                                                                                                                            |
| UNIT_SPACE_AVAILABLE    | BIGINT       | Space (in bytes) available on the unit for use.                                                                                                                                                                                                                                                                                                                                     |
| UNIT_SPACE_AVAILABLE_GB | BIGINT       | Space, in billions of bytes, available on the unit for use.                                                                                                                                                                                                                                                                                                                         |
| PERCENT_USED            | DECIMAL(7,3) | The percentage that the disk unit has been consumed.                                                                                                                                                                                                                                                                                                                                |
| UNIT_MEDIA_CAPACITY     | BIGINT       | Storage capacity (in bytes) of the unit.                                                                                                                                                                                                                                                                                                                                            |

Table 191. SYSDISKSTAT table function (continued)

| Column Name                | Data Type           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNIT_MEDIA_CAPACITY_GB     | BIGINT              | Storage capacity, in billions of bytes, of the unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| STORAGE_FOR_SYSTEM         | BIGINT              | The amount of auxiliary storage on the disk unit, in millions of bytes, reserved for use by the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| STORAGE_ALLOCATION_ALLOWED | VARCHAR(3)          | An indicator of whether new storage allocations are allowed on the disk unit.<br><br><b>NO</b> The disk unit does not allow new storage allocations.<br><b>YES</b> The disk unit allows new storage allocations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| PROTECTION_TYPE            | VARCHAR(8)          | The type of protection that has been assigned to this disk unit.<br><br><b>MIRRORED</b> The ASP is under system mirrored protection provided by the system software.<br><b>PARITY</b> This disk unit is part of a parity protection array.<br><br>Contains the null value if no storage protection has been set up for this disk unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| PROTECTION_STATUS          | VARCHAR(21)         | The disk protection status for the disk unit, when the unit is under device parity protection or the ASP is under mirrored protection.<br><br><b>ACTIVE</b> The disk unit is active.<br><b>BUSY</b> The disk unit is busy.<br><b>DEGRADED</b> There is a hardware failure within the disk subsystem that affects performance, but does not affect the function of the disk unit.<br><b>FAILED</b> The disk unit has failed.<br><b>HARDWARE FAILURE</b> There is a hardware failure within the disk subsystem that does not affect the function or performance of the disk unit.<br><b>NOT READY</b> The disk unit is not ready.<br><b>PARITY REBUILD</b> The disk unit's parity protection is being rebuilt.<br><b>POWER LOSS</b> The disk unit is not operational.<br><b>READ WRITE PROTECTED</b> The disk unit is read/write protected.<br><b>RESUME</b> The unit is part of a mirrored ASP and mirroring is in the process of being resumed on this unit.<br><b>RESUME PENDING</b> The unit is part of a mirrored independent ASP which is varied off. Mirror synchronization will resume when the independent ASP is varied on.<br><b>SUSPEND</b> The unit is part of a mirrored ASP and mirroring is suspended on this unit.<br><b>UNKNOWN</b> The disk unit has returned a status that is not recognizable by the system.<br><b>UNPROTECTED</b> Some other disk unit in the disk subsystem has failed.<br><b>WRITE PROTECTED</b> The disk unit is write protected.<br><br>Contains the null value if PROTECTION_TYPE is null. |
| RAID_TYPE                  | VARCHAR(6)          | The type of RAID protection that has been assigned to this disk unit.<br><br><b>RAID5</b> This disk unit has been set up with RAID 5 protection.<br><b>RAID6</b> This disk unit has been set up with RAID 6 protection.<br><b>RAID10</b> This disk unit has been set up with RAID 10 protection.<br><br>Contains the null value if PROTECTION_TYPE is not PARITY or no storage protection has been set up for this disk unit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| MIRRORED_SUBUNIT           | CHAR(1)<br>Nullable | Whether the disk unit is for subunit A or B of a mirrored pair.<br><br><b>A</b> This entry is for subunit A.<br><b>B</b> This entry is for subunit B.<br><br>Contains the null value if the unit is not a mirrored pair or if the information is not available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 191. SYSDISKSTAT table function (continued)

| Column Name                  | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LOGICAL_MIRRORED_PAIR_STATUS | CHAR(1)                   | Indicates the status of a mirrored pair of disks:<br><br><b>0</b> Indicates that one mirrored unit of a mirrored pair is not active.<br><b>1</b> Indicates that both mirrored units of a mirrored pair are active.<br><br>Contains the null value if PROTECTION_TYPE is not MIRRORED or no storage protection has been set up for this disk unit.                 |
| MIRRORED_UNIT_STATUS         | CHAR(1)                   | Indicates the status of a mirrored unit:<br><br><b>1</b> Indicates that this mirrored unit of a mirrored pair is active (online with current data).<br><b>2</b> Indicates that this mirrored unit is being synchronized.<br><b>3</b> Indicates that this mirrored unit is suspended.<br><br>Contains the null value if PROTECTION_TYPE is not MIRRORED.           |
| AVAILABILITY_PARITY_SET_UNIT | VARCHAR(3)                | Whether the disk unit is in an availability parity set.<br><br><b>NO</b> The disk unit is not in an availability parity set.<br><b>YES</b> The disk unit is in an availability parity set.                                                                                                                                                                        |
| HYPERSWAP                    | VARCHAR(3)                | Whether unit is using HyperSwap.<br><br><b>NO</b> Unit is not using HyperSwap.<br><b>YES</b> Unit is using HyperSwap.                                                                                                                                                                                                                                             |
| FIRMWARE_LEVEL               | VARCHAR(8)<br>Nullable    | The level of code running in the SSD device.<br>Contains the null value if this disk is not SSD or if the information is not available.                                                                                                                                                                                                                           |
| SSD_PART_NUMBER              | VARCHAR(12)<br>Nullable   | The part number as reported by the SSD device.<br>Contains the null value if this disk is not SSD or if the information is not available.                                                                                                                                                                                                                         |
| SSD_POWER_ON_DAYS            | BIGINT<br>Nullable        | The number of days that the SSD device has been active in a system.<br>Contains the null value if this disk is not SSD or if the information is not available.                                                                                                                                                                                                    |
| SSD_LIFE_REMAINING           | INTEGER<br>Nullable       | The percentage of the lifetime remaining for the SSD device. This estimates the percentage of usable function remaining for the drive before it should be replaced. Calculations for this percentage include more than just the number of bytes written and supported.<br>Contains the null value if this disk is not SSD or if the information is not available. |
| SSD_READ_WRITE_PROTECTED     | VARCHAR(3)<br>Nullable    | Whether the device is read/write protected.<br><br><b>NO</b> The SSD device is not read/write protected<br><b>YES</b> The SSD device is read/write protected<br><br>Contains the null value if this disk is not SSD or if the information is not available.                                                                                                       |
| SSD_BYTES_WRITTEN            | DECIMAL(20,0)<br>Nullable | The lifetime number of bytes, in gigabytes, that have been physically written to the NAND memory in this particular SSD disk unit. This is strongly related to bytes written by the applications using the drive, but will not match.<br>Contains the null value if this disk is not SSD or if the information is not available.                                  |
| SSD_SUPPORTED_BYTES_WRITTEN  | DECIMAL(20,0)<br>Nullable | The lifetime number of bytes, in gigabytes, that the SSD is expected to be able to physically write at a minimum. Additional writes beyond this number may start to fail due to the limited write endurance of a Read Intensive drive.<br>Contains the null value if this disk is not SSD or if the information is not available.                                 |
| SSD_PFA_WARNING              | VARCHAR(3)<br>Nullable    | Whether the Predictive Failure Analysis warning has been logged.<br><br><b>NO</b> The Predictive Failure Analysis warning has not been logged.<br><b>YES</b> The Predictive Failure Analysis warning has been logged.<br><br>Contains the null value if this disk is not SSD or if the information is not available.                                              |
| TOTAL_SAMPLE_COUNT           | BIGINT                    | The number of times the disk queue was checked to determine whether or not the queue is empty.                                                                                                                                                                                                                                                                    |



Table 191. SYSDISKSTAT table function (continued)

| Column Name                    | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOTAL_NOT_BUSY_COUNT           | BIGINT       | The number of times the disk queue was empty during the same time period that the sample count was taken. The busy count can be calculated as TOTAL_SAMPLE_COUNT - TOTAL_NOT_BUSY_COUNT.                                                                                                                                                                               |
| TOTAL_READ_REQUESTS            | BIGINT       | The number of input data transfer requests processed for the disk unit since the last IPL. This value is not directly related to the number of blocks transferred for the disk unit because the number of blocks to be transferred for a given transfer request can vary greatly. This value will wrap back to 1 when 2,147,483,647 is reached.                        |
| TOTAL_WRITE_REQUESTS           | BIGINT       | The number of output data transfer requests processed for the disk unit since the last IPL. This value is not directly related to the number of blocks transferred for the disk unit because the number of blocks to be transferred for a given transfer request can vary greatly. This value will wrap back to 1 when 2,147,483,647 is reached.                       |
| TOTAL_BLOCKS_READ              | BIGINT       | The number of 512-byte blocks transferred from the disk unit since the last IPL. This value will wrap back to 1 when 2,147,483,647 is reached.                                                                                                                                                                                                                         |
| TOTAL_BLOCKS_WRITTEN           | BIGINT       | The number of 512-byte blocks transferred to the disk unit since the last IPL. This value will wrap back to 1 when 2,147,483,647 is reached.                                                                                                                                                                                                                           |
| TOTAL_PERMANENT_BLOCKS_WRITTEN | BIGINT       | The number of 512-byte blocks of permanent storage transferred to the disk unit since the last IPL. This value will wrap back to 1 when 2,147,483,647 is reached.                                                                                                                                                                                                      |
| TOTAL_PERMANENT_WRITE_REQUESTS | BIGINT       | The number of output permanent data transfer requests processed for the disk unit since the last IPL. This value is not directly related to the permanent blocks transferred from main storage for the disk unit because the number of blocks transferred for a given transfer request can vary greatly. This value will wrap back to 1 when 2,147,483,647 is reached. |
| ELAPSED_TIME                   | INTEGER      | The time that has elapsed, in seconds, between the measurement start time and the current system time.                                                                                                                                                                                                                                                                 |
| ELAPSED_IO_REQUESTS            | DECIMAL(6,1) | The average number of I/O requests for read and write operations that occurred per second during the elapsed time.                                                                                                                                                                                                                                                     |
| ELAPSED_REQUEST_SIZE           | DECIMAL(6,1) | The average size of an I/O request in KB during the elapsed time.                                                                                                                                                                                                                                                                                                      |
| ELAPSED_READ_REQUESTS          | DECIMAL(6,1) | The average number of requests per second to transfer data from the disk unit during the elapsed time.                                                                                                                                                                                                                                                                 |
| ELAPSED_WRITE_REQUESTS         | DECIMAL(6,1) | The average number of requests per second to transfer data to the disk unit during the elapsed time.                                                                                                                                                                                                                                                                   |
| ELAPSED_DATA_READ              | DECIMAL(6,1) | The average amount of data, in KB, transferred from the disk unit, per request, during the elapsed time.                                                                                                                                                                                                                                                               |
| ELAPSED_DATA_WRITTEN           | DECIMAL(6,1) | The average amount of data, in KB, transferred to the disk unit, per request, during the elapsed time.                                                                                                                                                                                                                                                                 |
| ELAPSED_PERCENT_BUSY           | DECIMAL(4,1) | The estimated percentage of time the disk unit is being used during the elapsed time.                                                                                                                                                                                                                                                                                  |

## Example

- Return information about all disks, resetting the statistical information shown by the current job.

```
SELECT * FROM TABLE(QSYS2.SYSDISKSTAT(RESET_STATISTICS=>'YES'));
```

## SYSDISKSTAT view

The SYSDISKSTAT view contains information about spinning disk and solid-state drives (SSD).

The information returned is similar to the detail seen from the Work with Disk Status (WRKDSKSTS) command and from the Open List of ASPs (QYASPOL) API.

The view contains one or more rows for every disk unit on the system, including non-configured (unallocated) disk units. For non-configured units, the UNIT\_NUMBER is 0. For a disk which has multiple paths to the disk unit, there will be a row for each unique path to the disk unit. For such disks, the MULTIPLE\_PATH\_UNIT column will be YES and each RESOURCE\_NAME column will identify a different path to the disk unit.

**Authorization:** None required.

The following table describes the columns in the view. The system name is SYSDISKS. The schema is QSYS2.

Table 192. SYSDISKSTAT view

| Column Name                | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|--------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASP_NUMBER                 | ASP_NUMBER         | SMALLINT                 | Specifies the independent auxiliary storage pool (IASP) number.                                                                                                                                                                                                                                                                                                                     |
| DISK_TYPE                  | DISK_TYPE          | VARCHAR(4)               | Disk type number of the disk.                                                                                                                                                                                                                                                                                                                                                       |
| DISK_MODEL                 | DISK_MODEL         | VARCHAR(4)               | Model number of the disk.                                                                                                                                                                                                                                                                                                                                                           |
| UNIT_NUMBER                | UNITNBR            | SMALLINT                 | Unit number of the disk.                                                                                                                                                                                                                                                                                                                                                            |
| SERIAL_NUMBER              | SERIALNBR          | VARCHAR(15)              | The serial number of the disk unit.                                                                                                                                                                                                                                                                                                                                                 |
| RESOURCE_NAME              | RESOURCE           | VARCHAR(10)              | The unique system-assigned name of the disk unit.                                                                                                                                                                                                                                                                                                                                   |
| RESOURCE_STATUS            | PATHSTATUS         | VARCHAR(7)<br>Nullable   | The status of the resource.<br><br><b>ACTIVE</b> RESOURCE_NAME is active.<br><b>PASSIVE</b> RESOURCE_NAME is not active.<br><br>Contains the null value if the path status is not known.                                                                                                                                                                                            |
| MULTIPLE_PATH_UNIT         | MULTI_PATH         | VARCHAR(3)               | A disk unit may have multiple resource names. Each resource name represents a unique connection to the disk unit. All active connections are used to communicate with the disk unit. This attribute indicates whether the disk unit has more than one connection.<br><br><b>NO</b> The disk unit has only one connection.<br><b>YES</b> The disk unit has more than one connection. |
| UNIT_TYPE                  | UNIT_TYPE          | SMALLINT                 | Indicates the type of disk unit:<br><br><b>0</b> Not solid state disk<br><b>1</b> Solid state disk (SSD)                                                                                                                                                                                                                                                                            |
| UNIT_STORAGE_CAPACITY      | UNITSCAP           | BIGINT                   | Unit storage capacity has the same value as the unit media capacity for configured disk units. This value is 0 for non-configured units.                                                                                                                                                                                                                                            |
| UNIT_SPACE_AVAILABLE       | UNITSPACE          | BIGINT                   | Space (in bytes) available on the unit for use.                                                                                                                                                                                                                                                                                                                                     |
| UNIT_SPACE_AVAILABLE_GB    | UNITSPCGB          | BIGINT                   | Space, in billions of bytes, available on the unit for use.                                                                                                                                                                                                                                                                                                                         |
| PERCENT_USED               | PERCENTUSE         | DECIMAL(7,3)<br>Nullable | The percentage that the disk unit has been consumed.                                                                                                                                                                                                                                                                                                                                |
| UNIT_MEDIA_CAPACITY        | UNITMCAP           | BIGINT                   | Storage capacity (in bytes) of the unit.                                                                                                                                                                                                                                                                                                                                            |
| UNIT_MEDIA_CAPACITY_GB     | UNITMCAPGB         | BIGINT                   | Storage capacity, in billions of bytes, of the unit.                                                                                                                                                                                                                                                                                                                                |
| STORAGE_FOR_SYSTEM         | STORAGESYS         | BIGINT                   | The amount of auxiliary storage on the disk unit, in millions of bytes, reserved for use by the system.                                                                                                                                                                                                                                                                             |
| STORAGE_ALLOCATION_ALLOWED | NEW_ALLOC          | VARCHAR(3)               | An indicator of whether new storage allocations are allowed on the disk unit.<br><br><b>NO</b> The disk unit does not allow new storage allocations.<br><b>YES</b> The disk unit allows new storage allocations.                                                                                                                                                                    |
| PROTECTION_TYPE            | PROTECTION         | VARCHAR(8)<br>Nullable   | The type of protection that has been assigned to this disk unit.<br><br><b>MIRRORED</b> The ASP is under system mirrored protection provided by the system software.<br><b>PARITY</b> This disk unit is part of a parity protection array.<br><br>Contains the null value if no storage protection has been set up for this disk unit.                                              |

Table 192. SYSDISKSTAT view (continued)

| Column Name       | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PROTECTION_STATUS | STATUS             | VARCHAR(21)<br>Nullable | <p>The disk protection status for the disk unit, when the unit is under device parity protection or the ASP is under mirrored protection.</p> <p><b>ACTIVE</b> The disk unit is active.</p> <p><b>BUSY</b> The disk unit is busy.</p> <p><b>DEGRADED</b> There is a hardware failure within the disk subsystem that affects performance, but does not affect the function of the disk unit.</p> <p><b>FAILED</b> The disk unit has failed.</p> <p><b>HARDWARE FAILURE</b> There is a hardware failure within the disk subsystem that does not affect the function or performance of the disk unit.</p> <p><b>NOT READY</b> The disk unit is not ready.</p> <p><b>PARITY REBUILD</b> The disk unit's parity protection is being rebuilt.</p> <p><b>POWER LOSS</b> The disk unit is not operational.</p> <p><b>READ WRITE PROTECTED</b> The disk unit is read/write protected.</p> <p><b>RESUME</b> The unit is part of a mirrored ASP and mirroring is in the process of being resumed on this unit.</p> <p><b>RESUME PENDING</b> The unit is part of a mirrored independent ASP which is varied off. Mirror synchronization will resume when the independent ASP is varied on.</p> <p><b>SUSPEND</b> The unit is part of a mirrored ASP and mirroring is suspended on this unit.</p> <p><b>UNKNOWN</b> The disk unit has returned a status that is not recognizable by the system.</p> <p><b>UNPROTECTED</b> Some other disk unit in the disk subsystem has failed.</p> <p><b>WRITE PROTECTED</b> The disk unit is write protected.</p> <p>Contains the null value if PROTECTION_TYPE is null.</p> |
| RAID_TYPE         | RAID_TYPE          | VARCHAR(6)<br>Nullable  | <p>The type of RAID protection that has been assigned to this disk unit.</p> <p><b>RAID5</b> This disk unit has been set up with RAID 5 protection.</p> <p><b>RAID6</b> This disk unit has been set up with RAID 6 protection.</p> <p><b>RAID10</b> This disk unit has been set up with RAID 10 protection.</p> <p>Contains the null value if PROTECTION_TYPE is not PARITY or no storage protection has been set up for this disk unit.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| MIRRORED_SUBUNIT  | SUBUNIT            | CHAR(1)<br>Nullable     | <p>Whether the disk unit is for subunit A or B of a mirrored pair.</p> <p><b>A</b> This entry is for subunit A.</p> <p><b>B</b> This entry is for subunit B.</p> <p>Contains the null value if the unit is not a mirrored pair or if the information is not available.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 192. SYSDISKSTAT view (continued)

| Column Name                  | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------|--------------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LOGICAL_MIRRORED_PAIR_STATUS | MIRRORPS           | CHAR(1)<br>Nullable       | <p>Indicates the status of a mirrored pair of disks:</p> <ul style="list-style-type: none"> <li><b>0</b> Indicates that one mirrored unit of a mirrored pair is not active.</li> <li><b>1</b> Indicates that both mirrored units of a mirrored pair are active.</li> </ul> <p>Contains the null value if PROTECTION_TYPE is not MIRRORED or no storage protection has been set up for this disk unit.</p>             |
| MIRRORED_UNIT_STATUS         | MIRRORUS           | CHAR(1)<br>Nullable       | <p>Indicates the status of a mirrored unit:</p> <ul style="list-style-type: none"> <li><b>1</b> Indicates that this mirrored unit of a mirrored pair is active (online with current data).</li> <li><b>2</b> Indicates that this mirrored unit is being synchronized.</li> <li><b>3</b> Indicates that this mirrored unit is suspended.</li> </ul> <p>Contains the null value if PROTECTION_TYPE is not MIRRORED.</p> |
| AVAILABILITY_PARITY_SET_UNIT | PARITY             | VARCHAR(3)                | <p>Whether the disk unit is in an availability parity set.</p> <ul style="list-style-type: none"> <li><b>NO</b> The disk unit is not in an availability parity set.</li> <li><b>YES</b> The disk unit is in an availability parity set.</li> </ul>                                                                                                                                                                    |
| HYPERSWAP                    | HYPERSWAP          | VARCHAR(3)                | <p>Whether unit is using HyperSwap.</p> <ul style="list-style-type: none"> <li><b>NO</b> Unit is not using HyperSwap.</li> <li><b>YES</b> Unit is using HyperSwap.</li> </ul>                                                                                                                                                                                                                                         |
| FIRMWARE_LEVEL               | FIRMWARE           | VARCHAR(8)<br>Nullable    | <p>The level of code running in the SSD device.</p> <p>Contains the null value if this disk is not SSD or if the information is not available.</p>                                                                                                                                                                                                                                                                    |
| SSD_PART_NUMBER              | SSD_PART           | VARCHAR(12)<br>Nullable   | <p>The part number as reported by the SSD device.</p> <p>Contains the null value if this disk is not SSD or if the information is not available.</p>                                                                                                                                                                                                                                                                  |
| SSD_POWER_ON_DAYS            | SSD_DAYS           | BIGINT<br>Nullable        | <p>The number of days that the SSD device has been active in a system.</p> <p>Contains the null value if this disk is not SSD or if the information is not available.</p>                                                                                                                                                                                                                                             |
| SSD_LIFE_REMAINING           | SSD_LIFE           | INTEGER<br>Nullable       | <p>The percentage of the lifetime remaining for the SSD device. This estimates the percentage of usable function remaining for the drive before it should be replaced. Calculations for this percentage include more than just the number of bytes written and supported.</p> <p>Contains the null value if this disk is not SSD or if the information is not available.</p>                                          |
| SSD_READ_WRITE_PROTECTED     | SSD_PROT           | VARCHAR(3)<br>Nullable    | <p>Whether the device is read/write protected.</p> <ul style="list-style-type: none"> <li><b>NO</b> The SSD device is not read/write protected</li> <li><b>YES</b> The SSD device is read/write protected</li> </ul> <p>Contains the null value if this disk is not SSD or if the information is not available.</p>                                                                                                   |
| SSD_BYTES_WRITTEN            | SSD_WRITE          | DECIMAL(20,0)<br>Nullable | <p>The lifetime number of bytes, in gigabytes, that have been physically written to the NAND memory in this particular SSD disk unit. This is strongly related to bytes written by the applications using the drive, but will not match.</p> <p>Contains the null value if this disk is not SSD or if the information is not available.</p>                                                                           |

Table 192. SYSDISKSTAT view (continued)

| Column Name                    | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|--------------------|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SSD_SUPPORTED_BYTES_WRITTEN    | SSD_MAX_W          | DECIMAL(20,0)<br>Nullable | The lifetime number of bytes, in gigabytes, that the SSD is expected to be able to physically write at a minimum. Additional writes beyond this number may start to fail due to the limited write endurance of a Read Intensive drive.<br><br>Contains the null value if this disk is not SSD or if the information is not available.                                  |
| SSD_PFA_WARNING                | SSD_PFA            | VARCHAR(3)<br>Nullable    | Whether the Predictive Failure Analysis warning has been logged.<br><br><b>NO</b> The Predictive Failure Analysis warning has not been logged.<br><br><b>YES</b> The Predictive Failure Analysis warning has been logged.<br><br>Contains the null value if this disk is not SSD or if the information is not available.                                               |
| TOTAL_SAMPLE_COUNT             | SAMPLED            | BIGINT                    | The number of times the disk queue was checked to determine whether or not the queue is empty.                                                                                                                                                                                                                                                                         |
| TOTAL_NOT_BUSY_COUNT           | NOT_BUSY           | BIGINT                    | The number of times the disk queue was empty during the same time period that the sample count was taken. The busy count can be calculated as TOTAL_SAMPLE_COUNT - TOTAL_NOT_BUSY_COUNT.                                                                                                                                                                               |
| TOTAL_READ_REQUESTS            | REQ_IN             | BIGINT                    | The number of input data transfer requests processed for the disk unit since the last IPL. This value is not directly related to the number of blocks transferred for the disk unit because the number of blocks to be transferred for a given transfer request can vary greatly. This value will wrap back to 1 when 2,147,483,647 is reached.                        |
| TOTAL_WRITE_REQUESTS           | REQ_OUT            | BIGINT                    | The number of output data transfer requests processed for the disk unit since the last IPL. This value is not directly related to the number of blocks transferred for the disk unit because the number of blocks to be transferred for a given transfer request can vary greatly. This value will wrap back to 1 when 2,147,483,647 is reached.                       |
| TOTAL_BLOCKS_READ              | BLOCK_IN           | BIGINT                    | The number of 512-byte blocks transferred from the disk unit since the last IPL. This value will wrap back to 1 when 2,147,483,647 is reached.                                                                                                                                                                                                                         |
| TOTAL_BLOCKS_WRITTEN           | BLOCK_OUT          | BIGINT                    | The number of 512-byte blocks transferred to the disk unit since the last IPL. This value will wrap back to 1 when 2,147,483,647 is reached.                                                                                                                                                                                                                           |
| TOTAL_PERMANENT_BLOCKS_WRITTEN | BLOCK_PERM         | BIGINT                    | The number of 512-byte blocks of permanent storage transferred to the disk unit since the last IPL. This value will wrap back to 1 when 2,147,483,647 is reached.                                                                                                                                                                                                      |
| TOTAL_PERMANENT_WRITE_REQUESTS | REQ_PERM           | BIGINT                    | The number of output permanent data transfer requests processed for the disk unit since the last IPL. This value is not directly related to the permanent blocks transferred from main storage for the disk unit because the number of blocks transferred for a given transfer request can vary greatly. This value will wrap back to 1 when 2,147,483,647 is reached. |
| ELAPSED_TIME                   | ELAP_TIME          | INTEGER                   | The time that has elapsed, in seconds, between the measurement start time and the current system time.                                                                                                                                                                                                                                                                 |
| ELAPSED_IO_REQUESTS            | ELAP_IO            | DECIMAL(6,1)<br>Nullable  | The average number of I/O requests for read and write operations that occurred per second during the elapsed time.                                                                                                                                                                                                                                                     |
| ELAPSED_REQUEST_SIZE           | ELAP_SIZE          | DECIMAL(6,1)<br>Nullable  | The average size of an I/O request in KB during the elapsed time.                                                                                                                                                                                                                                                                                                      |
| ELAPSED_READ_REQUESTS          | ELAP_REQ_R         | DECIMAL(6,1)<br>Nullable  | The average number of requests per second to transfer data from the disk unit during the elapsed time.                                                                                                                                                                                                                                                                 |
| ELAPSED_WRITE_REQUESTS         | ELAP_REQ_W         | DECIMAL(6,1)<br>Nullable  | The average number of requests per second to transfer data to the disk unit during the elapsed time.                                                                                                                                                                                                                                                                   |

Table 192. SYSDISKSTAT view (continued)

| Column Name          | System Column Name | Data Type                | Description                                                                                              |
|----------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------|
| ELAPSED_DATA_READ    | ELAP_DTA_R         | DECIMAL(6,1)<br>Nullable | The average amount of data, in KB, transferred from the disk unit, per request, during the elapsed time. |
| ELAPSED_DATA_WRITTEN | ELAP_DTA_W         | DECIMAL(6,1)<br>Nullable | The average amount of data, in KB, transferred to the disk unit, per request, during the elapsed time.   |
| ELAPSED_PERCENT_BUSY | ELAP_BUSY          | DECIMAL(4,1)<br>Nullable | The estimated percentage of time the disk unit is being used during the elapsed time.                    |

## Notes

The values in the ELAPSED\_ columns are based on the TOTAL\_ columns. When an ELAPSED calculation notices that the ending value is less than the value at the start of the time interval, it adds 2,147,483,647 to the ending value for an accurate result. When this happens, a warning SQLSTATE '01687' is issued. It is recommended that the statistics get reset using the QSYS2.SYSDISKSTAT table function before the counters can wrap more than once. The frequency needed for this action depends on the size and activity of the disk units.

The ELAPSED\_ column information is derived from the values reported in the TOTAL\_ columns as shown in the following table. These formulas can be used to calculate identical statistics if you want to save historical disk statistics in a permanent table.

For clarity, the values prefixed by delta\_ indicate the difference between two rows in the corresponding TOTAL\_ columns. For example, delta\_READ\_REQUESTS means TOTAL\_READ\_REQUESTS(*time2*) - TOTAL\_READ\_REQUESTS(*time1*). The *delta\_time* value means the time in seconds between *time1* and *time2*.

Table 193. Calculating elapsed data

| Elapsed column name    | TOTAL_ columns used to calculate the elapsed value                                            | Notes                                                    |
|------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------|
| ELAPSED_IO_REQUESTS    | (delta_READ_REQUESTS + delta_WRITE_REQUESTS) / delta_time                                     |                                                          |
| ELAPSED_REQUEST_SIZE   | ((delta_BLOCKS_READ + delta_BLOCKS_WRITTEN) / 2) (delta_READ_REQUESTS + delta_WRITE_REQUESTS) | Divide by 2 to convert value from 512 byte blocks to KB. |
| ELAPSED_READ_REQUESTS  | delta_READ_REQUESTS / delta_time                                                              |                                                          |
| ELAPSED_WRITE_REQUESTS | delta_WRITE_REQUESTS / delta_time                                                             |                                                          |
| ELAPSED_DATA_READ      | (delta_BLOCKS_READ / 2) / delta_READ_REQUESTS                                                 | Divide by 2 to convert value from 512 byte blocks to KB. |
| ELAPSED_DATA_WRITTEN   | (delta_BLOCKS_WRITTEN / 2) / delta_WRITE_REQUESTS                                             | Divide by 2 to convert value from 512 byte blocks to KB. |
| ELAPSED_PERCENT_BUSY   | ((delta_SAMPLE_COUNT - delta_NOT_BUSY_COUNT) / delta_SAMPLE_COUNT) * 100                      |                                                          |

## Examples

- Return information about all disks.

```
SELECT * FROM QSYS2.SYSDISKSTAT
```

- Return information for all SSD units.

```
SELECT * FROM QSYS2.SYSDISKSTAT WHERE UNIT_TYPE = 1
```

## SYSTMPSTG view

The SYSTMPSTG view contains one row for every temporary storage bucket that is tracking some amount of temporary storage across the system.

Temporary storage is application working storage that does not persist across a restart of the operating system. Accounting for all the temporary storage being used on the system is implemented using the concept of temporary storage buckets.

There are two types of temporary storage buckets:

- global buckets that are used to track temporary storage that is scoped to all jobs on the system.
- job buckets that are used to track temporary storage that is scoped to a single job.

Each bucket has a bucket number. Global buckets managed by the licensed internal code have bucket numbers from 1 to 4095. Global buckets managed by IBM i Work Management have bucket numbers from 4096 to 65535. Job buckets have numbers greater than 65535.

A job temporary storage bucket is assigned when the job starts and does not change for the life of the job. A job temporary storage bucket will normally be empty after the associated job ends and all working storage for the job is deleted or freed. If the job temporary storage bucket is empty after the job ends, the bucket becomes available to be associated with a new job. If the job associated with the job buckets ends and some temporary objects tracked to that job are not deleted, the job bucket will show a status of \*ENDED as well as the date and time that the job ended. These job buckets identify jobs that are not deleting all of their temporary storage when the job ends.

Statistics for each job bucket indicate the current amount of storage (in bytes) used for temporary storage tracked by the bucket, the storage limit (in bytes) for disk storage used for temporary storage tracked by the bucket, and the peak amount of disk storage (in bytes) used for temporary storage tracked by the bucket. A job bucket does not include any temporary storage used for SQL query execution. For job buckets, the storage limit will reflect the MAXTMPSTG value of the class (\*CLS) object specified when the job was submitted; a null value is returned if the job has a MAXTMPSTG value of \*NOMAX.

The following table describes the columns in the view. The schema is QSYS2.

Table 194. SYSTMPSTG view

| Column Name         | System Column Name | Data Type                 | Description                                                                                                                                                           |
|---------------------|--------------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BUCKET_NUMBER       | BKTNBR             | INTEGER                   | Number that uniquely identifies the temporary storage bucket.                                                                                                         |
| GLOBAL_BUCKET_NAME  | GLBBKTNAME         | VARCHAR(30)<br>Nullable   | For global buckets, the name of the bucket.<br>For job buckets, contains the null value.                                                                              |
| JOB_NAME            | JOBNAME            | VARCHAR(10)<br>Nullable   | For job buckets, the job name.<br>For global buckets, contains the null value.                                                                                        |
| JOB_USER_NAME       | JOBUSRNAME         | VARCHAR(10)<br>Nullable   | For job buckets, the user profile under which the job is run.<br>For global buckets, contains the null value.                                                         |
| JOB_NUMBER          | JOBNBR             | CHAR(6)<br>Nullable       | For job buckets, the job number assigned by the system.<br>For global buckets, contains the null value.                                                               |
| BUCKET_CURRENT_SIZE | BKTCURSIZ          | DECIMAL(23,0)             | The current number of bytes of disk storage for this temporary storage bucket.                                                                                        |
| BUCKET_LIMIT_SIZE   | BKTLMTSIZ          | DECIMAL(23,0)<br>Nullable | The current limit, in bytes, for the amount of disk storage for this temporary storage bucket. If the temporary storage bucket has no limit, contains the null value. |

Table 194. SYSTMPSTG view (continued)

| Column Name      | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                               |
|------------------|--------------------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BUCKET_PEAK_SIZE | BKTPEAKSIZ         | DECIMAL(23,0)          | The largest number of bytes of disk storage for this temporary storage bucket. For global buckets, this is the peak amount of disk storage since the last restart of the operating system. For job buckets, this is the peak amount of disk storage since the job was started.                                                                                            |
| JOB_STATUS       | JOBSTS             | VARCHAR(7)<br>Nullable | For job buckets, indicates whether the bucket is associated with an active job or a job that ended without deleting all temporary objects associated with the job.<br><br><b>*ENDED</b> The job associated with this job bucket has ended.<br><b>*ACTIVE</b> The job associated with this job bucket is still active.<br><br>For global buckets, contains the null value. |
| JOB_ENDED_TIME   | JOBENDTIM          | TIMESTAMP<br>Nullable  | For job buckets associated with jobs that have ended, indicates the timestamp of when the associated job ended.<br><br>Contains the null value for global buckets and job buckets associated with active jobs.                                                                                                                                                            |

## USER\_STORAGE view

The USER\_STORAGE view contains details about storage by user profile.

The user storage consumption detail is determined by using Retrieve User Information (QSYRUSRI) API.

You must have \*OBJOPR and \*READ authority to a \*USRPRF or it will not be returned. To see information for independent ASPs (iASPs), the iASP must be varied on.

User storage is broken down by SYSBAS and iASPs.

The following table describes the columns in the view. The system name is USER\_STG. The schema is QSYS2.

Table 195. USER\_STORAGE view

| Column Name             | System Column Name | Data Type               | Description                                                                                                                                                                                            |
|-------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUTHORIZATION_NAME      | USER_NAME          | VARCHAR(10)<br>Nullable | User profile name.                                                                                                                                                                                     |
| ASPGRP                  | ASPGRP             | VARCHAR(10)<br>Nullable | Name of the independent ASP or *SYSBAS.                                                                                                                                                                |
| MAXIMUM_STORAGE_ALLOWED | MAXSTG             | BIGINT<br>Nullable      | The maximum amount of auxiliary storage (in kilobytes) that can be assigned to store permanent objects owned by the user. Contains null if the user does not have a maximum amount of allowed storage. |
| STORAGE_USED            | STGUSED            | BIGINT<br>Nullable      | The amount of auxiliary storage (in kilobytes) occupied by the user's owned objects for this ASPGRP.                                                                                                   |

## Example

Determine how much storage user SCOTTFF has consumed.

```
SELECT * FROM QSYS2.USER_STORAGE
WHERE USER_NAME = 'SCOTTFF'
```

## System Health Services

For the most important system resources, the IBM i operating system automatically tracks the highest consumption and consumers.

The IBM i operating system is comprised of many products and components. As an integrated operating system, not only do the products and components frequently rely upon each other, but common building



blocks and resources are used. Some of the resources are deemed to be critical because their proper use and consumption is directly related to achieving continued, normal operational behavior. The repository for this tracking lies within Db2 for i.

A table, a view, and global variables combine to provide information about limits on your system. Information about the important limits is logged in a Db2 for i supplied table named QSYS2.SYSLIMITBL. The QSYS2.SYSLIMITS view uses SYSLIMITBL and other Db2 resources to provide extended and formatted detail about these limits. You should generally work with the view rather than the underlying table. You can use Db2 for i provided global variables to control the number of rows kept for each type of limit in SYSLIMITBL.

The following tables list the limits that are tracked along with the corresponding limit ID, the system maximum value, the value that will cause the first row to be added to SYSLIMITBL (the floor), and the amount of change after this first row is written that causes subsequent rows to be added to the table (the increment).

| <i>Table 196. Database limits</i>                         |                 |                   |               |                  |
|-----------------------------------------------------------|-----------------|-------------------|---------------|------------------|
| <b>Limit description</b>                                  | <b>Limit ID</b> | <b>Maximum</b>    | <b>Floor</b>  | <b>Increment</b> |
| Maximum number of all rows in a partition                 | 15000           | 4,294,967,288     | 100,000       | 500,000          |
| Maximum number of valid rows in a partition               | 15001           | 4,294,967,288     | 100,000       | 500,000          |
| Maximum number of deleted rows in a partition             | 15002           | 4,294,967,288     | 10,000        | 100,000          |
| Maximum size of a table                                   | 15003           | 1,869,169,767,219 | 536,865,792   | 1,073,731,584    |
| Maximum number of overflow rows in a partition            | 15004           | 4,294,967,288     | 10,000        | 100,000          |
| Maximum number of variable-length segments                | 15104           | 65,533            | 100           | 100              |
| Maximum number of indexes over a partition                | 15106           | 15,000            | 20            | 100              |
| Maximum size of a *MAX4GB index                           | 15400           | 4,294,967,296     | 838,860,800   | 167,772,160      |
| Maximum size of a *MAX1TB index                           | 15401           | 1,869,166,411,776 | 8,388,608,000 | 8,388,608,000    |
| Maximum size of an encoded vector index                   | 15403           | 2,199,023,255,552 | 1,677,721,600 | 8,388,608,000    |
| Maximum number of members in a source physical file       | 16100           | 32,767            | 100           | 50               |
| Maximum number of rows locked in a unit of work           | 16200           | 500,000,000       | 10,000        | 100,000          |
| Maximum number of row change operations in a unit of work | 16201           | storage           | 10,000        | 100,000          |
| Maximum size of an extended dynamic package               | 16806           | 1,056,964,608     | 335,544,320   | 8,388,608        |

Table 197. Journal limits

| Limit description                                                        | Limit ID | Maximum                        | Floor          | Increment      |
|--------------------------------------------------------------------------|----------|--------------------------------|----------------|----------------|
| Maximum size of a journal receiver                                       | 18300    | 1,099,511,627,776              | 10,000,000,000 | 50,000,000,000 |
| Maximum number of objects that can be associated with a *MAX10M journal  | 18301    | 10,000,000                     | 10,000         | 200,000        |
| Maximum number of objects that can be associated with a *MAX250K journal | 18302    | 250,000                        | 10,000         | 50,000         |
| Maximum sequence number of a *MAXOPT3 journal                            | 18303    | 18,446,744,073<br>,709,551,600 | 10,000,000     | 100,000,000    |
| Maximum sequence number of a *MAXOPT1 or *MAXOPT2 journal                | 18304    | 9,999,999,999                  | 10,000,000     | 10,000,000     |

Table 198. File system limits

| Limit description                                           | Limit ID | Maximum           | Floor      | Increment |
|-------------------------------------------------------------|----------|-------------------|------------|-----------|
| Maximum number of object description entries in a library   | 18400    | 1,000,000         | 1,000      | 1,000     |
| Number of objects linked in a directory                     | 18402    | storage           | 100,000    | 10,000    |
| Maximum number of directories linked in a directory         | 18403    | 1,000,000         | 1,000      | 1,000     |
| Maximum number of file system objects in *SYSBAS ASPs       | 18404    | 2,147,483,647     | 100,000    | 10,000    |
| Maximum number of file system objects in an independent ASP | 18405    | 2,147,483,647     | 100,000    | 10,000    |
| Maximum number of document library objects in a folder      | 18406    | 65510             | 1,000      | 500       |
| Number of document library objects in the system ASP        | 18407    | storage           | 100,000    | 10,000    |
| Maximum number of document library objects in a user ASP    | 18408    | 1,000,000         | 100,000    | 10,000    |
| Maximum number of bytes in a stream file                    | 18409    | 1,099,511,627,776 | 16,777,216 | 1,048,576 |
| Maximum number of bytes in a document                       | 18410    | 2,147,483,647     | 16,777,216 | 1,048,576 |

Table 199. Work management limits

| Limit description                                       | Limit ID | Maximum    | Floor  | Increment |
|---------------------------------------------------------|----------|------------|--------|-----------|
| Maximum number of jobs                                  | 19000    | 970,000    | 1,000  | 400       |
| Maximum number of spool files                           | 19002    | 2,610,000  | 10,000 | 5,000     |
| Maximum number of spooled files in each independent ASP | 19003    | 10,000,000 | 10,000 | 5,000     |

## System limit alerts

Some system limits are instrumented by the IBM i operating system to send messages to the QSYSOPR message queue when a threshold value has been reached.

Once each day, the IBM i will look for any limits that have surpassed their alerting level. This happens when Collection Services is recycled, typically just past midnight. At this time, a call is made to the QSYS2.PROCESS\_SYSTEM\_LIMITS\_ALERTS procedure to identify and signal any alerts for the day.

The following limits are checked against their alerting level as part of this daily processing. If the level is exceeded, a severity 80 informational message SQL7062 is sent to the QSYSOPR message queue. Since these limits will prevent database or other system activity from continuing if they are reached, you should take action to get the object's percent used for the limit below the alerting level. Reducing data by archiving it is one example of an action that could be taken.

Table 200. System limits that send alerting messages

| Limit ID | Limit description                          | Maximum           | Default Alerting Level | Alerting Cadence |
|----------|--------------------------------------------|-------------------|------------------------|------------------|
| 15000    | Maximum number of all rows in a partition  | 4,294,967,288     | Greater than 90%       | Once per day     |
| 15003    | Maximum size of a table                    | 1,869,169,767,219 | Greater than 90%       | Once per day     |
| 15104    | Maximum number of variable-length segments | 65,533            | Greater than 90%       | Once per day     |
| 15400    | Maximum *MAX4GB Index Size                 | 4,294,967,296     | Greater than 90%       | Once per day     |
| 15401    | Maximum *MAX1TB Index Size                 | 1,869,166,411,776 | Greater than 90%       | Once per day     |
| 15403    | Maximum Encoded Vector Index Size          | 2,199,023,255,552 | Greater than 90%       | Once per day     |
| 19002    | Maximum number of spooled files            | 2,610,000         | Greater than 90%       | Once per day     |

The SQL7062 QSYSOPR message is formatted like this:

```
MYLIB/MYTABLE *FILE HAS CONSUMED MORE THAN 90% OF THE LIMIT:
15000-MAXIMUM NUMBER OF ALL ROWS (4008420999 OF 4294967288=93.33%).
REFER TO ibm.biz/DB2foriAlerts FOR MORE DETAIL.
```

## System limit alerts global variables

For each limit that is instrumented to send a system limit alert, a corresponding global variable is defined that can be used to modify the consumption level that causes the alert to be issued.

The following are the names of the global variables. The schema is SYSIBMADM.

| Limit ID | Global variable                           | Shipped limit |
|----------|-------------------------------------------|---------------|
| 15000    | QIBM_SYSTEM_LIMITS_ALERT_15000_PERCENTAGE | 90            |
| 15003    | QIBM_SYSTEM_LIMITS_ALERT_15003_PERCENTAGE | 90            |
| 15104    | QIBM_SYSTEM_LIMITS_ALERT_15104_PERCENTAGE | 90            |
| 15400    | QIBM_SYSTEM_LIMITS_ALERT_15400_PERCENTAGE | 90            |
| 15401    | QIBM_SYSTEM_LIMITS_ALERT_15401_PERCENTAGE | 90            |
| 15403    | QIBM_SYSTEM_LIMITS_ALERT_15403_PERCENTAGE | 90            |
| 19002    | QIBM_SYSTEM_LIMITS_ALERT_19002_PERCENTAGE | 90            |

You can redefine any of the global variable values to change the alerting percent on your system. The change will take effect the next time the altering levels are checked.

Use IBM i Access Client Solutions (ACS) to generate SQL for the global variable and use the OR REPLACE option to recreate it with a different default. For example, to send an alert when a file reaches 70% of its maximum size, use the following SQL statement:

```
CREATE OR REPLACE VARIABLE SYSIBMADM.QIBM_SYSTEM_LIMITS_ALERT_15000_PERCENTAGE
 INTEGER
 DEFAULT 70
```

## SYSLIMTBL table

The SYSLIMTBL table contains information about limits as they are being consumed. It is maintained by Db2 for i.

This table is not authorized or managed like a typical Db2 for i catalog. By default, all users have authority to view this table. If this table is removed or incompatibly altered, the IBM i operating system will automatically recreate it. The SYSLIMTBL table is designed to have as small a footprint as possible.

You can add AFTER INSERT or AFTER DELETE triggers to this table. This allows you to perform an action such as sending a notification when a limit is being logged to the table.

The following table describes the columns in the table. The schema is QSYS2.

Table 201. SYSLIMTBL table

| Column Name           | System Column Name | Data Type | Description                                                                                                                                                                                                                                                                                                                                   |
|-----------------------|--------------------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LAST_CHANGE_TIMESTAMP | LASTCHG            | TIMESTAMP | The timestamp when this row was last changed.                                                                                                                                                                                                                                                                                                 |
| LIMIT_CATEGORY        | CATEGORY           | SMALLINT  | The category of this limit. <ul style="list-style-type: none"> <li><b>0</b> Database</li> <li><b>1</b> Journal</li> <li><b>2</b> Security</li> <li><b>3</b> Miscellaneous</li> <li><b>4</b> Work management</li> <li><b>5</b> File system</li> <li><b>6</b> Save/restore</li> <li><b>7</b> Cluster</li> <li><b>8</b> Communication</li> </ul> |

Table 201. SYSLIMTBL table (continued)

| Column Name         | System Column Name | Data Type                              | Description                                                                                                                                                                                                                                                      |          |        |          |     |          |        |          |     |
|---------------------|--------------------|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|--------|----------|-----|----------|--------|----------|-----|
| LIMIT_TYPE          | LIMTYPE            | SMALLINT                               | The type of limit.<br><br><table border="0"> <tr> <td style="padding-right: 20px;"><b>1</b></td> <td>Object</td> </tr> <tr> <td><b>2</b></td> <td>Job</td> </tr> <tr> <td><b>3</b></td> <td>System</td> </tr> <tr> <td><b>4</b></td> <td>ASP</td> </tr> </table> | <b>1</b> | Object | <b>2</b> | Job | <b>3</b> | System | <b>4</b> | ASP |
| <b>1</b>            | Object             |                                        |                                                                                                                                                                                                                                                                  |          |        |          |     |          |        |          |     |
| <b>2</b>            | Job                |                                        |                                                                                                                                                                                                                                                                  |          |        |          |     |          |        |          |     |
| <b>3</b>            | System             |                                        |                                                                                                                                                                                                                                                                  |          |        |          |     |          |        |          |     |
| <b>4</b>            | ASP                |                                        |                                                                                                                                                                                                                                                                  |          |        |          |     |          |        |          |     |
| LIMIT_ID            | LIMIT_ID           | INTEGER                                | Unique identifier for this limit. Values are maintained in the SIZING_ID column in the QSYS2.SQL_SIZING table.                                                                                                                                                   |          |        |          |     |          |        |          |     |
| JOB_NAME            | JOB_NAME           | VARCHAR(28)                            | The name of the job that reported the current value.                                                                                                                                                                                                             |          |        |          |     |          |        |          |     |
| USER_NAME           | CURUSER            | VARCHAR(10)                            | The name of the user in effect when the current value was updated.                                                                                                                                                                                               |          |        |          |     |          |        |          |     |
| CURRENT_VALUE       | CURVAL             | BIGINT                                 | Reported value for this limit.                                                                                                                                                                                                                                   |          |        |          |     |          |        |          |     |
| SYSTEM_SCHEMA_NAME  | SYS_NAME           | VARCHAR(10)<br>Nullable                | The library name for the object. If no library name, contains the null value.                                                                                                                                                                                    |          |        |          |     |          |        |          |     |
| SYSTEM_OBJECT_NAME  | SYS_ONAME          | VARCHAR(30)<br>Nullable                | The object name for this row. If no object name, contains the null value.                                                                                                                                                                                        |          |        |          |     |          |        |          |     |
| SYSTEM_TABLE_MEMBER | SYS_MNAME          | VARCHAR(10)<br>Nullable                | The member name for an object limit specific to database members. Contains the null value if this row is not for a member limit.                                                                                                                                 |          |        |          |     |          |        |          |     |
| OBJECT_TYPE         | OBJTYPE            | VARCHAR(7)<br>Nullable                 | The IBM i object type when an object name has been logged in the SYSTEM_SCHEMA_NAME and SYSTEM_OBJECT_NAME columns. Contains the null value when no object name is specified.                                                                                    |          |        |          |     |          |        |          |     |
| ASP_NUMBER          | ASPNUM             | SMALLINT<br>Nullable                   | Contains the ASP number related to this row. Contains the null value if there is no ASP number.                                                                                                                                                                  |          |        |          |     |          |        |          |     |
| IFS_PATH_NAME       | PATHNAME           | DBCLOB(5000)<br>CCSID 1200<br>Nullable | IFS path for the object. Contains the null value if there is no path.                                                                                                                                                                                            |          |        |          |     |          |        |          |     |

## Example

Add a trigger to QSYS2.SYSLIMTBL to send a message when any table is approaching the maximum size. The trigger will be fired when any row is inserted into SYSLIMTBL. Within the trigger, it checks for the LIMIT\_ID indicating the maximum number of rows in a partition (15000) and the value when you want to be notified.

```

/* Force any pseudo closed cursors over SYSLIMTBL to be closed */
CL: ALCOBJ OBJ((QSYS2/SYSLIMTBL *FILE *EXCL)) CONFLICT(*QRSRLS) ;
CL: DLCOBJ OBJ((QSYS2/SYSLIMTBL *FILE *EXCL));

CREATE OR REPLACE TRIGGER MYLIB.SYSTEM_LIMITS_LARGE_FILE
AFTER INSERT ON QSYS2.SYSLIMTBL
REFERENCING NEW AS N FOR EACH ROW MODE DB2ROW
SET OPTION USRPRF=*OWNER, DYNUSRPRF=*OWNER
BEGIN ATOMIC
 DECLARE V_CMDSTMT VARCHAR(200) ;
 DECLARE ERROR INTEGER;
 DECLARE EXIT HANDLER FOR SQLEXCEPTION SET ERROR = 1;
 /* -----*/
 /* If a table is nearing the maximum size, alert the operator */
 /* -----*/
 IF (N.LIMIT_ID = 15000 AND
 N.CURRENT_VALUE > 3000000000) THEN
 SET V_CMDSTMT = 'SNDSMSG MSG(''Table: '
 CONCAT N.SYSTEM_SCHEMA_NAME CONCAT '/' CONCAT N.SYSTEM_OBJECT_NAME
 CONCAT ' (' CONCAT N.SYSTEM_TABLE_MEMBER CONCAT
 ') IS GETTING VERY LARGE - ROW COUNT = '
 CONCAT CURRENT_VALUE CONCAT ' ') TOUSR(*SYSOPR) MSGTYPE(*INFO) ' ;
 CALL QSYS2.QCMDXEC(V_CMDSTMT);

```

```
END IF;
END;
```

## SYSLIMITS view

The SYSLIMITS view contains information about limits. This view is built on the QSYS.SYSLIMTBL table along with other system information. If a job is still active, the view contains information about the job that logged the limit.

**Authorization:** For rows where the job is still active, the caller's user profile must be the same as the [job user identity](#) of the job for which the information is being returned, or must have \*JOBCTL user special authority, or QIBM\_DB\_SQLADM or QIBM\_DB\_SYSMON function usage authority. If the caller does not have sufficient authority, partial information is returned along with an SQL warning of '01548'.

For rows where LIMIT\_TYPE = 'OBJECT', additional authorization is required:

- If the user has \*EXECUTE authority to the library, and both \*OBJOPR and \*READ authority to an object, full details are returned.
- Otherwise, partial information is returned along with an SQL warning of '01548'.

The following table describes the columns in the view. The schema is QSYS2.

Table 202. SYSLIMITS view

| Column Name           | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                     |
|-----------------------|--------------------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LAST_CHANGE_TIMESTAMP | LASTCHG            | TIMESTAMP                 | The timestamp when this row was last changed.                                                                                                                                                                                                                                   |
| LIMIT_CATEGORY        | CATEGORY           | VARCHAR(15)               | The category for this limit. <ul style="list-style-type: none"> <li>• DATABASE</li> <li>• JOURNAL</li> <li>• SECURITY</li> <li>• MISCELLANEOUS</li> <li>• WORK MANAGEMENT</li> <li>• FILE SYSTEM</li> <li>• SAVE RESTORE</li> <li>• CLUSTER</li> <li>• COMMUNICATION</li> </ul> |
| LIMIT_TYPE            | TYPE               | VARCHAR(7)                | The type of limit. <ul style="list-style-type: none"> <li>• OBJECT</li> <li>• JOB</li> <li>• SYSTEM</li> <li>• ASP</li> </ul>                                                                                                                                                   |
| SIZING_NAME           | SIZING_NAM         | VARCHAR(128)              | Name that corresponds to the sizing ID.                                                                                                                                                                                                                                         |
| COMMENTS              | COMMENTS           | VARCHAR(2000)<br>Nullable | Description of the limit.                                                                                                                                                                                                                                                       |
| USER_NAME             | CURUSER            | VARCHAR(10)               | The name of the user in effect when this row was logged.                                                                                                                                                                                                                        |
| CURRENT_VALUE         | CURVAL             | BIGINT                    | Reported value for this limit.                                                                                                                                                                                                                                                  |
| MAXIMUM_VALUE         | MAXVAL             | DECIMAL(21,0)<br>Nullable | Maximum value allowed for this limit.                                                                                                                                                                                                                                           |
| JOB_NAME              | JOB_NAME           | VARCHAR(28)               | The name of the job when this row was logged.<br>Contains the null value if the job is no longer active.                                                                                                                                                                        |
| JOB_STATUS            | JOB_STATUS         | CHAR(10)<br>Nullable      | Status of the job.<br>Contains the null value if the job is no longer active.                                                                                                                                                                                                   |
| ACTIVE_JOB_STATUS     | AJSTATUS           | CHAR(4)<br>Nullable       | The active status of the initial thread of the job.<br>Contains the null value if the job is in transition or is no longer active.                                                                                                                                              |

Table 202. SYSLIMITS view (continued)

| Column Name          | System Column Name | Data Type                              | Description                                                                                                                                                                                               |
|----------------------|--------------------|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RUN_PRIORITY         | RUNPRI             | INTEGER<br>Nullable                    | The highest run priority allowed for any thread within this job.<br>Contains the null value if the job is no longer active.                                                                               |
| SBS_NAME             | SBS_NAME           | CHAR(10)<br>Nullable                   | Name of subsystem where job is running.<br>Contains the null value if the job is no longer active.                                                                                                        |
| CPU_USED             | CPU_USED           | BIGINT<br>Nullable                     | The amount of CPU time (in milliseconds) that has been currently used by this job.<br>Contains the null value if the job is no longer active.                                                             |
| TEMP_STORAGE_USED_MB | TEMPSTG            | INTEGER<br>Nullable                    | The amount of auxiliary storage (in megabytes) that is currently allocated to this job.<br>Contains the null value if the job is no longer active.                                                        |
| AUX_IO_REQUESTED     | AUXIO              | BIGINT<br>Nullable                     | The number of auxiliary I/O requests performed by the job across all routing steps. This includes both database and nondatabase paging.<br>Contains the null value if the job is no longer active.        |
| PAGE_FAULTS          | PAGEFAULT          | BIGINT<br>Nullable                     | The number of times an active program referenced an address that was not in main storage during the current routing step of the specified job.<br>Contains the null value if the job is no longer active. |
| CLIENT_WRKSTNNAME    | CLIENTWRK          | CHAR(255)<br>Nullable                  | Value of the SQL CLIENT_WRKSTNNAME special register.<br>Contains the null value if the job is no longer active.                                                                                           |
| CLIENT_APPLNAME      | CLIENTAPP          | CHAR(255)<br>Nullable                  | Value of the SQL CLIENT_APPLNAME special register.<br>Contains the null value if the job is no longer active.                                                                                             |
| CLIENT_ACCTNG        | CLIENTACT          | CHAR(255)<br>Nullable                  | Value of the SQL CLIENT_ACCTNG special register.<br>Contains the null value if the job is no longer active.                                                                                               |
| CLIENT_PROGRAMID     | CLIENTPGM          | CHAR(255)<br>Nullable                  | Value of the SQL CLIENT_PROGRAMID special register.<br>Contains the null value if the job is no longer active.                                                                                            |
| CLIENT_USERID        | CLIENTUSER         | CHAR(255)<br>Nullable                  | Value of the SQL CLIENT_USERID special register.<br>Contains the null value if the job is no longer active.                                                                                               |
| SQL_STATEMENT_TEXT   | SQLSTMT            | VARCHAR(10000)<br>Nullable             | Statement text of the last SQL statement to run or the SQL statement that is currently running.<br>Contains the null value if the job is no longer active.                                                |
| SCHEMA_NAME          | OBJ_SCHEMA         | VARCHAR(128)<br>Nullable               | The SQL schema name for this object.<br>Contains the null value if there is no schema name.                                                                                                               |
| OBJECT_NAME          | OBJ_NAME           | VARCHAR(128)<br>Nullable               | The SQL name for the object.<br>Contains the null value if there is no object name or if an SQL name could not be returned.                                                                               |
| SYSTEM_SCHEMA_NAME   | SYS_NAME           | VARCHAR(10)<br>Nullable                | The library name for the object.<br>Contains the null value if there is no library name.                                                                                                                  |
| SYSTEM_OBJECT_NAME   | SYS_ONAME          | VARCHAR(30)<br>Nullable                | The object name for this row.<br>Contains the null value if there is no object name.                                                                                                                      |
| SYSTEM_TABLE_MEMBER  | SYS_MNAME          | VARCHAR(10)<br>Nullable                | The member name for an object limit specific to database members.<br>Contains the null value if this row is not for a member limit.                                                                       |
| IFS_PATH_NAME        | PATHNAME           | DBCLOB(5000)<br>CCSID 1200<br>Nullable | IFS path for the object.<br>Contains the null value if there is no path.                                                                                                                                  |

Table 202. SYSLIMITS view (continued)

| Column Name     | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|--------------------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBJECT_TYPE     | OBJTYPE            | VARCHAR(7)<br>Nullable | The IBM i object type when an object name has been logged in the SYSTEM_SCHEMA_NAME and SYSTEM_OBJECT_NAME columns.<br><br>Contains the null value when no object name is specified.                                                                                                                                                                                                                                                                                                                                      |
| SQL_OBJECT_TYPE | SQLOBJTYPE         | VARCHAR(9)<br>Nullable | The SQL type of the object when an object name has been logged in the SYSTEM_SCHEMA_NAME and SYSTEM_OBJECT_NAME columns. Values are: <ul style="list-style-type: none"> <li>• ALIAS</li> <li>• FUNCTION</li> <li>• INDEX</li> <li>• PACKAGE</li> <li>• PROCEDURE</li> <li>• ROUTINE</li> <li>• SEQUENCE</li> <li>• TABLE</li> <li>• TRIGGER</li> <li>• TYPE</li> <li>• VARIABLE</li> <li>• VIEW</li> <li>• XSR</li> </ul> Contains the null value if the object is not an SQL object or when no object name is specified. |
| ASP_NUMBER      | ASPNUM             | SMALLINT<br>Nullable   | Contains the ASP number related to this row.<br><br>Contains the null value if there is no ASP number.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| LIMIT_ID        | LIMIT_ID           | INTEGER                | Unique identifier for this limit. Values are maintained in the SIZING_ID column in the QSYS2.SQL_SIZING table.                                                                                                                                                                                                                                                                                                                                                                                                            |

## Examples

- Find the 50 largest IFS stream files. Remove any duplicates from the result. Note that only stream files that have reached the documented floor and increment values will appear in SYSLIMITS.

```
SELECT IFS_PATH_NAME, MAX(CURRENT_VALUE) AS MAX_BYTE_SIZE
FROM QSYS2.SYSLIMITS
WHERE LIMIT_ID = 18409
GROUP BY IFS_PATH_NAME
ORDER BY MAX_BYTE_SIZE DESC LIMIT 50;
```

- Review the consumption of the 'Total number of jobs', relative to the QMAXJOB system value.

```
WITH TT(JOB_MAXIMUM)
AS (SELECT CURRENT_NUMERIC_VALUE
FROM QSYS2.SYSTEM_VALUE_INFO
WHERE SYSTEM_VALUE_NAME = 'QMAXJOB')
SELECT LAST_CHANGE_TIMESTAMP AS INCREMENT_TIME, CURRENT_VALUE AS JOB_COUNT,
TT.JOB_MAXIMUM,
DEC(DEC(CURRENT_VALUE,19,2) / DEC(TT.JOB_MAXIMUM,19,2) * 100,19,2)
AS PERCENT_CONSUMED
FROM QSYS2.SYSLIMITS, TT
WHERE LIMIT_ID = 19000 ORDER BY CURRENT_VALUE DESC;
```

## SYSLIMITS\_BASIC view

The SYSLIMITS\_BASIC view contains information about limits. This view is built on the QSYS.SYSLIMITBL table along with other system information. It does not return information about the job that logged the limit. This view returns less information than the SYSLIMITS view, but it requires less authorization and typically performs significantly better.

**Authorization:** None required.



The following table describes the columns in the view. The system name is SYSLIMIT\_B. The schema is QSYS2.

Table 203. SYSLIMITS\_BASIC view

| Column Name           | System Column Name | Data Type                              | Description                                                                                                                                                                                                                                                                     |
|-----------------------|--------------------|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LAST_CHANGE_TIMESTAMP | LASTCHG            | TIMESTAMP                              | The timestamp when this row was last changed.                                                                                                                                                                                                                                   |
| LIMIT_CATEGORY        | CATEGORY           | VARCHAR(15)                            | The category for this limit. <ul style="list-style-type: none"> <li>• DATABASE</li> <li>• JOURNAL</li> <li>• SECURITY</li> <li>• MISCELLANEOUS</li> <li>• WORK MANAGEMENT</li> <li>• FILE SYSTEM</li> <li>• SAVE RESTORE</li> <li>• CLUSTER</li> <li>• COMMUNICATION</li> </ul> |
| LIMIT_TYPE            | TYPE               | VARCHAR(7)                             | The type of limit. <ul style="list-style-type: none"> <li>• OBJECT</li> <li>• JOB</li> <li>• SYSTEM</li> <li>• ASP</li> </ul>                                                                                                                                                   |
| SIZING_NAME           | SIZING_NAM         | VARCHAR(128)                           | Name that corresponds to the sizing ID.                                                                                                                                                                                                                                         |
| COMMENTS              | COMMENTS           | VARCHAR(2000)<br>Nullable              | Description of the limit.                                                                                                                                                                                                                                                       |
| USER_NAME             | CURUSER            | VARCHAR(10)                            | The name of the user in effect when this row was logged.                                                                                                                                                                                                                        |
| CURRENT_VALUE         | CURVAL             | BIGINT                                 | Reported value for this limit.                                                                                                                                                                                                                                                  |
| MAXIMUM_VALUE         | MAXVAL             | DECIMAL(21,0)<br>Nullable              | Maximum value allowed for this limit.                                                                                                                                                                                                                                           |
| JOB_NAME              | JOB_NAME           | VARCHAR(28)                            | The name of the job when this row was logged.<br>Contains the null value if the job is no longer active.                                                                                                                                                                        |
| SYSTEM_SCHEMA_NAME    | SYS_NAME           | VARCHAR(10)<br>Nullable                | The library name for the object.<br>Contains the null value if there is no library name.                                                                                                                                                                                        |
| SYSTEM_OBJECT_NAME    | SYS_ONAME          | VARCHAR(30)<br>Nullable                | The object name for this row.<br>Contains the null value if there is no object name.                                                                                                                                                                                            |
| SYSTEM_TABLE_MEMBER   | SYS_MNAME          | VARCHAR(10)<br>Nullable                | The member name for an object limit specific to database members.<br>Contains the null value if this row is not for a member limit.                                                                                                                                             |
| IFS_PATH_NAME         | PATHNAME           | DBCLOB(5000)<br>CCSID 1200<br>Nullable | IFS path for the object.<br>Contains the null value if there is no path.                                                                                                                                                                                                        |
| OBJECT_TYPE           | OBJTYPE            | VARCHAR(7)<br>Nullable                 | The IBM i object type when an object name has been logged in the SYSTEM_SCHEMA_NAME and SYSTEM_OBJECT_NAME columns.<br>Contains the null value when no object name is specified.                                                                                                |
| ASP_NUMBER            | ASPNUM             | SMALLINT<br>Nullable                   | Contains the ASP number related to this row.<br>Contains the null value if there is no ASP number.                                                                                                                                                                              |
| LIMIT_ID              | LIMIT_ID           | INTEGER                                | Unique identifier for this limit. Values are maintained in the SIZING_ID column in the QSYS2.SQL_SIZING table.                                                                                                                                                                  |

## Examples

- Review the consumption of the 'Total number of jobs', relative to the QMAXJOB system value.

```
WITH TT(JOB_MAXIMUM)
AS (SELECT CURRENT_NUMERIC_VALUE
 FROM QSYS2.SYSTEM_VALUE_INFO
 WHERE SYSTEM_VALUE_NAME = 'QMAXJOB')
SELECT LAST_CHANGE_TIMESTAMP AS INCREMENT_TIME, CURRENT_VALUE AS JOB_COUNT,
 TT.JOB_MAXIMUM,
 DEC(DEC(CURRENT_VALUE,19,2) / DEC(TT.JOB_MAXIMUM,19,2) * 100,19,2)
 AS PERCENT_CONSUMED
FROM QSYS2.SYSLIMITS_BASIC, TT
WHERE LIMIT_ID = 19000 ORDER BY CURRENT_VALUE DESC;
```

## QIBM\_SYSTEM\_LIMITS global variables

To prevent excess storage consumption or retention of unnecessarily old system limits entries within the QSYS2/SYSLIMITBL table, Db2 for i will automatically delete (or prune) rows.

There are two ways the pruning is controlled. One method is by the number of days to keep a row. The other is by the maximum number of rows to keep for a specific limit. In each case, Db2 for i supplied global variables guide the pruning action.

Controls exist for deleting rows that have reached a certain age are handled with a set of Db2 for i provided global variables. These controls cause rows to be removed when they exceed the number of days.

The following are the names of the global variables that control pruning by number of days and the limit that is shipped for each one. The schema is SYSIBMADM.

| Global variable                   | Shipped limit |
|-----------------------------------|---------------|
| QIBM_SYSTEM_LIMITS_ASP_BY_DAYS    | 90            |
| QIBM_SYSTEM_LIMITS_JOB_BY_DAYS    | 90            |
| QIBM_SYSTEM_LIMITS_OBJECT_BY_DAYS | 90            |
| QIBM_SYSTEM_LIMITS_SYSTEM_BY_DAYS | 90            |

The second control for automatic deletion of rows is by the number of rows for a type of limit. For each type of limit, there are two global variables. The pruning variable is used to choose how many of the most recently logged entries should be retained. The high point variable is used to choose how many of the highest consumption value entries should be retained.

The following are the names of the global variables and the limit that is shipped for each one. The schema is SYSIBMADM.

| Global variable                               | Shipped limit |
|-----------------------------------------------|---------------|
| QIBM_SYSTEM_LIMITS_PRUNE_BY_ASP               | 20            |
| QIBM_SYSTEM_LIMITS_PRUNE_BY_JOB               | 20            |
| QIBM_SYSTEM_LIMITS_PRUNE_BY_OBJECT            | 20            |
| QIBM_SYSTEM_LIMITS_PRUNE_BY_SYSTEM            | 20            |
| QIBM_SYSTEM_LIMITS_SAVE_HIGH_POINTS_BY_ASP    | 25            |
| QIBM_SYSTEM_LIMITS_SAVE_HIGH_POINTS_BY_JOB    | 5             |
| QIBM_SYSTEM_LIMITS_SAVE_HIGH_POINTS_BY_OBJECT | 5             |
| QIBM_SYSTEM_LIMITS_SAVE_HIGH_POINTS_BY_SYSTEM | 25            |

You can adjust any of the global variable values to establish a custom behavior for the automatic deletion of system limits rows. The Db2 for i supplied global variables use the default value to guide an automatic row deletion process that runs nightly when Collection Services is recycled, which normally occurs just past midnight.

Use IBM i Access Client Solutions (ACS) to generate SQL for the global variable and use the OR REPLACE option to recreate it with a different default. For example, to remove all object limits older than 30 days, use the following SQL statement:

```
CREATE OR REPLACE VARIABLE SYSIBMADM.QIBM_SYSTEM_LIMITS_OBJECT_BY_DAYS
 INTEGER
 DEFAULT 30
```

## Work Management Services

These views and functions provide system value and job information.

### ACTIVE\_JOB\_INFO table function

The ACTIVE\_JOB\_INFO table function returns one row for every active job.

The information returned is similar to the detail seen from the Work with Active Jobs (WRKACTJOB) command and the List Job (QUSLJOB) API. The ACTIVE\_JOB\_INFO table function has two uses:

1. To see details for all, or a subset of, active jobs. A subset of active jobs can be requested by using the optional filter parameters.
2. To measure elapsed statistics for active jobs. You can use an optional parameter to reset statistics, similar to the WRKACTJOB command F10 Restart Statistics function. Measurements will be calculated based on this new starting point.

**Authorization:** None required to see general information or to see information for jobs where the caller's user profile is the same as the job user identity of the job for which the information is being returned.

For DETAILED\_INFO => NONE or DETAILED\_INFO => WORK:

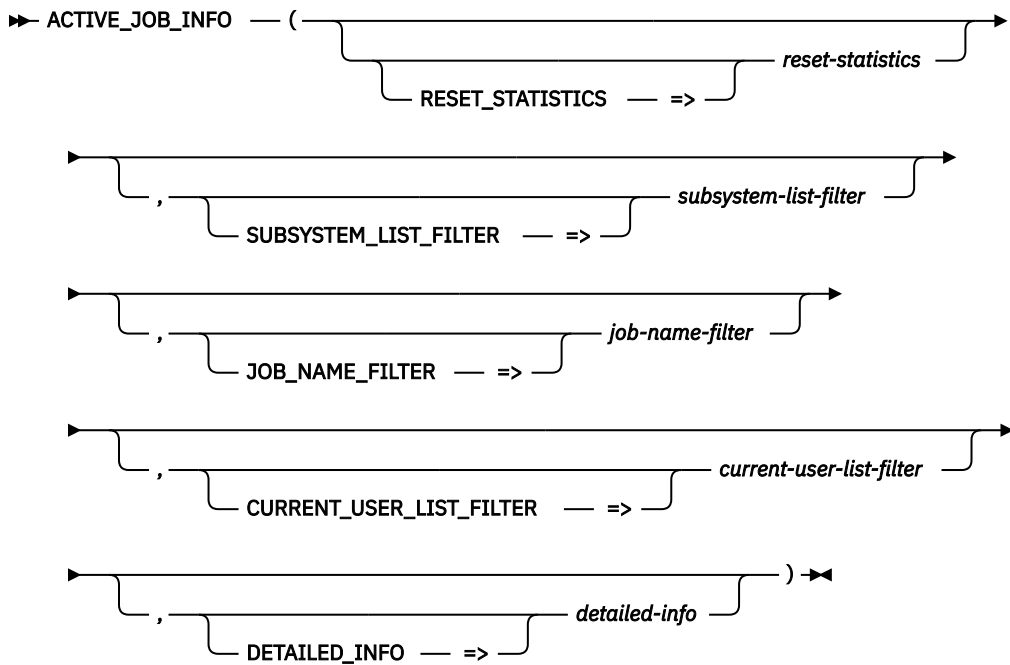
- None required.

For DETAILED\_INFO => QTEMP:

- The caller must have \*JOBCTL special authority.

For DETAILED\_INFO => ALL:

- All callers can see detailed column information for the columns that are included with DETAILED\_INFO => WORK. In addition, the CLIENT\_IP\_ADDRESS, PAGE\_FAULTS, PRESTART\_JOB\_REUSE\_COUNT, and PRESTART\_JOB\_MAX\_USE\_COUNT columns are returned.
- A caller with QIBM\_DB\_SQLADM or QIBM\_DB\_SYSMON function usage authority can see detailed column information that relates to SQL activity starting with the SQL\_STATEMENT\_TEXT column through the PSEUDO\_CLOSED\_CURSOR\_COUNT column.
- A caller with \*JOBCTL user special authority can see all detailed column information.



The schema is QSYS2.

**reset-statistics**

A character or graphic string expression that contains a value of YES or NO. If this parameter has a value of YES, statistics are reset such that the time of this query execution is used as the new baseline. Future invocations of ACTIVE\_JOB\_INFO within this connection will return statistical detail relative to the new baseline. If this parameter has a value of NO, statistics are not reset for the invocation unless the *subsystem-list-filter* or *job-name-filter* parameter values are different than the previous invocation. Changing the filter values will always cause statistics to be reset. If this parameter is not specified, the default is NO.

The first invocation of ACTIVE\_JOB\_INFO within a connection will always perform an implicit reset, regardless of whether a reset was explicitly requested.

**subsystem-list-filter**

A character or graphic string expression that contains a list of up to 25 subsystem names separated by exactly one comma. The filter determines which subsystems to use to return job information.

If this parameter is not specified, is an empty string, or is the null value, information for all subsystems is returned.

**job-name-filter**

A character or graphic string expression that contains an unqualified job name that determines the job information to be returned. The name can be a generic name.

The string can be one of the following special values:

- \* Only information for the current job is returned.
- \*ALL Information for all jobs is returned.
- \*CURRENT Information for all jobs with a job name that is the same as the current job is returned.
- \*SBS Information for all active subsystem monitors is returned.
- \*SYS Information for all active system jobs is returned. When using this value, the *subsystem-list-filter* must not be specified or must be the null value.

If this parameter is not specified, is an empty string, or is the null value, information for all jobs is returned.

**current-user-list-filter** A character or graphic string expression that contains a list of up to 10 user profile names separated by exactly one comma. The filter determines which current user values to use to return job information.

If this parameter is not specified, is an empty string, or is the null value, information for all users is returned.

**detailed-info** A character or graphic string expression that indicates the type of information to be returned.

**NONE** Only the general information is returned for active jobs. This is the information in the columns prior to the JOB\_DESCRIPTION\_LIBRARY column. This is the default.

**WORK** In addition to the general information for active jobs, additional work management information is returned.

**QTEMP** In addition to the general information for active jobs, the QTEMP\_SIZE column is returned.

**ALL** Information for all the columns is returned.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

The **DETAILED\_INFO option** column indicates which of the DETAILED\_INFO parameter values can return a non-null value for the corresponding result column.

Table 204. ACTIVE\_JOB\_INFO table function

| Column Name      | Data Type   | DETAILED_INFO option         | Description                                                                                                |
|------------------|-------------|------------------------------|------------------------------------------------------------------------------------------------------------|
| ORDINAL_POSITION | INTEGER     | NONE<br>WORK<br>QTEMP<br>ALL | A unique number for each row.                                                                              |
| JOB_NAME         | VARCHAR(28) | NONE<br>WORK<br>QTEMP<br>ALL | The qualified job name.                                                                                    |
| JOB_NAME_SHORT   | VARCHAR(10) | NONE<br>WORK<br>QTEMP<br>ALL | The name of the job.                                                                                       |
| JOB_USER         | VARCHAR(10) | NONE<br>WORK<br>QTEMP<br>ALL | The user profile that started the job.                                                                     |
| JOB_NUMBER       | VARCHAR(6)  | NONE<br>WORK<br>QTEMP<br>ALL | The job number of the job.                                                                                 |
| INTERNAL_JOB_ID  | BINARY(16)  | NONE<br>WORK<br>QTEMP<br>ALL | The internal job identifier.                                                                               |
| SUBSYSTEM        | VARCHAR(10) | NONE<br>WORK<br>QTEMP<br>ALL | The name of the subsystem where the job is running.<br>Contains the null value if the job is a system job. |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name            | Data Type   | DETAILED_INFO option         | Description                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------|-------------|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_LIBRARY_NAME | VARCHAR(10) | NONE<br>WORK<br>QTEMP<br>ALL | Library containing the subsystem description.<br>Contains the null value if the job is a system job.                                                                                                                                                                                                                                                                                                                              |
| AUTHORIZATION_NAME     | VARCHAR(10) | NONE<br>WORK<br>QTEMP<br>ALL | The user profile under which the initial thread is running at this time. For jobs that swap user profiles, this user profile name and the user profile that initiated the job can be different.                                                                                                                                                                                                                                   |
| JOB_TYPE               | VARCHAR(3)  | NONE<br>WORK<br>QTEMP<br>ALL | Type of active job.<br><br><b>ASJ</b> Autostart<br><b>BCH</b> Batch<br><b>BCI</b> Batch Immediate<br><b>EVK</b> Started by a procedure start request<br><b>INT</b> Interactive<br><b>M36</b> Advanced 36 server job<br><b>MRT</b> Multiple requester terminal<br><b>PDJ</b> Print driver job<br><b>PJ</b> Prestart job<br><b>RDR</b> Spool reader<br><b>SBS</b> Subsystem monitor<br><b>SYS</b> System<br><b>WTR</b> Spool writer |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name   | Data Type  | DETAILED_INFO option         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------|------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FUNCTION_TYPE | VARCHAR(3) | NONE<br>WORK<br>QTEMP<br>ALL | <p>The type of function described in the FUNCTION column.</p> <p><b>CMD</b> The FUNCTION column contains the name of the command being run.</p> <p><b>DLY</b> The initial thread of the job is processing a DLYJOB (Delay Job) command. The FUNCTION column contains a time that is the number of seconds the job is delayed (up to 999999 seconds), or the time when job is to resume processing (hh:mm:ss).</p> <p><b>GRP</b> The FUNCTION column contains the group name of a suspended group job.</p> <p><b>I/O</b> The job is a subsystem monitor that is performing input/output operations (I/O) to a work station for the sign-on display file. The FUNCTION column contains the name of the work station device.</p> <p><b>IDX</b> The FUNCTION column contains the name of the file associated with an index rebuild operation.</p> <p><b>JVM</b> The initial thread of the job is running a Java Virtual Machine. The FUNCTION column contains the name of the java class.</p> <p><b>LOG</b> The FUNCTION column contains <b>QHST</b> to indicate history information is being logged to a database file.</p> <p><b>MNU</b> The FUNCTION column contains the name of the menu.</p> <p><b>MRT</b> The job is either a multiple requester terminal (MRT) job if JOB_TYPE is BCH, or it is an interactive job attached to an MRT job if JOB_TYPE is INT.<br/>For an MRT job, the FUNCTION column contains information in the following format:</p> <ul style="list-style-type: none"> <li>• CHAR(2): The number of requesters currently attached to the MRT job.</li> <li>• CHAR(1): Contains a / (slash).</li> <li>• CHAR(2): The maximum number of requesters.</li> <li>• CHAR(1): Contains a blank.</li> <li>• CHAR(3): The never-ending program (NEP) indicator. A value of NEP indicates a never-ending program. A value of blanks indicates that it is not a never-ending program.</li> <li>• CHAR(1): Contains a blank.</li> </ul> <p>For an interactive job attached to an MRT, the FUNCTION column contains the name of the MRT procedure.</p> <p><b>PGM</b> The FUNCTION column contains the name of a program.</p> <p><b>PRC</b> The FUNCTION column contains the name of a procedure.</p> <p><b>USR</b> The FUNCTION column contains the user-specified function set with the Change Current Job (QWCCCJOB) API.</p> <p>Contains the null value if none of these values apply.</p> |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name | Data Type   | DETAILED_INFO option         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------|-------------|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FUNCTION    | VARCHAR(10) | NONE<br>WORK<br>QTEMP<br>ALL | <p>The last high-level function initiated by the initial thread.</p> <p>If FUNCTION_TYPE is not null, contains a value as defined by the FUNCTION_TYPE column. Otherwise, can contain one of the following values:</p> <p><b>ADLACTJOB</b> Auxiliary storage is being allocated for the number of active jobs specified in the QADLACTJ system value.</p> <p><b>ADLTOTJOB</b> Auxiliary storage is being allocated for the number of jobs specified in the QADLTOTJ system value.</p> <p><b>CMDENT</b> The command entry display is being used.</p> <p><b>COMMIT</b> The initial thread of the job is performing a commit operation.</p> <p><b>DIRSHD</b> This job is running under the directory shadowing function.</p> <p><b>DLTSPF</b> A spooled file is being deleted.</p> <p><b>DUMP</b> A dump is in process.</p> <p><b>JOBIDXRCY</b> A damaged job index is being recovered.</p> <p><b>JOBLOG</b> A job log is being produced.</p> <p><b>JOBLOGQRCY</b> The job log server queue is being recovered or rebuilt.</p> <p><b>PASSTHRU</b> The job is a pass-through job.</p> <p><b>RCLSPLSTG</b> Empty spooled database members are being deleted.</p> <p><b>ROLLBACK</b> The initial thread of the job is performing a rollback operation.</p> <p><b>SPLCLNUP</b> A cleanup of jobs on job queues and spooled files is being performed.</p> <p>Contains the null value if a logged function has not been performed.</p> |



Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name         | Data Type     | DETAILED_INFO option         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------|---------------|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_STATUS          | VARCHAR(4)    | NONE<br>WORK<br>QTEMP<br>ALL | <p>The status of the initial thread of the job. The following list contains some of the most common values. For a complete list of values, see <a href="#">Work Management API Attribute Descriptions in Application Programming Interfaces</a></p> <p><b>CMNW</b> Waiting for the completion of an I/O operation to a communications device.</p> <p><b>CNDW</b> Waiting on handle-based condition.</p> <p><b>DEQW</b> Waiting for completion of a dequeue operation.</p> <p><b>DLYW</b> Due to the Delay Job (DLYJOB) command, the initial thread of the job is delayed while it waits for a time interval to end, or for a specific delay end time.</p> <p><b>DSPW</b> Waiting for input from a work station display.</p> <p><b>END</b> The job has been ended with the *IMMED option, or its delay time has ended with the *CNTRLD option.</p> <p><b>EOJ</b> Ending for a reason other than running the End Job (ENDJOB) or End Subsystem (ENDSBS) command.</p> <p><b>EVTW</b> Waiting for an event.</p> <p><b>HLD</b> The job is being held.</p> <p><b>JVAW</b> Waiting for completion of a Java program operation.</p> <p><b>LCKW</b> Waiting for a lock.</p> <p><b>LSPW</b> Waiting for a lock space to be attached.</p> <p><b>MSGW</b> Waiting for a message from a message queue.</p> <p><b>MTXW</b> Waiting for a mutex.</p> <p><b>PSRW</b> A prestart job waiting for a program start request.</p> <p><b>RUN</b> Job is currently running.</p> <p><b>SEMW</b> Waiting for a semaphore.</p> <p><b>THDW</b> Waiting for another thread to complete an operation.</p> |
| MEMORY_POOL         | VARCHAR(9)    | NONE<br>WORK<br>QTEMP<br>ALL | The identifier of the system-related pool from which the job's main storage is allocated. This is the pool that the threads in the job start in.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| RUN_PRIORITY        | INTEGER       | NONE<br>WORK<br>QTEMP<br>ALL | The run priority of the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| THREAD_COUNT        | INTEGER       | NONE<br>WORK<br>QTEMP<br>ALL | The number of active threads in the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| TEMPORARY_STORAGE   | INTEGER       | NONE<br>WORK<br>QTEMP<br>ALL | The amount of temporary storage, in megabytes, that is currently allocated to this job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CPU_TIME            | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The total processing unit time used by the job, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| TOTAL_DISK_IO_COUNT | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The total number of disk I/O operations performed by the job across all routing steps. This is the sum of the asynchronous and synchronous disk I/O.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name                 | Data Type     | DETAILED_INFO option         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------|---------------|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ELAPSED_INTERACTION_COUNT   | INTEGER       | NONE<br>WORK<br>QTEMP<br>ALL | The number of interactions. This is the number of operator interactions during the measurement time interval.<br>Contains the null value if the job is not interactive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| ELAPSED_TOTAL_RESPONSE_TIME | INTEGER       | NONE<br>WORK<br>QTEMP<br>ALL | The total response time over the measurement time interval, in seconds.<br>Contains the null value if the job is not interactive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| ELAPSED_TOTAL_DISK_IO_COUNT | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The number of disk I/O operations performed by the job during the measurement time interval. This is the sum of the asynchronous and synchronous disk I/O.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| ELAPSED_ASYNC_DISK_IO_COUNT | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The number of asynchronous (physical) disk I/O operations performed by the job during the measurement time interval. This value is the sum of the asynchronous database and nondatabase reads and writes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| ELAPSED_SYNC_DISK_IO_COUNT  | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The number of synchronous (physical) disk I/O operations performed by the job during the measurement time interval. This value is the sum of the synchronous database and nondatabase reads and writes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| ELAPSED_CPU_PERCENTAGE      | DECIMAL(10,2) | NONE<br>WORK<br>QTEMP<br>ALL | The percent of processing unit time attributed to this job during the measurement time interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| ELAPSED_CPU_TIME            | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The total CPU time spent during the measurement time interval, in milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| ELAPSED_PAGE_FAULT_COUNT    | DECIMAL(20,0) | NONE<br>WORK<br>QTEMP<br>ALL | The number of times an active program referenced an address that is not in main storage for the specified job during the measurement time interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| JOB_END_REASON              | VARCHAR(60)   | NONE<br>WORK<br>QTEMP<br>ALL | Reason the job is ending. Contains one of the following values: <ul style="list-style-type: none"> <li>• JOB ENDED DUE TO A DEVICE ERROR</li> <li>• JOB ENDED DUE TO A SIGNAL</li> <li>• JOB ENDED DUE TO AN UNHANDLED ERROR</li> <li>• JOB ENDED DUE TO THE CPU LIMIT BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE DISCONNECT TIME INTERVAL BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE INACTIVITY TIME INTERVAL BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE MESSAGE SEVERITY LEVEL BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE STORAGE LIMIT BEING EXCEEDED</li> <li>• JOB ENDED WHILE IT WAS STILL ON A JOB QUEUE</li> <li>• JOB ENDING ABNORMALLY</li> <li>• JOB ENDING IMMEDIATELY</li> <li>• JOB ENDING IN NORMAL MANNER</li> <li>• JOB ENDING NORMALLY AFTER A CONTROLLED END WAS REQUESTED</li> <li>• SYSTEM ENDED ABNORMALLY</li> </ul> Contains the null value if job is not currently ending. |
| SERVER_TYPE                 | VARCHAR(30)   | NONE<br>WORK<br>QTEMP<br>ALL | The type of server represented by the job. See <a href="#">Server</a> table for a list of server type values.<br>Contains the null value if the job is not part of a server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name             | Data Type     | DETAILED_INFO option         | Description                                                                                                                                                                                                                                                          |
|-------------------------|---------------|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ELAPSED_TIME            | DECIMAL(20,3) | NONE<br>WORK<br>QTEMP<br>ALL | The time that has elapsed, in seconds, between the measurement start time and the current system time.                                                                                                                                                               |
| JOB_DESCRIPTION_LIBRARY | VARCHAR(10)   | WORK<br>ALL                  | The name of the library containing the job description.<br>Contains the null value if the job has no job description.                                                                                                                                                |
| JOB_DESCRIPTION         | VARCHAR(10)   | WORK<br>ALL                  | The name of the job description used for this job.<br>Contains the null value if the job has no job description.                                                                                                                                                     |
| JOB_QUEUE_LIBRARY       | VARCHAR(10)   | WORK<br>ALL                  | The name of the library containing the job queue.<br>Contains the null value if the job is not a batch job that was started from a job queue.                                                                                                                        |
| JOB_QUEUE               | VARCHAR(10)   | WORK<br>ALL                  | The name of the job queue that the job was on.<br>Contains the null value if the job is not a batch job that was started from a job queue.                                                                                                                           |
| OUTPUT_QUEUE_LIBRARY    | VARCHAR(10)   | WORK<br>ALL                  | The name of the library that contains the default output queue.<br>Contains the null value if the job has no default output queue.                                                                                                                                   |
| OUTPUT_QUEUE            | VARCHAR(10)   | WORK<br>ALL                  | The name of the default output queue that is used for spooled output produced by this job. The default output queue is only used by spooled printer files that specify *JOB for the output queue.<br>Contains the null value if the job has no default output queue. |
| WORKLOAD_GROUP          | VARCHAR(10)   | WORK<br>ALL                  | The name of the workload group to which the job belongs.<br>Contains the null value if the job is not part of a workload group.                                                                                                                                      |
| CCSID                   | INTEGER       | WORK<br>ALL                  | The coded character set identifier (CCSID) used for this job.                                                                                                                                                                                                        |
| DEFAULT_CCSID           | INTEGER       | WORK<br>ALL                  | The default coded character set identifier used for this job.                                                                                                                                                                                                        |
| SORT_SEQUENCE_LIBRARY   | VARCHAR(10)   | WORK<br>ALL                  | The name of the library that contains the sort sequence table.<br>Contains the null value if no sort sequence table is defined for this job or if SORT_SEQUENCE is a special value.                                                                                  |
| SORT_SEQUENCE           | VARCHAR(10)   | WORK<br>ALL                  | The name of the sort sequence table associated with this job.<br>Contains the null value if no sort sequence table is defined for this job.                                                                                                                          |
| LANGUAGE_ID             | CHAR(3)       | WORK<br>ALL                  | The language identifier associated with this job.                                                                                                                                                                                                                    |
| DATE_FORMAT             | CHAR(4)       | WORK<br>ALL                  | The date format used for this job.<br><b>*DMY</b> Day, month, year format.<br><b>*JUL</b> Julian format (year and day).<br><b>*MDY</b> Month, day, year format.<br><b>*YMD</b> Year, month, day format.                                                              |
| DATE_SEPARATOR          | CHAR(1)       | WORK<br>ALL                  | The date separator used for this job.                                                                                                                                                                                                                                |
| TIME_SEPARATOR          | CHAR(1)       | WORK<br>ALL                  | The time separator used for this job.                                                                                                                                                                                                                                |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name               | Data Type   | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|-------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DECIMAL_FORMAT            | VARCHAR(6)  | WORK<br>ALL          | <p>The decimal format used for this job.</p> <p><b>*BLANK</b> Uses a period for a decimal point, a comma for a 3-digit grouping character, and zero-suppress to the left of the decimal point.</p> <p><b>J</b> Uses a comma for a decimal point and a period for a 3-digit grouping character. The zero-suppression character is in the second position (rather than the first) to the left of the decimal notation. Balances with zero values to the left of the comma are written with one leading zero (0,04). The J entry also overrides any edit codes that might suppress the leading zero.</p> <p><b>I</b> Uses a comma for a decimal point, a period for a 3-digit grouping character, and zero-suppress to the left of the decimal point.</p> |
| TIMEZONE_DESCRIPTION      | VARCHAR(10) | ALL                  | The name of the time zone description that is used to calculate local job time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| TIMEZONE_CURRENT_OFFSET   | INTEGER     | ALL                  | The offset, in minutes, used to calculate local job time. This value has been adjusted for Daylight Saving Time, if necessary.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| TIMEZONE_FULL_NAME        | VARCHAR(50) | ALL                  | <p>The full, or long, name for the time zone. This column returns either the standard or Daylight Saving Time full name depending on whether or not Daylight Saving Time is in effect.</p> <p>Contains the null value if the time zone description uses a message to specify the current full name and the message cannot be retrieved.</p>                                                                                                                                                                                                                                                                                                                                                                                                            |
| TIMEZONE_ABBREVIATED_NAME | VARCHAR(10) | ALL                  | <p>The abbreviated, or short, name for the time zone. This column returns either the standard or Daylight Saving Time abbreviated name depending on whether or not Daylight Saving Time is in effect.</p> <p>Contains the null value if the time zone description uses a message to specify the current abbreviated name and the message cannot be retrieved.</p>                                                                                                                                                                                                                                                                                                                                                                                      |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name             | Data Type                                                   | DETAILED_INFO option | Description                                                                                                                                     |                                                                   |
|-------------------------|-------------------------------------------------------------|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| JOB_TYPE_ENHANCED       | VARCHAR(28)                                                 | WORK<br>ALL          | The combined job type and job subtype values.                                                                                                   |                                                                   |
|                         |                                                             |                      | <b>ALTERNATE_SPOOL_USER</b>                                                                                                                     | Batch - alternate spool user                                      |
|                         |                                                             |                      | <b>AUTOSTART</b>                                                                                                                                | Autostart job                                                     |
|                         |                                                             |                      | <b>BATCH</b>                                                                                                                                    | Batch job                                                         |
|                         |                                                             |                      | <b>BATCH_IMMEDIATE</b>                                                                                                                          | Batch immediate job                                               |
|                         |                                                             |                      | <b>BATCH_MRT</b>                                                                                                                                | Batch - System/36 multiple requester terminal (MRT) job           |
|                         |                                                             |                      | <b>COMM_PROCEDURE_START_REQUEST</b>                                                                                                             | Communications job - procedure start request job                  |
|                         |                                                             |                      | <b>INTERACTIVE</b>                                                                                                                              | Interactive job                                                   |
|                         |                                                             |                      | <b>INTERACTIVE_GROUP</b>                                                                                                                        | Interactive job - Part of group                                   |
|                         |                                                             |                      | <b>INTERACTIVE_SYSREQ</b>                                                                                                                       | Interactive job - Part of system request pair                     |
|                         |                                                             |                      | <b>INTERACTIVE_SYSREQ_AND_GROUP</b>                                                                                                             | Interactive job - Part of system request pair and part of a group |
|                         |                                                             |                      | <b>PRESTART</b>                                                                                                                                 | Prestart job                                                      |
|                         |                                                             |                      | <b>PRESTART_BATCH</b>                                                                                                                           | Prestart batch job                                                |
|                         |                                                             |                      | <b>PRESTART_COMM</b>                                                                                                                            | Prestart communications job                                       |
|                         |                                                             |                      | <b>READER</b>                                                                                                                                   | Reader job                                                        |
| <b>SUBSYSTEM</b>        | Subsystem job                                               |                      |                                                                                                                                                 |                                                                   |
| <b>SYSTEM</b>           | System job (all system jobs including SCPF)                 |                      |                                                                                                                                                 |                                                                   |
| <b>WRITER</b>           | Writer job (including both spool writers and print drivers) |                      |                                                                                                                                                 |                                                                   |
| JOB_ENTERED_SYSTEM_TIME | TIMESTAMP(0)                                                | WORK<br>ALL          | The timestamp for when the job was placed on the system.                                                                                        |                                                                   |
| JOB_ACTIVE_TIME         | TIMESTAMP(0)                                                | WORK<br>ALL          | The timestamp for when the job began to run on the system.                                                                                      |                                                                   |
| CLIENT_IP_ADDRESS       | VARCHAR(45)                                                 | ALL                  | Client IP address, in IPv4 format, being used by the job.<br>Contains the null value when no client IP address exists or the job is using IPv6. |                                                                   |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name               | Data Type   | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------|-------------|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_USER_IDENTITY_SETTING | VARCHAR(11) | ALL                  | <p>The method by which the job user identity was set.</p> <p><b>APPLICATION</b> The job user identity was explicitly set by an application using one of the Set Job User Identity APIs, QWTSJUID or QwtSetJuid(). The job may be running either single threaded or multithreaded.</p> <p><b>DEFAULT</b> The job is currently running single threaded and the job user identity is the name of the user profile under which the job is currently running.</p> <p><b>SYSTEM</b> The job is currently running multithreaded and the job user identity was implicitly set by the system when the job became multithreaded. It was set to the name of the user profile that the job was running under when it became multithreaded.</p> |
| JOB_USER_IDENTITY         | VARCHAR(10) | ALL                  | <p>The user profile name by which the job is known to other jobs on the system. The job user identity is used for authorization checks when other jobs on the system attempt to operate against the job.</p> <p>Contains the null value if the user profile no longer exists.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| DBCS_CAPABLE              | VARCHAR(3)  | ALL                  | <p>Whether the job is DBCS-capable.</p> <p><b>NO</b> The job is not DBCS-capable.</p> <p><b>YES</b> The job is DBCS-capable.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| SIGNAL_STATUS             | VARCHAR(3)  | ALL                  | <p>Whether the job is enabled to receive signals from another job or the system.</p> <p><b>NO</b> The job is not enabled for signals. This job cannot receive signals from another job or the system.</p> <p><b>YES</b> The job is enabled for signals. This job can receive signals from another job or the system.</p>                                                                                                                                                                                                                                                                                                                                                                                                           |
| MESSAGE_REPLY             | VARCHAR(3)  | ALL                  | <p>Whether the job is waiting for a reply to a specific message.</p> <p><b>NO</b> The job is not waiting for a reply to a message.</p> <p><b>YES</b> The job is waiting for a reply to a message.</p> <p>Contains the null value if the job is not in message wait status.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| END_STATUS                | VARCHAR(3)  | ALL                  | <p>Whether the system issued a controlled cancellation.</p> <p><b>NO</b> The system, subsystem, or job is not canceled.</p> <p><b>YES</b> The system, the subsystem in which the job is running, or the job itself is canceled.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| CANCEL_KEY                | VARCHAR(3)  | ALL                  | <p>Whether the user pressed the cancel key.</p> <p><b>NO</b> The user did not press the cancel key.</p> <p><b>YES</b> The user pressed the cancel key.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| EXIT_KEY                  | VARCHAR(3)  | ALL                  | <p>Whether the user pressed the exit key.</p> <p><b>NO</b> The user did not press the exit key.</p> <p><b>YES</b> The user pressed the exit key.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| MAXIMUM_ACTIVE_THREADS    | INTEGER     | ALL                  | <p>The maximum number of threads that a job can run with at any time. If multiple threads are initiated simultaneously, this value may be exceeded. If this maximum value is exceeded, the excess threads will be allowed to run to their normal completion. Initiation of additional threads will be inhibited until the maximum number of threads in the job drops below this maximum value.</p> <p>Contains the null value if there is no maximum.</p>                                                                                                                                                                                                                                                                          |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name                       | Data Type   | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------|-------------|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SYSTEM_POOL_ID                    | INTEGER     | ALL                  | <p>The identifier of the system-related pool from which main storage is currently being allocated for the job's initial thread. These identifiers are not the same as those specified in the subsystem description, but are the same as the system pool identifiers shown on the system status display. If a thread reaches its time-slice end, the pool the thread is running in can be switched based on the job's time-slice end pool value. The current system pool identifier returned will be the actual pool in which the initial thread of the job is running.</p> <p>Contains the null value if the value is not available.</p>                                                                                     |
| POOL_NAME                         | VARCHAR(10) | ALL                  | <p>The name of the memory pool in which the job started running. The name may be a number, in which case it is a private pool associated with a subsystem. Can contain one of the following special values:</p> <p><b>*BASE</b> This job is running in the base system pool, which can be shared with other subsystems.</p> <p><b>*INTERACT</b> This job is running in the shared pool used for interactive work.</p> <p><b>*MACHINE</b> This job is running in the machine pool.</p> <p><b>*SHRPOOL1 - *SHRPOOL60</b> This job is running in the identified shared pool.</p> <p><b>*SPOOL</b> This job is running in the shared pool for spooled writers.</p> <p>Contains the null value if the value is not available.</p> |
| QTEMP_SIZE                        | INTEGER     | QTEMP<br>ALL         | <p>The amount of storage, in megabytes, used by objects in the job's temporary library (QTEMP). Objects that are locked, damaged, or not authorized are not included.</p> <p>Contains the null value if the size cannot be returned.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| PEAK_TEMPORARY_STORAGE            | INTEGER     | ALL                  | <p>The maximum amount of auxiliary storage, in megabytes, that the job has used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| DEFAULT_WAIT                      | INTEGER     | ALL                  | <p>The default maximum time, in seconds, that a thread in the job waits for a system instruction, such as a LOCK machine interface (MI) instruction, to acquire a resource.</p> <p>Contains the null value if there is no maximum or if the value is not available.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| MAXIMUM_PROCESSING_TIME_ALLOWED   | INTEGER     | ALL                  | <p>The maximum processing unit time, in milliseconds, that the job can use. If the job consists of multiple routing steps, this is the maximum processing unit time that the current routing step can use. If the maximum time is exceeded, the job is held.</p> <p>Contains the null value if no maximum amount of processing unit time has been defined.</p>                                                                                                                                                                                                                                                                                                                                                               |
| MAXIMUM_TEMPORARY_STORAGE_ALLOWED | INTEGER     | ALL                  | <p>The maximum amount of auxiliary storage, in megabytes, that the job can use. If the job consists of multiple routing steps, this is the maximum temporary storage that the routing step can use. This temporary storage is used for storage required by the program itself and by implicitly created internal system objects used to support the routing step. (It does not include storage for objects in the QTEMP library.) If the maximum temporary storage is exceeded, the job is held. This does not apply to the use of permanent storage, which is controlled through the user profile.</p> <p>Contains the null value if no maximum amount of temporary storage has been defined.</p>                           |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name                     | Data Type      | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|----------------|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TIME_SLICE                      | INTEGER        | ALL                  | The maximum amount of processor time, in milliseconds, given to each thread in this job before other threads in this job and in other jobs are given the opportunity to run. The time slice establishes the amount of time needed by a thread in this job to accomplish a meaningful amount of processing. At the end of the time slice, the thread might be put in an inactive state so that other threads can become active in the storage pool. Values range from 8 through 9999999.<br><br>Contains the null value if the value is not available. |
| PAGE_FAULTS                     | BIGINT         | ALL                  | The number of times an active program referenced an address that was not in main storage during the current routing step of the specified job.                                                                                                                                                                                                                                                                                                                                                                                                        |
| TOTAL_RESPONSE_TIME             | BIGINT         | ALL                  | The total amount of response time for the initial thread, in milliseconds. This value does not include the time used by the machine, by the attached input/output (I/O) hardware, and by the transmission lines for sending and receiving data. Returns zero for jobs that have no interactions. A value of -1 is returned if the field is not large enough to hold the actual result.                                                                                                                                                                |
| INTERACTIVE_TRANSACTIONS        | INTEGER        | ALL                  | The count of operator interactions, such as pressing the Enter key or a function key. Returns zero for jobs that have no interactions.                                                                                                                                                                                                                                                                                                                                                                                                                |
| DATABASE_LOCK_WAITS             | INTEGER        | ALL                  | The number of times that the initial thread had to wait to obtain a database lock.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| NON_DATABASE_LOCK_WAITS         | INTEGER        | ALL                  | The number of times that the initial thread had to wait to obtain a nondatabase lock.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| INTERNAL_MACHINE_LOCK_WAITS     | INTEGER        | ALL                  | The number of times that the initial thread had to wait to obtain an internal machine lock.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| DATABASE_LOCK_WAIT_TIME         | INTEGER        | ALL                  | The cumulative amount of time, in milliseconds, that the initial thread has had to wait to obtain database locks.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| NON_DATABASE_LOCK_WAIT_TIME     | INTEGER        | ALL                  | The cumulative amount of time, in milliseconds, that the initial thread has had to wait to obtain nondatabase locks.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| INTERNAL_MACHINE_LOCK_WAIT_TIME | INTEGER        | ALL                  | The cumulative amount of time, in milliseconds, that the initial thread has had to wait to obtain internal machine locks.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| SQL_STATEMENT_TEXT              | VARCHAR(10000) | ALL                  | Statement text of the last SQL statement to run or the SQL statement that is currently running. The statement text will be truncated if it is longer than the column.<br><br>Contains the null value if no SQL statement has been run.                                                                                                                                                                                                                                                                                                                |
| SQL_STATEMENT_STATUS            | VARCHAR(8)     | ALL                  | The status of SQL within this job.<br><br><b>ACTIVE</b> An SQL statement is currently running<br><b>COMPLETE</b> At least one SQL statement has run and has completed<br><br>Contains the null value if no SQL statement has been run.                                                                                                                                                                                                                                                                                                                |
| SQL_STATEMENT_START_TIMESTAMP   | TIMESTAMP      | ALL                  | The timestamp of the execution start for an active SQL statement.<br><br>Contains the null value if there is no active SQL statement.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| SQL_STATEMENT_NAME              | VARCHAR(128)   | ALL                  | The name of the SQL statement.<br><br>Contains the null value when the SQL statement has no name.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| SQL_STATEMENT_LIBRARY_NAME      | VARCHAR(10)    | ALL                  | The library name for the SQL statement object.<br><br>Contains the null value when the SQL statement name is null or when the SQL statement does not exist within a permanent object.                                                                                                                                                                                                                                                                                                                                                                 |



Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name                | Data Type    | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|--------------|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SQL_STATEMENT_OBJECT_NAME  | VARCHAR(10)  | ALL                  | The name of the object which contains the last SQL statement executed in the job. When the current SQL statement belongs to an SQL function or an SQL procedure, the object name will be the external program name.<br><br>Contains the null value when the SQL statement name is null or when the SQL statement does not exist within a permanent object.          |
| SQL_STATEMENT_OBJECT_TYPE  | VARCHAR(7)   | ALL                  | The type of object containing the current SQL statement.<br><br><b>*PGM</b> The current SQL statement resides within a program.<br><b>*SQLPKG</b> The current SQL statement resides within an SQL package.<br><b>*SRVPGM</b> The current SQL statement resides within a service program.<br><br>Contains the null value when the SQL statement object name is null. |
| QUERY_OPTIONS_LIBRARY_NAME | VARCHAR(10)  | ALL                  | The name of the QAQQINI options library in use for this job.                                                                                                                                                                                                                                                                                                        |
| SQL_ACTIVATION_GROUP_COUNT | INTEGER      | ALL                  | The number of activation groups, current and ended, that have executed SQL statements for the job.<br><br>Contains the null value if no SQL statement has been run.                                                                                                                                                                                                 |
| SQL_DESCRIPTOR_COUNT       | BIGINT       | ALL                  | The number of SQL descriptors that are active for the job.<br><br>Contains the null value if no SQL descriptors are active for the job.                                                                                                                                                                                                                             |
| SQL_LOB_LOCATOR_COUNT      | INTEGER      | ALL                  | The number of LOB locators that are active for the job.<br><br>Contains the null value if no LOB locators are active for the job.                                                                                                                                                                                                                                   |
| CLI_HANDLE_COUNT           | BIGINT       | ALL                  | The number of SQL Call Level Interface (CLI) handles that are active for the job. This count includes CLI statement handles, descriptor handles, environment handles, and connection handles.<br><br>Contains the null value if no CLI handles are active for the job.                                                                                              |
| SQL_SERVER_MODE            | VARCHAR(3)   | ALL                  | Indicates whether the job is configured to use SQL Server Mode.<br><br><b>NO</b> The job is not configured to use SQL Server Mode.<br><b>YES</b> The job is configured to use SQL Server Mode.                                                                                                                                                                      |
| CLIENT_ACCTNG              | VARCHAR(255) | ALL                  | Value of the SQL CURRENT CLIENT_ACCTNG special register. The value can be null. For more information, see <a href="#">CURRENT CLIENT_ACCTNG</a> .                                                                                                                                                                                                                   |
| CLIENT_APPLNAME            | VARCHAR(255) | ALL                  | Value of the SQL CURRENT CLIENT_APPLNAME special register. The value can be null. For more information, see <a href="#">CURRENT CLIENT_APPLNAME</a> .                                                                                                                                                                                                               |
| CLIENT_PROGRAMID           | VARCHAR(255) | ALL                  | Value of the SQL CURRENT CLIENT_PROGRAMID special register. The value can be null. For more information, see <a href="#">CURRENT CLIENT_PROGRAMID</a> .                                                                                                                                                                                                             |
| CLIENT_USERID              | VARCHAR(255) | ALL                  | Value of the SQL CURRENT CLIENT_USERID special register. The value can be null. For more information, see <a href="#">CURRENT CLIENT_USERID</a> .                                                                                                                                                                                                                   |
| CLIENT_WRKSTNNAME          | VARCHAR(255) | ALL                  | Value of the SQL CURRENT CLIENT_WRKSTNNAME special register. The value can be null. For more information, see <a href="#">CURRENT CLIENT_WRKSTNNAME</a> .                                                                                                                                                                                                           |
| ROUTINE_TYPE               | CHAR(1)      | ALL                  | For a routine defined using SQL, the type of the currently executing routine.<br><br><b>F</b> Function<br><b>P</b> Procedure<br><br>Contains the null value if there is no SQL routine currently executing.                                                                                                                                                         |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name                   | Data Type    | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                |
|-------------------------------|--------------|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ROUTINE_SCHEMA                | VARCHAR(128) | ALL                  | For a routine defined using SQL, the schema name of the currently executing routine.<br>Contains the null value if there is no SQL routine currently executing.                                                                                                                                                            |
| ROUTINE_SPECIFIC_NAME         | VARCHAR(128) | ALL                  | For a routine defined using SQL, the name of the currently executing routine.<br>Contains the null value if there is no SQL routine currently executing.                                                                                                                                                                   |
| CLIENT_PORT                   | INTEGER      | ALL                  | The port number used by the current client to communicate with the server.<br>Contains the null value if the target job does not correspond to a connection formed using the TCP/IP protocol.                                                                                                                              |
| CLIENT_HOST                   | VARCHAR(255) | ALL                  | The host name used by the current client to communicate with the server.<br>Contains the null value if the target job does not correspond to a connection formed using the TCP/IP protocol.                                                                                                                                |
| INTERFACE_NAME                | VARCHAR(127) | ALL                  | The client database interface name.<br>Contains the null value if there is no client database interface name.                                                                                                                                                                                                              |
| INTERFACE_TYPE                | VARCHAR(63)  | ALL                  | The client database interface type.<br>Contains the null value if there is no client database interface type.                                                                                                                                                                                                              |
| INTERFACE_LEVEL               | VARCHAR(63)  | ALL                  | The client database interface level in the following form: "VVRMMFP". VV - Version RR - Release MM - Modification level FP - Fix pack level (only applicable for certain interfaces).<br>Contains the null value if there is no client database interface level.                                                           |
| SERVER_MODE_CONNECTING_JOB    | VARCHAR(28)  | ALL                  | The qualified job name of the job that established the SQL Server Mode connection. If the job name is QSQRVR, then the qualified job name of the connecting job is returned.<br>Contains the null value if the job name is not QSQRVR or JOB_STATUS is PSRW.                                                               |
| SERVER_MODE_CONNECTING_THREAD | BIGINT       | ALL                  | If the job name is QSQRVR and the server mode job is in use, the thread identifier of the last thread to use this connection is returned. When SQL_STATEMENT_STATUS is COMPLETE, this application thread identifier might no longer exist.<br>Contains the null value if the job name is not QSQRVR or JOB_STATUS is PSRW. |
| PRESTART_JOB_REUSE_COUNT      | INTEGER      | ALL                  | The number of times the prestart job has been used. The prestart job reuse count is incremented when a disconnect is processed for a prestart job. When the prestart job reuse count exceeds the prestart job maximum number of uses, the job is ended.<br>Contains the null value if the job is not a prestart job.       |
| PRESTART_JOB_MAX_USE_COUNT    | INTEGER      | ALL                  | The maximum number of times the prestart job can be used before it is ended. A value of -1 is returned for *NOMAX.<br>Contains the null value if the job is not a prestart job.                                                                                                                                            |
| AVAILABLE_RESULT_SETS         | INTEGER      | ALL                  | The current count of unconsumed SQL result sets for the job.<br>Contains the null value if the job has no unconsumed SQL result sets.                                                                                                                                                                                      |
| UNCONSUMED_RESULT_SETS        | INTEGER      | ALL                  | The cumulative count of unconsumed SQL result sets that were discarded for the job.<br>Contains the null value if the job has no unconsumed SQL result sets that have been discarded.                                                                                                                                      |
| OPEN_CURSOR_COUNT             | INTEGER      | ALL                  | The number of SQL cursors that are currently open for the job.<br>Contains the null value if no SQL cursors are currently open for the job.                                                                                                                                                                                |

Table 204. ACTIVE\_JOB\_INFO table function (continued)

| Column Name                | Data Type  | DETAILED_INFO option | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------|------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FULL_OPEN_CURSOR_COUNT     | BIGINT     | ALL                  | The total number of SQL cursors that have been full opened for the life of the job.<br>Contains the null value if no SQL cursors have been full opened during the life of the job.                                                                                                                                                                                                                                                                                                                                          |
| PSEUDO_OPEN_CURSOR_COUNT   | BIGINT     | ALL                  | The total number of SQL cursors that have been pseudo opened for the life of the job. Pseudo opens are also known as reused SQL cursors.<br>Contains the null value if no SQL cursors have been pseudo opened during the life of the job.                                                                                                                                                                                                                                                                                   |
| PSEUDO_CLOSED_CURSOR_COUNT | INTEGER    | ALL                  | The active number of pseudo closed SQL cursors within the job. Pseudo closed cursors are cursors that have been closed by the application, but remain open within the database. A pseudo closed cursor may be reused when the same query is executed many times, resulting in a performance improvement on the open. Conversely, accumulating too many pseudo closed cursors within the job can have a negative impact on the storage footprint of the job.<br>Contains the null value if no SQL cursors are pseudo closed. |
| CQE_CURSOR_COUNT           | INTEGER    | ALL                  | The number of cursors using CQE for this job. This includes SQL cursors (both fully opened and pseudo closed) and cursors used to implement native database queries.<br>Contains the null value if no cursors have used CQE for this job.                                                                                                                                                                                                                                                                                   |
| CQE_CURSOR_STORAGE         | INTEGER    | ALL                  | The amount of storage, in megabytes, used by cursors using CQE for this job.<br>Contains the null value if no cursors have used CQE for this job.                                                                                                                                                                                                                                                                                                                                                                           |
| SQE_CURSOR_COUNT           | INTEGER    | ALL                  | The number of cursors using SQE for this job. This includes SQL cursors (both fully opened and pseudo closed) and cursors used to implement native database queries.<br>Contains the null value if no cursors have used SQE for this job.                                                                                                                                                                                                                                                                                   |
| SQE_CURSOR_STORAGE         | INTEGER    | ALL                  | The amount of storage, in megabytes, used by cursors using SQE for this job.<br>Contains the null value if no cursors have used SQE for this job.                                                                                                                                                                                                                                                                                                                                                                           |
| LARGEST_QUERY_SIZE         | INTEGER    | ALL                  | The amount of storage, in megabytes, used by the SQE cursor that used the most storage for this job.<br>Contains the null value if no cursors have used SQE for this job.                                                                                                                                                                                                                                                                                                                                                   |
| QRO_HASH                   | VARCHAR(8) | ALL                  | An internally generated identifier for the SQE query referred to in the LARGEST_QUERY_SIZE column. The QRO hash surfaces within Visual Explain and from Show Statements exploration of the SQL Plan Cache and SQL Plan Cache Snapshots.<br>Contains the null value if no cursors have used SQE for this job.                                                                                                                                                                                                                |
| OPEN_FILES                 | INTEGER    | ALL                  | The number of open files (*FILE objects) for this job. For details about the types of files and their usage, use the QSYS2.OPEN_FILES table function.                                                                                                                                                                                                                                                                                                                                                                       |

## Examples

- **Example 1:** Looking at only QZDASOINIT jobs, find the top 10 consumers of Elapsed I/O.

```
SELECT JOB_NAME, AUTHORIZATION_NAME, ELAPSED_TOTAL_DISK_IO_COUNT,
 ELAPSED_CPU_PERCENTAGE
FROM TABLE(QSYS2.ACTIVE_JOB_INFO(
 JOB_NAME_FILTER => 'QZDASOINIT',
 SUBSYSTEM_LIST_FILTER => 'QUSRWRK')) X
ORDER BY ELAPSED_TOTAL_DISK_IO_COUNT DESC
FETCH FIRST 10 ROWS ONLY;
```

Note: The data in the ELAPSED\_xxx columns is updated upon each re-execution of the query. Elapsed data will not get returned the first time a query is run for ACTIVE\_JOB\_INFO for a connection. See the *reset-statistics* parameter for details.

- **Example 2:** Find the active jobs using the most temporary storage. Include the most recently executed SQL statement for each target job.

```
SELECT JOB_NAME, AUTHORIZATION_NAME, TEMPORARY_STORAGE, SQL_STATEMENT_TEXT
FROM TABLE (QSYS2.ACTIVE_JOB_INFO(DETAILED_INFO => 'ALL')) X
WHERE JOB_TYPE <> 'SYS'
ORDER BY TEMPORARY_STORAGE DESC;
```

## AUTOSTART\_JOB\_INFO view

The AUTOSTART\_JOB\_INFO view returns information about autostart jobs.

The values returned for the columns in the view are closely related to the values returned by the Display Autostart Job Entries panel accessed through the DSPSBSD (Display Subsystem Description) CL command and by the List Subsystem Entries (QWDLSE) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is AUTOJ\_INFO. The schema is QSYS2.

Table 205. AUTOSTART\_JOB\_INFO view

| Column Name                   | System Column Name | Data Type   | Description                                                                                        |
|-------------------------------|--------------------|-------------|----------------------------------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10) | The name of the library in which the subsystem description resides.                                |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10) | The name of the subsystem about which information is being returned.                               |
| AUTOSTART_JOB_NAME            | AJ_NAME            | VARCHAR(10) | The simple name of the job that is automatically started when the associated subsystem is started. |
| JOB_DESCRIPTION_LIBRARY       | JOBDLIB            | VARCHAR(10) | The name of the library in which the job description for the autostart job entry resides.          |
| JOB_DESCRIPTION               | JOBID              | VARCHAR(10) | The name of the job description for the autostart job entry.                                       |

## Example

- List all the autostart job entries in the QUSRWRK subsystem.

```
SELECT AUTOSTART_JOB_NAME, JOB_DESCRIPTION_LIBRARY, JOB_DESCRIPTION
FROM QSYS2.AUTOSTART_JOB_INFO
WHERE SUBSYSTEM_DESCRIPTION_LIBRARY = 'QSYS' AND
 SUBSYSTEM_DESCRIPTION = 'QSYSWRK'
ORDER BY 1, 2, 3;
```

## COMMUNICATIONS\_ENTRY\_INFO view

The COMMUNICATIONS\_ENTRY\_INFO view returns information about subsystem communications entries.

The values returned for the columns in the view are closely related to the values returned by the Display Communications Entries and Display Remote Location Name Entries panels accessed through the DSPSBSD (Display Subsystem Description) CL command and by the List Subsystem Entries (QWDLSE) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is COMM\_INFO. The schema is QSYS2.

Table 206. COMMUNICATIONS\_ENTRY\_INFO view

| Column Name                   | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10)             | The name of the library in which the subsystem description resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10)             | The name of the subsystem about which information is being returned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| DEVICE                        | DEVICE             | VARCHAR(10)<br>Nullable | <p>The name of the device description or the type of the device being used with this communications entry. Can contain the following special values:</p> <p><b>*ALL</b> All communications device types are used with this communications entry.</p> <p><b>*APPC</b> All advanced program-to-program communications devices can be used with this communications entry.</p> <p><b>*ASYNC</b> All asynchronous communications devices can be used with this communications entry.</p> <p><b>*BSCCL</b> All bisynchronous equivalency link communications devices can be used with this communications entry.</p> <p><b>*FINANCE</b> All finance communications devices can be used with this communications entry.</p> <p><b>*INTRA</b> All intrasystem communications devices can be used with this communications entry.</p> <p><b>*RETAIL</b> All retail communications devices can be used with this communications entry.</p> <p><b>*SNUF</b> All SNA upline facility communications devices can be used with this communications entry.</p> <p>Contains the null value if this entry was defined using a remote location name.</p> |
| REMOTE_LOCATION               | RMT_LOC            | VARCHAR(8)<br>Nullable  | <p>The name of the remote location for this entry.</p> <p>Contains the null value if this entry was defined using a device description name.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| MODE                          | MODE               | VARCHAR(8)              | <p>The mode name of the communications device. Can contain the following special value:</p> <p><b>*ANY</b> Any available modes defined to the communications device are allocated to the subsystem. If the communications device does not have defined modes associated with it, the communications device itself is allocated to the subsystem.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| JOB_DESCRIPTION_LIBRARY       | JOBDLIB            | VARCHAR(10)<br>Nullable | <p>The name of the library in which the communications entry job description resides.</p> <p>Contains the null value if JOB_DESCRIPTION is *USRPRF.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 206. COMMUNICATIONS\_ENTRY\_INFO view (continued)

| Column Name         | System Column Name | Data Type                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------|--------------------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_DESCRIPTION     | JOBID              | VARCHAR(10)                 | The name of the job description used when a job is started as a result of receiving a program start request and processed through this communications entry. Can contain the following special value:<br><br><b>*USRPRF</b> The job description name that is specified in the user profile of the user that made the program start request is used for jobs that are processed through this communications entry.                                                                                                                                                                   |
| DEFAULT_USER        | DFT_USER           | VARCHAR(10)<br><br>Nullable | The name of the default user profile used for evoke requests that enter the subsystem through this entry and contain no security information. Can contain the following special value:<br><br><b>*SYS</b> All user program start requests are treated the same as if no user profile is specified as the default. For program start requests that are sent by system functions, the request runs under a predetermined user profile if a user profile is not specified on the program start request.<br><br>Contains the null value if no user profile is specified as the default. |
| MAXIMUM_ACTIVE_JOBS | MAX_ACTIVE         | INTEGER<br><br>Nullable     | The maximum number of jobs that can be active at the same time through this entry.<br><br>Contains the null value if the entry specifies *NOMAX, indicating that there is no maximum.                                                                                                                                                                                                                                                                                                                                                                                               |

## Example

- List all the communications entries defined for the QCMN subsystem.

```
SELECT *
FROM QSYS2.COMMUNICATIONS_ENTRY_INFO
WHERE SUBSYSTEM_DESCRIPTION_LIBRARY = 'QSYS' AND
SUBSYSTEM_DESCRIPTION = 'QCMN';
```

## GET\_JOB\_INFO table function

The GET\_JOB\_INFO table function returns one row containing the information about a specific job.

**Authorization:** None required to for a job where the caller's user profile is the same as the [job user identity](#) of the job for which the information is being returned.

Otherwise, the caller must have either \*JOBCTL special authority, or QIBM\_DB\_SQLADM or QIBM\_DB\_SYSMON function usage authority.

```
➔➔ GET_JOB_INFO ((job-name ➔
 V_JOB_NAME =>)
, (ignore-errors ➔
 V_IGNORE_ERRORS =>)) ➔➔
```

The schema is QSYS2.

**job-name** A character or graphic string expression that identifies the name of a job. Two forms of a job name are supported.

1. A fully qualified job name in the form *job-number/job-user/job-name*.
2. The first 10 characters are the job name, the second 10 characters are the user name, and the last 6 characters are the job number.

The special value of '\*' indicates the current job.

**ignore-errors**

A character or graphic string expression that identifies what to do when an error is encountered.

**NO** An error is returned. This is the default.

**YES** A warning is returned. No row is returned when an error is encountered.

The result of the function is a table containing a single row with the format shown in the following table. All the columns are nullable.

Table 207. GET\_JOB\_INFO table function

| Column Name            | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V_JOB_STATUS           | CHAR(10)    | Status of the job.<br><br><b>*ACTIVE</b> Job is active. It could be a group job, system request job, or disconnected job.<br><b>*JOBQ</b> Job is currently on job queue.<br><b>*OUTQ</b> Job has completed running but has output on an output queue or the job log has not yet been written.                                                                                                                                                                                                                              |
| V_ACTIVE_JOB_STATUS    | CHAR(4)     | The active status of the initial thread of the job.<br>For the list of values see <a href="#">Work Management API Attribute Descriptions in Application Programming Interfaces</a> and search on "Active job status".<br>Returns null if the job is in transition or is not active.                                                                                                                                                                                                                                        |
| V_ACTIVE_JOB_TYPE      | VARCHAR(3)  | Type of active job.<br><br><b>ASJ</b> Autostart<br><b>BCH</b> Batch<br><b>BCI</b> Batch Immediate<br><b>EVK</b> Started by a procedure start request<br><b>INT</b> Interactive<br><b>M36</b> Advanced 36 server job<br><b>MRT</b> Multiple requester terminal<br><b>PDJ</b> Print driver job<br><b>PJ</b> Prestart job<br><b>RDR</b> Spool reader<br><b>SBS</b> Subsystem monitor<br><b>SYS</b> System<br><b>WTR</b> Spool writer<br><br>Returns null if the job type is not available for jobs that are no longer active. |
| V_RUN_PRIORITY         | INTEGER     | The highest run priority allowed for any thread within this job.                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| V_AUTHORIZATION_NAME   | VARCHAR(10) | The user profile under which the initial thread is running at this time. For jobs that swap user profiles, this user profile name and the user profile that initiated the job can be different.                                                                                                                                                                                                                                                                                                                            |
| V_SBS_NAME             | CHAR(10)    | Name of subsystem where job is running.<br>Returns null if the job is not active.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| V_CPU_USED             | BIGINT      | The amount of CPU time (in milliseconds) that has been currently used by this job.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| V_TEMP_STORAGE_USED_MB | INTEGER     | The amount of auxiliary storage (in megabytes) that is currently allocated to this job.                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 207. GET\_JOB\_INFO table function (continued)

| Column Name                | Data Type      | Description                                                                                                                                                                                                                                                                                                                           |
|----------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| V_AUX_IO_REQUESTED         | BIGINT         | The number of auxiliary I/O requests performed by the job across all routing steps. This includes both database and nondatabase paging.                                                                                                                                                                                               |
| V_PAGE_FAULTS              | BIGINT         | The number of times an active program referenced an address that was not in main storage during the current routing step of the specified job.                                                                                                                                                                                        |
| V_CLIENT_WRKSTNNAME        | CHAR(255)      | Value of the SQL CLIENT_WRKSTNNAME special register.<br>Returns null if the job is not active.                                                                                                                                                                                                                                        |
| V_CLIENT_APPLNAME          | CHAR(255)      | Value of the SQL CLIENT_APPLNAME special register.<br>Returns null if the job is not active.                                                                                                                                                                                                                                          |
| V_CLIENT_ACCTNG            | CHAR(255)      | Value of the SQL CLIENT_ACCTNG special register.<br>Returns null if the job is not active.                                                                                                                                                                                                                                            |
| V_CLIENT_PROGRAMID         | CHAR(255)      | Value of the SQL CLIENT_PROGRAMID special register.<br>Returns null if the job is not active.                                                                                                                                                                                                                                         |
| V_CLIENT_USERID            | CHAR(255)      | Value of the SQL CLIENT_USERID special register.<br>Returns null if the job is not active.                                                                                                                                                                                                                                            |
| V_SQL_STATEMENT_TEXT       | VARCHAR(10000) | Statement text of the last SQL statement to run or the SQL statement that is currently running.<br>Returns null if the job is not active.                                                                                                                                                                                             |
| V_SQL_STMT_STATUS          | CHAR(8)        | The status of SQL within this job.<br><br><b>ACTIVE</b> An SQL statement is currently running<br><b>COMPLETE</b> At least one SQL statement has run and has completed<br><b>UNKNOWN</b> The SQL status is not known<br>Returns null if no SQL statement has been run.                                                                 |
| V_SQL_STMT_START_TIMESTAMP | TIMESTAMP      | The timestamp of the execution start for an active SQL statement. If there is no active SQL statement, the null value is returned.                                                                                                                                                                                                    |
| V_QUERY_OPTIONS_LIB_NAME   | CHAR(10)       | The name of the QAQQINI options library in use for this job.                                                                                                                                                                                                                                                                          |
| V_CLIENT_IP_ADDRESS        | VARCHAR(45)    | Client IP address being used by the job.<br>Returns null when no client IP address exists or the job is using IPv6.                                                                                                                                                                                                                   |
| V_PJ_REUSE_COUNT           | INTEGER        | The number of times the prestart job has been used. The prestart job reuse count is incremented when a disconnect is processed for a prestart job. When the prestart job reuse count exceeds the prestart job maximum number of uses, the job is ended.<br>Returns null if the job is not active or if the job is not a prestart job. |
| V_PJ_MAXUSE_COUNT          | INTEGER        | The maximum number of times the prestart job can be used before it is ended. A value of -1 is returned for *NOMAX.<br>Returns null if the job is not active or if the job is not a prestart job.                                                                                                                                      |

## Example

Return information about job 347117/QUSER/QZDASOINIT.

```
SELECT * FROM TABLE(QSYS2.GET_JOB_INFO('347117/QUSER/QZDASOINIT'));
```

## JOB\_DESCRIPTION\_INFO view

The JOB\_DESCRIPTION\_INFO view returns information about job descriptions.

The values returned for the columns in the view are closely related to the values returned by the Display Job Description (DSPJOBDD) CL command and the Retrieve Job Description Information (QWDRJOBDD) API.

**Authorization:** The caller must have:

- \*EXECUTE authority to the library containing the job description, and



- \*OBJOPR and \*READ authorities to the job description.

The following table describes the columns in the view. The system name is JOBD\_INFO. The schema is QSYS2.

Table 208. JOB\_DESCRIPTION\_INFO view

| Column Name             | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|--------------------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_DESCRIPTION_LIBRARY | JOBDLIB            | VARCHAR(10)               | The name of the library in which the job description resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| JOB_DESCRIPTION         | JOBID              | VARCHAR(10)               | The name of the job description about which information is being returned.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| AUTHORIZATION_NAME      | USER_NAME          | VARCHAR(10)               | The name of the user profile associated with this job description. Can contain the following special value:<br><br><b>*RQD</b> A user name is required to use the job description.                                                                                                                                                                                                                                                                                                                            |
| JOB_DATE                | JOB_DATE           | DATE<br>Nullable          | The date that will be assigned to jobs using this job description when they are started.<br><br>Contains the null value if this job will use the QDATE system value.                                                                                                                                                                                                                                                                                                                                          |
| ACCOUNTING_CODE         | ACGCDE             | VARCHAR(15)               | An identifier assigned to jobs that use this job description. This code is used to collect system resource use information. Can contain the following special value:<br><br><b>*USRPRF</b> The accounting code used for jobs using this job description is obtained from the job's user profile.                                                                                                                                                                                                              |
| ROUTING_DATA            | RTGDTA             | VARCHAR(80)               | The routing data that is used with this job description to start jobs. Can contain one of the following special values:<br><br><b>QCMDI</b> The default routing data QCMDI is used by the IBM-supplied interactive subsystem to route the job to the IBM-supplied control language processor QCMD in the QSYS library.<br><br><b>*RQSDTA</b> Up to the first 80 characters of the request data specified in the request data field are used as the routing data for the job.                                  |
| REQUEST_DATA            | RQSDTA             | VARCHAR(256)<br>Nullable  | The request data that is placed as the last entry in the job's message queue for jobs that use this job description. Can contain the following special value:<br><br><b>*RTGDTA</b> The data specified in the routing data parameter is placed as the last entry in the job's message queue.<br><br>Contains the null value if no request data is placed in the job's message queue.                                                                                                                          |
| LIBRARY_LIST_COUNT      | LIBL_COUNT         | INTEGER                   | The number of libraries in the user portion of the initial library list.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| LIBRARY_LIST            | LIBL               | VARCHAR(2750)<br>Nullable | The initial library list that is used for jobs that use this job description. Only the libraries in the user portion of the library list are included. The list is an array of 11 character entries. Each entry contains a ten character name followed by one blank. Can contain the following special value:<br><br><b>*SYSVAL</b> The jobs using this job description will use the library list specified by the QUSRLIBL system value.<br><br>Contains the null value is there is no initial library list. |
| JOB_SWITCHES            | SWITCHES           | CHAR(8)                   | The initial settings for a group of eight job switches used by jobs that use this job description. These switches can be set or tested in a program and used to control a program's flow. The possible values are '0' (off) and '1' (on).                                                                                                                                                                                                                                                                     |
| TEXT_DESCRIPTION        | TEXT               | VARCHAR(50)<br>Nullable   | The user text, if any, used to briefly describe the job description. Contains the null value is there is no descriptive text.                                                                                                                                                                                                                                                                                                                                                                                 |
| JOB_QUEUE_LIBRARY       | JOBQLIB            | VARCHAR(10)               | The library of the job queue into which batch jobs using this job description are placed.                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 208. JOB\_DESCRIPTION\_INFO view (continued)

| Column Name           | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------|--------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_QUEUE             | JOBQ               | VARCHAR(10)             | The name of the job queue into which batch jobs using this job description are placed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| JOB_QUEUE_PRIORITY    | JOBQ_PRI           | SMALLINT                | The scheduling priority of each job that uses this job description. The highest priority is 1 and the lowest priority is 9.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| HOLD_ON_JOB_QUEUE     | JOBQ_HOLD          | VARCHAR(4)              | Whether jobs using this job description are put on the job queue with a status of held.<br><br><b>*NO</b> Jobs using this job description are not put on the job queue as held.<br><br><b>*YES</b> Jobs using this job description are put on the job queue as held.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| OUTPUT_QUEUE_LIBRARY  | OUTQLIB            | VARCHAR(10)<br>Nullable | The name of the library in which the output queue resides. Contains the null value if OUTPUT_QUEUE is a special value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| OUTPUT_QUEUE          | OUTQ               | VARCHAR(10)             | The name of the default output queue that is used for spooled output produced by jobs that use this job description. Can contain one of the following special values:<br><br><b>*DEV</b> The output queue with the same name as the printer device for this job description is used.<br><br><b>*USRPRF</b> The output queue name for jobs using this job description is obtained from the user profile of the job at the time the job is started.<br><br><b>*WRKSTN</b> The output queue name is obtained from the device description from which this job is started.                                                                                                                                                                                                                                                                                             |
| OUTPUT_QUEUE_PRIORITY | OUTQ_PRI           | SMALLINT                | The output priority for spooled files that are produced by jobs using this job description. The highest priority is 1, and the lowest priority is 9.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| SPOOLED_FILE_ACTION   | SPOOL_ACT          | VARCHAR(7)              | Specifies whether spooled files can be accessed through job interfaces once a job has completed its normal activity.<br><br><b>*DETACH</b> Spooled files are detached from the job when the job completes its activity.<br><br><b>*KEEP</b> When the job completes its activity, as long as at least one spooled file for the job exists in the system auxiliary storage pool (ASP 1) or in a basic user ASP (ASPs 2-32), the spooled files are kept with the job and the status of the job is updated to indicate that the job has completed. If all remaining spooled files for the job are in independent ASPs (ASPs 33-255), the spooled files will be detached from the job and the job will be removed from the system.<br><br><b>*SYSVAL</b> The jobs using this job description will take the spooled file action specified by the QSPLFACN system value. |
| PRINTER_DEVICE        | DEV_NAME           | VARCHAR(10)             | The name of the printer device that is used for all spooled files created by jobs that use this job description. Can contain one of the following special values:<br><br><b>*SYSVAL</b> The value in the system value QPRTDEV at the time the job is started is used as the printer device name.<br><br><b>*USRPRF</b> The printer device name is obtained from the user profile of the job at the time the job is started.<br><br><b>*WRKSTN</b> The printer device name is obtained from the work station where the job was started.                                                                                                                                                                                                                                                                                                                            |

Table 208. JOB\_DESCRIPTION\_INFO view (continued)

| Column Name                      | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRINT_TEXT                       | PRINT_TEXT         | VARCHAR(30)<br>Nullable | <p>The line of text that is printed at the bottom of each page of printed output for jobs using this job description. Can contain the following special value:</p> <p><b>*SYSVAL</b> The value in the system value QPRTTXT is used for jobs using this job description.</p> <p>Contains the null value if there is no text to print.</p>                                                                                                                                                                                                                                                                                                     |
| JOB_MESSAGE_QUEUE_MAXIMIMUM_SIZE | MSGQ_MAX           | SMALLINT<br>Nullable    | <p>The maximum size (in megabytes) of the job message queue. The possible values are 2 to 64.</p> <p>Contains the null value if the maximum size is set by system value QJOBMSGQMX at the time the job is started.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| JOB_MESSAGE_QUEUE_FULL_ACTION    | MSGQ_FULL          | VARCHAR(8)              | <p>The action taken when the job message queue becomes full.</p> <p><b>*NOWRAP</b> When the message queue becomes full, do not wrap. This action will cause the job to end.</p> <p><b>*PRTWRAP</b> When the message queue becomes full, wrap the job queue and print the messages that are being overlaid.</p> <p><b>*SYSVAL</b> The value is specified by the system value QJOBMSGQFL.</p> <p><b>*WRAP</b> When the message queue becomes full, wrap to the beginning and start filling again.</p>                                                                                                                                          |
| SYNTAX_CHECK_SEVERITY            | SYNTAX             | SMALLINT<br>Nullable    | <p>Whether requests placed on the job's message queue are checked for syntax as CL commands, and the message severity that causes a syntax error to end processing of a job. The possible values are:</p> <p><b>0-99</b> Specifies the lowest message severity that causes a running job to end. The request data is checked for syntax as CL commands, and, if a syntax error occurs that is greater than or equal to the error message severity specified here, the running of the job that contains the erroneous command is suppressed.</p> <p>Contains the null value if the request data is not checked for syntax as CL commands.</p> |
| JOB_END_SEVERITY                 | JOB_ENDSEV         | SMALLINT                | <p>The message severity level of escape messages that can cause a batch job to end. The batch job ends when a request in the batch input stream sends an escape message whose severity is equal to or greater than this value to the request processing program. The possible values are from 0 through 99.</p>                                                                                                                                                                                                                                                                                                                              |

Table 208. JOB\_DESCRIPTION\_INFO view (continued)

| Column Name           | System Column Name | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------|--------------------|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOBLOG_OUTPUT         | JOBLOG_OUT         | VARCHAR(10) | <p>How the job log will be produced when the job completes. This does not affect job logs produced when the message queue is full and the job message queue full action specifies *PRTWRAP. Messages in the job message queue are written to a spooled file, from which the job log can be printed, unless the Control Job Log Output (QMHCTLJL) API was used in the job to specify that the messages in the job log are to be written to a database file.</p> <p>The job log output value can be changed at any time until the job log has been produced or removed. To change the job log output value for a job, use the Change Job (QWTCGJJB) API or the Change Job (CHGJOB) command.</p> <p>The job log can be displayed at any time until the job log has been produced or removed. To display the job log, use the Display Job Log (DSPJOBLOG) command.</p> <p>The job log can be removed when the job has completed and the job log has not yet been produced or removed. To remove the job log, use the Remove Pending Job Log (QWTRMVJL) API or the End Job (ENDJOB) command.</p> <p>The possible values are:</p> <p><b>*JOBEND</b> The job log will be produced by the job itself. If the job cannot produce its own job log, the job log will be produced by a job log server. For example, a job does not produce its own job log when the system is processing a Power Down System (PWRDWNSYS) command.</p> <p><b>*JOBLOGSVR</b> The job log will be produced by a job log server. For more information about job log servers, refer to the Start Job Log Server (STRLOGSVR) command.</p> <p><b>*PND</b> The job log will not be produced. The job log remains pending until removed.</p> <p><b>*SYSVAL</b> The value is specified by the QLOGOUTPUT system value.</p> |
| INQUIRY_MESSAGE_REPLY | INQ_REPLY          | VARCHAR(8)  | <p>How inquiry messages are answered for jobs that use this job description.</p> <p><b>*DFT</b> The system uses the default message reply to answer any inquiry messages issued while the job is running. The default reply is either defined in the message description or is the default system reply.</p> <p><b>*RQD</b> The job requires an answer for any inquiry messages that occur while the job is running.</p> <p><b>*SYSRPYL</b> The system reply list is checked to see if there is an entry for an inquiry message issued while the job is running. If a match occurs, the system uses the reply value for that entry. If no entry exists for that message, the system uses an inquiry message.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

Table 208. JOB\_DESCRIPTION\_INFO view (continued)

| Column Name              | System Column Name | Data Type  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|--------------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MESSAGE_LOGGING_LEVEL    | LOG_LEVEL          | SMALLINT   | <p>The type of information logged.</p> <p><b>0</b> No messages are logged.</p> <p><b>1</b> All messages sent to the job's external message queue with a severity greater than or equal to the message logging severity are logged. This includes the indication of job start, job end, and job completion status.</p> <p><b>2</b> The following information is logged:</p> <ul style="list-style-type: none"> <li>Level 1 information.</li> <li>Request messages that result in a high-level message with a severity code greater than or equal to the logging severity cause the request message and all associated messages to be logged.</li> </ul> <p><b>Note:</b> A high-level message is one that is sent to the program message queue of the program that receives the request message. For example, QCMD is an IBM-supplied request processing program that receives request messages.</p> <p><b>3</b> The following information is logged:</p> <ul style="list-style-type: none"> <li>Level 1 and 2 information.</li> <li>All request messages.</li> <li>Commands run by a CL program are logged if it is allowed by the logging of CL programs job attribute and the log attribute of the CL program.</li> </ul> <p><b>4</b> The following information is logged:</p> <ul style="list-style-type: none"> <li>All request messages and all messages with a severity greater than or equal to the message logging severity, including trace messages.</li> <li>Commands run by a CL program are logged if it is allowed by the logging of CL programs job attribute and the log attribute of the CL program.</li> </ul> |
| MESSAGE_LOGGING_SEVERITY | LOG_SEV            | SMALLINT   | <p>The severity level that is used in conjunction with the logging level to determine which error messages are logged in the job log. The possible values are from 0 through 99.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| MESSAGE_LOGGING_TEXT     | LOG_TEXT           | VARCHAR(7) | <p>The level of message text that is written in the job log when a message is logged according to the logging level and logging severity.</p> <p><b>*MSG</b> Only the message text is written to the job log.</p> <p><b>*NOLIST</b> If the job ends normally, no job log is produced. If the job ends abnormally (if the job end code is 20 or higher), a job log is produced. The messages that appear in the job log contain both the message text and the message help.</p> <p><b>*SECLVL</b> Both the message text and the message help (cause and recovery) of the error message are written to the job log.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| LOG_CL_PROGRAM_COMMANDS  | LOG_CL             | VARCHAR(4) | <p>Whether or not commands are logged for CL programs that are run.</p> <p><b>*NO</b> CL programs are not logged.</p> <p><b>*YES</b> CL programs are logged.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 208. JOB\_DESCRIPTION\_INFO view (continued)

| Column Name            | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DEVICE_RECOVERY_ACTION | DEVRECOVER         | VARCHAR(13)             | <p>The action to take when an I/O error occurs for the interactive job's requesting program device.</p> <p><b>*DSCENDRQS</b> Disconnects the job when an I/O error occurs. When the job reconnects, the system sends the End Request (ENDRQS) command to return control to the previous request level.</p> <p><b>*DSCMSG</b> Disconnects the job when an I/O error occurs. When the job reconnects, the system sends a message to the application program indicating the job has reconnected and that the workstation device has recovered.</p> <p><b>*ENDJOB</b> Ends the job when an I/O error occurs. A message is sent to the job's log and to the history log (QHST). This message indicates that the job ended because of a device error.</p> <p><b>*ENDJOBNO LIST</b> Ends the job when an I/O error occurs. There is no job log produced for the job. The system sends a message to the history log (QHST). This message indicates that the job ended because of a device error.</p> <p><b>*MSG</b> Signals the I/O error message to the application and lets the application program perform error recovery.</p> <p><b>*SYSVAL</b> The value in the system value QDEVRCYACN at the time the job is started is used as the device recovery action for this job description.</p> |
| TIME_SLICE_END_POOL    | TIME_SLICE         | VARCHAR(7)              | <p>Whether interactive jobs using this job description should be moved to another main storage pool when they reach time-slice end.</p> <p><b>*BASE</b> The job is moved to the base pool when it reaches time-slice end.</p> <p><b>*NONE</b> The job is not moved when it reaches time-slice end.</p> <p><b>*SYSVAL</b> The system value is used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| ALLOW_MULTIPLE_THREADS | ALWMLTTHD          | VARCHAR(4)              | <p>Whether or not the job is allowed to run with multiple user threads. This attribute does not prevent the operating system from creating system threads in the job. This attribute is not allowed to be changed once a job starts. This attribute applies to autostart jobs, prestart jobs, batch jobs submitted from job schedule entries, and jobs started by using the Submit Job (SBMJOB) and Batch Job (BCHJOB) commands. This attribute is ignored when starting all other types of jobs. This attribute should be set to *YES only in job descriptions that are used exclusively with functions that create multiple user threads.</p> <p><b>*NO</b> The job is not allowed to run with multiple user threads.</p> <p><b>*YES</b> The job is allowed to run with multiple user threads.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| ASPGRP                 | ASPGRP             | VARCHAR(10)<br>Nullable | <p>The name of the ASP group. This is the name of the primary ASP device in an ASP group or the name of an ASP device description. This specifies the initial ASP group setting for jobs using this job description.</p> <p>Contains the null value if jobs using this job description do not have an initial ASP group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Table 208. JOB\_DESCRIPTION\_INFO view (continued)

| Column Name      | System Column Name | Data Type  | Description                                                                                                                                                                                                                                                                                                                  |
|------------------|--------------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DDM_CONVERSATION | DDM_CONV           | VARCHAR(5) | Whether the Distributed Data Management conversations are kept or dropped when they are not being used. The possible values are:<br><br><b>*DROP</b> The system ends a DDM-allocated conversation when there are no users.<br><br><b>*KEEP</b> The system keeps DDM conversation connections active when there are no users. |

## Examples

- Review information about the job queues associated with each job description.

```
SELECT JOB_DESCRIPTION_LIBRARY, JOB_DESCRIPTION,
 JOB_QUEUE_LIBRARY, JOB_QUEUE, JOB_QUEUE_PRIORITY
FROM QSYS2.JOB_DESCRIPTION_INFO;
```

- Find the job descriptions that have APPLIB1 in their library list

```
SELECT JOB_DESCRIPTION_LIBRARY, JOB_DESCRIPTION, LIBRARY_LIST
FROM QSYS2.JOB_DESCRIPTION_INFO
WHERE LIBRARY_LIST LIKE '%APPLIB1%';
```

- Examine the library lists for every job description.

Since the library list column returns a character string containing a list of libraries, to see the individual library names it needs to be broken apart. To do this, you can create a table function that takes the library list string and returns a list of library names.

```
CREATE OR REPLACE FUNCTION QGPL.GET_LIB_NAMES(JOBD_LIBL VARCHAR(2750),
 JOBD_LIBL_CNT INT)
RETURNS TABLE(LIBL_POSITION INT, LIBRARY_NAME VARCHAR(10))
BEGIN
 DECLARE IN_POS INT;
 DECLARE LIB_CNT INT;
 SET IN_POS = 1;
 SET LIB_CNT = 1;
 WHILE LIB_CNT <= JOBD_LIBL_CNT
 DO
 PIPE (LIB_CNT, RTRIM((SUBSTR(JOBD_LIBL, IN_POS, 10))));
 SET IN_POS = IN_POS + 11;
 SET LIB_CNT = LIB_CNT + 1;
 END WHILE;
 RETURN;
END;
```

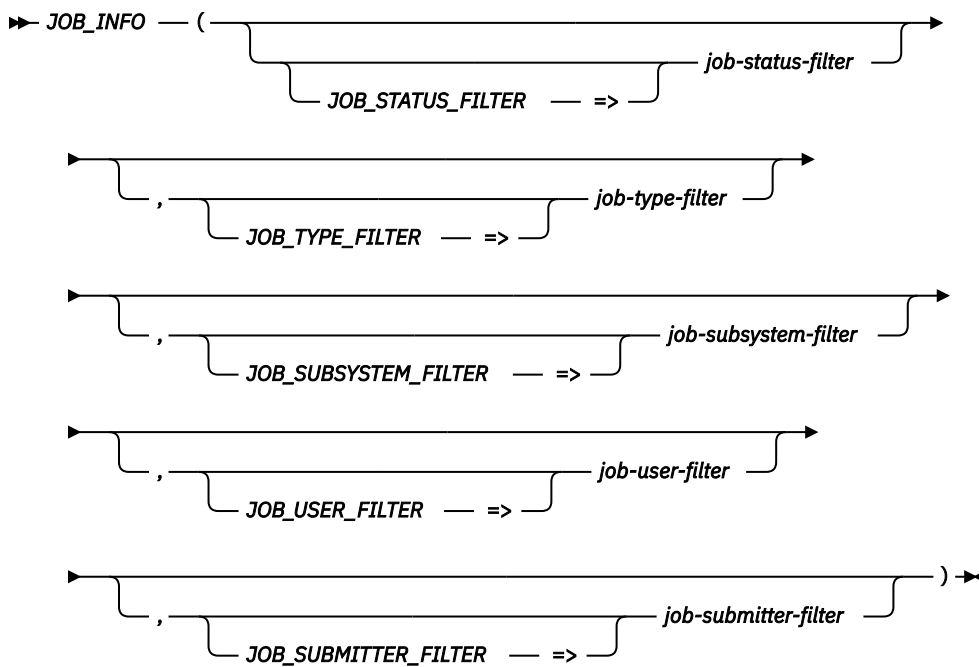
Now this function can be used to return the list of library names.

```
SELECT JOB_DESCRIPTION, JOB_DESCRIPTION_LIBRARY, LIBL_POSITION, LIBRARY_NAME
FROM QSYS2.JOB_DESCRIPTION_INFO,
TABLE (QGPL.GET_LIB_NAMES(LIBRARY_LIST, LIBRARY_LIST_COUNT)) X;
```

## JOB\_INFO table function

The JOB\_INFO table function returns one row for each job meeting the selection criteria. It returns information similar to what is returned by the Work with User Jobs (WRKUSRJOB), Work with Subsystem Jobs (WRKSBSJOB), and Work with Submitted Jobs (WRKSBMJOB) CL commands and the List Job (QUSLJOB) API.

**Authorization:** None required to see information for jobs where the caller's user profile is the same as the job user identity of the job for which the information is being returned. Otherwise, the caller must have \*JOBCTL special authority.



The schema is QSYS2.

***job-status-filter*** A character or graphic string expression that specifies the value to use as the job status filtering criteria. The string must be one of the following special values:

- \*ALL** Jobs of any status including jobs on job queues, active jobs, and jobs on an output queue.
- \*ACTIVE** Jobs that are active. You can use the `QSYS2.ACTIVE_JOB_INFO` table function to get additional details for these jobs.
- \*JOBQ** Jobs that are not active because they are waiting on a job queue.
- \*OUTQ** Jobs that have completed execution and have output on an output queue.

If this parameter is not provided, a value of **\*ALL** is used.

***job-type-filter*** A character or graphic string expression that specifies the value to use as the job type filtering criteria. The string must be one of the following special values:

- \*ALL** All types of user jobs, including interactive jobs and batch jobs.
- \*BATCH** Only batch user jobs, including prestart jobs, batch immediate jobs, and autostart jobs.
- \*INTERACT** Only interactive user jobs.

If this parameter is not provided, a value of **\*ALL** is used.

***job-subsystem-filter*** A character or graphic string expression that specifies the subsystem value to use as the job subsystem filtering criteria. The string can be a subsystem name or the following value:

- \*ALL** All jobs in all subsystems, including jobs that are on job queues and on output queues.

If a subsystem name is provided, only active jobs are found.

If this parameter is not provided, a value of **\*ALL** is used.



**job-user-filter** The USER special register or a character or graphic string expression that specifies the user profile name to use as the job user filtering criteria.

The string can be a user name or one of the following special values:

**\*ALL** All jobs being processed under all user names.

**\*USER** The user part of the qualified job name.

The USER special register is specified as a non-string value. It represents the current user of the job invoking the function.

If this parameter is not provided, the value of the USER special register is used.

**job-submitter-filter**

A character or graphic string expression that specifies the type of submitted jobs to return. The string must be one of the following values:

**\*ALL** All submitted jobs.

**\*JOB** Jobs that were submitted from the same job that is invoking this function.

**\*USER** Jobs that were submitted from a job having the same user profile as the job invoking this function.

**\*WRKSTN** Jobs that were submitted from the same work station as the job invoking this function.

If this parameter is not provided, a value of **\*ALL** is used.

**Restrictions:**

- Only one of these filters can have a value other than **\*ALL**: *job-subsystem-filter* and *job-submitter-filter*.
- If a value other than **\*ALL** is specified for *job-submitter-filter*, you must specify **\*ALL** for *job-user-filter*.

**Notes:**

- Jobs submitted with **\*NO** specified for the Allow display by WRKSBMJOB (DSPSBMJOB) parameter of the SBMJOB command are not returned by this table function.

For each of the WRKSBMJOB, WRKSBSJOB, and WRKUSRJOB CL commands shown below, the corresponding invocation of JOB\_INFO will return the same list of jobs. Note that to get exact equivalence, predicates must be added to some queries to achieve equivalent results:

- For equivalence with WRKUSRJOB, a query must always include the predicate WHERE JOB\_TYPE NOT IN ('SBS','SYS','RDR','WTR')
- For equivalence with WRKSBSJOB SBS(\*OUTQ) or WRKSBSJOB SBS(\*ALL), a query must always include the predicate WHERE JOB\_TYPE NOT IN ('SBS','SYS')

| CL Command | CL Parameters    | JOB_INFO invocation                                                                                            |
|------------|------------------|----------------------------------------------------------------------------------------------------------------|
| WRKSBMJOB  | SBMFROM(*USER)   | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_SUBMITTER_FILTER => '*USER',<br>JOB_USER_FILTER => '*ALL'<br>)) X   |
|            | SBMFROM(*WRKSTN) | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_SUBMITTER_FILTER => '*WRKSTN',<br>JOB_USER_FILTER => '*ALL'<br>)) X |
|            | SBMFROM(*JOB)    | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_SUBMITTER_FILTER => '*JOB',<br>JOB_USER_FILTER => '*ALL'<br>)) X    |

Table 209. Equivalent CL command and JOB\_INFO invocations (continued)

| CL Command | CL Parameters                                     | JOB_INFO invocation                                                                                                                                                                             |
|------------|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WRKSBSJOB  | SBS(QBATCH)<br>USER(*ALL)                         | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_SUBSYSTEM_FILTER => 'QBATCH',<br>JOB_USER_FILTER => '*ALL'<br>)) X                                                                                   |
|            | SBS(*JOBQ)<br>USER(*ALL)                          | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_STATUS_FILTER => '*JOBQ',<br>JOB_USER_FILTER => '*ALL'<br>)) X                                                                                       |
|            | SBS(*OUTQ)<br>USER(JOEUSER)                       | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_STATUS_FILTER => '*OUTQ',<br>JOB_USER_FILTER => 'JOEUSER'<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS')                                            |
|            | SBS(*ALL)<br>USER(*ALL)                           | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_STATUS_FILTER => '*ALL',<br>JOB_USER_FILTER => '*ALL'<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS')                                                |
| WRKUSRJOB  | USER(*)<br>STATUS(*ALL)<br>JOBTYPE(*ALL)          | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS', 'RDR', 'WTR')                                                                                               |
|            | USER(*)<br>STATUS(*ALL)<br>JOBTYPE(*INTERACT)     | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_TYPE_FILTER => '*INTERACT'<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS', 'RDR', 'WTR')                                                             |
|            | USER(JOEUSER)<br>STATUS(*ACTIVE)<br>JOBTYPE(*ALL) | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_USER_FILTER => 'JOEUSER',<br>JOB_STATUS_FILTER => '*ACTIVE'<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS', 'RDR', 'WTR')                            |
|            | USER(*)<br>STATUS(*OUTQ)<br>JOBTYPE(*ALL)         | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_STATUS_FILTER => '*OUTQ'<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS', 'RDR', 'WTR')                                                               |
|            | USER(*ALL)<br>STATUS(*JOBQ)<br>JOBTYPE(*BATCH)    | SELECT * FROM TABLE(QSYS2.JOB_INFO(<br>JOB_USER_FILTER => '*ALL',<br>JOB_STATUS_FILTER => '*JOBQ',<br>JOB_TYPE_FILTER => '*BATCH'<br>)) X<br>WHERE JOB_TYPE NOT IN ('SBS', 'SYS', 'RDR', 'WTR') |

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 210. JOB\_INFO table function

| Column Name    | Data Type   | Description             |
|----------------|-------------|-------------------------|
| JOB_NAME       | VARCHAR(28) | The qualified job name. |
| JOB_NAME_SHORT | VARCHAR(10) | The name of the job.    |

Table 210. JOB\_INFO table function (continued)

| Column Name     | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_USER        | VARCHAR(10) | The user profile that started the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| JOB_NUMBER      | VARCHAR(6)  | The job number of the job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| JOB_INFORMATION | VARCHAR(12) | <p>Indicates whether information is available for the job.</p> <p><b>NO</b> The information is not available because the job was not accessible.</p> <p><b>YES</b> The information is available.</p> <p>When this value is NO, all columns other than JOB_NAME return the null value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| JOB_STATUS      | VARCHAR(6)  | <p>The status of the job.</p> <p><b>ACTIVE</b> The job has started, and it can use system resources (processing unit, main storage, and so on). This does not guarantee that the job is currently running, however. For example, an active job may be in one of the following states where it is not in a position to use system resources:</p> <ul style="list-style-type: none"> <li>• The Hold Job (HLDJOB) command holds the job; the Release Job (RLSJOB) command allows the job to run again.</li> <li>• The Transfer Group Job (TFRGRPJOB) or Transfer Secondary Job (TFRSECJOB) command suspends the job. When control returns to the job, the job can run again.</li> <li>• The job is disconnected using the Disconnect Job (DSCJOB) command. When the interactive user signs back on, thereby connecting back into the job, the job can run again.</li> <li>• The job is waiting for any reason. For example, when the job receives the reply for an inquiry message, the job can start running again.</li> </ul> <p><b>JOBQ</b> The job is currently on a job queue. The job possibly was previously active and was placed back on the job queue because of the Transfer Job (TFRJOB) or Transfer Batch Job (TFRBCHJOB) command, or the job was never active because it was just submitted.</p> <p><b>OUTQ</b> The job has completed running and has spooled output that has not yet printed or the job's job log has not yet been written.</p> |
| JOB_TYPE        | VARCHAR(3)  | <p>The type of job.</p> <p><b>ASJ</b> Autostart</p> <p><b>BCH</b> Batch</p> <p><b>BCI</b> Batch Immediate</p> <p><b>EVK</b> Started by a procedure start request</p> <p><b>INT</b> Interactive</p> <p><b>M36</b> Advanced 36 server job</p> <p><b>MRT</b> Multiple requester terminal</p> <p><b>PDJ</b> Print driver job</p> <p><b>PJ</b> Prestart job</p> <p><b>RDR</b> Spool reader</p> <p><b>SBS</b> Subsystem monitor</p> <p><b>SYS</b> System</p> <p><b>WTR</b> Spool writer</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 210. JOB\_INFO table function (continued)

| Column Name                                                               | Data Type   | Description                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_TYPE_ENHANCED                                                         | VARCHAR(28) | The combined job type and job subtype values.                                                                                                                                                                                                                                                                 |
|                                                                           |             | <b>ALTERNATE_SPOOL_USER</b> Batch - alternate spool user                                                                                                                                                                                                                                                      |
|                                                                           |             | <b>AUTOSTART</b> Autostart job                                                                                                                                                                                                                                                                                |
|                                                                           |             | <b>BATCH</b> Batch job                                                                                                                                                                                                                                                                                        |
|                                                                           |             | <b>BATCH_IMMEDIATE</b> Batch immediate job                                                                                                                                                                                                                                                                    |
|                                                                           |             | <b>BATCH_MRT</b> Batch - System/36 multiple requester terminal (MRT) job                                                                                                                                                                                                                                      |
|                                                                           |             | <b>COMM_PROCEDURE_START_REQUEST</b> Communications job - procedure start request job                                                                                                                                                                                                                          |
|                                                                           |             | <b>INTERACTIVE</b> Interactive job                                                                                                                                                                                                                                                                            |
|                                                                           |             | <b>INTERACTIVE_GROUP</b> Interactive job - Part of group                                                                                                                                                                                                                                                      |
|                                                                           |             | <b>INTERACTIVE_SYSREQ</b> Interactive job - Part of system request pair                                                                                                                                                                                                                                       |
|                                                                           |             | <b>INTERACTIVE_SYSREQ_AND_GROUP</b> Interactive job - Part of system request pair and part of a group                                                                                                                                                                                                         |
|                                                                           |             | <b>PRESTART</b> Prestart job                                                                                                                                                                                                                                                                                  |
|                                                                           |             | <b>PRESTART_BATCH</b> Prestart batch job                                                                                                                                                                                                                                                                      |
|                                                                           |             | <b>PRESTART_COMM</b> Prestart communications job                                                                                                                                                                                                                                                              |
|                                                                           |             | <b>READER</b> Reader job                                                                                                                                                                                                                                                                                      |
|                                                                           |             | <b>SUBSYSTEM</b> Subsystem job                                                                                                                                                                                                                                                                                |
| <b>SYSTEM</b> System job (all system jobs including SCPF)                 |             |                                                                                                                                                                                                                                                                                                               |
| <b>WRITER</b> Writer job (including both spool writers and print drivers) |             |                                                                                                                                                                                                                                                                                                               |
| JOB_SUBSYSTEM                                                             | VARCHAR(10) | The name of the subsystem for the job.<br>Contains the null value if JOB_TYPE is SYS, JOB_STATUS is JOBQ or OUTQ, or if the job has no subsystem.                                                                                                                                                             |
| JOB_DATE                                                                  | VARCHAR(10) | The date that is assigned to the job, in *ISO format. The job date remains the same for the duration of the job unless it is changed by the user. Can also contain the following special value:<br><br><b>SYSVAL</b> This job will use the system date.<br><br>Contains the null value if JOB_STATUS is OUTQ. |
| JOB_DESCRIPTION_LIBRARY                                                   | VARCHAR(10) | The name of the library containing the job description.<br>Contains the null value if JOB_DESCRIPTION is null.                                                                                                                                                                                                |
| JOB_DESCRIPTION                                                           | VARCHAR(10) | The name of the job description used for this job.<br>Contains the null value if the job has no job description.                                                                                                                                                                                              |
| JOB_ACCOUNTING_CODE                                                       | VARCHAR(15) | An identifier assigned to the job by the system to collect resource use information for the job when job accounting is active.<br>Contains the null value if the job has no accounting code.                                                                                                                  |
| SUBMITTER_JOB_NAME                                                        | VARCHAR(28) | The qualified job name of the submitter's job.<br>Contains the null value if the job has no submitter.                                                                                                                                                                                                        |
| SUBMITTER_MESSAGE_QUEUE_LIBRARY                                           | VARCHAR(10) | The name of the library containing the message queue.<br>Contains the null value if the job has no submitter.                                                                                                                                                                                                 |

Table 210. JOB\_INFO table function (continued)

| Column Name             | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBMITTER_MESSAGE_QUEUE | VARCHAR(10)  | The name of the message queue where the system sends a completion message when a batch job ends.<br>Contains the null value if the job has no submitter.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| SERVER_TYPE             | VARCHAR(30)  | The type of server represented by the job. See <a href="#">Server table</a> for a list of server type values.<br>Contains the null value if the job is not part of a server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| JOB_ENTERED_SYSTEM_TIME | TIMESTAMP(0) | The timestamp for when the job was placed on the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| JOB_SCHEDULED_TIME      | TIMESTAMP(0) | The timestamp for when the job is scheduled to become active.<br>Contains the null value if this is not a scheduled job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| JOB_ACTIVE_TIME         | TIMESTAMP(0) | The time the job began to run on the system.<br>Contains the null value if the job did not become active.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| JOB_END_TIME            | TIMESTAMP(0) | The timestamp for when the job completed running on the system.<br>Contains the null value if the job has not ended.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| JOB_END_SEVERITY        | SMALLINT     | The message severity level of escape messages that can cause a batch job to end. The batch job ends when a request in the batch input stream sends an escape message, whose severity is equal to or greater than this value, to the request processing program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| COMPLETION_STATUS       | VARCHAR(8)   | The completion status of the job.<br><br><b>ABNORMAL</b> The job completed abnormally.<br><b>NORMAL</b> The job completed normally.<br>Contains the null value if this the job has not completed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| JOB_END_REASON          | VARCHAR(60)  | The most recent action that caused the job to end. Contains one of the following values: <ul style="list-style-type: none"> <li>• JOB ENDED DUE TO A DEVICE ERROR</li> <li>• JOB ENDED DUE TO A SIGNAL</li> <li>• JOB ENDED DUE TO AN UNHANDLED ERROR</li> <li>• JOB ENDED DUE TO THE CPU LIMIT BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE DISCONNECT TIME INTERVAL BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE INACTIVITY TIME INTERVAL BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE MESSAGE SEVERITY LEVEL BEING EXCEEDED</li> <li>• JOB ENDED DUE TO THE STORAGE LIMIT BEING EXCEEDED</li> <li>• JOB ENDED WHILE IT WAS STILL ON A JOB QUEUE</li> <li>• JOB ENDING ABNORMALLY</li> <li>• JOB ENDING IMMEDIATELY</li> <li>• JOB ENDING IN NORMAL MANNER</li> <li>• JOB ENDING NORMALLY AFTER A CONTROLLED END WAS REQUESTED</li> <li>• SYSTEM ENDED ABNORMALLY</li> </ul> Contains the null value if job is not currently ending. |
| JOB_QUEUE_LIBRARY       | VARCHAR(10)  | The name of the library containing the job queue.<br>Contains the null value if JOB_STATUS is OUTQ or if job is not on a job queue and the job is not a batch job that was started from a job queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| JOB_QUEUE_NAME          | VARCHAR(10)  | The name of the job queue that the job is currently on, or that the job was on if it is currently active.<br>Contains the null value if JOB_STATUS is OUTQ or if job is not on a job queue and the job is not a batch job that was started from a job queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 210. JOB\_INFO table function (continued)

| Column Name                       | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_QUEUE_STATUS                  | VARCHAR(9)   | <p>The status of this job on the job queue.</p> <p><b>HELD</b> This job is being held on the job queue.</p> <p><b>RELEASED</b> This job is ready to be selected.</p> <p><b>SCHEDULED</b> This job will run as scheduled.</p> <p>Contains the null value if the job is not on a job queue.</p>                                                                                                                                                                                                                                                                                                                                                                                                                               |
| JOB_QUEUE_PRIORITY                | SMALLINT     | <p>The scheduling priority of the job compared to other jobs on the same job queue. The highest priority is 0 and the lowest is 9.</p> <p>Contains the null value if JOB_STATUS is not JOBQ.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| JOB_QUEUE_TIME                    | TIMESTAMP(0) | <p>The timestamp when the job was put on the job queue.</p> <p>Contains the null value if this the job is not on a job queue.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| JOB_MESSAGE_QUEUE_MAXIMUM_SIZE    | SMALLINT     | <p>The maximum size, in megabytes, that the job message queue can become. The range is 2 to 64.</p> <p>Contains the null value if JOB_QUEUE_NAME is null.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| JOB_MESSAGE_QUEUE_FULL_ACTION     | VARCHAR(8)   | <p>The action to take when the message queue is full.</p> <p><b>*NOWRAP</b> When the job message queue is full, do not wrap. This action causes the job to end.</p> <p><b>*PRTWRAP</b> When the job message queue is full, wrap the message queue and print the messages that are being overlaid because of the wrapping.</p> <p><b>*WRAP</b> When the job message queue is full, wrap to the beginning and start filling again.</p> <p>Contains the null value if JOB_QUEUE_NAME is null.</p>                                                                                                                                                                                                                              |
| ALLOW_MULTIPLE_THREADS            | VARCHAR(3)   | <p>Indicates whether this job allows multiple user threads. This attribute does not prevent the operating system from creating system threads in the job.</p> <p><b>NO</b> This job does not allow multiple user threads.</p> <p><b>YES</b> This job allows multiple user threads.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| PEAK_TEMPORARY_STORAGE            | INTEGER      | <p>The maximum amount of auxiliary storage, in megabytes, that the job has used.</p> <p>Contains the null value if JOB_STATUS is OUTQ or for a job on a job queue if a value has not been set for the job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| DEFAULT_WAIT                      | INTEGER      | <p>The default maximum time, in seconds, that a thread in the job waits for a system instruction, such as a LOCK machine interface (MI) instruction, to acquire a resource.</p> <p>Contains the null value if there is no maximum, if JOB_STATUS is OUTQ, or for a job on a job queue if a value has not been set for the job.</p>                                                                                                                                                                                                                                                                                                                                                                                          |
| MAXIMUM_PROCESSING_TIME_ALLOWED   | INTEGER      | <p>The maximum processing unit time, in milliseconds, that the job can use. If the job consists of multiple routing steps, this is the maximum processing unit time that the current routing step can use. If the maximum time is exceeded, the job is held.</p> <p>Contains the null value if JOB_STATUS is OUTQ or if no maximum amount of processing unit time has been defined.</p>                                                                                                                                                                                                                                                                                                                                     |
| MAXIMUM_TEMPORARY_STORAGE_ALLOWED | INTEGER      | <p>The maximum amount of auxiliary storage, in megabytes, that the job can use. If the job consists of multiple routing steps, this is the maximum temporary storage that the routing step can use. This temporary storage is used for storage required by the program itself and by implicitly created internal system objects used to support the routing step. (It does not include storage for objects in the QTEMP library.) If the maximum temporary storage is exceeded, the job is held. This does not apply to the use of permanent storage, which is controlled through the user profile.</p> <p>Contains the null value if JOB_STATUS is OUTQ or if no maximum amount of temporary storage has been defined.</p> |

Table 210. JOB\_INFO table function (continued)

| Column Name                  | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TIME_SLICE                   | INTEGER     | <p>The maximum amount of processor time, in milliseconds, given to each thread in this job before other threads in this job and in other jobs are given the opportunity to run. The time slice establishes the amount of time needed by a thread in this job to accomplish a meaningful amount of processing. At the end of the time slice, the thread might be put in an inactive state so that other threads can become active in the storage pool. Values range from 8 through 9999999.</p> <p>Contains the null value if JOB_STATUS is OUTQ or for a job on a job queue if a value has not been set for the job.</p>                                                                                                     |
| JOB_SWITCHES                 | CHAR(8)     | <p>The current setting of the job switches used by this job.</p> <p>Contains the null value no job switches are set.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| ROUTING_DATA                 | VARCHAR(80) | <p>The routing data that is used to determine the routing entry that identifies the program to start for the routing step.</p> <p>Contains the null value if there is no routing data for this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CCSID                        | INTEGER     | <p>The coded character set identifier (CCSID) used for this job.</p> <p>Contains the null value if no CCSID is defined for this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| CHARACTER_IDENTIFIER_CONTROL | VARCHAR(9)  | <p>The character identifier control for the job. This attribute controls the type of CCSID conversion that occurs for display files, printer files, and panel groups. The *CHRIDCTL special value must be specified on the CHRID command parameter on the create, change, or override command for display files, printer files, and panel groups before this attribute will be used.</p> <p><b>*DEV D</b> The *DEV D special value performs the same function as on the CHRID command parameter for display files, printer files, and panel groups.</p> <p><b>*JOBCCSID</b> The *JOBCCSID special value performs the same function as on the CHRID command parameter for display files, printer files, and panel groups.</p> |
| SORT_SEQUENCE_LIBRARY        | VARCHAR(10) | <p>The name or the library that contains the sort sequence table.</p> <p>Contains the null value if no sort sequence table is defined for this job or if SORT_SEQUENCE_NAME is a special value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| SORT_SEQUENCE_NAME           | VARCHAR(10) | <p>The name of the sort sequence table associated with this job.</p> <p>Contains the null value if no sort sequence table is defined for this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| LANGUAGE_ID                  | CHAR(3)     | <p>The language identifier associated with this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| COUNTRY_ID                   | CHAR(2)     | <p>The country or region identifier associated with this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| DATE_FORMAT                  | CHAR(4)     | <p>The date format used for this job.</p> <p><b>*DMY</b> Day, month, year format.</p> <p><b>*JUL</b> Julian format (year and day).</p> <p><b>*MDY</b> Month, day, year format.</p> <p><b>*YMD</b> Year, month, day format.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| DATE_SEPARATOR               | CHAR(1)     | <p>The date separator used for this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| TIME_SEPARATOR               | CHAR(1)     | <p>The time separator used for this job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

Table 210. JOB\_INFO table function (continued)

| Column Name                | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DECIMAL_FORMAT             | VARCHAR(6)  | <p>The decimal format used for this job.</p> <p><b>*BLANK</b> Uses a period for a decimal point, a comma for a 3-digit grouping character, and zero-suppress to the left of the decimal point.</p> <p><b>J</b> Uses a comma for a decimal point and a period for a 3-digit grouping character. The zero-suppression character is in the second position (rather than the first) to the left of the decimal notation. Balances with zero values to the left of the comma are written with one leading zero (0,04). The J entry also overrides any edit codes that might suppress the leading zero.</p> <p><b>I</b> Uses a comma for a decimal point, a period for a 3-digit grouping character, and zero-suppress to the left of the decimal point.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| TIME_ZONE_DESCRIPTION_NAME | VARCHAR(10) | The name of the time zone description that is used to calculate local job time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| MESSAGE_LOGGING_LEVEL      | SMALLINT    | <p>The type of information that is logged.</p> <p><b>0</b> No messages are logged.</p> <p><b>1</b> All messages sent to the job's external message queue with a severity greater than or equal to the message logging severity are logged. This includes the indication of job start, job end and job completion status.</p> <p><b>2</b> The following information is logged:</p> <ul style="list-style-type: none"> <li>• Level 1 information</li> <li>• Request messages that result in a high-level message with a severity code greater than or equal to the logging severity cause the request message and all associated messages to be logged.</li> </ul> <p>Note: A high-level message is one that is sent to the program message queue of the program that receives the request message. For example, QCMD is an IBM-supplied request processing program that receives request messages.</p> <p><b>3</b> The following information is logged:</p> <ul style="list-style-type: none"> <li>• Level 1 and 2 information</li> <li>• All request messages</li> <li>• Commands run by a CL program are logged if it is allowed by the logging of CL programs job attribute and the log attribute of the CL program.</li> </ul> <p><b>4</b> The following information is logged:</p> <ul style="list-style-type: none"> <li>• All request messages and all messages with a severity greater than or equal to the message logging severity, including trace messages.</li> <li>• Commands run by a CL program are logged if it is allowed by the logging of CL programs job attribute and the log attribute of the CL program.</li> </ul> |
| MESSAGE_LOGGING_SEVERITY   | SMALLINT    | The severity level that is used in conjunction with the logging level to determine which error messages are logged in the job log. The values range from 0 through 99.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| MESSAGE_LOGGING_TEXT       | VARCHAR(7)  | <p>The level of message text that is written in the job log when a message is logged according to the logging level and logging severity.</p> <p><b>*MSG</b> Only the message text is written to the job log.</p> <p><b>*NOLIST</b> If the job ends normally, no job log is produced. If the job ends abnormally (the job end code is 20 or higher), a job log is produced. The messages that appear in the job log contain both the message text and the message help.</p> <p><b>*SECLVL</b> Both the message text and the message help (cause and recovery) of the error message are written to the job log.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |



Table 210. JOB\_INFO table function (continued)

| Column Name             | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LOG_CL_PROGRAM_COMMANDS | VARCHAR(4)  | Specifies whether or not commands are logged for CL programs that are run.<br><br><b>*NO</b> Commands are not logged.<br><b>*YES</b> Commands are logged.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| STATUS_MESSAGE          | VARCHAR(7)  | Specifies whether status messages are displayed for this job.<br><br><b>*NONE</b> This job does not display status messages.<br><b>*NORMAL</b> This job displays status messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| INQUIRY_MESSAGE_REPLY   | VARCHAR(8)  | Specifies how the job answers inquiry messages.<br><br><b>*RQD</b> The job requires an answer for any inquiry messages that occur while this job is running.<br><b>*DFT</b> The system uses the default message reply to answer any inquiry messages issued while this job is running. The default reply is either defined in the message description or is the default system reply.<br><b>*SYSRPLY</b> The system reply list is checked to see if there is an entry for an inquiry message issued while this job is running. If a match occurs, the system uses the reply value for that entry. If no entry exists for that message, the system uses an inquiry message. |
| BREAK_MESSAGE           | VARCHAR(7)  | Specifies how this job handles break messages.<br><br><b>*HOLD</b> The message queue holds break messages until a user or program requests them. The work station user uses the Display Message (DSPMSG) command to display the messages; a program must issue a Receive Message (RCVMSG) command to receive a message and handle it.<br><b>*NORMAL</b> The message queue status determines break message handling.<br><b>*NOTIFY</b> The system notifies the job's message queue when a message arrives. For interactive jobs, the audible alarm sounds if there is one, and the message-waiting light comes on.                                                          |
| JOB_LOG_OUTPUT          | VARCHAR(10) | Specifies how the job log will be produced when the job completes.<br><br><b>*JOBEND</b> The job log will be produced by the job itself. If the job cannot produce its own job log, the job log will be produced by a job log server. For example, a job does not produce its own job log when the system is processing a Power Down System (PWRDWN SYS) command.<br><b>*JOBLOGSVR</b> The job log will be produced by a job log server. For more information about job log servers, refer to the Start Job Log Server (STRLOGSVR) command.<br><b>*PND</b> The job log will not be produced. The job log remains pending until removed.                                    |
| JOB_LOG_PENDING         | VARCHAR(3)  | Specifies whether there is a job log that has not yet been written. The writing of the job log may become pending based on the value of the job log output job attribute when the job completes its activity.<br><br><b>NO</b> Job log is not pending.<br><b>YES</b> Job log is pending.                                                                                                                                                                                                                                                                                                                                                                                   |
| OUTPUT_QUEUE_PRIORITY   | SMALLINT    | The output priority for spooled output files that this job produces. The highest priority is 0, and the lowest is 9.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| OUTPUT_QUEUE_LIBRARY    | VARCHAR(10) | The name of the library that contains the default output queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| OUTPUT_QUEUE_NAME       | VARCHAR(10) | The name of the default output queue that is used for spooled output produced by this job and the name of the library that contains the output queue. The default output queue is only for spooled printer files that specify *JOB for the output queue.                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 210. JOB\_INFO table function (continued)

| Column Name            | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SPOOLED_FILE_ACTION    | VARCHAR(7)  | <p>Specifies whether spooled files are accessed through job interfaces after the job has completed is normal activity.</p> <p><b>*DETACH</b> The spooled files are detached from the job when the job completes its activity.</p> <p><b>*KEEP</b> When the job completes its activity, as long as at least one spooled file for the job exists in the system auxiliary storage pool (ASP 1) or in a basic user ASP (ASPs 2-32), the spooled files are kept with the job and the status of the job is updated to indicate that the job has completed. If all remaining spooled files for the job are in independent ASPs (ASPs 33-255), the spooled files will be detached from the job and the job will be removed from the system.</p>                                                                                                                                                                                                                                                                                                                                                                                                                       |
| PRINTER_DEVICE_NAME    | VARCHAR(10) | The printer device used for printing output from this job.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| PRINT_KEY_FORMAT       | VARCHAR(7)  | <p>Specifies whether border and header information is provided when the Print key is pressed.</p> <p><b>*NONE</b> The border and header information is not included with output from the Print key.</p> <p><b>*PRTBDR</b> The border information is included with output from the Print key.</p> <p><b>*PRTHDR</b> The header information is included with output from the Print key.</p> <p><b>*PRTALL</b> The border and header information is included with output from the Print key.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| PRINT_TEXT             | VARCHAR(30) | <p>The line of text that is printed at the bottom of each page of printed output for the job.</p> <p>Contains the null value if there is no text defined to print at the bottom of each page.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| DEVICE_NAME            | VARCHAR(10) | <p>The name of the device as identified to the system. For an interactive job it is the device where the job started.</p> <p>Contains the null value if this is not an interactive job.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| DEVICE_RECOVERY_ACTION | VARCHAR(13) | <p>The action taken for interactive jobs when an I/O error occurs for the job's requesting program device.</p> <p><b>*DSCENDRQS</b> Disconnects the job when an I/O error occurs. When the job reconnects, the system sends the End Request (ENDRQS) command to return control to the previous request level.</p> <p><b>*DSCMSG</b> Disconnects the job when an I/O error occurs. When the job reconnects, the system sends an error message to the application program, indicating the job has reconnected and that the work station device has recovered.</p> <p><b>*ENDJOB</b> Ends the job when an I/O error occurs. A message is sent to the job's log and to the history log (QHST) indicating the job ended because of a device error.</p> <p><b>*ENDJOBNO LIST</b> Ends the job when an I/O error occurs. There is no job log produced for the job. The system sends a message to the QHST log indicating the job ended because of a device error.</p> <p><b>*MSG</b> Signals the I/O error message to the application and lets the application program perform error recovery.</p> <p>Contains the null value if this is not an interactive job.</p> |

Table 210. JOB\_INFO table function (continued)

| Column Name      | Data Type  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DDM_CONVERSATION | VARCHAR(5) | <p>Specifies whether connections using distributed data management (DDM) protocols remain active when they are not being used. The connections include APPC conversations, active TCP/IP connections or Opti-Connect connections.</p> <p><b>*DROP</b> The system ends a DDM connection when there are no users. Examples include when an application closes a DDM file, or when a DRDA application runs a SQL DISCONNECT statement.</p> <p><b>*KEEP</b> The system keeps DDM connections active when there are no users, except for the following:</p> <ul style="list-style-type: none"> <li>• The routing step ends on the source system. The routing step ends when the job ends or when the job is rerouted to another routing step.</li> <li>• The Reclaim Distributed Data Management Conversation (RCLDDMCNV) command or the Reclaim Resources (RCLRSC) command runs.</li> <li>• A communications failure or an internal failure occurs.</li> <li>• A DRDA connection to an application server not running on the system ends.</li> </ul> |
| MODE_NAME        | VARCHAR(8) | <p>The mode name of the advanced program-to-program communications device that started the job. The following special value may be returned:</p> <p><b>*BLANK</b> The mode name is a blank name.</p> <p>Contains the null value if the job is not using advanced program-to-program communications (APPC).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| UNIT_OF_WORK_ID  | CHAR(24)   | <p>The unit of work ID is used to track jobs across multiple systems.</p> <p>Contains the null value if the job is not associated with a source or target system using advanced program-to-program communications (APPC).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| INTERNAL_JOB_ID  | BINARY(16) | The internal job identifier.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## Examples

- Find all interactive jobs.

```
SELECT * FROM TABLE(QSYS2.JOB_INFO(JOB_TYPE_FILTER => '*INTERACT')) X;
```

- Find jobs submitted by SCOTTF that have not been started.


```
SELECT * FROM TABLE(QSYS2.JOB_INFO(JOB_USER_FILTER => 'SCOTTF',
JOB_STATUS_FILTER => '*JOBQ')) X;
```

## JOB\_LOCK\_INFO table function

The JOB\_LOCK\_INFO table function returns a list of objects that have been locked or have lower level locks acquired by the specified job. If the job is not active, no rows are returned.

This information is similar to what is returned by the Retrieve Job Locks (QWCRJBLK) API.

**Authorization:** None required to see information for jobs where the caller's user profile is the same as the job user identity of the job for which the information is being returned. Otherwise, the caller must have \*JOBCTL special authority.



The schema is QSYS2.

**job-name** A character or graphic string expression that identifies the qualified name of a job. The special value of '\*' indicates the current job.

**internal-objects** A character or graphic string expression that indicates whether rows are returned for internal objects.

**NO** Only external objects are returned. This is the default.

**YES** Rows for internal objects and internal space objects are returned in addition to external objects.

**ignore-errors** A character or graphic string expression that identifies what to do when an error is encountered.

**NO** An error is returned.

**YES** A warning is returned.

No row is returned when an error is encountered. This is the default.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 211. JOB\_LOCK\_INFO table function

| Column Name                                            | Data Type   | Description                                                                                                                                                                           |
|--------------------------------------------------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LOCK_CATEGORY                                          | VARCHAR(23) | The type of entity to which this row of lock information applies.                                                                                                                     |
|                                                        |             | <b>EXTERNAL</b> IBM i external object                                                                                                                                                 |
|                                                        |             | <b>EXTERNAL SPACE LOCATION</b> IBM i external object space location                                                                                                                   |
|                                                        |             | <b>INTERNAL</b> Internal system object<br>This type of row is only returned if the <i>internal-objects</i> parameter value is YES.                                                    |
|                                                        |             | <b>INTERNAL SPACE LOCATION</b> Internal system object space location<br>This type of row is only returned if the <i>internal-objects</i> parameter value is YES.                      |
|                                                        |             | <b>LOCK SPACE</b> Lock space object                                                                                                                                                   |
|                                                        |             | <b>MEMBER</b> Member object                                                                                                                                                           |
| Contains the null value if LOCK_CATEGORY is not known. |             |                                                                                                                                                                                       |
| OBJECT_LIBRARY                                         | VARCHAR(10) | The name of the library containing the locked object.<br>Contains the null value if LOCK_CATEGORY is INTERNAL or INTERNAL SPACE LOCATION or if the library name cannot be determined. |

Table 211. JOB\_LOCK\_INFO table function (continued)

| Column Name      | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBJECT_NAME      | VARCHAR(30) | <p>The name of the object that is locked.</p> <p>If LOCK_CATEGORY is MEMBER, the object name is the name of the file that owns the member. If the MEMBER_LOCK_TYPE is MEMBER or ACCESS PATH, the file that owns the member may be either a physical file or a logical file. If the MEMBER_LOCK_TYPE is DATA, the file that owns the member will be a physical file.</p> <p>If LOCK_CATEGORY is LOCK SPACE, the name is the lock space id for that lock space.</p> <p>Contains the null value if LOCK_CATEGORY is INTERNAL or INTERNAL SPACE LOCATION and the user does not have *JOBCTL special authority or if the object name cannot be determined.</p> |
| OBJECT_TYPE      | VARCHAR(7)  | The object type. If the lock is on a database member the object type will be *FILE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| OBJECT_ATTRIBUTE | VARCHAR(10) | <p>The extended attribute value for this object's type.</p> <p>Contains the null value if there is no extended attribute associated with the object type.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| LOCK_STATE       | VARCHAR(7)  | <p>The lock condition for the lock request. Lower level locks are returned and can occur when a member of a file is locked but the file itself is not locked.</p> <p><b>*EXCL</b> Lock exclusive, no read allowed.</p> <p><b>*EXCLRD</b> Lock exclusive, read allowed.</p> <p><b>*SHRNUP</b> Lock shared, no update.</p> <p><b>*SHRRD</b> Lock shared for read.</p> <p><b>*SHRUPD</b> Lock shared for update.</p> <p>Contains the null value if the object is not locked but there are locks on lower level objects.</p>                                                                                                                                  |
| LOCK_STATUS      | VARCHAR(17) | <p>The status of the lock request.</p> <p><b>ASYNCHRONOUS WAIT</b> The job or thread has a lock request outstanding for this object (asynchronous). The lock may be a single request or part of a multiple lock request for which some other object specified in the request has been identified as unavailable.</p> <p><b>HELD</b> The lock on this object currently is held by the job or thread.</p> <p><b>SYNCHRONOUS WAIT</b> The job or thread is waiting to get the lock on this object (synchronous).</p> <p>Contains the null value if the object is not locked but there are locks on lower level objects.</p>                                  |
| LOCK_COUNT       | INTEGER     | The number of identical locks on this entity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| LOCK_SCOPE       | VARCHAR(10) | <p>The scope of the lock. Lower level locks are returned and can occur when a member of a file is locked but the file itself is not locked.</p> <p><b>JOB</b> The lock is scoped to the job.</p> <p><b>LOCK SPACE</b> The lock is scoped to a lock space.</p> <p><b>THREAD</b> The lock is scoped to a thread.</p> <p>Contains the null value if the object is not locked but there are locks on lower level objects.</p>                                                                                                                                                                                                                                 |
| MEMBER_LOCKS     | INTEGER     | <p>The number of member locks for a database file.</p> <p>Contains the null value if the object is not a database file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| MEMBER_NAME      | VARCHAR(10) | <p>The name of the member that has a lock held or waiting on it. Contains *N if the member name cannot be determined.</p> <p>Contains the null value if LOCK_CATEGORY is not MEMBER.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

Table 211. JOB\_LOCK\_INFO table function (continued)

| Column Name                | Data Type   | Description                                                                                                                                                                                                                                                                               |
|----------------------------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MEMBER_LOCK_TYPE           | VARCHAR(11) | The type of member lock.                                                                                                                                                                                                                                                                  |
|                            |             | <b>ACCESS PATH</b> The lock is an access path lock.                                                                                                                                                                                                                                       |
|                            |             | <b>DATA</b> The lock is a data lock.                                                                                                                                                                                                                                                      |
|                            |             | <b>MEMBER</b> The lock is a member lock.                                                                                                                                                                                                                                                  |
|                            |             | Contains the null value if LOCK_CATEGORY is not MEMBER.                                                                                                                                                                                                                                   |
| SPACE_LOCATION_LOCK_OFFSET | BIGINT      | A value in bytes to the location in the space that is locked.<br>Contains the null value if LOCK_CATEGORY is not EXTERNAL SPACE LOCATION or INTERNAL SPACE LOCATION.                                                                                                                      |
| LOCK_SPACE_ID              | BINARY(20)  | The identifier of the lock space for which the lock is being waited on.<br>Contains the null value if LOCK_SCOPE is not LOCK SPACE.                                                                                                                                                       |
| THREAD_ID                  | BIGINT      | A value which uniquely identifies a thread within a job holding a thread scope lock or the thread waiting for a lock.<br>Contain the null value if this lock is not thread scoped.                                                                                                        |
| THREAD_HANDLE              | INTEGER     | A value which addresses a particular thread within a job holding a thread scope lock or the thread waiting for a lock.<br>Contain the null value if this lock is not thread scoped.                                                                                                       |
| ASP_NAME                   | VARCHAR(10) | The name of the Auxiliary Storage Pool (ASP) that contains the object that is locked.<br>Contains the null value if the name of the ASP device cannot be determined.                                                                                                                      |
| ASP_NUMBER                 | INTEGER     | The numeric identifier of the ASP containing the locked object.<br>Contains the null value if the name of the ASP device cannot be determined.                                                                                                                                            |
| OBJECT_LOCK_HANDLE         | BINARY(64)  | An identifier that can be input to Retrieve Lock Information (QWCRLOCKI) API to find additional information about other holders of locks on this object.<br>Contains the null value if additional information cannot be retrieved.                                                        |
| LOCK_REQUEST_HANDLE        | BINARY(64)  | A handle to lock request information. Using the Retrieve Lock Request Information (QWCRQRQI) API and passing in this handle you can retrieve additional information about the program that requested this lock.<br>Contains the null value if additional information cannot be retrieved. |

### Example

- Retrieve information about locks for the current job.

```
SELECT * FROM TABLE(QSYS2.JOB_LOCK_INFO('*'));
```

## JOB\_QUEUE\_INFO view

The JOB\_QUEUE\_INFO view returns one row for each job queue.

The values returned for the columns in the view are similar to the values returned by the Work with Job Queue (WRKJOBQ) CL command and the Retrieve Job Queue Information (QSPRJOBQ) API.

**Authorization:** Rows will be returned for job queues when the caller has:

- Execute authority to the job queue library and
  - Read authority to the job queue, or
  - \*JOBCTL special authority and the job queue has OPRCTL(\*YES), or
  - \*SPLCTL special authority

The following table describes the columns in the view. The system name is JOBQ\_INFO. The schema is QSYS2.

Table 212. JOB\_QUEUE\_INFO view

| Column Name            | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------|--------------------|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_QUEUE_NAME         | JOBQ               | VARCHAR(10)             | The name of the job queue.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| JOB_QUEUE_LIBRARY      | JOBQ_LIB           | VARCHAR(10)             | The name of the library that contains the job queue.                                                                                                                                                                                                                                                                                                                                                                                                                   |
| JOB_QUEUE_STATUS       | STATUS             | VARCHAR(8)              | The status of the job queue.<br><br><b>HELD</b> The queue is held.<br><b>RELEASED</b> The queue is released.                                                                                                                                                                                                                                                                                                                                                           |
| NUMBER_OF_JOBS         | JOBS               | INTEGER                 | The number of jobs in the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| SUBSYSTEM_NAME         | SUB_NAME           | VARCHAR(10)<br>Nullable | The name of the subsystem that can receive jobs from this job queue.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                                                                                                                                                      |
| SUBSYSTEM_LIBRARY_NAME | SUBLIB_NAM         | VARCHAR(10)<br>Nullable | The library in which the subsystem description resides.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                                                                                                                                                                   |
| SEQUENCE_NUMBER        | SEQNO              | INTEGER<br>Nullable     | The job queue entry sequence number. The subsystem uses this number to determine the order in which job queues are processed. Jobs from the queue with the lowest sequence number are processed first.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                    |
| MAXIMUM_ACTIVE_JOBS    | MAX_JOBS           | INTEGER<br>Nullable     | The maximum number of jobs that can be active at the same time through this job queue entry. A value of -1 indicates *NOMAX, no maximum number of jobs is defined.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                                                        |
| ACTIVE_JOBS            | ACT_JOBS           | INTEGER<br>Nullable     | The current number of jobs that are active that came through this job queue entry.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                                                                                                                                        |
| HELD_JOBS              | HELD_JOBS          | INTEGER                 | The current number of jobs that are in *HELD status. This is the sum of the 10 HELD_JOBS_PRIORITY_n columns.                                                                                                                                                                                                                                                                                                                                                           |
| RELEASED_JOBS          | RLS_JOBS           | INTEGER                 | The current number of jobs that are in *RELEASED status. This is the sum of the 10 RELEASED_JOBS_PRIORITY_n columns.                                                                                                                                                                                                                                                                                                                                                   |
| SCHEDULED_JOBS         | SCHED_JOBS         | INTEGER                 | The current number of jobs that are in *SCHEDULED status. This is the sum of the 10 SCHEDULED_JOBS_PRIORITY_n columns.                                                                                                                                                                                                                                                                                                                                                 |
| TEXT_DESCRIPTION       | TEXT               | VARCHAR(50)<br>Nullable | Text that describes the job queue.<br><br>Contains the null value if there is no text description for the job queue.                                                                                                                                                                                                                                                                                                                                                   |
| OPERATOR_CONTROLLED    | OPR_CTRL           | VARCHAR(4)              | Whether users with job control authority are allowed to control this job queue and manage the jobs on the queue. Users have job control authority if SPCAUT(*JOBCTL) is specified in their user profile.<br><br><b>*NO</b> This queue and its jobs cannot be controlled by users with job control authority unless they also have other special authority.<br><br><b>*YES</b> Users with job control authority can control the queue and manage the jobs on the queue. |

Table 212. JOB\_QUEUE\_INFO view (continued)

| Column Name                    | System Column Name | Data Type           | Description                                                                                                                                                                                                                                                                                                |
|--------------------------------|--------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUTHORITY_TO_CHECK             | ALL_AUTH           | VARCHAR(7)          | Whether the user must be the owner of the queue in order to control the queue by holding or releasing the queue.<br><br><b>*DTAAUT</b> Any user with *READ, *ADD, or *DELETE authority to the job queue can control the queue.<br><br><b>*OWNER</b> Only the owner of the job queue can control the queue. |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_1 | MAXIMUM1           | INTEGER<br>Nullable | The maximum number of priority 1 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_2 | MAXIMUM2           | INTEGER<br>Nullable | The maximum number of priority 2 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_3 | MAXIMUM3           | INTEGER<br>Nullable | The maximum number of priority 3 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_4 | MAXIMUM4           | INTEGER<br>Nullable | The maximum number of priority 4 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_5 | MAXIMUM5           | INTEGER<br>Nullable | The maximum number of priority 5 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_6 | MAXIMUM6           | INTEGER<br>Nullable | The maximum number of priority 6 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_7 | MAXIMUM7           | INTEGER<br>Nullable | The maximum number of priority 7 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_8 | MAXIMUM8           | INTEGER<br>Nullable | The maximum number of priority 8 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| MAXIMUM_ACTIVE_JOBS_PRIORITY_9 | MAXIMUM9           | INTEGER<br>Nullable | The maximum number of priority 9 jobs that can be active at the same time. A value of -1 indicates *NOMAX, no maximum number of jobs.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                         |
| ACTIVE_JOBS_PRIORITY_0         | ACTIVE0            | INTEGER<br>Nullable | The number of priority 0 jobs that are active.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                |
| ACTIVE_JOBS_PRIORITY_1         | ACTIVE1            | INTEGER<br>Nullable | The number of priority 1 jobs that are active.<br><br>Contains the null value if this job queue is not associated with an active subsystem.                                                                                                                                                                |



Table 212. JOB\_QUEUE\_INFO view (continued)

| Column Name               | System Column Name | Data Type           | Description                                                                                                                             |
|---------------------------|--------------------|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| ACTIVE_JOBS_PRIORITY_2    | ACTIVE2            | INTEGER<br>Nullable | The number of priority 2 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_3    | ACTIVE3            | INTEGER<br>Nullable | The number of priority 3 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_4    | ACTIVE4            | INTEGER<br>Nullable | The number of priority 4 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_5    | ACTIVE5            | INTEGER<br>Nullable | The number of priority 5 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_6    | ACTIVE6            | INTEGER<br>Nullable | The number of priority 6 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_7    | ACTIVE7            | INTEGER<br>Nullable | The number of priority 7 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_8    | ACTIVE8            | INTEGER<br>Nullable | The number of priority 8 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| ACTIVE_JOBS_PRIORITY_9    | ACTIVE9            | INTEGER<br>Nullable | The number of priority 9 jobs that are active.<br>Contains the null value if this job queue is not associated with an active subsystem. |
| RELEASED_JOBS_PRIORITY_0  | RELEASED0          | INTEGER             | The number of priority 0 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_1  | RELEASED1          | INTEGER             | The number of priority 1 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_2  | RELEASED2          | INTEGER             | The number of priority 2 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_3  | RELEASED3          | INTEGER             | The number of priority 3 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_4  | RELEASED4          | INTEGER             | The number of priority 4 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_5  | RELEASED5          | INTEGER             | The number of priority 5 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_6  | RELEASED6          | INTEGER             | The number of priority 6 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_7  | RELEASED7          | INTEGER             | The number of priority 7 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_8  | RELEASED8          | INTEGER             | The number of priority 8 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| RELEASED_JOBS_PRIORITY_9  | RELEASED9          | INTEGER             | The number of priority 9 jobs currently sitting on the job queue in *RELEASED status.                                                   |
| SCHEDULED_JOBS_PRIORITY_0 | SCHEDULED0         | INTEGER             | The number of priority 0 jobs currently sitting on the job queue in *SCHEDULED status.                                                  |
| SCHEDULED_JOBS_PRIORITY_1 | SCHEDULED1         | INTEGER             | The number of priority 1 jobs currently sitting on the job queue in *SCHEDULED status.                                                  |
| SCHEDULED_JOBS_PRIORITY_2 | SCHEDULED2         | INTEGER             | The number of priority 2 jobs currently sitting on the job queue in *SCHEDULED status.                                                  |
| SCHEDULED_JOBS_PRIORITY_3 | SCHEDULED3         | INTEGER             | The number of priority 3 jobs currently sitting on the job queue in *SCHEDULED status.                                                  |
| SCHEDULED_JOBS_PRIORITY_4 | SCHEDULED4         | INTEGER             | The number of priority 4 jobs currently sitting on the job queue in *SCHEDULED status.                                                  |

Table 212. JOB\_QUEUE\_INFO view (continued)

| Column Name               | System Column Name | Data Type | Description                                                                            |
|---------------------------|--------------------|-----------|----------------------------------------------------------------------------------------|
| SCHEDULED_JOBS_PRIORITY_5 | SCHEDULED5         | INTEGER   | The number of priority 5 jobs currently sitting on the job queue in *SCHEDULED status. |
| SCHEDULED_JOBS_PRIORITY_6 | SCHEDULED6         | INTEGER   | The number of priority 6 jobs currently sitting on the job queue in *SCHEDULED status. |
| SCHEDULED_JOBS_PRIORITY_7 | SCHEDULED7         | INTEGER   | The number of priority 7 jobs currently sitting on the job queue in *SCHEDULED status. |
| SCHEDULED_JOBS_PRIORITY_8 | SCHEDULED8         | INTEGER   | The number of priority 8 jobs currently sitting on the job queue in *SCHEDULED status. |
| SCHEDULED_JOBS_PRIORITY_9 | SCHEDULED9         | INTEGER   | The number of priority 9 jobs currently sitting on the job queue in *SCHEDULED status. |
| HELD_JOBS_PRIORITY_0      | HELD0              | INTEGER   | The number of priority 0 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_1      | HELD1              | INTEGER   | The number of priority 1 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_2      | HELD2              | INTEGER   | The number of priority 2 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_3      | HELD3              | INTEGER   | The number of priority 3 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_4      | HELD4              | INTEGER   | The number of priority 4 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_5      | HELD5              | INTEGER   | The number of priority 5 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_6      | HELD6              | INTEGER   | The number of priority 6 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_7      | HELD7              | INTEGER   | The number of priority 7 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_8      | HELD8              | INTEGER   | The number of priority 8 jobs currently sitting on the job queue in *HELD status.      |
| HELD_JOBS_PRIORITY_9      | HELD9              | INTEGER   | The number of priority 9 jobs currently sitting on the job queue in *HELD status.      |

## Example

- Examine the job queues with the largest number of active jobs

```
SELECT * FROM QSYS2.JOB_QUEUE_INFO
WHERE ACTIVE_JOBS IS NOT NULL
ORDER BY NUMBER_OF_JOBS DESC;
```

## MEMORY\_POOL table function

The MEMORY\_POOL table function returns one row for every pool.

The information returned is similar to the detail seen from the Work System Status (WRKSYSSTS) command.

**Authorization:** None required.

```
MEMORY_POOL ((reset_statistics)) =>
```

The schema is QSYS2.

**reset\_statistics** A character or graphic string expression that contains a value of YES or NO.

If this parameter has a value of YES, statistics are reset such that the time of this query execution is used as the new baseline. The columns that contain this statistical data have names that are prefixed with ELAPSED\_. Future invocations of MEMORY\_POOL within this connection will return statistical detail relative to the new baseline. If this parameter has a value of NO, statistics are not reset for the invocation. If this parameter is not specified, the default is NO.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 213. MEMORY\_POOL table function

| Column Name                | Data Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SYSTEM_POOL_ID             | INTEGER       | The system-related pool identifier for each of the system storage pools that currently has main storage allocated to it.                                                                                                                                                                                                                                                                                                                                       |
| POOL_NAME                  | VARCHAR(10)   | The name of this storage pool. The name may be a number, in which case it is a private pool associated with a subsystem, or one of the following special values.<br><br><b>*MACHINE</b> The machine pool.<br><b>*BASE</b> The base system pool, which can be shared with other subsystems.<br><b>*INTERACT</b> The shared pool used for the QINTER subsystem.<br><b>*SPOOL</b> The shared pool used for spooled writers.<br><b>*SHRPOOLx</b> A shared pool.    |
| CURRENT_SIZE               | DECIMAL(20,2) | The amount of main storage, in megabytes, in the pool.                                                                                                                                                                                                                                                                                                                                                                                                         |
| RESERVED_SIZE              | DECIMAL(10,2) | The amount of storage, in megabytes, in the pool reserved for system use (for example, for save/restore operations).                                                                                                                                                                                                                                                                                                                                           |
| DEFINED_SIZE               | DECIMAL(20,2) | The size of the pool, in megabytes, as defined in the shared pool, subsystem description, or system value QMCHPOOL. Contains the null value for a pool without a defined size.                                                                                                                                                                                                                                                                                 |
| MAXIMUM_ACTIVE_THREADS     | INTEGER       | The maximum number of threads that can be active in the pool at any one time.                                                                                                                                                                                                                                                                                                                                                                                  |
| CURRENT_THREADS            | INTEGER       | The number of threads currently using the pool.                                                                                                                                                                                                                                                                                                                                                                                                                |
| CURRENT_INELIGIBLE_THREADS | INTEGER       | The number of ineligible threads in the pool.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| STATUS                     | VARCHAR(8)    | The status of the pool.<br><br><b>ACTIVE</b> Pool is currently active.<br><b>INACTIVE</b> Pool is currently not active.                                                                                                                                                                                                                                                                                                                                        |
| SUBSYSTEM_LIBRARY_NAME     | VARCHAR(10)   | The library containing the subsystem name. Contains the null value for shared pools.                                                                                                                                                                                                                                                                                                                                                                           |
| SUBSYSTEM_NAME             | VARCHAR(10)   | The subsystem with which this storage pool is associated. Contains the null value for shared pools.                                                                                                                                                                                                                                                                                                                                                            |
| DESCRIPTION                | VARCHAR(50)   | The description of the shared pool. Contains the null value for private pools or if a description does not exist for a shared pool.                                                                                                                                                                                                                                                                                                                            |
| PAGING_OPTION              | VARCHAR(10)   | Whether the system will dynamically adjust the paging characteristics of the storage pool for optimum performance.<br><br><b>*FIXED</b> The system does not dynamically adjust the paging characteristics.<br><b>*CALC</b> The system dynamically adjusts the paging characteristics.<br><b>USRDFN</b> The system does not dynamically adjust the paging characteristics for the storage pool but uses values that have been defined through the QWCCHGTN API. |
| ELAPSED_TIME               | INTEGER       | The time, in seconds, since the measurement start time.                                                                                                                                                                                                                                                                                                                                                                                                        |
| ELAPSED_DATABASE_FAULTS    | DECIMAL(10,1) | The rate, in page faults per second, of database page faults against pages containing either database access paths or data.                                                                                                                                                                                                                                                                                                                                    |

Table 213. MEMORY\_POOL table function (continued)

| Column Name                  | Data Type     | Description                                                                                                                                                                                                                    |
|------------------------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ELAPSED_NON_DATABASE_FAULTS  | DECIMAL(10,1) | The rate, in page faults per second, of nondatabase page faults against pages other than those designated as database pages.                                                                                                   |
| ELAPSED_TOTAL_FAULTS         | DECIMAL(10,1) | The rate, in page faults per second, of database faults and non-database faults.                                                                                                                                               |
| ELAPSED_DATABASE_PAGES       | DECIMAL(10,1) | The rate, in pages per second, at which database pages are brought into the storage pool.                                                                                                                                      |
| ELAPSED_NON_DATABASE_PAGES   | DECIMAL(10,1) | The rate in pages per second at which nondatabase pages are brought into the storage pool.                                                                                                                                     |
| ELAPSED_ACTIVE_TO_WAIT       | DECIMAL(10,1) | The rate, in transitions per minute, of transitions of threads from an active condition to a waiting condition.                                                                                                                |
| ELAPSED_WAIT_TO_INELIGIBLE   | DECIMAL(10,1) | The rate, in transitions per minute, of transitions of threads from a waiting condition to an ineligible condition.                                                                                                            |
| ELAPSED_ACTIVE_TO_INELIGIBLE | DECIMAL(10,1) | The rate, in transitions per minute, of transitions of threads from an active condition to an ineligible condition.                                                                                                            |
| TUNING_PRIORITY              | INTEGER       | The priority of the shared storage pool used by the system when making automatic performance adjustments. Contains the null value for private pools defined in subsystem descriptions.                                         |
| TUNING_MINIMUM_SIZE          | DECIMAL(10,2) | The minimum amount of storage to allocate to the shared storage pool (as a percentage of total main storage). Contains the null value for private pools defined in subsystem descriptions.                                     |
| TUNING_MAXIMUM_SIZE          | DECIMAL(10,2) | The maximum amount of storage to allocate to the shared storage pool (as a percentage of total main storage). Contains the null value for private pools defined in subsystem descriptions.                                     |
| TUNING_MINIMUM_FAULTS        | DECIMAL(10,2) | The maximum page faults per second to use as a guideline for the shared storage pool. Contains the null value for private pools defined in subsystem descriptions.                                                             |
| TUNING_MAXIMUM_FAULTS        | DECIMAL(10,2) | The minimum page faults per second to use as a guideline for the shared storage pool. Contains the null value for private pools defined in subsystem descriptions.                                                             |
| TUNING_THREAD_FAULTS         | DECIMAL(10,2) | The page faults per second for each active thread to use as a guideline for the shared storage pool. Contains the null value for private pools defined in subsystem descriptions.                                              |
| TUNING_MINIMUM_ACTIVITY      | DECIMAL(10,2) | The minimum value that the shared pool's activity level can be set to by the performance adjuster when the QPFRADJ system value is set to 2 or 3. Contains the null value for private pools defined in subsystem descriptions. |
| TUNING_MAXIMUM_ACTIVITY      | DECIMAL(10,2) | The maximum value that the shared pool's activity level can be set to by the performance adjuster when the QPFRADJ system value is set to 2 or 3. Contains the null value for private pools defined in subsystem descriptions. |

## Example

Return all available pool information, both private and shared, active and inactive. Specify to reset all the elapsed values to 0.

```
SELECT * FROM TABLE(QSYS2.MEMORY_POOL(RESET_STATISTICS=>'YES')) X;
```

## MEMORY\_POOL\_INFO view

The MEMORY\_POOL\_INFO view returns one row for every active pool.

The information returned is similar to the detail seen from the Work System Status (WRKSYSSTS) command. It does not reset the statistical columns; to do this, use the associated table function, [MEMORY\\_POOL](#).

**Authorization:** None required.

The following table describes the columns in the view. The system name is POOL\_INFO. The schema is QSYS2.

Table 214. MEMORY\_POOL\_INFO view

| Column Name                 | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SYSTEM_POOL_ID              | POOL_ID            | INTEGER                 | The system-related pool identifier for each of the system storage pools that currently has main storage allocated to it.                                                                                                                                                                                                                                                                                                                                       |
| POOL_NAME                   | POOL_NAME          | VARCHAR(10)             | The name of this storage pool. The name may be a number, in which case it is a private pool associated with a subsystem, or one of the following special values.<br><br><b>*MACHINE</b> The machine pool.<br><b>*BASE</b> The base system pool, which can be shared with other subsystems.<br><b>*INTERACT</b> The shared pool used for the QINTER subsystem.<br><b>*SPOOL</b> The shared pool used for spooled writers.<br><b>*SHRPOOLx</b> A shared pool.    |
| CURRENT_SIZE                | CURR_SIZE          | DECIMAL(20,2)           | The amount of main storage, in megabytes, in the pool.                                                                                                                                                                                                                                                                                                                                                                                                         |
| RESERVED_SIZE               | RSVD_SIZE          | DECIMAL(10,2)           | The amount of storage, in megabytes, in the pool reserved for system use (for example, for save/restore operations).                                                                                                                                                                                                                                                                                                                                           |
| DEFINED_SIZE                | DFND_SIZE          | DECIMAL(20,2)           | The size of the pool, in megabytes, as defined in the shared pool, subsystem description, or system value QMCHPOOL. Contains the null value for a pool without a defined size.                                                                                                                                                                                                                                                                                 |
| MAXIMUM_ACTIVE_THREADS      | MAX_THREAD         | INTEGER                 | The maximum number of threads that can be active in the pool at any one time.                                                                                                                                                                                                                                                                                                                                                                                  |
| CURRENT_THREADS             | CURR_THRD          | INTEGER                 | The number of threads currently using the pool.                                                                                                                                                                                                                                                                                                                                                                                                                |
| CURRENT_INELIGIBLE_THREADS  | INEL_THRD          | INTEGER                 | The number of ineligible threads in the pool.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| SUBSYSTEM_LIBRARY_NAME      | SUBLIB_NAM         | VARCHAR(10)<br>Nullable | The library containing the subsystem name. Contains the null value for shared pools.                                                                                                                                                                                                                                                                                                                                                                           |
| SUBSYSTEM_NAME              | SUB_NAME           | VARCHAR(10)<br>Nullable | The subsystem with which this storage pool is associated. Contains the null value for shared pools.                                                                                                                                                                                                                                                                                                                                                            |
| DESCRIPTION                 | DESC               | VARCHAR(50)<br>Nullable | The description of the shared pool. Contains the null value for private pools or if a description does not exist for a shared pool.                                                                                                                                                                                                                                                                                                                            |
| PAGING_OPTION               | PAGE_OPT           | VARCHAR(10)             | Whether the system will dynamically adjust the paging characteristics of the storage pool for optimum performance.<br><br><b>*FIXED</b> The system does not dynamically adjust the paging characteristics.<br><b>*CALC</b> The system dynamically adjusts the paging characteristics.<br><b>USRDFN</b> The system does not dynamically adjust the paging characteristics for the storage pool but uses values that have been defined through the QWCCHGTN API. |
| ELAPSED_TIME                | ELAP_TIME          | INTEGER                 | The time, in seconds, since the measurement start time.                                                                                                                                                                                                                                                                                                                                                                                                        |
| ELAPSED_DATABASE_FAULTS     | ELAP_DBF           | DECIMAL(10,1)           | The rate, in page faults per second, of database page faults against pages containing either database access paths or data.                                                                                                                                                                                                                                                                                                                                    |
| ELAPSED_NON_DATABASE_FAULTS | ELAP_NDBF          | DECIMAL(10,1)           | The rate, in page faults per second, of nondatabase page faults against pages other than those designated as database pages.                                                                                                                                                                                                                                                                                                                                   |
| ELAPSED_TOTAL_FAULTS        | ELAP_TOTF          | DECIMAL(10,1)           | The rate, in page faults per second, of database faults and non-database faults.                                                                                                                                                                                                                                                                                                                                                                               |
| ELAPSED_DATABASE_PAGES      | ELAP_DBP           | DECIMAL(10,1)           | The rate, in pages per second, at which database pages are brought into the storage pool.                                                                                                                                                                                                                                                                                                                                                                      |
| ELAPSED_NON_DATABASE_PAGES  | ELAP_NDBP          | DECIMAL(10,1)           | The rate in pages per second at which nondatabase pages are brought into the storage pool.                                                                                                                                                                                                                                                                                                                                                                     |
| ELAPSED_ACTIVE_TO_WAIT      | ELAP_ATW           | DECIMAL(10,1)           | The rate, in transitions per minute, of transitions of threads from an active condition to a waiting condition.                                                                                                                                                                                                                                                                                                                                                |

Table 214. MEMORY\_POOL\_INFO view (continued)

| Column Name                  | System Column Name | Data Type                 | Description                                                                                                                                                                                                                    |
|------------------------------|--------------------|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ELAPSED_WAIT_TO_INELIGIBLE   | ELAP_WTI           | DECIMAL(10,1)             | The rate, in transitions per minute, of transitions of threads from a waiting condition to an ineligible condition.                                                                                                            |
| ELAPSED_ACTIVE_TO_INELIGIBLE | ELAP_ATI           | DECIMAL(10,1)             | The rate, in transitions per minute, of transitions of threads from an active condition to an ineligible condition.                                                                                                            |
| TUNING_PRIORITY              | TUN_PRIOR          | INTEGER<br>Nullable       | The priority of the shared storage pool used by the system when making automatic performance adjustments. Contains the null value for private pools defined in subsystem descriptions.                                         |
| TUNING_MINIMUM_SIZE          | TUN_MIN_SZ         | DECIMAL(10,2)<br>Nullable | The minimum amount of storage to allocate to the shared storage pool (as a percentage of total main storage). Contains the null value for private pools defined in subsystem descriptions.                                     |
| TUNING_MAXIMUM_SIZE          | TUN_MAX_SZ         | DECIMAL(10,2)<br>Nullable | The maximum amount of storage to allocate to the shared storage pool (as a percentage of total main storage). Contains the null value for private pools defined in subsystem descriptions.                                     |
| TUNING_MINIMUM_FAULTS        | TUN_MIN_FT         | DECIMAL(10,2)<br>Nullable | The maximum page faults per second to use as a guideline for the shared storage pool. Contains the null value for private pools defined in subsystem descriptions.                                                             |
| TUNING_MAXIMUM_FAULTS        | TUN_MAX_FT         | DECIMAL(10,2)<br>Nullable | The minimum page faults per second to use as a guideline for the shared storage pool. Contains the null value for private pools defined in subsystem descriptions.                                                             |
| TUNING_THREAD_FAULTS         | TUN_THR_FT         | DECIMAL(10,2)<br>Nullable | The page faults per second for each active thread to use as a guideline for the shared storage pool. Contains the null value for private pools defined in subsystem descriptions.                                              |
| TUNING_MINIMUM_ACTIVITY      | TUN_MIN_AC         | DECIMAL(10,2)<br>Nullable | The minimum value that the shared pool's activity level can be set to by the performance adjuster when the QPFRADJ system value is set to 2 or 3. Contains the null value for private pools defined in subsystem descriptions. |
| TUNING_MAXIMUM_ACTIVITY      | TUN_MAX_AC         | DECIMAL(10,2)<br>Nullable | The maximum value that the shared pool's activity level can be set to by the performance adjuster when the QPFRADJ system value is set to 2 or 3. Contains the null value for private pools defined in subsystem descriptions. |

## Example

Return all active pool information.

```
SELECT * FROM QSYS2.MEMORY_POOL_INFO;
```

## OBJECT\_LOCK\_INFO view

The OBJECT\_LOCK\_INFO view returns one row for every lock held for every object on the partition in \*SYSBAS and in the current thread's ASP group.

The values returned for the columns in the view are closely related to the values returned by [Retrieve Lock Information API](#) and [Retrieve Lock Request Information API](#). Refer to the APIs for more detailed information.

**Authorization:** The caller must have:

- \*EXECUTE authority to the library containing the object, and
- \*OBJOPR and \*READ authority to the database file

The following table describes the columns in the view. The system name is OBJ\_LOCK. The schema is QSYS2.

Table 215. OBJECT\_LOCK\_INFO view

| Column Name   | System Column Name | Data Type    | Description                                   |
|---------------|--------------------|--------------|-----------------------------------------------|
| OBJECT_SCHEMA | OSHEMA             | VARCHAR(128) | The name of the schema containing the object. |

Table 215. OBJECT\_LOCK\_INFO view (continued)

| Column Name          | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|--------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OBJECT_NAME          | NAME               | VARCHAR(128)            | The name of the object.                                                                                                                                                                                                                                                                                                                                                                     |
| SYSTEM_OBJECT_SCHEMA | SYS_DNAME          | VARCHAR(10)             | The system library name of the object.                                                                                                                                                                                                                                                                                                                                                      |
| SYSTEM_OBJECT_NAME   | SYS_ONAME          | VARCHAR(10)             | The system name of the object                                                                                                                                                                                                                                                                                                                                                               |
| SYSTEM_TABLE_MEMBER  | SYS_MNAME          | VARCHAR(10)<br>Nullable | The name of the member that is locked in the file.<br>Contains the null value if the lock information is not for a member.                                                                                                                                                                                                                                                                  |
| OBJECT_TYPE          | OBJTYPE            | VARCHAR(8)              | The system object type of the locked object.                                                                                                                                                                                                                                                                                                                                                |
| SQL_OBJECT_TYPE      | SQLTYPE            | VARCHAR(9)<br>Nullable  | The SQL type of the object. Values are:<br><ul style="list-style-type: none"> <li>• ALIAS</li> <li>• FUNCTION</li> <li>• INDEX</li> <li>• PACKAGE</li> <li>• PROCEDURE</li> <li>• ROUTINE</li> <li>• SEQUENCE</li> <li>• TABLE</li> <li>• TRIGGER</li> <li>• TYPE</li> <li>• VARIABLE</li> <li>• VIEW</li> <li>• XSR</li> </ul> Contains the null value if the object is not an SQL object. |
| ASP_NUMBER           | ASPNUM             | INTEGER                 | The numeric identifier of the ASP containing the object that is locked.                                                                                                                                                                                                                                                                                                                     |
| ASPGRP               | ASPGRP             | VARCHAR(10)             | The name of the ASP device containing the object that is locked. Can contain the special value of *SYSBAS.                                                                                                                                                                                                                                                                                  |
| MEMBER_LOCK_TYPE     | LOCK_TYPE          | VARCHAR(10)<br>Nullable | The type of lock that is held.<br><b>ACCESSPATH</b> Lock on the access path used to access the member's data.<br><b>DATA</b> Lock on the actual data within the member.<br><b>MEMBER</b> Lock on the member control block.<br>Contains the null value if the lock information is not for a member.                                                                                          |
| LOCK_STATE           | LOCK_STATE         | VARCHAR(7)              | The lock condition for the object or member.<br><b>*EXCL</b> Lock exclusive no read.<br><b>*EXCLRD</b> Lock exclusive allow read.<br><b>*SHRNUP</b> Lock shared no update.<br><b>*SHRRD</b> Lock shared for read.<br><b>*SHRUPD</b> Lock shared for update.                                                                                                                                 |
| LOCK_STATUS          | STATUS             | VARCHAR(9)              | The status of the lock.<br><b>HELD</b> The lock is currently held by the job.<br><b>REQUESTED</b> The job has a lock request outstanding for the object.<br><b>WAITING</b> The job is waiting for the lock.                                                                                                                                                                                 |
| LOCK_SCOPE           | LOCK_SCOPE         | VARCHAR(10)             | The scope of the lock. Values are:<br><ul style="list-style-type: none"> <li>• JOB</li> <li>• LOCK SPACE</li> <li>• THREAD</li> </ul>                                                                                                                                                                                                                                                       |

Table 215. OBJECT\_LOCK\_INFO view (continued)

| Column Name          | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------|--------------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_NAME             | JOB_NAME           | VARCHAR(28)<br>Nullable   | The qualified job name.<br><br>Contains the null value when the LOCK_SCOPE column value is LOCK SPACE.                                                                                                                                                                                                                                                                                                                                                                                                                |
| THREAD_ID            | THREAD_ID          | BIGINT<br>Nullable        | The thread that is associated with the lock.<br><br><ul style="list-style-type: none"> <li>If a held lock is job scoped, returns the null value. If a held lock is thread scoped, contains the identifier for the thread holding the lock.</li> <li>If the scope of the lock is to the lock space and the lock is not held, contains the identifier of the thread requesting the lock.</li> <li>If the lock is requested but not yet available, contains the identifier of the thread requesting the lock.</li> </ul> |
| LOCK_SPACE_ID        | LOCKID             | BINARY(20)<br>Nullable    | When the LOCK_SCOPE column value is LOCK SPACE and the lock is held, contains the lock space ID value of the lock space that holds the lock. If the lock is being waited on by a thread, contains the lock space ID value for which the lock is being waited on.<br><br>Otherwise, contains the null value.                                                                                                                                                                                                           |
| LOCK_COUNT           | LOCK_COUNT         | INTEGER                   | The number of identical locks held.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| PROGRAM_LIBRARY_NAME | PROGLIB            | VARCHAR(10)<br>Nullable   | The name of the library containing the program or service program.<br><br>Contains the null value if the lock holder information is not available.                                                                                                                                                                                                                                                                                                                                                                    |
| PROGRAM_NAME         | PROGNAME           | VARCHAR(10)<br>Nullable   | The name of the program holding the lock. This can be any type of program object, including objects of type *PGM and *SRVPGM.<br><br>Contains the null value if the lock holder information is not available.                                                                                                                                                                                                                                                                                                         |
| MODULE_NAME_LIBRARY  | MODLIB             | VARCHAR(10)<br>Nullable   | The library containing the module.<br><br>Contains the null value if the lock holder information is not available or if the program is not an ILE program.                                                                                                                                                                                                                                                                                                                                                            |
| MODULE_NAME          | MODNAME            | VARCHAR(10)<br>Nullable   | The module containing the ILE procedure.<br><br>Contains the null value if the lock holder information is not available or if the program is not an ILE program.                                                                                                                                                                                                                                                                                                                                                      |
| PROCEDURE_NAME       | PROCNAME           | VARCHAR(4096)<br>Nullable | The name of the procedure.<br><br>Contains the null value if the lock holder information is not available.                                                                                                                                                                                                                                                                                                                                                                                                            |
| STATEMENT_ID         | STMTID             | CHAR(10)<br>Nullable      | The high-level language statement identifier. For a character representation of a number, the number is right-adjusted and padded on the left with zeros (for example, '000000246').<br><br>Contains the null value if the lock holder information is not available.                                                                                                                                                                                                                                                  |
| MACHINE_INSTRUCTION  | INSTRUCT           | INTEGER<br>Nullable       | The current machine instruction number in the program.<br><br>Contains the null value if the lock holder information is not available or if it is an ILE procedure.                                                                                                                                                                                                                                                                                                                                                   |

## Example

Find all the jobs holding object locks over the SALES table:

```
SELECT * FROM QSYS2.OBJECT_LOCK_INFO
WHERE SYSTEM_OBJECT_NAME = 'SALES'
```

## OPEN\_FILES table function

The OPEN\_FILES table function returns a list of files (\*FILE objects) that are open in all threads for a job.

This information is similar to what can be accessed through the Display Job (DSPJOB) CL command and by the List Open Files (QDMLOPNF) API.



**Authorization:** None required to see information for jobs where the caller's user profile is the same as the job user identity of the job for which the information is being returned. To see information for other jobs, the caller must have \*JOBCTL special authority.

➤ OPEN\_FILES — ( — JOB\_NAME — => — ) job-name — ) ➤

The schema is QSYS2.

**job-name** A character or graphic string expression that identifies the qualified name of a job. The special value of '\*' indicates the current job.

The result of the function is a table containing rows with the format shown in the following table. All the columns are nullable.

Table 216. OPEN\_FILES table function

| Column Name           | Data Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LIBRARY_NAME          | VARCHAR(10) | The name of the library that contains the open file.<br>Contains the null value if the file is an inline data file.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| FILE_NAME             | VARCHAR(10) | The name of the file that is open. For an unnamed inline data file, contains the value QINLINE.                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| FILE_TYPE             | VARCHAR(7)  | The type of file that is open.<br><b>BSCF</b> Binary Synchronous Communications (BSC) file<br><b>CMNF</b> Communications file<br><b>DDMF</b> Distributed Data Management file<br><b>DKTF</b> Diskette file (spooled and non-spooled)<br><b>DSPF</b> Display file<br><b>ICFF</b> Intersystem Communications Function file<br><b>LF</b> Logical file<br><b>MXDF</b> Mixed file<br><b>PF</b> Physical file<br><b>PRTF</b> Printer file (spooled and non-spooled)<br><b>SAVF</b> Save file<br><b>TAPF</b> Tape file<br><b>*INLINE</b> Inline data file |
| MEMBER_NAME           | VARCHAR(10) | If FILE_TYPE is physical (PF) or logical (LF), the name of the database member. If multiple member processing is being performed, the value is *ALL.<br>Contains the null value for a DDM file, an inline data file, and a device file.                                                                                                                                                                                                                                                                                                            |
| DEVICE_NAME           | VARCHAR(10) | The name of the last program device used for an I/O operation if FILE_TYPE is a device file (BSCF, CMNF, DKTF, DSPF, ICFF, MXDF, PRTF, SAVF, or TAPF). If the file is a spooled file, the value is *SPOOL.<br>Contains the null value for a device file when no I/O operation has been performed, a database physical or logical file, a DDM file, and an inline data file.                                                                                                                                                                        |
| RECORD_FORMAT         | VARCHAR(10) | The name of the last record format that was used for an I/O operation to the file.<br>Contains the null value if no record format name was used or no I/O operations have been performed.                                                                                                                                                                                                                                                                                                                                                          |
| ACTIVATION_GROUP_NAME | VARCHAR(10) | The name of the activation group to which the open file is scoped.<br><b>*DFTACTGRP</b> The file is scoped to the default activation group.<br><b>*NEW</b> The file is scoped to a *NEW activation group.<br>Contains the null value for a file scoped to the job, not a specific activation group.                                                                                                                                                                                                                                                |

Table 216. OPEN\_FILES table function (continued)

| Column Name             | Data Type   | Description                                                                                                                                                                                                                                                                                                    |
|-------------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ACTIVATION_GROUP_NUMBER | INTEGER     | The number of the activation group to which the open file is scoped.<br>Contains the null value for a file scoped to the job, not a specific activation group.                                                                                                                                                 |
| THREAD_ID               | BIGINT      | The thread handle assigned by the system which identifies the thread in which the file was opened.                                                                                                                                                                                                             |
| OPEN_OPTION             | VARCHAR(6)  | The type of open operation that was performed.<br><br><b>ALL</b> The file was opened for all operations (input, output, update, and delete).<br><b>INPUT</b> The file was opened for input operations only.<br><b>OUTPUT</b> The file was opened for output operations only.                                   |
| SHARED_OPENS            | BIGINT      | The number of times the file was opened for shared processing.<br>Contains the null value for open operations that are not shared.                                                                                                                                                                             |
| WRITE_COUNT             | BIGINT      | The number of successful write operations. If record blocking is not in effect for the file, this is the number of records. If record blocking is in effect for the file, this is the number of record blocks.                                                                                                 |
| READ_COUNT              | BIGINT      | Number of successful read operations. If record blocking is not in effect for the file, this is the number of records. If record blocking is in effect for the file, this is the number of record blocks.                                                                                                      |
| WRITE_READ_COUNT        | BIGINT      | The number of successful write/read operations.                                                                                                                                                                                                                                                                |
| OTHER_IO_COUNT          | BIGINT      | Number of successful I/O operations of the following types: <ul style="list-style-type: none"> <li>• update</li> <li>• delete</li> <li>• change end-of-data</li> <li>• force end-of-data</li> <li>• force end-of-volume</li> <li>• release record lock</li> <li>• acquire or release program device</li> </ul> |
| RELATIVE_RECORD_NUMBER  | BIGINT      | Relative record number of the last record referred to by an I/O or open operation for database files.<br>Contains the null value for nondatabase files and database files on which no I/O operations have been performed.                                                                                      |
| IASP_NAME               | VARCHAR(10) | The auxiliary storage pool (ASP) device name where the library is stored. If the library is in the system ASP or one of the basic user ASPs, contains *SYSBAS.<br><br>Contains the null value if the file is an inline data file or if the name of the ASP device cannot be determined.                        |
| IASP_NUMBER             | INTEGER     | The number of the auxiliary storage pool (ASP) from which the system allocates storage for the library.<br><br>Contains the null value if the file is an inline data file or if the number of the ASP device cannot be determined.                                                                             |

### Example

- List all the database files that are open for job 429467/QUSER/QZDASOINIT.

```
SELECT * FROM TABLE(QSYS2.OPEN_FILES('429467/QUSER/QZDASOINIT'))
WHERE FILE_TYPE IN ('PF', 'LF');
```

### PRESTART\_JOB\_INFO view

The PRESTART\_JOB\_INFO view returns information about prestart jobs.

The values returned for the columns in the view are closely related to the values returned by the Display Prestart Job Entries panel accessed through the DSPSBSD (Display Subsystem Description) CL command and by the List Subsystem Entries (QWDLSE) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is PREJ\_INFO. The schema is QSYS2.

Table 217. PRESTART\_JOB\_INFO view

| Column Name                   | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|--------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10)             | The name of the library in which the subsystem description resides.                                                                                                                                                                                                                                                                                     |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10)             | The name of the subsystem about which information is being returned.                                                                                                                                                                                                                                                                                    |
| PRESTART_JOB_PROGRAM_LIBRARY  | PJ_PGM_LIB         | VARCHAR(10)             | The name of the library in which the prestart job program resides.                                                                                                                                                                                                                                                                                      |
| PRESTART_JOB_PROGRAM          | PJ_PGM             | VARCHAR(10)             | The program name that is used to match an incoming request with an available prestart job.                                                                                                                                                                                                                                                              |
| SUBSYSTEM_ACTIVE              | SBS_ACTIVE         | VARCHAR(3)              | Whether the subsystem for this prestart job is active.<br><br><b>NO</b> The status of the subsystem is not ACTIVE.<br><br><b>YES</b> The status of the subsystem is ACTIVE.                                                                                                                                                                             |
| USER_PROFILE                  | USRPRF             | VARCHAR(10)             | The name of the user profile under which the prestart job runs.                                                                                                                                                                                                                                                                                         |
| PRESTART_JOB_NAME             | PJ_NAME            | VARCHAR(10)             | The name of the prestart job.                                                                                                                                                                                                                                                                                                                           |
| JOB_DESCRIPTION_LIBRARY       | JOBDLIB            | VARCHAR(10)<br>Nullable | The name of the library in which the job description for the prestart job entry resides.<br>Contains the null value if JOB_DESCRIPTION is *USRPRF.                                                                                                                                                                                                      |
| JOB_DESCRIPTION               | JOB                | VARCHAR(10)             | The name of the job description that is used for the prestart job entry. Can contain the following special value:<br><br><b>*USRPRF</b> The job description that has the same name as the user profile that is used.                                                                                                                                    |
| START_JOBS                    | START_JOBS         | VARCHAR(3)              | Whether the prestart jobs are started at the time the subsystem is started.<br><br><b>NO</b> The prestart jobs are not started at the time the subsystem is started. The Start Prestart Jobs (STRPJ) command is used to start these prestart jobs.<br><br><b>YES</b> The prestart jobs are started at the time the subsystem is started.                |
| INITIAL_JOBS                  | INIT_JOBS          | INTEGER                 | The initial number of prestart jobs that are started when the subsystem is started.                                                                                                                                                                                                                                                                     |
| THRESHOLD                     | THRESHOLD          | INTEGER                 | The number at which additional prestart jobs are started. When the pool of available jobs (jobs available to service a program start request) is reduced below this number, more jobs (specified by ADDITIONAL_JOBS) are started and added to the available pool. This number is checked after a prestart job is attached to a procedure start request. |
| ADDITIONAL_JOBS               | ADD_JOBS           | INTEGER                 | The additional number of prestart jobs that are started when the number of prestart jobs drops below the value of THRESHOLD.                                                                                                                                                                                                                            |
| MAXIMUM_JOBS                  | MAX_JOBS           | INTEGER                 | The maximum number of prestart jobs that can be active at the same time for this prestart job entry. A value of -1 indicates *NOMAX.                                                                                                                                                                                                                    |

Table 217. PRESTART\_JOB\_INFO view (continued)

| Column Name    | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MAXIMUM_USES   | MAX_USES           | INTEGER                 | The maximum number of requests that can be handled by each prestart job in the pool before the job is ended. A value of -1 indicates *NOMAX.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| WAIT_FOR_JOB   | WAIT               | VARCHAR(3)              | Whether requests wait for a prestart job to become available or are rejected if a prestart job is not immediately available when the request is received.<br><br><b>NO</b> Requests are rejected if a prestart job is not immediately available when the request is received.<br><br><b>YES</b> Requests wait until there is an available prestart job, or until a prestart job is started, to handle the request.                                                                                                                                                                                                                                                                                                                                                                                                   |
| POOL_ID        | POOL_ID            | INTEGER                 | The subsystem pool in which the prestart jobs will run.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| CLASS_LIBRARY  | CLASS1_LIB         | VARCHAR(10)             | The name of the library in which the first class resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| CLASS          | CLASS1             | VARCHAR(10)             | The name of the first class that the prestart jobs run under. Jobs start by using the first class that is specified until the number of jobs specified for the first class is reached. After the number of jobs that are specified for the first class is reached, then jobs are started by using the second class.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| CLASS_JOBS     | CLASS1_JOB         | INTEGER                 | The maximum number of jobs to run using the first class. Can contain the following special values:<br><br><b>-3</b> The system calculates how many prestart jobs use this class.<br><br>If only one class is specified and -3 is specified, all of the jobs use that class.<br><br>If two classes are specified and -3 is specified for both, the first class is the value of MAXIMUM_JOBS divided by two, and the second class is the value of MAXIMUM_JOBS minus the value that is calculated for the first class.<br><br>If a specific number of jobs is specified for either class and -3 is specified for the other class, the system calculates the difference between MAXIMUM_JOBS and the specific number of jobs for the -3 designation.<br><br><b>-4</b> All of the prestart jobs use the specified class. |
| CLASS2_LIBRARY | CLASS2_LIB         | VARCHAR(10)<br>Nullable | The name of the library in which the second class resides.<br><br>Contains the null value if only one class is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| CLASS2         | CLASS2             | VARCHAR(10)<br>Nullable | The second class that the prestart jobs run under.<br><br>Contains the null value if only one class is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| CLASS2_JOBS    | CLASS2_JOB         | INTEGER<br>Nullable     | The maximum number of jobs to run using the second class. Can contain the following special values:<br><br><b>-3</b> The system calculates how many prestart jobs use this class.<br><br><b>-4</b> All of the prestart jobs use the specified class.<br><br>Contains the null value if only one class is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

Table 217. PRESTART\_JOB\_INFO view (continued)

| Column Name                     | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|--------------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RESOURCES_AFFINITY_GROUP        | RAG                | VARCHAR(3)             | <p>Specifies whether the prestart jobs started by this entry are grouped together having affinity to the same set of processors and memory.</p> <p><b>NO</b> Prestart jobs will not be grouped together. They will be spread across all the available system resources.</p> <p><b>YES</b> Prestart jobs will be grouped together such that they will have affinity to the same system resources.</p>                                                                                                                                                                                                                                                                                            |
| THREAD_RESOURCES_AFFINITY_GROUP | T_RAG              | VARCHAR(7)             | <p>Specifies whether secondary threads running in the prestart jobs are grouped together with the initial thread, or spread across the system resources.</p> <p><b>GROUP</b> Secondary threads running in the prestart job will all have affinity to the same set of processors and memory as the initial thread.</p> <p><b>NOGROUP</b> Secondary threads running in the prestart job will not necessarily have affinity to the same set of processors and memory as the initial thread. They will be spread across all the available system resources.</p> <p><b>SYSVAL</b> The thread resources affinity group and level in the QTHDRSCAFN system value will be used when the job starts.</p> |
| THREAD_RESOURCES_AFFINITY_LEVEL | T_RAL              | VARCHAR(6)<br>Nullable | <p>The degree to which the system tries to maintain the affinity between threads and system resources.</p> <p><b>HIGH</b> A thread will only use the resources it has affinity to, and will wait until they become available if necessary.</p> <p><b>NORMAL</b> A thread will use any processor or memory in the system if the resources it has affinity to are not readily available.</p> <p>Contains the null value when THREAD_RESOURCES_AFFINITY_GROUP is SYSVAL.</p>                                                                                                                                                                                                                       |

## Example

- List all the prestart job entries in the QUSRWRK subsystem.

```
SELECT *
FROM QSYS2.PRESTART_JOB_INFO
WHERE SUBSYSTEM_DESCRIPTION_LIBRARY = 'QSYS' AND
SUBSYSTEM_DESCRIPTION = 'QUSRWRK';
```


## PRESTART\_JOB\_STATISTICS table function

The PRESTART\_JOB\_STATISTICS table function returns statistics and performance information for an active prestart job entry in an active subsystem. The information is collected from the time the reset key is pressed on the Display Active Prestart Jobs (DSPACTPJ) command, from the time the prestart job entry is started, or from the time RESET\_STATISTICS => 'YES' is run for this table function. The prestart job entry is started when the subsystem starts or when the Start Prestart Jobs (STRPJ) command is used.

This information is similar to what is returned by the Display Active Prestart Jobs (DSPACTPJ) command and the Retrieve Active Prestart Jobs Status (QWTRAPJS) API.

**Authorization:** The caller must have:

- \*EXECUTE authority to the prestart job's program library.
- \*USE authority to the ASP device description, if the subsystem description specifies an ASP group name.



The schema is QSYS2.

|                                     |                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>subsystem-name</b>               | A character or graphic string expression that contains the name of the active subsystem that contains the prestart job entry.                                                                                                                                                                                                                                                                             |
| <b>prestart-job-program-library</b> | A character or graphic string expression that contains the name of the library where <i>prestart-job-program</i> is located. Can contain the following special values: <ul style="list-style-type: none"> <li><b>*CURLIB</b> The thread's current library is used to locate <i>prestart-job-program</i>.</li> <li><b>*LIBL</b> The library list is used to locate <i>prestart-job-program</i>.</li> </ul> |
| <b>prestart-job-program</b>         | A character or graphic string expression that contains the name of the program that identifies the active prestart job entry.                                                                                                                                                                                                                                                                             |
| <b>reset-statistics</b>             | A character or graphic string expression that contains a value of YES or NO to indicate whether the statistical information should be reset. If this parameter is not specified, the default is NO.<br><br>If the value is YES, the statistics are reset after the information is gathered. The reset applies to the statistics shown by the Display Active Prestart Jobs (DSPACTPJ) command as well.     |
| <b>ignore-errors</b>                | A character or graphic string expression that identifies what to do when an error is encountered. <ul style="list-style-type: none"> <li><b>NO</b> An error is returned.</li> <li><b>YES</b> A warning is returned.<br/>No row is returned when an error is encountered. This is the default.</li> </ul>                                                                                                  |

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 218. PRESTART\_JOB\_STATISTICS view

| Column Name                  | Data Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_NAME               | VARCHAR(10)   | The name of the active subsystem that contains the prestart job entry.                                                                                                                                                                                                                                                                                                                                                                                                |
| PRESTART_JOB_PROGRAM_LIBRARY | VARCHAR(10)   | The library where the prestart job program is located.                                                                                                                                                                                                                                                                                                                                                                                                                |
| PRESTART_JOB_PROGRAM         | VARCHAR(10)   | The program that identifies the active prestart job entry.                                                                                                                                                                                                                                                                                                                                                                                                            |
| STATUS_TIMESTAMP             | TIMESTAMP     | The date and time this active prestart jobs status information was collected.                                                                                                                                                                                                                                                                                                                                                                                         |
| ELAPSED_TIME                 | INTEGER       | The time that has elapsed, in seconds, since the last reset date/time.<br>The value will not reflect elapsed time beyond 10,000 hours. A value of 0 might mean that the elapsed time could not be calculated.                                                                                                                                                                                                                                                         |
| CURRENT_JOBS                 | INTEGER       | The current number of active jobs associated with the prestart job entry.                                                                                                                                                                                                                                                                                                                                                                                             |
| AVERAGE_JOBS                 | DECIMAL(10,1) | The average number of jobs that have been active since the reset date/time. This value is based on calculations involving time intervals and is inaccurate if the system clock was changed while information is being collected.<br>This value will be recalculated after statistics are reset.                                                                                                                                                                       |
| PEAK_JOBS                    | INTEGER       | The maximum number of jobs that have been active since the reset date/time.<br>This value will be recalculated after statistics are reset.                                                                                                                                                                                                                                                                                                                            |
| CURRENT_INUSE_JOBS           | INTEGER       | The current number of jobs in use.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| AVERAGE_INUSE_JOBS           | DECIMAL(10,1) | The average number of jobs in use since the reset date/time. This value is based on calculations involving time intervals and is inaccurate if the system clock was changed while information is being collected.<br>This value will be recalculated after statistics are reset.                                                                                                                                                                                      |
| PEAK_INUSE_JOBS              | INTEGER       | The maximum number of jobs in use since the reset date/time.<br>This value will be recalculated after statistics are reset.                                                                                                                                                                                                                                                                                                                                           |
| CURRENT_WAIT_REQUESTS        | INTEGER       | The number of requests that are waiting for a prestart job.                                                                                                                                                                                                                                                                                                                                                                                                           |
| AVERAGE_WAIT_REQUESTS        | DECIMAL(10,1) | The average number of requests that have waited for a prestart job since the reset date/time. This value is based on calculations involving time intervals and is inaccurate if the system clock was changed while information is being collected.<br>This value will be recalculated after statistics are reset.                                                                                                                                                     |
| PEAK_WAIT_REQUESTS           | INTEGER       | The maximum number of requests that have waited at one time for an available prestart job since the reset date/time.<br>This value will be recalculated after statistics are reset.                                                                                                                                                                                                                                                                                   |
| AVERAGE_WAIT_TIME            | DECIMAL(10,1) | The average time, in seconds, that requests have waited to be attached to a prestart job since the reset date/time. This includes requests that wait and requests that do not wait. This value is based on calculations involving time intervals and is inaccurate if the system clock was changed while information is being collected.<br>This value will be recalculated after statistics are reset.                                                               |
| ACCEPTED_REQUESTS            | INTEGER       | The total number of accepted requests since the reset date/time. An accepted request is one that is either attached immediately to a prestart job, or is queued because a prestart job is not available.<br>This value will be recalculated after statistics are reset.                                                                                                                                                                                               |
| REJECTED_REQUESTS            | INTEGER       | The total number of rejected requests since the reset date/time. A request is rejected if <ul style="list-style-type: none"> <li>*NO is specified for the WAIT parameter on the prestart job entry and there are no jobs currently available to handle the request.</li> <li>*YES is specified for the WAIT parameter and there are more program start requests than the maximum jobs allowed.</li> </ul> This value will be recalculated after statistics are reset. |

## Example

List the current statistics for all prestart jobs in the QUSRWRK subsystem. A warning will be issued for each prestart job returned from PRESTART\_JOB\_INFO that is not active.

```
SELECT PJ.PRESTART_JOB_NAME, STAT.*
FROM QSYS2.PRESTART_JOB_INFO PJ,
 LATERAL (
 SELECT * FROM TABLE
 (QSYS2.PRESTART_JOB_STATISTICS(PJ.SUBSYSTEM_DESCRIPTION,
 PJ.PRESTART_JOB_PROGRAM_LIBRARY,
 PJ.PRESTART_JOB_PROGRAM))) AS STAT
WHERE PJ.SUBSYSTEM_ACTIVE = 'YES' AND
 PJ.SUBSYSTEM_DESCRIPTION = 'QUSRWRK';
```

## RECORD\_LOCK\_INFO view

The RECORD\_LOCK\_INFO view returns one row for every record lock for the partition.

The values returned for the columns in the view are closely related to the values returned by [Retrieve Record Locks API](#). Refer to the APIs for more detailed information.

**Authorization:** The caller must have:

- \*EXECUTE authority to the library containing the database file, and
- \*OBJOPR and \*READ authority to the database file

The following table describes the columns in the view. The system name is RCD\_LOCK. The schema is QSYS2.

Table 219. RECORD\_LOCK\_INFO view

| Column Name            | System Column Name | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------|--------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TABLE_SCHEMA           | TABSCHEMA          | VARCHAR(128) | Name of the schema.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| TABLE_NAME             | TABNAME            | VARCHAR(128) | Name of the table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| TABLE_PARTITION        | TABPART            | VARCHAR(128) | Name of the table partition or member that contains the locked record.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| SYSTEM_TABLE_SCHEMA    | SYS_DNAME          | VARCHAR(10)  | System name of the schema.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| SYSTEM_TABLE_NAME      | SYS_TNAME          | VARCHAR(10)  | System name of the table                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| SYSTEM_TABLE_MEMBER    | SYS_MNAME          | VARCHAR(10)  | The name of the member that contains the locked record.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| RELATIVE_RECORD_NUMBER | RRN                | BIGINT       | The relative record number (RRN) of the record that is locked.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| LOCK_STATE             | LOCK_STATE         | VARCHAR(8)   | The lock condition for the record.<br><br><b>READ</b> The record is locked for read. Another job may read the same record but cannot lock the record for update intent. The record cannot be changed by another job as long as one job holds a read lock on the record.<br><br><b>UPDATE</b> The record is locked for update intent. Another job may read the record but may not obtain a read or update lock on it until the lock is released.<br><br><b>INTERNAL</b> The row is locked internally for read. For a short time the operating system holds an internal lock to access the row. Another job may read the same row and may even have the row locked for update intent. However, if another job does have the row locked for update intent, the actual change of the row will not proceed until the internal lock is released. |
| LOCK_STATUS            | STATUS             | VARCHAR(9)   | The status of the lock.<br><br><b>HELD</b> The lock is currently held by the job.<br><b>WAITING</b> The job is waiting for the lock.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |



Table 219. RECORD\_LOCK\_INFO view (continued)

| Column Name   | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|--------------------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LOCK_SCOPE    | LOCK_SCOPE         | VARCHAR(10)            | The scope of the lock. Values are: <ul style="list-style-type: none"> <li>• JOB</li> <li>• THREAD</li> <li>• LOCK SPACE</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                   |
| JOB_NAME      | JOB_NAME           | VARCHAR(28)            | The qualified job name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| THREAD_ID     | THREAD_ID          | BIGINT<br>Nullable     | The thread that is associated with the lock. <ul style="list-style-type: none"> <li>• If a held lock is job scoped, returns the null value. If a held lock is thread scoped, contains the identifier for the thread holding the lock.</li> <li>• If the scope of the lock is to the lock space and the lock is not held, contains the identifier of the thread requesting the lock.</li> <li>• If the lock is requested but not yet available, contains the identifier of the thread requesting the lock.</li> </ul> |
| LOCK_SPACE_ID | LOCKID             | BINARY(20)<br>Nullable | When the LOCK_SCOPE column value is LOCK SPACE and the lock is being waited on by a thread, contains the lock space ID value for which the lock is being waited on.<br><br>Otherwise, contains the null value.                                                                                                                                                                                                                                                                                                       |

## Example

Review the jobs that are updating the SALES table:

```
SELECT JOB_NAME, COUNT(*) AS ROWS_UPDATING
FROM QSYS2.RECORD_LOCK_INFO
WHERE SYSTEM_TABLE_NAME = 'SALES' AND
 SYSTEM_TABLE_SCHEMA = 'TOYSTORE' AND
 LOCK_STATE = 'UPDATE'
GROUP BY JOB_NAME
ORDER BY ROWS_UPDATING DESC
```

## ROUTING\_ENTRY\_INFO view

The ROUTING\_ENTRY\_INFO view returns information about routing entries.

The values returned for the columns in the view are closely related to the values returned by the Display Routing Entries panel accessed through the DSPSBSD (Display Subsystem Description) CL command and by the List Subsystem Entries (QWDLSE) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is RTG\_INFO. The schema is QSYS2.

Table 220. ROUTING\_ENTRY\_INFO view

| Column Name                   | System Column Name | Data Type               | Description                                                                                                                    |
|-------------------------------|--------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10)             | The name of the library in which the subsystem description resides.                                                            |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10)             | The name of the subsystem about which information is being returned.                                                           |
| SEQUENCE_NUMBER               | SEQNO              | INTEGER                 | The sequence number of the routing entry.                                                                                      |
| PROGRAM_LIBRARY               | PGM_LIB            | VARCHAR(10)<br>Nullable | The name of the library in which the routing entry program resides.<br><br>Contains the null value if PROGRAM_NAME is *RTGDTA. |

Table 220. ROUTING\_ENTRY\_INFO view (continued)

| Column Name              | System Column Name | Data Type           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|--------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PROGRAM_NAME             | PGM_NAME           | VARCHAR(10)         | <p>The name of the program that is started when a routing step is started through this routing entry. Can contain the following special value:</p> <p><b>*RTGDTA</b> The program name is taken from the routing data that was supplied and matched against this entry. The qualified program name will be taken from the routing data in this case, where the program name is specified in positions 37 through 46 and the library name is taken from positions 47 through 56.</p> |
| CLASS_LIBRARY            | CLASS_LIB          | VARCHAR(10)         | The name of the library in which the routing entry class resides.                                                                                                                                                                                                                                                                                                                                                                                                                  |
| CLASS                    | CLASS              | VARCHAR(10)         | The name of the class that is used when a routing step is started through this routing entry.                                                                                                                                                                                                                                                                                                                                                                                      |
| MAXIMUM_STEPS            | MAX_STEPS          | INTEGER<br>Nullable | <p>The maximum number of routing steps (jobs) that can be active at the same time through this routing entry.</p> <p>Contains the null value if the routing entry specifies *NOMAX, indicating that there is no maximum.</p>                                                                                                                                                                                                                                                       |
| POOL_ID                  | POOL_ID            | INTEGER             | The pool identifier of the storage pool in which the routing entry program is run.                                                                                                                                                                                                                                                                                                                                                                                                 |
| COMPARISON_DATA          | COMPDATA           | VARCHAR(80)         | <p>A value that is compared with the routing data to determine whether this is the routing entry that is used for starting a routing step. Can contain the following special value:</p> <p><b>*ANY</b> Any routing data is considered a match.</p>                                                                                                                                                                                                                                 |
| COMPARISON_START         | COMP_START         | INTEGER<br>Nullable | <p>The starting position for the routing data comparison. The comparison between the compare value and the routing data begins with this position in the routing data character string, and the last character position compared must be less than or equal to the length of the routing data used in the comparison.</p> <p>Contains the null value when COMPARISON_DATA is *ANY.</p>                                                                                             |
| RESOURCES_AFFINITY_GROUP | RAG                | VARCHAR(3)          | <p>Specifies whether jobs using this routing entry are grouped together having affinity to the same set of processors and memory.</p> <p><b>NO</b> The jobs will not be grouped together. They will be spread across all the available system resources.</p> <p><b>YES</b> The jobs will be grouped together such that they will have affinity to the same system resources.</p>                                                                                                   |

Table 220. ROUTING\_ENTRY\_INFO view (continued)

| Column Name                     | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| THREAD_RESOURCES_AFFINITY_GROUP | T_RAG              | VARCHAR(7)             | <p>Specifies whether secondary threads running in jobs that started through this routing entry are grouped together with the initial thread, or spread across the system resources.</p> <p><b>GROUP</b> Secondary threads will all have affinity to the same set of processors and memory as the initial thread.</p> <p><b>NOGROUP</b> Secondary threads will not necessarily have affinity to the same set of processors and memory as the initial thread. They will be spread across all the available system resources.</p> <p><b>SYSVAL</b> The thread resources affinity group and level in the QTHDRSCAFN system value will be used when the job starts.</p> |
| THREAD_RESOURCES_AFFINITY_LEVEL | T_RAL              | VARCHAR(6)<br>Nullable | <p>The degree to which the system tries to maintain the affinity between threads and system resources.</p> <p><b>HIGH</b> A thread will only use the resources it has affinity to, and will wait until they become available if necessary.</p> <p><b>NORMAL</b> A thread will use any processor or memory in the system if the resources it has affinity to are not readily available.</p> <p>Contains the null value when THREAD_RESOURCES_AFFINITY_GROUP is SYSVAL.</p>                                                                                                                                                                                          |

## Example

- List all the routing entries defined for the QBATCH subsystem.

```
SELECT *
FROM QSYS2.ROUTING_ENTRY_INFO
WHERE SUBSYSTEM_DESCRIPTION_LIBRARY = 'QSYS' AND
SUBSYSTEM_DESCRIPTION = 'QBATCH';
```

## SCHEDULED\_JOB\_INFO view

The SCHEDULED\_JOB\_INFO view returns information that can also be seen through the Work with Job Schedule Entries (WRKJOBSCDE) command interface. Each job schedule entry contains the information to automatically submit a batch job once or at regularly scheduled intervals.

**Authorization:** No authority is required to access scheduled job rows, but some columns return NULL if you don't have the required authority. You must have \*JOBCTL special authority or be the user profile listed in the SCHEDULED\_BY column to see the data in all columns.

The following table describes the columns in the view. The system name is SCHED\_JOB. The schema is QSYS2.

Table 221. SCHEDULED\_JOB\_INFO view

| Column Name                | System Column Name | Data Type | Description                                                                                |
|----------------------------|--------------------|-----------|--------------------------------------------------------------------------------------------|
| SCHEDULED_JOB_ENTRY_NUMBER | ENTRYNO            | INTEGER   | The number assigned to the job schedule entry when the entry is added to the job schedule. |

Table 221. SCHEDULED\_JOB\_INFO view (continued)

| Column Name            | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SCHEDULED_JOB_NAME     | SCDJOBNAME         | VARCHAR(10)             | The name of the job schedule entry.<br><br>This is the simple job name portion of the fully qualified job name used when the job is submitted. It is also used to identify the job schedule entry through change, hold, release and remove functions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| SCHEDULED_DATE_VALUE   | SCDDATEV           | VARCHAR(14)             | Indicates the date on which the job is scheduled to be submitted.<br><br><b>SCHEDULED_DATE</b> The date in the SCHEDULED_DATE column is used<br><br><b>SCHEDULED_DAYS</b> The days in the SCHEDULED_DAYS column are used<br><br><b>*MONTHSTR</b> The first day of the month is used.<br><br><b>*MONTHEND</b> The last day of the month is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| SCHEDULED_DATE         | SCDDATE            | DATE<br>Nullable        | The date on which the job is scheduled to be submitted.<br><br>Contains the null value if the SCHEDULED_DATE_VALUE column is not SCHEDULED_DATE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| SCHEDULED_TIME         | SCDTIME            | TIME                    | The time when the job is scheduled to be submitted on the scheduled date.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| SCHEDULED_DAYS         | SCDDAYS            | VARCHAR(34)<br>Nullable | The days on which the job is submitted if a specific date is not specified.<br><br>The value is a comma separated string with any or all of the values: *MON *TUE *WED *THU *FRI *SAT *SUN. The single value of *ALL can be returned to represent all seven values.<br><br>Contains the null value if SCHEDULED_DATE_VALUE is not SCHEDULED_DAYS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| FREQUENCY              | FREQUENCY          | VARCHAR(8)              | How often the job is to be submitted.<br><br><b>*ONCE</b> The job is scheduled to be submitted a single time.<br><br><b>*WEEKLY</b> The job is scheduled to be submitted on the same day or days of each week at the scheduled time.<br><br><b>*MONTHLY</b> The job is scheduled to be submitted on the same day or days of each month at the scheduled time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| RELATIVE_DAYS_OF_MONTH | RELDAYSMON         | VARCHAR(13)<br>Nullable | Specifies which occurrence during the month (for the days listed in the SCHEDULED_DAYS column) the job is scheduled to be run. The value is a comma separated string with up to five of the following values:<br><br><b>1</b> The job is scheduled for the first occurrence of the day or days (SCHEDULED_DAYS: *MON and *WED for example) of the month.<br><br><b>2</b> The job is scheduled for the second occurrence of the day or days of the month.<br><br><b>3</b> The job is scheduled for the third occurrence of the day or days of the month.<br><br><b>4</b> The job is scheduled for the fourth occurrence of the day or days of the month.<br><br><b>5</b> The job is scheduled for the fifth occurrence of the day or days of the month.<br><br><b>*LAST</b> The job is scheduled for the last occurrence of the day or days of the month.<br><br>Contains the null value if the FREQUENCY column does not have a value of MONTHLY or SCHEDULED_DAYS is null. |
| RECOVERY_ACTION        | RECOVERY           | VARCHAR(7)              | The recovery action taken when the system is powered down or in the restricted state at the time a job is scheduled to be submitted.<br><br><b>*SBMRLS</b> Submit a job to the job queue as a released job.<br><br><b>*SBMHLD</b> Submit a job to the job queue as a held job.<br><br><b>*NOSBM</b> Do not submit a job to the job queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 221. SCHEDULED\_JOB\_INFO view (continued)

| Column Name                    | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------|--------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NEXT_SUBMISSION_DATE           | NXTSUBDATE         | DATE<br>Nullable         | The next date that the job scheduling process is scheduled to submit this job.<br><br>Contains the null value if the job is not scheduled to be submitted again.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| STATUS                         | STATUS             | VARCHAR(9)               | The status of the job schedule entry.<br><br><b>HELD</b> The entry is held. If an entry has a status of HELD at the scheduled date and time, a job is not submitted.<br><br><b>SAVED</b> The entry is defined with a frequency of ONCE and a save value of *YES at a time later than the scheduled date and time.<br><br><b>SCHEDULED</b> The entry is waiting until the scheduled date and time for a job to be submitted.                                                                                                                                                                                                                                |
| JOB_QUEUE_NAME                 | JOBQ               | VARCHAR(10)              | The job queue to which the job is scheduled to be submitted. Can contain the special value of *JOBQ, meaning that the job is submitted to the job queue specified in the job description listed in the JOB_DESCRIPTION_NAME and JOB_DESCRIPTION_LIBRARY_NAME columns.                                                                                                                                                                                                                                                                                                                                                                                      |
| JOB_QUEUE_LIBRARY_NAME         | JOBQLIB            | VARCHAR(10)<br>Nullable  | The library containing the job queue.<br><br>Contains the null value if JOB_QUEUE_NAME is *JOBQ.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| JOB_QUEUE_STATUS               | JOBQSTATUS         | VARCHAR(10)<br>Nullable  | The status of the job queue.<br><br><b>HLD</b> The job queue is held, but not attached to an active subsystem.<br><br><b>HLD/SBS</b> The job queue is held and attached to an active subsystem.<br><br><b>LOCKED</b> The status of the job queue could not be determined because a lock could not be obtained on the job queue.<br><br><b>RLS</b> The job queue is released, but not attached to an active subsystem.<br><br><b>RLS/SBS</b> The job queue is released and attached to an active subsystem.<br><br>Contains the null value if JOB_QUEUE_NAME is *JOBQ, if the job queue is not found or is damaged, or if the information is not available. |
| DATES_OMITTED                  | OMITDATES          | VARCHAR(219)<br>Nullable | A comma separated string with up to 20 dates in *ISO format indicating dates when the job will not be scheduled to run.<br><br>Contains the null value if no dates were specified to omit or if the information is not available.                                                                                                                                                                                                                                                                                                                                                                                                                          |
| SCHEDULED_BY                   | CREATEDBY          | VARCHAR(10)              | The user profile of the job which added the entry to the job schedule.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| DESCRIPTION                    | TEXT               | VARCHAR(50)<br>Nullable  | The descriptive text for the job schedule entry.<br><br>Contains the null value if the job schedule entry has no description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| COMMAND_STRING                 | COMMAND            | VARCHAR(512)<br>Nullable | The command that is run in the submitted job.<br><br>Contains the null value if the information is not available.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| USER_PROFILE_FOR_SUBMITTED_JOB | SBMJOBUSR          | VARCHAR(10)<br>Nullable  | The user profile to be used when the job is submitted. Can contain the special value *JOBQ to indicate that the user profile from the job description is used.<br><br>Contains the null value if the information is not available.                                                                                                                                                                                                                                                                                                                                                                                                                         |
| JOB_DESCRIPTION_NAME           | JOBQ               | VARCHAR(10)<br>Nullable  | The job description used when the job is submitted. Can contain the special value of *USRPRF to indicate that the job description specified in the user profile under which the submitted job runs is used.<br><br>Contains the null value if the information is not available.                                                                                                                                                                                                                                                                                                                                                                            |

Table 221. SCHEDULED\_JOB\_INFO view (continued)

| Column Name                          | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------|--------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOB_DESCRIPTION_LIBRARY_NAME         | JOBDLIB            | VARCHAR(10)<br>Nullable  | The library containing the job description.<br><br>Contains the null value if JOB_DESCRIPTION_NAME has a value of *USRPRF or if the information is not available.                                                                                                                                                                                                                                                                                                                                                                              |
| MESSAGE_QUEUE_NAME                   | MSGQ               | VARCHAR(10)<br>Nullable  | The name of the message queue where the messages for this job schedule entry are sent. Can contain the special value *USRPRF to indicate that the message queue specified in the user profile under which the submitted job runs is used.<br><br>Contains the null value if no specific message queue is associated with this job schedule entry or if the information is not available.                                                                                                                                                       |
| MESSAGE_QUEUE_LIBRARY_NAME           | MSGQLIB            | VARCHAR(10)<br>Nullable  | The library containing the message queue.<br><br>Contains the null value if MESSAGE_QUEUE_NAME is null, contains the special value of *USRPRF, or if the information is not available.                                                                                                                                                                                                                                                                                                                                                         |
| LAST_SUCCESSFUL_SUBMISSION_TIMESTAMP | SBMTIMSTMP         | TIMESTAMP(0)<br>Nullable | The timestamp when a batch job was last successfully submitted for the job schedule entry.<br><br>Contains the null value if the job schedule entry has not been used to submit a job.                                                                                                                                                                                                                                                                                                                                                         |
| LAST_SUCCESSFUL_SUBMISSION_JOB       | LASTSBMJOB         | VARCHAR(28)<br>Nullable  | The qualified job name used when this scheduled job was last submitted.<br><br>Contains the null value if the scheduled job has never been submitted or if the information is not available.                                                                                                                                                                                                                                                                                                                                                   |
| LAST_ATTEMPTED_SUBMISSION_TIMESTAMP  | ATTSBMTIM          | TIMESTAMP(0)<br>Nullable | The timestamp when this scheduled job was last submitted.<br><br>Contains the null value if the scheduled job has never been submitted or if the information is not available.                                                                                                                                                                                                                                                                                                                                                                 |
| LAST_ATTEMPTED_SUBMISSION_STATUS     | SBMJOBSTS          | VARCHAR(68)<br>Nullable  | The status from when this scheduled job was last submitted. Values are: <ul style="list-style-type: none"> <li>• JOB SUCCESSFULLY SUBMITTED</li> <li>• LAST JOB SUBMISSION FAILED, CHECK THE JOB MESSAGE QUEUE FOR DETAILS</li> <li>• JOB NOT SUBMITTED DUE TO HELD STATUS</li> <li>• JOB SUBMITTED AFTER SCHEDULED TIME AS SPECIFIED BY RECOVERY ACTION</li> <li>• JOB NOT SUBMITTED AS SPECIFIED BY RECOVERY ACTION</li> </ul> Contains the null value if the scheduled job has never been submitted or if the information is not available. |
| KEEP_ENTRY                           | KEEP               | VARCHAR(3)<br>Nullable   | Whether the job schedule entry is kept or removed after the job has been submitted. <p><b>YES</b>     The job schedule entry is kept.</p> <p><b>NO</b>        The job schedule entry is removed.</p> Contains the null value when the FREQUENCY column does not contain *ONCE or if the information is not available.                                                                                                                                                                                                                          |

## Example

Review the job scheduled entries which are no longer in effect, either because they were explicitly held or because they were scheduled to run a single time and the scheduled date and time has passed.

```
SELECT * FROM QSYS2.SCHEDULED_JOB_INFO WHERE STATUS IN ('HELD', 'SAVED')
ORDER BY SCHEDULED_BY;
```

## SUBSYSTEM\_INFO view

The SUBSYSTEM\_INFO view returns information about all subsystems.

The values returned for the columns in the view are closely related to the values returned by the WRKSBS (Work with Subsystems) CL command and by the Retrieve Subsystem Information (QWDRSBSD) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is SBS\_INFO. The schema is QSYS2.

Table 222. SUBSYSTEM\_INFO view

| Column Name                   | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10)             | The name of the library in which the subsystem description resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10)             | The name of the subsystem about which information is being returned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| STATUS                        | STATUS             | VARCHAR(10)             | The status of the subsystem.<br><br><b>ACTIVE</b> The subsystem is running.<br><b>ENDING</b> An ENDSBS command has been issued for the subsystem or an ENDSYS command has been issued, but the subsystem is still running.<br><b>INACTIVE</b> The subsystem is not running.<br><b>RESTRICTED</b> An ENDSBS command for the controlling subsystem, an ENDSYS *ALL command, or an ENDSYS command has placed the controlling subsystem in a restricted condition.<br><b>STARTING</b> A STRSBS command has been issued for the subsystem, but it is still in the process of being started. |
| MAXIMUM_ACTIVE_JOBS           | MAX_ACT            | INTEGER<br>Nullable     | The maximum number of jobs that can run or use resources in the subsystem at one time.<br><br>Contains the null value if the subsystem description specifies *NOMAX, indicating that there is no maximum.                                                                                                                                                                                                                                                                                                                                                                              |
| CURRENT_ACTIVE_JOBS           | CUR_ACT            | INTEGER                 | The number of jobs currently active in the subsystem. This number includes held jobs but excludes jobs that are disconnected or suspended because of a transfer secondary job or a transfer group job. If STATUS is INACTIVE, returns 0.                                                                                                                                                                                                                                                                                                                                               |
| SUBSYSTEM_MONITOR_JOB         | SBSMONJOB          | VARCHAR(28)<br>Nullable | The qualified job name for the subsystem monitor job as identified to the system.<br><br>Contains the null value if STATUS is INACTIVE.                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| TEXT_DESCRIPTION              | TEXT               | VARCHAR(50)<br>Nullable | The text description of the subsystem description.<br><br>Contains the null value if there is no text description.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| CONTROLLING_SUBSYSTEM         | CTL_SBS            | VARCHAR(3)              | Whether this subsystem is the controlling subsystem.<br><br><b>NO</b> This subsystem is not the controlling subsystem.<br><b>YES</b> This subsystem is the controlling subsystem.                                                                                                                                                                                                                                                                                                                                                                                                      |
| WORKLOAD_GROUP                | WRK_GROUP          | VARCHAR(10)<br>Nullable | The name of the workload group used for jobs started in this subsystem.<br><br>Contains the null value if there is no workload group defined for the subsystem.                                                                                                                                                                                                                                                                                                                                                                                                                        |
| SIGNON_DEVICE_FILE_LIBRARY    | DEVFILELIB         | VARCHAR(10)             | The name of the library in which the sign-on device file resides.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| SIGNON_DEVICE_FILE            | DEVFILE            | VARCHAR(10)             | The name of the sign-on device file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 222. SUBSYSTEM\_INFO view (continued)

| Column Name                | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|--------------------|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SECONDARY_LANGUAGE_LIBRARY | LANG_LIB           | VARCHAR(10)<br>Nullable | The name of the subsystem's secondary language library.<br><br>Contains the null value if there is no secondary language library defined for the subsystem.                                                                                                                                                                                                                                                                                              |
| IASP_NAME                  | IASP_NAME          | VARCHAR(10)<br>Nullable | The name of the auxiliary storage pool (ASP) group associated with the subsystem monitor job. This is the name of the primary ASP device in an ASP group or the name of an ASP device description.<br><br>Contains the null value if the subsystem does not have an auxiliary storage pool (ASP) group associated with it. Only the libraries in the system ASP and any basic user ASPs will be in the library name space for the subsystem monitor job. |

## Example

- Show information for all active subsystems.

```
SELECT *
FROM QSYS2.SUBSYSTEM_INFO
WHERE STATUS = 'ACTIVE' ;
```

## SUBSYSTEM\_POOL\_INFO view

The SUBSYSTEM\_POOL\_INFO view returns information about storage pools defined for subsystems.

The values returned for the columns in the view are closely related to the values returned by the WRKSBS (Work with Subsystems) CL command and by the Retrieve Subsystem Information (QWDRSBS) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is SBS\_POOL. The schema is QSYS2.

Table 223. SUBSYSTEM\_POOL\_INFO view

| Column Name                   | System Column Name | Data Type   | Description                                                               |
|-------------------------------|--------------------|-------------|---------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10) | The name of the library in which the subsystem description resides.       |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10) | The name of the subsystem about which pool information is being returned. |
| POOL_ID                       | POOL_ID            | INTEGER     | The pool ID for the subsystem pool.                                       |



Table 223. SUBSYSTEM\_POOL\_INFO view (continued)

| Column Name         | System Column Name | Data Type                 | Description                                                                                                                                                               |                                                                                                                                                                                                                                                                |
|---------------------|--------------------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| POOL_NAME           | POOL_NAME          | VARCHAR(10)               | The name of the subsystem pool.                                                                                                                                           |                                                                                                                                                                                                                                                                |
|                     |                    |                           | <b>*BASE</b>                                                                                                                                                              | The system base pool, which can be shared with other subsystems. The QBASPOOL system value defines the base pool's minimum size. The base pool contains all main storage not allocated to other pools. The QBASACTLVL system value defines its activity level. |
|                     |                    |                           | <b>*INTERACT</b>                                                                                                                                                          | The shared pool used for interactive work.                                                                                                                                                                                                                     |
|                     |                    |                           | <b>*NOSTG</b>                                                                                                                                                             | No storage size or activity level is assigned to this storage pool.                                                                                                                                                                                            |
|                     |                    |                           | <b>*SHRPOOL1-<br/>*SHRPOOL60</b>                                                                                                                                          | Shared pools.                                                                                                                                                                                                                                                  |
|                     |                    |                           | <b>*SPOOL</b>                                                                                                                                                             | The shared pool for spooling writers.                                                                                                                                                                                                                          |
|                     | <b>*USERPOOL</b>   | The pool is user-defined. |                                                                                                                                                                           |                                                                                                                                                                                                                                                                |
| MAXIMUM_ACTIVE_JOBS | MAX_JOBS           | INTEGER<br>Nullable       | If the pool name is *USERPOOL, the maximum number of jobs that can be active in the pool at one time.<br><br>Contains the null value if the pool name is not *USERPOOL.   |                                                                                                                                                                                                                                                                |
| POOL_SIZE           | POOL_SIZE          | DECIMAL(20,2)<br>Nullable | If the pool name is *USERPOOL, the amount of storage, in megabytes, that the pool attempts to allocate.<br><br>Contains the null value if the pool name is not *USERPOOL. |                                                                                                                                                                                                                                                                |

## Example

- List all the system storage pools in the QBATCH subsystem.

```
SELECT *
FROM QSYS2.SUBSYSTEM_POOL_INFO
WHERE SUBSYSTEM_DESCRIPTION_LIBRARY = 'QSYS' AND
SUBSYSTEM_DESCRIPTION = 'QBATCH' AND
POOL_NAME <> '*USERPOOL';
```

## SYSTEM\_ACTIVITY\_INFO table function

The SYSTEM\_ACTIVITY\_INFO table function returns a single row containing statistical information about CPU usage.

The information returned is similar to the detail provided by the Work with System Activity (WRKSYSACT) CL command.

**Authorization:** The caller must have \*JOBCTL special authority

►► SYSTEM\_ACTIVITY\_INFO ( ( \_\_\_\_\_ ) ) ◀◀  
└──────────────────────────┘ *delay-seconds*  
└───────────┘ DELAY\_SECONDS — =>

The schema is QSYS2.

**delay-seconds** An integer value that specifies an interval of time to wait between statistical readings. Detail gathered from two points in time is used to calculate the table function results.

**integer** A delay of *integer* seconds separates the collection of details. *integer* must be a value of 1 or greater. One second is the default.

No response from the table function is returned until at least *integer* seconds have elapsed.

The result of the function is a table containing a single row with the format shown in the following table. All the columns are nullable.

Table 224. SYSTEM\_ACTIVITY\_INFO table function

| Column Name             | Data Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AVERAGE_CPU_RATE        | DECIMAL(20,2) | The average CPU rate expressed as a percentage where 100% indicates the processor is running at its nominal frequency. A value above or below 100% indicates how much the processor has been slowed down (throttled) or speeded up (turbo) relative to the nominal frequency for the processor model. For instance, a value of 120% indicates the processor is running 20% faster against its nominal speed. |
| AVERAGE_CPU_UTILIZATION | DECIMAL(20,2) | The average CPU utilization for all the active processors.                                                                                                                                                                                                                                                                                                                                                   |
| MINIMUM_CPU_UTILIZATION | DECIMAL(20,2) | The CPU utilization of the processor that reported the minimum amount of CPU utilization.                                                                                                                                                                                                                                                                                                                    |
| MAXIMUM_CPU_UTILIZATION | DECIMAL(20,2) | The CPU utilization of the processor that reported the maximum amount of CPU utilization.                                                                                                                                                                                                                                                                                                                    |

### Example

Return statistical CPU information. Use the default one second interval.

```
SELECT * FROM TABLE(QSYS2.SYSTEM_ACTIVITY_INFO());
```


### SYSTEM\_STATUS table function

The SYSTEM\_STATUS table function returns a single row containing details about the current partition.

The information returned is similar to the detail seen from the Work with System Status (WRKSYSSTS) and the Work with System Activity (WRKSYSACT) commands, and information available in the HMC and with the Retrieve Partition Information (dlpar\_get\_info) API.

**Authorization:** Non-null values are returned for the columns from PARTITION\_NAME through UNUSED\_CPU\_TIME\_SHARED\_POOL for callers with \*USE authority on the QSYS/QPMLPMGT service program.

To return values for TEMPORARY\_JOB\_STRUCTURES\_AVAILABLE and PERMANENT\_JOB\_STRUCTURES\_AVAILABLE, callers must have \*JOBCTL special authority and \*USE authority on the QSYS/QWTCTJBS program.



The schema is QSYS2.

**reset-statistics** A character or graphic string expression that contains a value of YES or NO.

If this parameter has a value of YES, statistics are reset such that the time of this query execution is used as the new baseline. The columns that contain this statistical data have names that are prefixed with ELAPSED\_. Future invocations of SYSTEM\_STATUS within this connection will return statistical detail relative to the new baseline. If this parameter has a value of NO, statistics are not reset for the invocation. If this parameter is not specified, the default is NO.

**detailed\_info** A character or graphic string expression that indicates the type of information to be returned.

**ALL** All available system information is returned.

**BASIC** Only the basic system information is returned. This is all the columns except for the 7 job table columns. This is the default.

The result of the function is a table containing multiple rows with the format shown in the following table. All the columns are nullable.

Table 225. SYSTEM\_STATUS table function

| Column Name                   | Data Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOTAL_JOBS_IN_SYSTEM          | INTEGER       | The total number of user and system jobs that are currently in the system. The total includes: <ul style="list-style-type: none"> <li>All jobs on job queues waiting to be processed.</li> <li>All jobs currently active (being processed).</li> <li>All jobs that have completed running but still have output on output queues to be produced.</li> </ul>                                                                                          |
| MAXIMUM_JOBS_IN_SYSTEM        | INTEGER       | The maximum number of jobs that are allowed on the system. When the number of jobs reaches this maximum, you can no longer submit or start more jobs on the system. The total includes: <ul style="list-style-type: none"> <li>All jobs on job queues waiting to be processed.</li> <li>All jobs currently active (being processed).</li> <li>All jobs that have completed running but still have output on output queues to be produced.</li> </ul> |
| ACTIVE_JOBS_IN_SYSTEM         | INTEGER       | The number of jobs active in the system (jobs that have been started, but have not yet ended), including both user and system jobs.                                                                                                                                                                                                                                                                                                                  |
| INTERACTIVE_JOBS_IN_SYSTEM    | DECIMAL(10,2) | The percentage of interactive performance assigned to this logical partition. This value is a percentage of the total interactive performance available to the entire physical system.                                                                                                                                                                                                                                                               |
| ELAPSED_TIME                  | INTEGER       | The time that has elapsed, in seconds, between the measurement start time and the current system time.                                                                                                                                                                                                                                                                                                                                               |
| ELAPSED_CPU_USED              | DECIMAL(10,2) | The average of the elapsed time during which the processing units were in use.                                                                                                                                                                                                                                                                                                                                                                       |
| ELAPSED_CPU_SHARED            | DECIMAL(10,2) | The percentage of the total shared processor pool capacity used by all partitions using the pool during the elapsed time.<br>Returns null if this is a dedicated partition.                                                                                                                                                                                                                                                                          |
| ELAPSED_CPU_UNCAPPED_CAPACITY | DECIMAL(10,2) | The percentage of the uncapped shared processing capacity for the partition used since the last time statistics were reset.<br>Returns null if this partition cannot use more than its configured processing capacity.                                                                                                                                                                                                                               |
| CONFIGURED_CPUS               | INTEGER       | Total number of configured CPUs for the partition.                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 225. SYSTEM\_STATUS table function (continued)

| Column Name                    | Data Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU_SHARING_ATTRIBUTE          | VARCHAR(8)    | <p>This attribute indicates whether this partition is sharing processors. If the value indicates the partition does not share physical processors, then this partition uses only dedicated processors. If the value indicates the partition shares physical processors, then this partition uses physical processors from a shared pool of physical processors.</p> <p><b>CAPPED</b> Partition shares processors. The partition is limited to using its configured capacity.</p> <p><b>UNCAPPED</b> Partition shares processors. The partition can use more than its configured capacity.</p> <p>Contains the null value if this is a dedicated partition.</p> |
| CURRENT_CPU_CAPACITY           | DECIMAL(10,2) | The current processing capacity specifies the processor units that are being used in the partition. For a partition sharing physical processors, the current processing capacity represents the share of the physical processors in the pool it is running. For a partition using dedicated processors, the current processing capacity represents the number of virtual processors that are currently active in the partition.                                                                                                                                                                                                                                |
| AVERAGE_CPU_RATE               | DECIMAL(20,2) | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: <a href="#">“SYSTEM_ACTIVITY_INFO table function” on page 863.</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| AVERAGE_CPU_UTILIZATION        | DECIMAL(20,2) | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: <a href="#">“SYSTEM_ACTIVITY_INFO table function” on page 863.</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MINIMUM_CPU_UTILIZATION        | DECIMAL(20,2) | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: <a href="#">“SYSTEM_ACTIVITY_INFO table function” on page 863.</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAXIMUM_CPU_UTILIZATION        | DECIMAL(20,2) | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: <a href="#">“SYSTEM_ACTIVITY_INFO table function” on page 863.</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| SQL_CPU_UTILIZATION            | DECIMAL(10,2) | Always contains the null value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAIN_STORAGE_SIZE              | BIGINT        | The amount of main storage, in kilobytes, in the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| SYSTEM_ASP_STORAGE             | BIGINT        | The storage capacity of the system auxiliary storage pool (ASP number 1) in millions of bytes. This value represents the amount of space available for storage of both permanent and temporary objects.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| TOTAL_AUXILIARY_STORAGE        | BIGINT        | The total auxiliary storage, in millions of bytes, on the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| SYSTEM_ASP_USED                | DECIMAL(10,2) | The percentage of the system storage pool (ASP number 1) currently in use.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| CURRENT_TEMPORARY_STORAGE      | INTEGER       | The current amount of storage, in millions of bytes, in use for temporary objects.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MAXIMUM_TEMPORARY_STORAGE_USED | INTEGER       | The largest amount of storage, in millions of bytes, used for temporary objects at any one time since the last IPL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| PERMANENT_ADDRESS_RATE         | DECIMAL(6,3)  | The percentage of the maximum possible addresses for permanent objects that have been used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 225. SYSTEM\_STATUS table function (continued)

| Column Name                        | Data Type     | Description                                                                                                                                                                                                                                                     |
|------------------------------------|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TEMPORARY_ADDRESS_RATE             | DECIMAL(6,3)  | The percentage of the maximum possible addresses for temporary objects that have been used.                                                                                                                                                                     |
| TEMPORARY_256MB_SEGMENTS           | DECIMAL(10,2) | The percentage of the maximum possible temporary 256MB segments that have been used.                                                                                                                                                                            |
| TEMPORARY_4GB_SEGMENTS             | DECIMAL(10,2) | The percentage of the maximum possible temporary 4GB segments that have been used.                                                                                                                                                                              |
| PERMANENT_256MB_SEGMENTS           | DECIMAL(10,2) | The percentage of the maximum possible permanent 256MB segments that have been used.                                                                                                                                                                            |
| PERMANENT_4GB_SEGMENTS             | DECIMAL(10,2) | The percentage of the maximum possible permanent 4GB segments that have been used.                                                                                                                                                                              |
| TEMPORARY_JOB_STRUCTURES_AVAILABLE | INTEGER       | The number of temporary job structures that currently exist on the system that are not in use.<br><br>Returns the null value if the user does not have *USE authority on the QSYS/QWTCTJBS program.                                                             |
| PERMANENT_JOB_STRUCTURES_AVAILABLE | INTEGER       | The number of permanent job structures that currently exist on the system that are not in use.<br><br>Returns the null value if the user does not have *USE authority on the QSYS/QWTCTJBS program.                                                             |
| TOTAL_JOB_TABLE_ENTRIES            | INTEGER       | The total number of job table entries. This includes AVAILABLE_JOB_TABLE_ENTRIES and IN_USE_JOB_TABLE_ENTRIES.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC.                                                                        |
| AVAILABLE_JOB_TABLE_ENTRIES        | INTEGER       | The total number of job table entries that are available.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC.                                                                                                                             |
| IN_USE_JOB_TABLE_ENTRIES           | INTEGER       | The total number of job table entries that are in use. This includes ACTIVE_JOB_TABLE_ENTRIES, JOBQ_JOB_TABLE_ENTRIES, OUTQ_JOB_TABLE_ENTRIES, and JOBLLOG_PENDING_JOB_TABLE_ENTRIES.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC. |
| ACTIVE_JOB_TABLE_ENTRIES           | INTEGER       | The total number of entries that are in use by active jobs.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC.                                                                                                                           |
| JOBQ_JOB_TABLE_ENTRIES             | INTEGER       | The total number of entries that are in use by jobs on a JOBQ.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC.                                                                                                                        |
| OUTQ_JOB_TABLE_ENTRIES             | INTEGER       | The total number of entries that are in use by jobs that have ended but have spooled output still attached to the job.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC.                                                                |
| JOBLLOG_PENDING_JOB_TABLE_ENTRIES  | INTEGER       | The total number of entries that are in use by jobs that have ended but have a pending job log.<br><br>Contains the null value when the DETAILED_INFO parameter is BASIC.                                                                                       |
| HOST_NAME                          | VARCHAR(255)  | Name of the system where this information was generated. This is the name set by CHGNETA.                                                                                                                                                                       |
| PARTITION_ID                       | INTEGER       | The identifier for the partition in which this view is being run.                                                                                                                                                                                               |
| NUMBER_OF_PARTITIONS               | INTEGER       | The number of partitions on the physical machine. This includes partitions that are currently powered on (running) and partitions that are powered off.                                                                                                         |

Table 225. SYSTEM\_STATUS table function (continued)

| Column Name                                                                                                                                                                                           | Data Type                  | Description                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ACTIVE_THREADS_IN_SYSTEM                                                                                                                                                                              | INTEGER                    | The number of initial and secondary threads in the system (threads that have been started, but have not yet ended), including both user and system threads.                                                                                                                                                |
| RESTRICTED_STATE                                                                                                                                                                                      | VARCHAR(3)                 | Whether the system is in restricted state.<br><br><b>NO</b> System is not in restricted state.<br><b>YES</b> System is in restricted state.                                                                                                                                                                |
| <b>The columns from PARTITION_NAME through UNUSED_CPU_TIME_SHARED_POOL require the user to have *USE authority on the QSYS/QPMLPMGT service program. Otherwise, they will contain the null value.</b> |                            |                                                                                                                                                                                                                                                                                                            |
| PARTITION_NAME                                                                                                                                                                                        | VARGRAPHIC(255) CCSID 1200 | The name of the partition as it is known to the HMC.                                                                                                                                                                                                                                                       |
| PARTITION_GROUP_ID                                                                                                                                                                                    | INTEGER                    | The LPAR group for this partition.                                                                                                                                                                                                                                                                         |
| SHARED_PROCESSOR_POOL_ID                                                                                                                                                                              | INTEGER                    | The shared processor pool this partition is a member of. A shared processor pool is a set of physical processors on the physical machine that is used to run a set of shared processor partitions. A value of 0 indicates the default pool.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES. |
| DEFINED_MEMORY                                                                                                                                                                                        | BIGINT                     | The amount of memory (in megabytes) that was configured for this partition through the HMC.                                                                                                                                                                                                                |
| MINIMUM_MEMORY                                                                                                                                                                                        | BIGINT                     | The minimum amount of main storage (in megabytes) that can be assigned to this partition.                                                                                                                                                                                                                  |
| MAXIMUM_MEMORY                                                                                                                                                                                        | BIGINT                     | The maximum amount of main storage (in megabytes) that can be assigned to this partition.                                                                                                                                                                                                                  |
| MEMORY_INCREMENT                                                                                                                                                                                      | BIGINT                     | The smallest amount of main storage (in megabytes) that can be added to or removed from this partition's memory.                                                                                                                                                                                           |
| DEDICATED_PROCESSORS                                                                                                                                                                                  | VARCHAR(3)                 | Whether the partition uses dedicated processors.<br><br><b>NO</b> The partition does not use dedicated processors.<br><b>YES</b> The partition uses dedicated processors.                                                                                                                                  |
| PHYSICAL_PROCESSORS                                                                                                                                                                                   | INTEGER                    | The number of physical processors in this physical machine that are available for use. This does not include processors on demand that have not been turned on.                                                                                                                                            |
| PHYSICAL_PROCESSORS_SHARED_POOL                                                                                                                                                                       | INTEGER                    | The number of physical processors that are allocated to the shared processor pool used by this partition.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES.                                                                                                                                   |
| MAXIMUM_PHYSICAL_PROCESSORS                                                                                                                                                                           | INTEGER                    | The maximum number of physical processors that can be active in this physical machine without installing additional processors. This value includes currently active processors and any standby (on demand) processors that are present in this physical machine.                                          |
| DEFINED_VIRTUAL_PROCESSORS                                                                                                                                                                            | INTEGER                    | The number of virtual processors configured for this partition through the HMC.                                                                                                                                                                                                                            |
| VIRTUAL_PROCESSORS                                                                                                                                                                                    | INTEGER                    | The number of virtual processors currently used by this partition.                                                                                                                                                                                                                                         |
| MINIMUM_VIRTUAL_PROCESSORS                                                                                                                                                                            | INTEGER                    | The minimum number of virtual processors that can be configured for this partition.                                                                                                                                                                                                                        |
| MAXIMUM_VIRTUAL_PROCESSORS                                                                                                                                                                            | INTEGER                    | The maximum number of virtual processors that can be configured for this partition.                                                                                                                                                                                                                        |
| DEFINED_PROCESSING_CAPACITY                                                                                                                                                                           | DECIMAL(5,2)               | The amount of processing capacity that was configured for this partition through the HMC.                                                                                                                                                                                                                  |

Table 225. SYSTEM\_STATUS table function (continued)

| Column Name                          | Data Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PROCESSING_CAPACITY                  | DECIMAL(5,2) | The current (usable) amount of processing capacity available to the partition (also known as partition's entitled capacity).                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| UNALLOCATED_PROCESSING_CAPACITY      | DECIMAL(5,2) | The amount of processing capacity in the partition group this partition belongs to, which is not allocated to any partition and is available for allocation.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| MINIMUM_REQUIRED_PROCESSING_CAPACITY | DECIMAL(5,2) | The minimum amount of processing capacity that the operating system in this partition requires for its operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| MAXIMUM_LICENSED_PROCESSING_CAPACITY | DECIMAL(5,2) | The current limit on processing capacity of this partition imposed by the operating system software license for this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| MINIMUM_PROCESSING_CAPACITY          | DECIMAL(5,2) | The minimum amount of processing capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| MAXIMUM_PROCESSING_CAPACITY          | DECIMAL(5,2) | The maximum amount of processing capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| PROCESSING_CAPACITY_INCREMENT        | DECIMAL(5,2) | The smallest capacity that can be added to or removed from this partition's processing capacity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| DEFINED_INTERACTIVE_CAPACITY         | DECIMAL(5,2) | The amount of interactive capacity that was configured for this partition through the HMC. A partition's interactive capacity is defined as this partition's portion of total interactive capacity of the physical machine.                                                                                                                                                                                                                                                                                                                                                                            |
| INTERACTIVE_CAPACITY                 | DECIMAL(5,2) | This partition's current (usable) portion of the physical machine interactive capacity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| INTERACTIVE_THRESHOLD                | DECIMAL(5,2) | The maximum interactive CPU utilization which can be sustained in this partition, without causing a disproportionate increase in system overhead.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| UNALLOCATED_INTERACTIVE_CAPACITY     | DECIMAL(5,2) | The amount of interactive capacity in the partition group this partition belongs to, which is not allocated to any partition and is available for allocation. Interactive capacity is defined as the portion of total interactive capacity of the physical machine.                                                                                                                                                                                                                                                                                                                                    |
| MINIMUM_INTERACTIVE_CAPACITY         | DECIMAL(5,2) | The minimum portion of the physical machine's interactive capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MAXIMUM_INTERACTIVE_CAPACITY         | DECIMAL(5,2) | The maximum portion of the physical machine's interactive capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| DEFINED_VARIABLE_CAPACITY_WEIGHT     | INTEGER      | The weighting factor that was configured for this partition through the HMC. Variable capacity weight is used for uncapped partitions when they compete for unused CPU cycles in the shared pool. Variable capacity weight can be in the range of 0 - 255. The larger the weight, the more the chance this partition will get additional CPU cycles from the shared pool.<br><br>Contains the null value if the DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED.                                                                                                                     |
| VARIABLE_CAPACITY_WEIGHT             | INTEGER      | The weighting factor that is used to assign additional unused CPU cycles (from the shared processor pool) to the partition. Variable capacity weight is used for uncapped partitions when they compete for unused CPU cycles in the shared pool. This factor is in the range of 0 - 255. The larger the weight, the greater the chance this partition will get additional CPU cycles from the pool. A value of 0 effectively caps this partition at its current (usable) processing capacity.<br><br>Contains the null value if the DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED. |

Table 225. SYSTEM\_STATUS table function (continued)

| Column Name                          | Data Type     | Description                                                                                                                                                                                                                                                                                                          |
|--------------------------------------|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNALLOCATED_VARIABLE_CAPACITY_WEIGHT | INTEGER       | The amount of capacity weight that is available for allocation to the partition's variable capacity weight.<br><br>Contains the null value if the DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED.                                                                                                 |
| HARDWARE_MULTITHREADING              | VARCHAR(3)    | Indicates whether hardware multi-threading is enabled.<br><br><b>NO</b> Hardware multi-threading is not enabled.<br><b>YES</b> Hardware multi-threading is enabled.                                                                                                                                                  |
| BOUND_HARDWARE_THREADS               | VARCHAR(3)    | Whether hardware threads are bound.<br><br><b>NO</b> Hardware threads are not bound.<br><b>YES</b> Hardware threads are bound.                                                                                                                                                                                       |
| THREADS_PER_PROCESSOR                | INTEGER       | The number of hardware threads per processor when hardware multi-threading is enabled.<br><br>Contains the null value if HARDWARE_MULTITHREADING is NO.                                                                                                                                                              |
| DISPATCH_LATENCY                     | DECIMAL(20,0) | The maximum time in nanoseconds between dispatches of this partition on a physical processor.                                                                                                                                                                                                                        |
| DISPATCH_WHEEL_ROTATION_TIME         | DECIMAL(20,0) | The number of nanoseconds in the hypervisor's scheduling window. Each virtual processor will be given the opportunity to execute on a physical processor some time during this period. The amount of time each virtual processor is able to use a physical processor is determined by partition processing capacity. |
| TOTAL_CPU_TIME                       | DECIMAL(20,0) | The number of nanoseconds of CPU time used by this partition since IPL.                                                                                                                                                                                                                                              |
| INTERACTIVE_CPU_TIME                 | DECIMAL(20,0) | The amount of CPU time, in nanoseconds, used by interactive processes in this partition since partition IPL. An interactive process is any process doing 5250 display device I/O.                                                                                                                                    |
| INTERACTIVE_CPU_TIME_ABOVE_THRESHOLD | DECIMAL(20,0) | The amount of CPU time, in nanoseconds, used by interactive processes while exceeding the interactive threshold. This is a total since IPL.                                                                                                                                                                          |
| UNUSED_CPU_TIME_SHARED_POOL          | DECIMAL(20,0) | The number of nanoseconds of CPU time that the physical processors in a shared processor pool have been idle since system IPL.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if the partition is not authorized to retrieve shared pool data.                                                     |
| MACHINE_TYPE                         | CHAR(4)       | The machine type.                                                                                                                                                                                                                                                                                                    |
| MACHINE_MODEL                        | CHAR(4)       | The machine model.                                                                                                                                                                                                                                                                                                   |
| SERIAL_NUMBER                        | CHAR(8)       | The machine serial number.                                                                                                                                                                                                                                                                                           |
| ATTENTION_LIGHT                      | VARCHAR(3)    | The status of the system attention light.<br><br><b>OFF</b> The light is off.<br><b>ON</b> The light is on.                                                                                                                                                                                                          |



Table 225. SYSTEM\_STATUS table function (continued)

| Column Name             | Data Type  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPL_MODE                | VARCHAR(9) | <p>The current IPL mode setting.</p> <p><b>AUTOMATIC</b> Used for automatic remote IPL, automatic IPL by date and time, and automatic IPL after a power failure.</p> <p><b>MANUAL</b> An operator uses the control panel to direct the system for special needs.</p> <p><b>NORMAL</b> Requires no operator intervention during the IPL.</p> <p><b>SECURE</b> Prevents use of the control panel to perform an IPL.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| IPL_TYPE                | CHAR(1)    | <p>Type of IPL performed.</p> <p><b>A</b> Used for special work, such as applying fixes (PTFs) and diagnostic work.</p> <p><b>B</b> Used for routine work and when directed by a PTF procedure.</p> <p><b>C</b> Reserved for system support.</p> <p><b>D</b> Used for special work, such as installing and reloading programs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| JOURNAL_RECOVERY_COUNT  | INTEGER    | <p>Specifies the system wide default journal recovery count. The journal recovery count allows you to choose between faster runtime processing of changes to journaled objects and faster IPL or vary on recovery after an abnormal shutdown. The value specified influences the frequency with which journaled objects are forced to auxiliary storage as those objects are changed. The specified journal recovery count indicates the approximate number of journaled changes that would need to be recovered during journal synchronization for this journal in the event of an IPL or vary on after an abnormal shutdown. A smaller value decreases the number of changes that would need to be recovered from this journal by increasing the frequency with which changed objects are forced to disk. A larger value increases the runtime processing of changes to journaled objects by decreasing the frequency with which changed objects are forced to disk. Changing this value may affect overall system performance as it affects the utilization of auxiliary storage devices.</p> <p>This value can be changed with the Change Journal Attributes (CHGJRNA) CL command.</p> <p>The system default for this value is 250,000.</p> |
| JOURNAL_CACHE_WAIT_TIME | INTEGER    | <p>The cache wait time, in seconds, for journal environments with caching enabled. The cache wait time is the maximum number of seconds that the system will wait before writing any lingering journal entries from main memory to disk.</p> <p>This value can be changed with the Change Journal Attributes (CHGJRNA) CL command.</p> <p>The system default for this value is 30 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Example

Return storage and CPU status for the partition. Specify to reset all the elapsed values to 0.

```
SELECT * FROM TABLE(QSYS2.SYSTEM_STATUS(RESET_STATISTICS=>'YES')) X;
```

## SYSTEM\_STATUS\_INFO view

The SYSTEM\_STATUS\_INFO view returns a single row containing details about the current partition. This view uses the QSYS2.SYSTEM\_STATUS table function with DETAILED\_INFO => 'ALL'.

For better performance, use SYSTEM\_STATUS\_INFO\_BASIC except when the additional job table columns returned by SYSTEM\_STATUS\_INFO are needed by the query.

The information returned is similar to the detail seen from the Work with System Status (WRKSYSSTS) and the Work with System Activity (WRKSYSACT) commands, information available in the HMC and with the Retrieve Partition Information (dlpar\_get\_info) API, and journal attribute values. It does not reset the statistical columns; to do this, use the associated table function, "[SYSTEM\\_STATUS table function](#)" on page 864.

**Authorization:** Non-null values are returned for the columns from PARTITION\_NAME through UNUSED\_CPU\_TIME\_SHARED\_POOL for callers with \*USE authority on the QSYS/QPMLPMGT service program.

To return values for TEMPORARY\_JOB\_STRUCTURES\_AVAILABLE and PERMANENT\_JOB\_STRUCTURES\_AVAILABLE, callers must have \*JOBCTL special authority and \*USE authority on the QSYS/QWTCTJBS program.

The following table describes the columns in the view. The system name is SYS\_STATUS. The schema is QSYS2.

Table 226. SYSTEM\_STATUS\_INFO view

| Column Name                   | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------|--------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOTAL_JOBS_IN_SYSTEM          | TOTAL_JOBS         | INTEGER                  | The total number of user and system jobs that are currently in the system. The total includes: <ul style="list-style-type: none"> <li>All jobs on job queues waiting to be processed.</li> <li>All jobs currently active (being processed).</li> <li>All jobs that have completed running but still have output on output queues to be produced.</li> </ul>                                                                                          |
| MAXIMUM_JOBS_IN_SYSTEM        | MAX_JOBS           | INTEGER                  | The maximum number of jobs that are allowed on the system. When the number of jobs reaches this maximum, you can no longer submit or start more jobs on the system. The total includes: <ul style="list-style-type: none"> <li>All jobs on job queues waiting to be processed.</li> <li>All jobs currently active (being processed).</li> <li>All jobs that have completed running but still have output on output queues to be produced.</li> </ul> |
| ACTIVE_JOBS_IN_SYSTEM         | ACT_JOBS           | INTEGER                  | The number of jobs active in the system (jobs that have been started, but have not yet ended), including both user and system jobs.                                                                                                                                                                                                                                                                                                                  |
| INTERACTIVE_JOBS_IN_SYSTEM    | INTER_JOBS         | DECIMAL(5,2)             | The percentage of interactive performance assigned to this logical partition. This value is a percentage of the total interactive performance available to the entire physical system.                                                                                                                                                                                                                                                               |
| ELAPSED_TIME                  | ELAP_TIME          | INTEGER                  | The time that has elapsed, in seconds, between the measurement start time and the current system time.                                                                                                                                                                                                                                                                                                                                               |
| ELAPSED_CPU_USED              | ELAP_USED          | DECIMAL(5,2)             | The average of the elapsed time during which the processing units were in use.                                                                                                                                                                                                                                                                                                                                                                       |
| ELAPSED_CPU_SHARED            | ELAP_SHARE         | DECIMAL(5,2)<br>Nullable | The percentage of the total shared processor pool capacity used by all partitions using the pool during the elapsed time.<br>Returns null if this is a dedicated partition.                                                                                                                                                                                                                                                                          |
| ELAPSED_CPU_UNCAPPED_CAPACITY | ELAP_UNCAP         | DECIMAL(5,2)<br>Nullable | The percentage of the uncapped shared processing capacity for the partition used since the last time statistics were reset.<br>Returns null if this partition cannot use more than its configured processing capacity.                                                                                                                                                                                                                               |
| CONFIGURED_CPUS               | CONFIGCPUS         | INTEGER                  | Total number of configured CPUs for the partition.                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 226. SYSTEM\_STATUS\_INFO view (continued)

| Column Name                    | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CPU_SHARING_ATTRIBUTE          | CPU_SHARE          | VARCHAR(8)<br>Nullable   | <p>This attribute indicates whether this partition is sharing processors. If the value indicates the partition does not share physical processors, then this partition uses only dedicated processors. If the value indicates the partition shares physical processors, then this partition uses physical processors from a shared pool of physical processors.</p> <p><b>CAPPED</b> Partition shares processors. The partition is limited to using its configured capacity.</p> <p><b>UNCAPPED</b> Partition shares processors. The partition can use more than its configured capacity.</p> <p>Contains the null value if this is a dedicated partition.</p> |
| CURRENT_CPU_CAPACITY           | CPU_CAP            | DECIMAL(5,2)             | The current processing capacity specifies the processor units that are being used in the partition. For a partition sharing physical processors, the current processing capacity represents the share of the physical processors in the pool it is running. For a partition using dedicated processors, the current processing capacity represents the number of virtual processors that are currently active in the partition.                                                                                                                                                                                                                                |
| AVERAGE_CPU_RATE               | CPU_RATE           | DECIMAL(5,2)             | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: "SYSTEM_ACTIVITY_INFO table function" on page 863.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| AVERAGE_CPU_UTILIZATION        | CPU_AVG            | DECIMAL(5,2)             | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: "SYSTEM_ACTIVITY_INFO table function" on page 863.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MINIMUM_CPU_UTILIZATION        | CPU_MIN            | DECIMAL(5,2)             | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: "SYSTEM_ACTIVITY_INFO table function" on page 863.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAXIMUM_CPU_UTILIZATION        | CPU_MAX            | DECIMAL(5,2)             | <p>Always returns 0.</p> <p>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: "SYSTEM_ACTIVITY_INFO table function" on page 863.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| SQL_CPU_UTILIZATION            | CPU_SQL            | DECIMAL(5,2)<br>Nullable | Always contains the null value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| MAIN_STORAGE_SIZE              | MAIN_STG           | BIGINT                   | The amount of main storage, in kilobytes, in the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| SYSTEM_ASP_STORAGE             | SYS_STG            | BIGINT                   | The storage capacity of the system auxiliary storage pool (ASP number 1) in millions of bytes. This value represents the amount of space available for storage of both permanent and temporary objects.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| TOTAL_AUXILIARY_STORAGE        | AUX_STG            | BIGINT                   | The total auxiliary storage, in millions of bytes, on the system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| SYSTEM_ASP_USED                | SYS_RATE           | DECIMAL(5,2)             | The percentage of the system storage pool (ASP number 1) currently in use.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| CURRENT_TEMPORARY_STORAGE      | TEMP_CUR           | INTEGER                  | The current amount of storage, in millions of bytes, in use for temporary objects.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| MAXIMUM_TEMPORARY_STORAGE_USED | TEMP_MAX           | INTEGER                  | The largest amount of storage, in millions of bytes, used for temporary objects at any one time since the last IPL.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| PERMANENT_ADDRESS_RATE         | PERM_RATE          | DECIMAL(6,3)             | The percentage of the maximum possible addresses for permanent objects that have been used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| TEMPORARY_ADDRESS_RATE         | TEMP_RATE          | DECIMAL(6,3)             | The percentage of the maximum possible addresses for temporary objects that have been used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| TEMPORARY_256MB_SEGMENTS       | TEMP_256MB         | DECIMAL(5,2)             | The percentage of the maximum possible temporary 256MB segments that have been used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| TEMPORARY_4GB_SEGMENTS         | TEMP_4GB           | DECIMAL(5,2)             | The percentage of the maximum possible temporary 4GB segments that have been used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

Table 226. SYSTEM\_STATUS\_INFO view (continued)

| Column Name                                                                                                                                                                                           | System Column Name | Data Type                              | Description                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PERMANENT_256MB_SEGMENTS                                                                                                                                                                              | PERM_256MB         | DECIMAL(5,2)                           | The percentage of the maximum possible permanent 256MB segments that have been used.                                                                                                                                                                                                                       |
| PERMANENT_4GB_SEGMENTS                                                                                                                                                                                | PERM_4GB           | DECIMAL(5,2)                           | The percentage of the maximum possible permanent 4GB segments that have been used.                                                                                                                                                                                                                         |
| TEMPORARY_JOB_STRUCTURES_AVAILABLE                                                                                                                                                                    | TEMP_JS            | INTEGER<br>Nullable                    | The number of temporary job structures that currently exist on the system that are not in use.<br><br>Returns the null value if the user does not have *USE authority on the QSYS/QWTCTJBS program.                                                                                                        |
| PERMANENT_JOB_STRUCTURES_AVAILABLE                                                                                                                                                                    | PERM_JS            | INTEGER<br>Nullable                    | The number of permanent job structures that currently exist on the system that are not in use.<br><br>Returns the null value if the user does not have *USE authority on the QSYS/QWTCTJBS program.                                                                                                        |
| TOTAL_JOB_TABLE_ENTRIES                                                                                                                                                                               | TOTAL_JOB_T        | INTEGER                                | The total number of job table entries.<br>This includes AVAILABLE_JOB_TABLE_ENTRIES and IN_USE_JOB_TABLE_ENTRIES.                                                                                                                                                                                          |
| AVAILABLE_JOB_TABLE_ENTRIES                                                                                                                                                                           | AVAIL_JOB_T        | INTEGER                                | The total number of job table entries that are available.                                                                                                                                                                                                                                                  |
| IN_USE_JOB_TABLE_ENTRIES                                                                                                                                                                              | INUSE_JOB_T        | INTEGER                                | The total number of job table entries that are in use. This includes ACTIVE_JOB_TABLE_ENTRIES, JOBQ_JOB_TABLE_ENTRIES, OUTQ_JOB_TABLE_ENTRIES, and JOBLG_PENDING_JOB_TABLE_ENTRIES.                                                                                                                        |
| ACTIVE_JOB_TABLE_ENTRIES                                                                                                                                                                              | ACT_JOB_T          | INTEGER                                | The total number of entries that are in use by active jobs.                                                                                                                                                                                                                                                |
| JOBQ_JOB_TABLE_ENTRIES                                                                                                                                                                                | JOBQ_JOB_T         | INTEGER                                | The total number of entries that are in use by jobs on a JOBQ.                                                                                                                                                                                                                                             |
| OUTQ_JOB_TABLE_ENTRIES                                                                                                                                                                                | OUTQ_JOB_T         | INTEGER                                | The total number of entries that are in use by jobs that have ended but have spooled output still attached to the job.                                                                                                                                                                                     |
| JOBLG_PENDING_JOB_TABLE_ENTRIES                                                                                                                                                                       | PEND_JOB_T         | INTEGER                                | The total number of entries that are in use by jobs that have ended but have a pending job log.                                                                                                                                                                                                            |
| HOST_NAME                                                                                                                                                                                             | HOST_NAME          | VARCHAR(255)                           | Name of the system where this information was generated. This is the name set by CHGNETA.                                                                                                                                                                                                                  |
| PARTITION_ID                                                                                                                                                                                          | PART_ID            | INTEGER                                | The identifier for the partition in which this view is being run.                                                                                                                                                                                                                                          |
| NUMBER_OF_PARTITIONS                                                                                                                                                                                  | NUM_PART           | INTEGER                                | The number of partitions on the physical machine. This includes partitions that are currently powered on (running) and partitions that are powered off.                                                                                                                                                    |
| ACTIVE_THREADS_IN_SYSTEM                                                                                                                                                                              | ACT_THREAD         | INTEGER                                | The number of initial and secondary threads in the system (threads that have been started, but have not yet ended), including both user and system threads.                                                                                                                                                |
| RESTRICTED_STATE                                                                                                                                                                                      | REST_STATE         | VARCHAR(3)                             | Whether the system is in restricted state.<br><br><b>NO</b> System is not in restricted state.<br><b>YES</b> System is in restricted state.                                                                                                                                                                |
| <b>The columns from PARTITION_NAME through UNUSED_CPU_TIME_SHARED_POOL require the user to have *USE authority on the QSYS/QPMLPMGT service program. Otherwise, they will contain the null value.</b> |                    |                                        |                                                                                                                                                                                                                                                                                                            |
| PARTITION_NAME                                                                                                                                                                                        | PART_NAME          | VARGRAPHIC(255) CCSID 1200<br>Nullable | The name of the partition as it is known to the HMC.                                                                                                                                                                                                                                                       |
| PARTITION_GROUP_ID                                                                                                                                                                                    | PART_GROUP         | INTEGER<br>Nullable                    | The LPAR group for this partition.                                                                                                                                                                                                                                                                         |
| SHARED_PROCESSOR_POOL_ID                                                                                                                                                                              | POOL_ID            | INTEGER<br>Nullable                    | The shared processor pool this partition is a member of. A shared processor pool is a set of physical processors on the physical machine that is used to run a set of shared processor partitions. A value of 0 indicates the default pool.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES. |
| DEFINED_MEMORY                                                                                                                                                                                        | DEF_MEM            | BIGINT<br>Nullable                     | The amount of memory (in megabytes) that was configured for this partition through the HMC.                                                                                                                                                                                                                |

Table 226. SYSTEM\_STATUS\_INFO view (continued)

| Column Name                          | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                       |
|--------------------------------------|--------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MINIMUM_MEMORY                       | MIN_MEM            | BIGINT<br>Nullable       | The minimum amount of main storage (in megabytes) that can be assigned to this partition.                                                                                                                                                                         |
| MAXIMUM_MEMORY                       | MAX_MEM            | BIGINT<br>Nullable       | The maximum amount of main storage (in megabytes) that can be assigned to this partition.                                                                                                                                                                         |
| MEMORY_INCREMENT                     | MEM_INCR           | BIGINT<br>Nullable       | The smallest amount of main storage (in megabytes) that can be added to or removed from this partition's memory.                                                                                                                                                  |
| DEDICATED_PROCESSORS                 | DED_PROC           | VARCHAR(3)<br>Nullable   | Whether the partition uses dedicated processors.<br><b>NO</b> The partition does not use dedicated processors.<br><b>YES</b> The partition uses dedicated processors.                                                                                             |
| PHYSICAL_PROCESSORS                  | PHY_PROC           | INTEGER<br>Nullable      | The number of physical processors in this physical machine that are available for use. This does not include processors on demand that have not been turned on.                                                                                                   |
| PHYSICAL_PROCESSORS_SHARED_POOL      | PHY_SHARE          | INTEGER<br>Nullable      | The number of physical processors that are allocated to the shared processor pool used by this partition.<br>Contains the null value if DEDICATED_PROCESSORS is YES.                                                                                              |
| MAXIMUM_PHYSICAL_PROCESSORS          | MAX_PHY            | INTEGER<br>Nullable      | The maximum number of physical processors that can be active in this physical machine without installing additional processors. This value includes currently active processors and any standby (on demand) processors that are present in this physical machine. |
| DEFINED_VIRTUAL_PROCESSORS           | DEF_VIRT           | INTEGER<br>Nullable      | The number of virtual processors configured for this partition through the HMC.                                                                                                                                                                                   |
| VIRTUAL_PROCESSORS                   | VIRT_PROC          | INTEGER<br>Nullable      | The number of virtual processors currently used by this partition.                                                                                                                                                                                                |
| MINIMUM_VIRTUAL_PROCESSORS           | MIN_VIRT           | INTEGER<br>Nullable      | The minimum number of virtual processors that can be configured for this partition.                                                                                                                                                                               |
| MAXIMUM_VIRTUAL_PROCESSORS           | MAX_VIRT           | INTEGER<br>Nullable      | The maximum number of virtual processors that can be configured for this partition.                                                                                                                                                                               |
| DEFINED_PROCESSING_CAPACITY          | DEF_CAP            | DECIMAL(5,2)<br>Nullable | The amount of processing capacity that was configured for this partition through the HMC.                                                                                                                                                                         |
| PROCESSING_CAPACITY                  | CAPACITY           | DECIMAL(5,2)<br>Nullable | The current (usable) amount of processing capacity available to the partition (also known as partition's entitled capacity).                                                                                                                                      |
| UNALLOCATED_PROCESSING_CAPACITY      | AVAIL_CAP          | DECIMAL(5,2)<br>Nullable | The amount of processing capacity in the partition group this partition belongs to, which is not allocated to any partition and is available for allocation.                                                                                                      |
| MINIMUM_REQUIRED_PROCESSING_CAPACITY | MIN_REQCAP         | DECIMAL(5,2)<br>Nullable | The minimum amount of processing capacity that the operating system in this partition requires for its operation.                                                                                                                                                 |
| MAXIMUM_LICENSED_PROCESSING_CAPACITY | MAX_LICCAP         | DECIMAL(5,2)<br>Nullable | The current limit on processing capacity of this partition imposed by the operating system software license for this partition.                                                                                                                                   |
| MINIMUM_PROCESSING_CAPACITY          | MIN_CAP            | DECIMAL(5,2)<br>Nullable | The minimum amount of processing capacity that can be assigned to this partition.                                                                                                                                                                                 |
| MAXIMUM_PROCESSING_CAPACITY          | MAX_CAP            | DECIMAL(5,2)<br>Nullable | The maximum amount of processing capacity that can be assigned to this partition.                                                                                                                                                                                 |
| PROCESSING_CAPACITY_INCREMENT        | CAP_INCR           | DECIMAL(5,2)<br>Nullable | The smallest capacity that can be added to or removed from this partition's processing capacity.                                                                                                                                                                  |

Table 226. SYSTEM\_STATUS\_INFO view (continued)

| Column Name                          | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------|--------------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DEFINED_INTERACTIVE_CAPACITY         | DEF_INTCAP         | DECIMAL(5,2)<br>Nullable  | The amount of interactive capacity that was configured for this partition through the HMC. A partition's interactive capacity is defined as this partition's portion of total interactive capacity of the physical machine.                                                                                                                                                                                                                                                                                                                                                                        |
| INTERACTIVE_CAPACITY                 | INT_CAP            | DECIMAL(5,2)<br>Nullable  | This partition's current (usable) portion of the physical machine interactive capacity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| INTERACTIVE_THRESHOLD                | INT_THRESH         | DECIMAL(5,2)<br>Nullable  | The maximum interactive CPU utilization which can be sustained in this partition, without causing a disproportionate increase in system overhead.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| UNALLOCATED_INTERACTIVE_CAPACITY     | AVL_INTCAP         | DECIMAL(5,2)<br>Nullable  | The amount of interactive capacity in the partition group this partition belongs to, which is not allocated to any partition and is available for allocation. Interactive capacity is defined as the portion of total interactive capacity of the physical machine.                                                                                                                                                                                                                                                                                                                                |
| MINIMUM_INTERACTIVE_CAPACITY         | MIN_INTCAP         | DECIMAL(5,2)<br>Nullable  | The minimum portion of the physical machine's interactive capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| MAXIMUM_INTERACTIVE_CAPACITY         | MAX_INTCAP         | DECIMAL(5,2)<br>Nullable  | The maximum portion of the physical machine's interactive capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| DEFINED_VARIABLE_CAPACITY_WEIGHT     | DEF_CAPW           | INTEGER<br>Nullable       | The weighting factor that was configured for this partition through the HMC. Variable capacity weight is used for uncapped partitions when they compete for unused CPU cycles in the shared pool. Variable capacity weight can be in the range of 0 - 255. The larger the weight, the more the chance this partition will get additional CPU cycles from the shared pool.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED.                                                                                                                     |
| VARIABLE_CAPACITY_WEIGHT             | VAR_CAPW           | INTEGER<br>Nullable       | The weighting factor that is used to assign additional unused CPU cycles (from the shared processor pool) to the partition. Variable capacity weight is used for uncapped partitions when they compete for unused CPU cycles in the shared pool. This factor is in the range of 0 - 255. The larger the weight, the greater the chance this partition will get additional CPU cycles from the pool. A value of 0 effectively caps this partition at its current (usable) processing capacity.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED. |
| UNALLOCATED_VARIABLE_CAPACITY_WEIGHT | AVAIL_CAPW         | INTEGER<br>Nullable       | The amount of capacity weight that is available for allocation to the partition's variable capacity weight.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED.                                                                                                                                                                                                                                                                                                                                                                                   |
| HARDWARE_MULTITHREADING              | HW_MLT_THR         | VARCHAR(3)<br>Nullable    | Indicates whether hardware multi-threading is enabled.<br><br><b>NO</b> Hardware multi-threading is not enabled.<br><b>YES</b> Hardware multi-threading is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| BOUND_HARDWARE_THREADS               | HW_BND_THR         | VARCHAR(3)<br>Nullable    | Whether hardware threads are bound.<br><br><b>NO</b> Hardware threads are not bound.<br><b>YES</b> Hardware threads are bound.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| THREADS_PER_PROCESSOR                | THREADS_PP         | INTEGER<br>Nullable       | The number of hardware threads per processor when hardware multi-threading is enabled.<br><br>Contains the null value if HARDWARE_MULTITHREADING is NO.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| DISPATCH_LATENCY                     | LATENCY            | DECIMAL(20,0)<br>Nullable | The maximum time in nanoseconds between dispatches of this partition on a physical processor.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| DISPATCH_WHEEL_ROTATION_TIME         | DISPATCH_T         | DECIMAL(20,0)<br>Nullable | The number of nanoseconds in the hypervisor's scheduling window. Each virtual processor will be given the opportunity to execute on a physical processor some time during this period. The amount of time each virtual processor is able to use a physical processor is determined by partition processing capacity.                                                                                                                                                                                                                                                                               |

Table 226. SYSTEM\_STATUS\_INFO view (continued)

| Column Name                          | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------|--------------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOTAL_CPU_TIME                       | CPU_TOTAL          | DECIMAL(20,0)<br>Nullable | The number of nanoseconds of CPU time used by this partition since IPL.                                                                                                                                                                                                                                                                                                                                        |
| INTERACTIVE_CPU_TIME                 | CPU_INTER          | DECIMAL(20,0)<br>Nullable | The amount of CPU time, in nanoseconds, used by interactive processes in this partition since partition IPL. An interactive process is any process doing 5250 display device I/O.                                                                                                                                                                                                                              |
| INTERACTIVE_CPU_TIME_ABOVE_THRESHOLD | CPU_THRESH         | DECIMAL(20,0)<br>Nullable | The amount of CPU time, in nanoseconds, used by interactive processes while exceeding the interactive threshold. This is a total since IPL.                                                                                                                                                                                                                                                                    |
| UNUSED_CPU_TIME_SHARED_POOL          | CPU_UNUSED         | DECIMAL(20,0)<br>Nullable | The number of nanoseconds of CPU time that the physical processors in a shared processor pool have been idle since system IPL.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if the partition is not authorized to retrieve shared pool data.                                                                                                                                               |
| MACHINE_TYPE                         | MACH_TYPE          | CHAR(4)                   | The machine type.                                                                                                                                                                                                                                                                                                                                                                                              |
| MACHINE_MODEL                        | MACH_MOD           | CHAR(4)                   | The machine model.                                                                                                                                                                                                                                                                                                                                                                                             |
| SERIAL_NUMBER                        | SERIAL             | CHAR(8)                   | The machine serial number.                                                                                                                                                                                                                                                                                                                                                                                     |
| ATTENTION_LIGHT                      | ATTN_LIGHT         | VARCHAR(3)                | The status of the system attention light.<br><br><b>OFF</b> The light is off.<br><b>ON</b> The light is on.                                                                                                                                                                                                                                                                                                    |
| IPL_MODE                             | IPL_MODE           | VARCHAR(9)                | The current IPL mode setting.<br><br><b>AUTOMATIC</b> Used for automatic remote IPL, automatic IPL by date and time, and automatic IPL after a power failure.<br><br><b>MANUAL</b> An operator uses the control panel to direct the system for special needs.<br><br><b>NORMAL</b> Requires no operator intervention during the IPL.<br><br><b>SECURE</b> Prevents use of the control panel to perform an IPL. |
| IPL_TYPE                             | IPL_TYPE           | CHAR(1)                   | Type of IPL performed.<br><br><b>A</b> Used for special work, such as applying fixes (PTFs) and diagnostic work.<br><b>B</b> Used for routine work and when directed by a PTF procedure.<br><b>C</b> Reserved for system support.<br><b>D</b> Used for special work, such as installing and reloading programs.                                                                                                |

Table 226. SYSTEM\_STATUS\_INFO view (continued)

| Column Name             | System Column Name | Data Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|--------------------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| JOURNAL_RECOVERY_COUNT  | JRNRCYCNT          | INTEGER   | <p>Specifies the system wide default journal recovery count. The journal recovery count allows you to choose between faster runtime processing of changes to journaled objects and faster IPL or vary on recovery after an abnormal shutdown. The value specified influences the frequency with which journaled objects are forced to auxiliary storage as those objects are changed. The specified journal recovery count indicates the approximate number of journaled changes that would need to be recovered during journal synchronization for this journal in the event of an IPL or vary on after an abnormal shutdown. A smaller value decreases the number of changes that would need to be recovered from this journal by increasing the frequency with which changed objects are forced to disk. A larger value increases the runtime processing of changes to journaled objects by decreasing the frequency with which changed objects are forced to disk. Changing this value may affect overall system performance as it affects the utilization of auxiliary storage devices.</p> <p>This value can be changed with the Change Journal Attributes (CHGJRNA) CL command.</p> <p>The system default for this value is 250,000.</p> |
| JOURNAL_CACHE_WAIT_TIME | CACHEWAIT          | INTEGER   | <p>The cache wait time, in seconds, for journal environments with caching enabled. The cache wait time is the maximum number of seconds that the system will wait before writing any lingering journal entries from main memory to disk.</p> <p>This value can be changed with the Change Journal Attributes (CHGJRNA) CL command.</p> <p>The system default for this value is 30 seconds.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Examples

- Review the state of the job table entries for the partition.

```
SELECT TOTAL_JOB_TABLE_ENTRIES, AVAILABLE_JOB_TABLE_ENTRIES,
 IN_USE_JOB_TABLE_ENTRIES, ACTIVE_JOB_TABLE_ENTRIES,
 JOBQ_JOB_TABLE_ENTRIES, OUTQ_JOB_TABLE_ENTRIES,
 JOBLLOG_PENDING_JOB_TABLE_ENTRIES
FROM QSYS2.SYSTEM_STATUS_INFO;
```

## SYSTEM\_STATUS\_INFO\_BASIC view

The SYSTEM\_STATUS\_INFO\_BASIC view returns a single row containing details about the current partition. This view uses the QSYS2.SYSTEM\_STATUS table function with DETAILED\_INFO => 'BASIC'.

This view contains the same information as SYSTEM\_STATUS\_INFO except it excludes the job table columns.

The information returned is similar to the detail seen from the Work with System Status (WRKSYSSTS) and the Work with System Activity (WRKSYSACT) commands, information available in the HMC and with the Retrieve Partition Information (dlpar\_get\_info) API, and journal attribute values. It does not reset the statistical columns; to do this, use the associated table function, "[SYSTEM\\_STATUS table function](#)" on page 864.

**Authorization:** Non-null values are returned for the columns from PARTITION\_NAME through UNUSED\_CPU\_TIME\_SHARED\_POOL for callers with \*USE authority on the QSYS/QPMLPMGT service program.

To return values for TEMPORARY\_JOB\_STRUCTURES\_AVAILABLE and PERMANENT\_JOB\_STRUCTURES\_AVAILABLE, callers must have \*JOBCTL special authority and \*USE authority on the QSYS/QWTCTJBS program.

The following table describes the columns in the view. The system name is SYS\_STAT\_B. The schema is QSYS2.



Table 227. SYSTEM\_STATUS\_INFO\_BASIC view

| Column Name                   | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|--------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TOTAL_JOBS_IN_SYSTEM          | TOTAL_JOBS         | INTEGER                  | The total number of user and system jobs that are currently in the system. The total includes: <ul style="list-style-type: none"> <li>All jobs on job queues waiting to be processed.</li> <li>All jobs currently active (being processed).</li> <li>All jobs that have completed running but still have output on output queues to be produced.</li> </ul>                                                                                                                                                                                                                                                                                             |
| MAXIMUM_JOBS_IN_SYSTEM        | MAX_JOBS           | INTEGER                  | The maximum number of jobs that are allowed on the system. When the number of jobs reaches this maximum, you can no longer submit or start more jobs on the system. The total includes: <ul style="list-style-type: none"> <li>All jobs on job queues waiting to be processed.</li> <li>All jobs currently active (being processed).</li> <li>All jobs that have completed running but still have output on output queues to be produced.</li> </ul>                                                                                                                                                                                                    |
| ACTIVE_JOBS_IN_SYSTEM         | ACT_JOBS           | INTEGER                  | The number of jobs active in the system (jobs that have been started, but have not yet ended), including both user and system jobs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| INTERACTIVE_JOBS_IN_SYSTEM    | INTER_JOBS         | DECIMAL(5,2)             | The percentage of interactive performance assigned to this logical partition. This value is a percentage of the total interactive performance available to the entire physical system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| ELAPSED_TIME                  | ELAP_TIME          | INTEGER                  | The time that has elapsed, in seconds, between the measurement start time and the current system time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| ELAPSED_CPU_USED              | ELAP_USED          | DECIMAL(5,2)             | The average of the elapsed time during which the processing units were in use.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| ELAPSED_CPU_SHARED            | ELAP_SHARE         | DECIMAL(5,2)<br>Nullable | The percentage of the total shared processor pool capacity used by all partitions using the pool during the elapsed time.<br>Returns null if this is a dedicated partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| ELAPSED_CPU_UNCAPPED_CAPACITY | ELAP_UNCAP         | DECIMAL(5,2)<br>Nullable | The percentage of the uncapped shared processing capacity for the partition used since the last time statistics were reset. Returns null if this partition cannot use more than its configured processing capacity.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| CONFIGURED_CPUS               | CONFIGCPUS         | INTEGER                  | Total number of configured CPUs for the partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| CPU_SHARING_ATTRIBUTE         | CPU_SHARE          | VARCHAR(8)<br>Nullable   | This attribute indicates whether this partition is sharing processors. If the value indicates the partition does not share physical processors, then this partition uses only dedicated processors. If the value indicates the partition shares physical processors, then this partition uses physical processors from a shared pool of physical processors.<br><br><b>CAPPED</b> Partition shares processors. The partition is limited to using its configured capacity.<br><br><b>UNCAPPED</b> Partition shares processors. The partition can use more than its configured capacity.<br><br>Contains the null value if this is a dedicated partition. |
| CURRENT_CPU_CAPACITY          | CPU_CAP            | DECIMAL(5,2)             | The current processing capacity specifies the processor units that are being used in the partition. For a partition sharing physical processors, the current processing capacity represents the share of the physical processors in the pool it is running. For a partition using dedicated processors, the current processing capacity represents the number of virtual processors that are currently active in the partition.                                                                                                                                                                                                                         |
| AVERAGE_CPU_RATE              | CPU_RATE           | DECIMAL(5,2)             | Always returns 0.<br><br>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: <a href="#">“SYSTEM_ACTIVITY_INFO table function” on page 863.</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| AVERAGE_CPU_UTILIZATION       | CPU_AVG            | DECIMAL(5,2)             | Always returns 0.<br><br>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: <a href="#">“SYSTEM_ACTIVITY_INFO table function” on page 863.</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

Table 227. SYSTEM\_STATUS\_INFO\_BASIC view (continued)

| Column Name                        | System Column Name | Data Type                | Description                                                                                                                                                                                             |
|------------------------------------|--------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MINIMUM_CPU_UTILIZATION            | CPU_MIN            | DECIMAL(5,2)             | Always returns 0.<br>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: "SYSTEM_ACTIVITY_INFO table function" on page 863.                                                    |
| MAXIMUM_CPU_UTILIZATION            | CPU_MAX            | DECIMAL(5,2)             | Always returns 0.<br>This information has moved to the QSYS2.SYSTEM_ACTIVITY_INFO table function: "SYSTEM_ACTIVITY_INFO table function" on page 863.                                                    |
| SQL_CPU_UTILIZATION                | CPU_SQL            | DECIMAL(5,2)<br>Nullable | Always contains the null value.                                                                                                                                                                         |
| MAIN_STORAGE_SIZE                  | MAIN_STG           | BIGINT                   | The amount of main storage, in kilobytes, in the system.                                                                                                                                                |
| SYSTEM_ASP_STORAGE                 | SYS_STG            | BIGINT                   | The storage capacity of the system auxiliary storage pool (ASP number 1) in millions of bytes. This value represents the amount of space available for storage of both permanent and temporary objects. |
| TOTAL_AUXILIARY_STORAGE            | AUX_STG            | BIGINT                   | The total auxiliary storage, in millions of bytes, on the system.                                                                                                                                       |
| SYSTEM_ASP_USED                    | SYS_RATE           | DECIMAL(5,2)             | The percentage of the system storage pool (ASP number 1) currently in use.                                                                                                                              |
| CURRENT_TEMPORARY_STORAGE          | TEMP_CUR           | INTEGER                  | The current amount of storage, in millions of bytes, in use for temporary objects.                                                                                                                      |
| MAXIMUM_TEMPORARY_STORAGE_USED     | TEMP_MAX           | INTEGER                  | The largest amount of storage, in millions of bytes, used for temporary objects at any one time since the last IPL.                                                                                     |
| PERMANENT_ADDRESS_RATE             | PERM_RATE          | DECIMAL(6,3)             | The percentage of the maximum possible addresses for permanent objects that have been used.                                                                                                             |
| TEMPORARY_ADDRESS_RATE             | TEMP_RATE          | DECIMAL(6,3)             | The percentage of the maximum possible addresses for temporary objects that have been used.                                                                                                             |
| TEMPORARY_256MB_SEGMENTS           | TEMP_256MB         | DECIMAL(5,2)             | The percentage of the maximum possible temporary 256MB segments that have been used.                                                                                                                    |
| TEMPORARY_4GB_SEGMENTS             | TEMP_4GB           | DECIMAL(5,2)             | The percentage of the maximum possible temporary 4GB segments that have been used.                                                                                                                      |
| PERMANENT_256MB_SEGMENTS           | PERM_256MB         | DECIMAL(5,2)             | The percentage of the maximum possible permanent 256MB segments that have been used.                                                                                                                    |
| PERMANENT_4GB_SEGMENTS             | PERM_4GB           | DECIMAL(5,2)             | The percentage of the maximum possible permanent 4GB segments that have been used.                                                                                                                      |
| TEMPORARY_JOB_STRUCTURES_AVAILABLE | TEMP_JS            | INTEGER<br>Nullable      | The number of temporary job structures that currently exist on the system that are not in use.<br>Returns the null value if the user does not have *USE authority on the QSYS/QWTCTJBS program.         |
| PERMANENT_JOB_STRUCTURES_AVAILABLE | PERM_JS            | INTEGER<br>Nullable      | The number of permanent job structures that currently exist on the system that are not in use.<br>Returns the null value if the user does not have *USE authority on the QSYS/QWTCTJBS program.         |
| HOST_NAME                          | HOST_NAME          | VARCHAR(255)             | Name of the system where this information was generated. This is the name set by CHGNETA.                                                                                                               |
| PARTITION_ID                       | PART_ID            | INTEGER                  | The identifier for the partition in which this view is being run.                                                                                                                                       |
| NUMBER_OF_PARTITIONS               | NUM_PART           | INTEGER                  | The number of partitions on the physical machine. This includes partitions that are currently powered on (running) and partitions that are powered off.                                                 |
| ACTIVE_THREADS_IN_SYSTEM           | ACT_THREAD         | INTEGER                  | The number of initial and secondary threads in the system (threads that have been started, but have not yet ended), including both user and system threads.                                             |
| RESTRICTED_STATE                   | REST_STATE         | VARCHAR(3)               | Whether the system is in restricted state.<br><b>NO</b> System is not in restricted state.<br><b>YES</b> System is in restricted state.                                                                 |

Table 227. SYSTEM\_STATUS\_INFO\_BASIC view (continued)

| Column Name                                                                                                                                                                                           | System Column Name | Data Type                                 | Description                                                                                                                                                                                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>The columns from PARTITION_NAME through UNUSED_CPU_TIME_SHARED_POOL require the user to have *USE authority on the QSYS/QPMLPMGT service program. Otherwise, they will contain the null value.</p> |                    |                                           |                                                                                                                                                                                                                                                                                                        |
| PARTITION_NAME                                                                                                                                                                                        | PART_NAME          | VARGRAPHIC(255)<br>CCSID 1200<br>Nullable | The name of the partition as it is known to the HMC.                                                                                                                                                                                                                                                   |
| PARTITION_GROUP_ID                                                                                                                                                                                    | PART_GROUP         | INTEGER<br>Nullable                       | The LPAR group for this partition.                                                                                                                                                                                                                                                                     |
| SHARED_PROCESSOR_POOL_ID                                                                                                                                                                              | POOL_ID            | INTEGER<br>Nullable                       | The shared processor pool this partition is a member of. A shared processor pool is a set of physical processors on the physical machine that is used to run a set of shared processor partitions. A value of 0 indicates the default pool.<br>Contains the null value if DEDICATED_PROCESSORS is YES. |
| DEFINED_MEMORY                                                                                                                                                                                        | DEF_MEM            | BIGINT<br>Nullable                        | The amount of memory (in megabytes) that was configured for this partition through the HMC.                                                                                                                                                                                                            |
| MINIMUM_MEMORY                                                                                                                                                                                        | MIN_MEM            | BIGINT<br>Nullable                        | The minimum amount of main storage (in megabytes) that can be assigned to this partition.                                                                                                                                                                                                              |
| MAXIMUM_MEMORY                                                                                                                                                                                        | MAX_MEM            | BIGINT<br>Nullable                        | The maximum amount of main storage (in megabytes) that can be assigned to this partition.                                                                                                                                                                                                              |
| MEMORY_INCREMENT                                                                                                                                                                                      | MEM_INCR           | BIGINT<br>Nullable                        | The smallest amount of main storage (in megabytes) that can be added to or removed from this partition's memory.                                                                                                                                                                                       |
| DEDICATED_PROCESSORS                                                                                                                                                                                  | DED_PROC           | VARCHAR(3)<br>Nullable                    | Whether the partition uses dedicated processors.<br><b>NO</b> The partition does not use dedicated processors.<br><b>YES</b> The partition uses dedicated processors.                                                                                                                                  |
| PHYSICAL_PROCESSORS                                                                                                                                                                                   | PHY_PROC           | INTEGER<br>Nullable                       | The number of physical processors in this physical machine that are available for use. This does not include processors on demand that have not been turned on.                                                                                                                                        |
| PHYSICAL_PROCESSORS_SHARED_POOL                                                                                                                                                                       | PHY_SHARE          | INTEGER<br>Nullable                       | The number of physical processors that are allocated to the shared processor pool used by this partition.<br>Contains the null value if DEDICATED_PROCESSORS is YES.                                                                                                                                   |
| MAXIMUM_PHYSICAL_PROCESSORS                                                                                                                                                                           | MAX_PHY            | INTEGER<br>Nullable                       | The maximum number of physical processors that can be active in this physical machine without installing additional processors. This value includes currently active processors and any standby (on demand) processors that are present in this physical machine.                                      |
| DEFINED_VIRTUAL_PROCESSORS                                                                                                                                                                            | DEF_VIRT           | INTEGER<br>Nullable                       | The number of virtual processors configured for this partition through the HMC.                                                                                                                                                                                                                        |
| VIRTUAL_PROCESSORS                                                                                                                                                                                    | VIRT_PROC          | INTEGER<br>Nullable                       | The number of virtual processors currently used by this partition.                                                                                                                                                                                                                                     |
| MINIMUM_VIRTUAL_PROCESSORS                                                                                                                                                                            | MIN_VIRT           | INTEGER<br>Nullable                       | The minimum number of virtual processors that can be configured for this partition.                                                                                                                                                                                                                    |
| MAXIMUM_VIRTUAL_PROCESSORS                                                                                                                                                                            | MAX_VIRT           | INTEGER<br>Nullable                       | The maximum number of virtual processors that can be configured for this partition.                                                                                                                                                                                                                    |
| DEFINED_PROCESSING_CAPACITY                                                                                                                                                                           | DEF_CAP            | DECIMAL(5,2)<br>Nullable                  | The amount of processing capacity that was configured for this partition through the HMC.                                                                                                                                                                                                              |
| PROCESSING_CAPACITY                                                                                                                                                                                   | CAPACITY           | DECIMAL(5,2)<br>Nullable                  | The current (usable) amount of processing capacity available to the partition (also known as partition's entitled capacity).                                                                                                                                                                           |

Table 227. SYSTEM\_STATUS\_INFO\_BASIC view (continued)

| Column Name                          | System Column Name | Data Type                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------|--------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| UNALLOCATED_PROCESSING_CAPACITY      | AVAIL_CAP          | DECIMAL(5,2)<br>Nullable | The amount of processing capacity in the partition group this partition belongs to, which is not allocated to any partition and is available for allocation.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| MINIMUM_REQUIRED_PROCESSING_CAPACITY | MIN_REQCAP         | DECIMAL(5,2)<br>Nullable | The minimum amount of processing capacity that the operating system in this partition requires for its operation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| MAXIMUM_LICENSED_PROCESSING_CAPACITY | MAX_LICCAP         | DECIMAL(5,2)<br>Nullable | The current limit on processing capacity of this partition imposed by the operating system software license for this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| MINIMUM_PROCESSING_CAPACITY          | MIN_CAP            | DECIMAL(5,2)<br>Nullable | The minimum amount of processing capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| MAXIMUM_PROCESSING_CAPACITY          | MAX_CAP            | DECIMAL(5,2)<br>Nullable | The maximum amount of processing capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| PROCESSING_CAPACITY_INCREMENT        | CAP_INCR           | DECIMAL(5,2)<br>Nullable | The smallest capacity that can be added to or removed from this partition's processing capacity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| DEFINED_INTERACTIVE_CAPACITY         | DEF_INTCAP         | DECIMAL(5,2)<br>Nullable | The amount of interactive capacity that was configured for this partition through the HMC. A partition's interactive capacity is defined as this partition's portion of total interactive capacity of the physical machine.                                                                                                                                                                                                                                                                                                                                                                        |
| INTERACTIVE_CAPACITY                 | INT_CAP            | DECIMAL(5,2)<br>Nullable | This partition's current (usable) portion of the physical machine interactive capacity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| INTERACTIVE_THRESHOLD                | INT_THRESH         | DECIMAL(5,2)<br>Nullable | The maximum interactive CPU utilization which can be sustained in this partition, without causing a disproportionate increase in system overhead.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| UNALLOCATED_INTERACTIVE_CAPACITY     | AVL_INTCAP         | DECIMAL(5,2)<br>Nullable | The amount of interactive capacity in the partition group this partition belongs to, which is not allocated to any partition and is available for allocation. Interactive capacity is defined as the portion of total interactive capacity of the physical machine.                                                                                                                                                                                                                                                                                                                                |
| MINIMUM_INTERACTIVE_CAPACITY         | MIN_INTCAP         | DECIMAL(5,2)<br>Nullable | The minimum portion of the physical machine's interactive capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| MAXIMUM_INTERACTIVE_CAPACITY         | MAX_INTCAP         | DECIMAL(5,2)<br>Nullable | The maximum portion of the physical machine's interactive capacity that can be assigned to this partition.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| DEFINED_VARIABLE_CAPACITY_WEIGHT     | DEF_CAPW           | INTEGER<br>Nullable      | The weighting factor that was configured for this partition through the HMC. Variable capacity weight is used for uncapped partitions when they compete for unused CPU cycles in the shared pool. Variable capacity weight can be in the range of 0 - 255. The larger the weight, the more the chance this partition will get additional CPU cycles from the shared pool.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED.                                                                                                                     |
| VARIABLE_CAPACITY_WEIGHT             | VAR_CAPW           | INTEGER<br>Nullable      | The weighting factor that is used to assign additional unused CPU cycles (from the shared processor pool) to the partition. Variable capacity weight is used for uncapped partitions when they compete for unused CPU cycles in the shared pool. This factor is in the range of 0 - 255. The larger the weight, the greater the chance this partition will get additional CPU cycles from the pool. A value of 0 effectively caps this partition at its current (usable) processing capacity.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED. |
| UNALLOCATED_VARIABLE_CAPACITY_WEIGHT | AVAIL_CAPW         | INTEGER<br>Nullable      | The amount of capacity weight that is available for allocation to the partition's variable capacity weight.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if CPU_SHARING_ATTRIBUTE is CAPPED.                                                                                                                                                                                                                                                                                                                                                                                   |

Table 227. SYSTEM\_STATUS\_INFO\_BASIC view (continued)

| Column Name                          | System Column Name | Data Type                 | Description                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------|--------------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HARDWARE_MULTITHREADING              | HW_MLT_THR         | VARCHAR(3)<br>Nullable    | Indicates whether hardware multi-threading is enabled.<br><br><b>NO</b> Hardware multi-threading is not enabled.<br><b>YES</b> Hardware multi-threading is enabled.                                                                                                                                                                                                                                |
| BOUND_HARDWARE_THREADS               | HW_BND_THR         | VARCHAR(3)<br>Nullable    | Whether hardware threads are bound.<br><br><b>NO</b> Hardware threads are not bound.<br><b>YES</b> Hardware threads are bound.                                                                                                                                                                                                                                                                     |
| THREADS_PER_PROCESSOR                | THREADS_PP         | INTEGER<br>Nullable       | The number of hardware threads per processor when hardware multi-threading is enabled.<br>Contains the null value if HARDWARE_MULTITHREADING is NO.                                                                                                                                                                                                                                                |
| DISPATCH_LATENCY                     | LATENCY            | DECIMAL(20,0)<br>Nullable | The maximum time in nanoseconds between dispatches of this partition on a physical processor.                                                                                                                                                                                                                                                                                                      |
| DISPATCH_WHEEL_ROTATION_TIME         | DISPATCH_T         | DECIMAL(20,0)<br>Nullable | The number of nanoseconds in the hypervisor's scheduling window. Each virtual processor will be given the opportunity to execute on a physical processor some time during this period. The amount of time each virtual processor is able to use a physical processor is determined by partition processing capacity.                                                                               |
| TOTAL_CPU_TIME                       | CPU_TOTAL          | DECIMAL(20,0)<br>Nullable | The number of nanoseconds of CPU time used by this partition since IPL.                                                                                                                                                                                                                                                                                                                            |
| INTERACTIVE_CPU_TIME                 | CPU_INTER          | DECIMAL(20,0)<br>Nullable | The amount of CPU time, in nanoseconds, used by interactive processes in this partition since partition IPL. An interactive process is any process doing 5250 display device I/O.                                                                                                                                                                                                                  |
| INTERACTIVE_CPU_TIME_ABOVE_THRESHOLD | CPU_THRESH         | DECIMAL(20,0)<br>Nullable | The amount of CPU time, in nanoseconds, used by interactive processes while exceeding the interactive threshold. This is a total since IPL.                                                                                                                                                                                                                                                        |
| UNUSED_CPU_TIME_SHARED_POOL          | CPU_UNUSED         | DECIMAL(20,0)<br>Nullable | The number of nanoseconds of CPU time that the physical processors in a shared processor pool have been idle since system IPL.<br><br>Contains the null value if DEDICATED_PROCESSORS is YES or if the partition is not authorized to retrieve shared pool data.                                                                                                                                   |
| MACHINE_TYPE                         | MACH_TYPE          | CHAR(4)                   | The machine type.                                                                                                                                                                                                                                                                                                                                                                                  |
| MACHINE_MODEL                        | MACH_MOD           | CHAR(4)                   | The machine model.                                                                                                                                                                                                                                                                                                                                                                                 |
| SERIAL_NUMBER                        | SERIAL             | CHAR(8)                   | The machine serial number.                                                                                                                                                                                                                                                                                                                                                                         |
| ATTENTION_LIGHT                      | ATTN_LIGHT         | VARCHAR(3)                | The status of the system attention light.<br><br><b>OFF</b> The light is off.<br><b>ON</b> The light is on.                                                                                                                                                                                                                                                                                        |
| IPL_MODE                             | IPL_MODE           | VARCHAR(9)                | The current IPL mode setting.<br><br><b>AUTOMATIC</b> Used for automatic remote IPL, automatic IPL by date and time, and automatic IPL after a power failure.<br><b>MANUAL</b> An operator uses the control panel to direct the system for special needs.<br><b>NORMAL</b> Requires no operator intervention during the IPL.<br><b>SECURE</b> Prevents use of the control panel to perform an IPL. |

Table 227. SYSTEM\_STATUS\_INFO\_BASIC view (continued)

| Column Name             | System Column Name | Data Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------|--------------------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPL_TYPE                | IPL_TYPE           | CHAR(1)   | Type of IPL performed.<br><br><b>A</b> Used for special work, such as applying fixes (PTFs) and diagnostic work.<br><b>B</b> Used for routine work and when directed by a PTF procedure.<br><b>C</b> Reserved for system support.<br><b>D</b> Used for special work, such as installing and reloading programs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| JOURNAL_RECOVERY_COUNT  | JRNRCYNT           | INTEGER   | Specifies the system wide default journal recovery count. The journal recovery count allows you to choose between faster runtime processing of changes to journaled objects and faster IPL or vary on recovery after an abnormal shutdown. The value specified influences the frequency with which journaled objects are forced to auxiliary storage as those objects are changed. The specified journal recovery count indicates the approximate number of journaled changes that would need to be recovered during journal synchronization for this journal in the event of an IPL or vary on after an abnormal shutdown. A smaller value decreases the number of changes that would need to be recovered from this journal by increasing the frequency with which changed objects are forced to disk. A larger value increases the runtime processing of changes to journaled objects by decreasing the frequency with which changed objects are forced to disk. Changing this value may affect overall system performance as it affects the utilization of auxiliary storage devices.<br><br>This value can be changed with the Change Journal Attributes (CHGJRNA) CL command.<br>The system default for this value is 250,000. |
| JOURNAL_CACHE_WAIT_TIME | CACHEWAIT          | INTEGER   | The cache wait time, in seconds, for journal environments with caching enabled. The cache wait time is the maximum number of seconds that the system will wait before writing any lingering journal entries from main memory to disk.<br><br>This value can be changed with the Change Journal Attributes (CHGJRNA) CL command.<br>The system default for this value is 30 seconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Examples

- Review the storage and CPU status for the partition.

```
SELECT * FROM QSYS2.SYSTEM_STATUS_INFO_BASIC;
```

- Return job structure information.

```
SELECT
 (SELECT CURRENT_NUMERIC_VALUE FROM QSYS2.SYSTEM_VALUE_INFO
 WHERE SYSTEM_VALUE_NAME = 'QTOTJOB') AS INITIAL_PERM_JOB_STRUCTURES,
 (SELECT CURRENT_NUMERIC_VALUE FROM QSYS2.SYSTEM_VALUE_INFO
 WHERE SYSTEM_VALUE_NAME = 'QADLTOTJ') AS ADDITIONAL_PERM_JOB_STRUCTURES,
 PERMANENT_JOB_STRUCTURES_AVAILABLE AS AVAILABLE_PERM_JOB_STRUCTURES,
 PERMANENT_JOB_STRUCTURES_AVAILABLE + TOTAL_JOBS_IN_SYSTEM AS TOTAL_PERM_JOB_STRUCTURES,
 (SELECT CURRENT_NUMERIC_VALUE FROM QSYS2.SYSTEM_VALUE_INFO
 WHERE SYSTEM_VALUE_NAME = 'QACTJOB') AS INITIAL_TEMP_JOB_STRUCTURES,
 (SELECT CURRENT_NUMERIC_VALUE FROM QSYS2.SYSTEM_VALUE_INFO
 WHERE SYSTEM_VALUE_NAME = 'QADLACTJ') AS ADDITIONAL_TEMP_JOB_STRUCTURES,
 TEMPORARY_JOB_STRUCTURES_AVAILABLE AS AVAILABLE_TEMP_JOB_STRUCTURES,
 (SELECT BUCKET_CURRENT_SIZE FROM QSYS2.SYSTMPSTG
 WHERE GLOBAL_BUCKET_NAME = '*ACTJOB') AS TOTAL_TEMP_STORAGE_USED
FROM QSYS2.SYSTEM_STATUS_INFO_BASIC;
```

## SYSTEM\_VALUE\_INFO view

The SYSTEM\_VALUE\_INFO view contains information about system values.

This view returns the names of system values and their values. The list of system values can be found in [Retrieve System Values \(QWCRSVAL\) API](#).

\*ALLOBJ or \*AUDIT special authority is required to retrieve the values for QAUDCTL, QAUDENDACN, QAUDFRCLVL, QAUDLVL, QAUDLVL2, and QCRTOBJAUD. The current value column will contain '\*NOTAVL' or -1 when accessed by an unauthorized user.

The following table describes the columns in the view. The system name is SYSVALINFO. The schema is QSYS2.

Table 228. SYSTEM\_VALUE\_INFO view

| Column Name             | System Column Name | Data Type                       | Description                                                                                   |
|-------------------------|--------------------|---------------------------------|-----------------------------------------------------------------------------------------------|
| SYSTEM_VALUE_NAME       | SYSVALNAME         | VARCHAR(10)                     | Name of the system value.                                                                     |
| CURRENT_NUMERIC_VALUE   | CURNUMVAL          | BIGINT                          | Contains the value if the system value is numeric data. Otherwise, contains the null value.   |
| CURRENT_CHARACTER_VALUE | CURCHARVAL         | VARGRAPHIC(1280)<br>CCSID(1200) | Contains the value if the system value is character data. Otherwise, contains the null value. |

## Example

Look at the system values related to maximums.

```
SELECT * FROM QSYS2.SYSTEM_VALUE_INFO
WHERE SYSTEM_VALUE_NAME LIKE '%MAX%'
```

returns

| SYSTEM_VALUE_NAME | CURRENT_NUMERIC_VALUE | CURRENT_CHARACTER_VALUE |
|-------------------|-----------------------|-------------------------|
| QMAXACTLVL        | 32,767                | -                       |
| QMAXSIGN          | -                     | 000005                  |
| QPWDMAXLEN        | 8                     | -                       |
| QMAXSGNACN        | -                     | 3                       |
| QMAXJOB           | 163,520               | -                       |
| QMAXSPLF          | 9,999                 | -                       |

## WORKLOAD\_GROUP\_INFO view

The WORKLOAD\_GROUP\_INFO view returns information about workload groups including their product entries.

The values returned for the columns in the view are closely related to the detail generated by the DSPWLCGRP CL command and the Retrieve Workload Groups Information (QLZRTVWC) API.

**Authorization:** The caller must have \*USE authority on the QSYS/QLZRTVWC program.

The following table describes the columns in the view. The system name is WLG\_INFO. The schema is QSYS2.

Table 229. WORKLOAD\_GROUP\_INFO view

| Column Name        | System Column Name | Data Type   | Description                                                                                                                                        |
|--------------------|--------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| WORKLOAD_GROUP     | WL_GRP             | VARCHAR(10) | The workload group name.                                                                                                                           |
| PROCESSOR_LIMIT    | PROC_LIMIT         | INTEGER     | The processor limit for jobs and threads associated with this workload group. Values are 1-256.                                                    |
| NUMBER_OF_PRODUCTS | PROD_COUNT         | INTEGER     | Number of products defined for this workload group. The view returns one row for each of the products defined for a specific WORKLOAD_GROUP value. |

Table 229. WORKLOAD\_GROUP\_INFO view (continued)

| Column Name  | System Column Name | Data Type              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------|--------------------|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PRODUCT_ID   | LICPGM             | VARCHAR(7)<br>Nullable | The identifier of the product.<br><br>Contains the null value if NUMBER_OF_PRODUCTS is 0.                                                                                                                                                                                                                                                                                                                                                    |
| LICENSE_TERM | LIC_TERM           | VARCHAR(6)<br>Nullable | The license term associated with this workload group, in one of the following formats:<br><br><b>Vx, VxRy, or VxRyMz</b> where x and y are a number from 0 through 9, and z is a number 0 through 9 or a letter A through Z<br><br><b>vv, vvr, or vvrmm</b> where vv and rr are a number from 00 through 35, and mm is a number 00 through 09 or a letter sequence 0A through 0Z.<br><br>Contains the null value if NUMBER_OF_PRODUCTS is 0. |
| FEATURE_ID   | FEATURE            | VARCHAR(4)<br>Nullable | The feature number of the product associated with this workload group. Values are 5001 to 9999.<br><br>Contains the null value if NUMBER_OF_PRODUCTS is 0.                                                                                                                                                                                                                                                                                   |

## Example

- Return information about all workload groups.

```
SELECT * FROM QSYS2.WORKLOAD_GROUP_INFO;
```

## WORKSTATION\_INFO view

The WORKSTATION\_INFO view returns information about workstation entries.

The values returned for the columns in the view are closely related to the values returned by the Display Work Station Name Entries and the Display Work Station Type Entries panels accessed through the DSPSBSD (Display Subsystem Description) CL command and by the List Subsystem Entries (QWDLSE) API.

**Authorization:** The caller must have:

- \*USE authority to the subsystem description, and
- \*EXECUTE authority to the library containing the subsystem description.

The following table describes the columns in the view. The system name is WRKST\_INFO. The schema is QSYS2.

Table 230. WORKSTATION\_INFO view

| Column Name                   | System Column Name | Data Type               | Description                                                                                                                                                                 |
|-------------------------------|--------------------|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SUBSYSTEM_DESCRIPTION_LIBRARY | SBSD_LIB           | VARCHAR(10)             | The name of the library in which the subsystem description resides.                                                                                                         |
| SUBSYSTEM_DESCRIPTION         | SBSD               | VARCHAR(10)             | The name of the subsystem about which information is being returned.                                                                                                        |
| WORKSTATION_NAME              | WS_NAME            | VARCHAR(10)<br>Nullable | The name of the workstation that is used by the subsystem. It can be a generic workstation entry like DSP*.<br><br>Contains the null value if WORKSTATION_TYPE is not null. |



Table 230. WORKSTATION\_INFO view (continued)

| Column Name             | System Column Name | Data Type               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------|--------------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WORKSTATION_TYPE        | WS_TYPE            | VARCHAR(10)<br>Nullable | <p>The display device type. Can contain the following special values:</p> <p><b>*ALL</b> All workstation devices. This includes devices with 5250, ASCII, and 327x device types.</p> <p><b>*ASCII</b> ASCII display station.</p> <p><b>*CONS</b> System console display. This entry overrides a device type entry that specifies the same device type as the device being used as the console.</p> <p><b>*NONASCII</b> All workstation devices that use a 5250 data stream. This includes the 327x device types.</p> <p>Contains the null value if WORKSTATION_NAME is not null.</p>                                                                                                                                                                                                        |
| JOB_DESCRIPTION_LIBRARY | JOBDLIB            | VARCHAR(10)<br>Nullable | <p>The name of the library in which the job description resides.</p> <p>Contains the null value if JOB_DESCRIPTION is *USRPRF.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| JOB_DESCRIPTION         | JOBID              | VARCHAR(10)             | <p>The name of the job description that is used for jobs started through this workstation entry. Can contain the following special value:</p> <p><b>*USRPRF</b> The job description named in the user profile of the user that signs on at this type of workstation is used for jobs started through this entry.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| ALLOCATION              | ALLOCATION         | VARCHAR(7)              | <p>How the workstations that are associated with this entry are allocated.</p> <p><b>*ENTER</b> The workstations are not allocated when the subsystem is started. However, the interactive jobs that are associated with the workstations are allowed to enter this subsystem through the Transfer Job (TFRJOB) command.</p> <p><b>*SIGNON</b> The workstations are allocated when the subsystem is started if the workstation is not already in use (signed on) in another subsystem. A sign-on prompt is displayed at each workstation that is associated with this work entry. If a workstation becomes allocated to a different subsystem, interactive jobs that are associated with the workstation are allowed to enter this subsystem through the Transfer Job (TFRJOB) command.</p> |
| MAXIMUM_ACTIVE_JOBS     | MAX_ACTIVE         | INTEGER<br>Nullable     | <p>The maximum number of jobs that can be active at the same time through this entry.</p> <p>Contains the null value if the entry specifies *NOMAX, indicating that there is no maximum.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Example

- List all the workstation entries in the QINTER subsystem.

```
SELECT *
FROM QSYS2.WORKSTATION_INFO
WHERE SUBSYSTEM_DESCRIPTION_LIBRARY = 'QSYS' AND
 SUBSYSTEM_DESCRIPTION = 'QINTER';
```

# SYSTOOLS

SYSTOOLS is a set of Db2 for IBM i supplied examples and tools.

SYSTOOLS is the name of a Database supplied schema (library). SYSTOOLS differs from other Db2 for i supplied schemas (QSYS, QSYS2, SYSIBM, and SYSIBMADM) in that it is not part of the default system path. As general purpose useful tools or examples are built by IBM, they are considered for inclusion within SYSTOOLS. SYSTOOLS provides a wider audience with the opportunity to extract value from the tools.

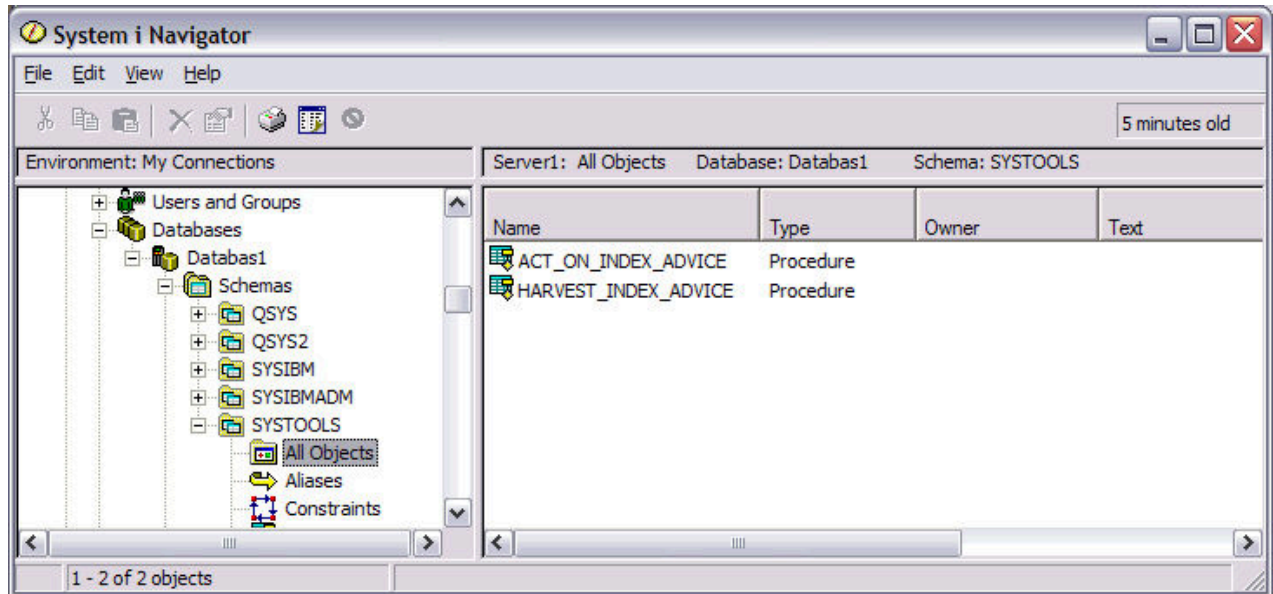
It is the intention of IBM to add content dynamically to SYSTOOLS, either on base releases or through PTFs for field releases. A best practice for customers who are interested in such tools would be to periodically review the contents of SYSTOOLS.

## Using SYSTOOLS

You can generate the sample SQL procedures, learn how to call the procedures, and understand the outcome that is expected. You can also modify the procedure source to customize an example into your business operations.

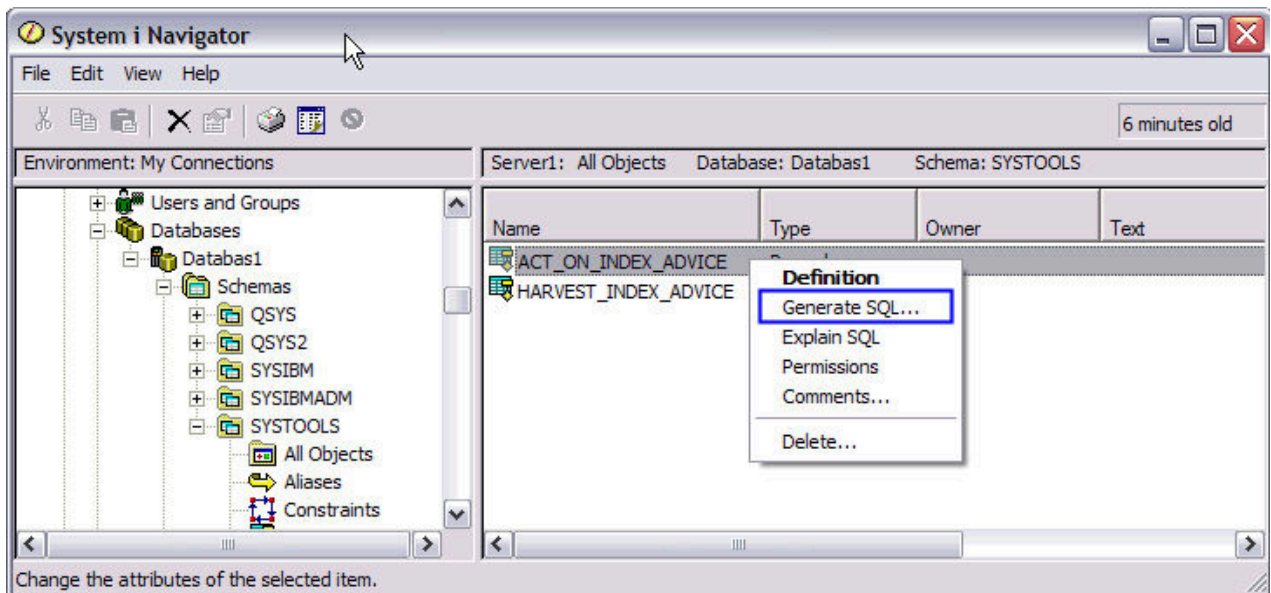
Use System i Navigator, as shown in Figure 1.

Figure 1. System i Navigator schema view of SYSTOOLS:



Start with the Generate SQL action, as shown in Figure 2, to discover and learn within SYSTOOLS. This action utilizes the Generate Data Definition Language (QSQGNDL) API to produce the CREATE PROCEDURE (SQL) statement. This statement is needed to create a replica of the IBM supplied procedure.

Figure 2. Launching Generate SQL from System i Navigator:



After the Generate SQL action completes, as shown in figure 3, you will have a Run SQL Scripts window active, allowing you to do the following:

1. Scroll down and read the procedure prolog.
2. Understand how to call the procedure and the outcome that is expected.
3. Modify the procedure source, including the procedure name and schema. This capability could be the most useful aspect of SYSTOOLS, allowing you to quickly claim and customize an IBM supplied example into your business operations.

Figure 3. Run SQL Scripts view of the generated SQL:

```

--
-- Example: PROCEDURE SYSTOOLS.ACT_ON_INDEX_ADVICE
--
-- Created on July 16, 2009 by Bob Smith (IBMer@us.ibm.com)
--
-- Disclaimer:
-- This example is being provided by IBM to allow IBM i users to understand how index advice
-- could be consumed to improve an index strategy. The creation of indexes can be time consuming
-- and having seldom used indexes may result in a performance degradation. As with any index
-- strategy, it is recommended that you carefully consider the performance characteristics of
-- your application prior to creating new indexes and that you evaluate the index usage
-- statistics.
--
-- While efforts were made to verify the completeness and accuracy of this sample procedure,
-- this sample is provided 'as is' without any warranty whatsoever and to the maximum extent permitted,
-- IBM disclaims all implied warranties.
--
-- Parameters:
-- 1) P_LIBRARY - The system name of the library which contains the file (table)
-- 2) P_FILE - The system name of the file (table) which may have advised indexes
-- 3) P_TIMES_ADVISED - The number of times an index should be advised before creation of a permanent index.
-- Pass in 1 if you don't want to limit the index creation by times advised.
-- 4) P_MTI_USED - The number of times an Maintained Temporary Index has been used since a matching
-- permanent index did not exist.
-- Pass in 0 if you don't want to limit the index creation by MTI used.
-- 5) P_AVERAGE_QUERY_ESTIMATE - The average estimated number of seconds needed to execute the query
-- which drove the index advice.
-- Pass in 1 if you don't want to limit the index creation by average query estimate.
--
-- Other columns which might be useful parameters to gauge index advisor consumption are:
-- LAST_ADVISED
-- ESTIMATED_CREATION_TIME
-- MOST_EXPENSIVE_QUERY
-- TABLE_SIZE
-- MTI_CREATED
-- LAST_MTI_USED
--
-- QSYS2.SYSIXADV documentation can be found here:
-- http://publib.boulder.ibm.com/infocenter/iserics/v6r1m0/topic/rzajq/rzajqindexcols.htm
--
-- Note: This procedure is hard-coded to work against index advice with Sort Sequence = '*HEX';
-- If you want it to work against index advice with sort sequence tables, search for *HEX and replace
-- with the NLSS schema and library. This restriction is necessary because the procedure should be
-- running with a sort sequence which matches the sort sequence of the index advice.
--
-- DECLARE V_DYNSTMT VARCHAR (30000) ;
-- DECLARE V_INDEXNAME VARCHAR (128) ;
-- DECLARE V_COUNT INTEGER DEFAULT 0 ;
-- DECLARE V1 INT DEFAULT 0 ;

```

The IBM maintenance of SYSTOOLS includes periodically dropping and recreating the IBM supplied objects. Customers are allowed to create their own objects within SYSTOOLS. However, if your user created objects conflict with the IBM supplied objects, your objects might be deleted. The tools and examples within SYSTOOLS are considered ready for use. However, they are not subject to IBM Service and Support as they are not considered part of any IBM product.

## HTTP function overview

The HTTP functions in SYSTOOLS are used to make HTTP requests to web services. These functions allow the SQL programmer to use Representational State Transfer (RESTful) via SQL, including Embedded SQL.

An alternate version of the HTTP functions are shipped in QSYS2. These are implemented using the HTTP Transport support provided by the AXISC APIs instead of Java. See [HTTP\\_DELETE](#), [HTTP\\_GET](#), [HTTP\\_POST](#), [HTTP\\_PUT](#), [HTTP\\_DELETE\\_VERBOSE](#), [HTTP\\_GET\\_VERBOSE](#), [HTTP\\_POST\\_VERBOSE](#), [HTTP\\_PUT\\_VERBOSE](#)

## Environmental considerations

Because the HTTP functions use functionality provided by Java, a Java environment is created in the current job. This requires the following conditions.

1. A JVM must not already exist in the job (with the exception of a JVM created by the Java stored procedures support).
2. The job CCSID cannot be 65535.
3. PASE must be installed and operational. The CHKPRDOPT PRDID(5770SS1) OPTION(33) CL command can be used to verify that PASE is installed.

## Foundational HTTP functions

The foundational functions are named according to the three dimensions used when making HTTP requests.

- The first dimension is the HTTP operation. There are 4 different HTTP operations: GET, PUT, POST, and DELETE.
- The second dimension is the data type used for returning data. There are two alternatives: BLOB and CLOB.
- The third dimension is if the verbose version of the function should be used. Non-verbose functions are scalar functions. The VERBOSE functions are table functions that return the return header information that is sent from the HTTP server.

The names of the functions reflect these three dimensions. For example, the HTTPGETCLOBVERBOSE uses the GET operation from the first dimension, the CLOB data type from the second dimension, and the VERBOSE setting from the third dimension.

The following table summarizes the 16 functions that are named according to the three dimensions.

| <b>make HTTP requests</b><br><b>Function</b> | <b>Description</b>                                                     | <b>Documentation link</b>                                                                     |
|----------------------------------------------|------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| HTTPDELETEBLOB                               | Makes an HTTP DELETE request and return BLOB data.                     | <a href="#">“HTTPDELETEBLOB and HTTPDELETECLOB scalar functions” on page 898</a>              |
| HTTPDELETEBLOBVERBOSE                        | Makes an HTTP DELETE request and return BLOB data and response header. | <a href="#">“HTTPDELETEBLOBVERBOSE and HTTPDELETECLOBVERBOSE table functions” on page 899</a> |
| HTTPDELETECLOB                               | Makes an HTTP DELETE request and return CLOB data.                     | <a href="#">“HTTPDELETEBLOB and HTTPDELETECLOB scalar functions” on page 898</a>              |
| HTTPDELETECLOBVERBOSE                        | Makes an HTTP DELETE request and return CLOB data and response header. | <a href="#">“HTTPDELETEBLOBVERBOSE and HTTPDELETECLOBVERBOSE table functions” on page 899</a> |
| HTTPGETBLOB                                  | Makes an HTTP GET request and return BLOB data.                        | <a href="#">“HTTPGETBLOB and HTTPGETCLOB scalar functions” on page 899</a>                    |
| HTTPGETBLOBVERBOSE                           | Makes an HTTP GET request and return BLOB data and response header.    | <a href="#">“HTTPGETBLOBVERBOSE and HTTPGETCLOBVERBOSE table functions” on page 900</a>       |
| HTTPGETCLOB                                  | Makes an HTTP GET request and return CLOB data.                        | <a href="#">“HTTPGETBLOB and HTTPGETCLOB scalar functions” on page 899</a>                    |
| HTTPGETCLOBVERBOSE                           | Makes an HTTP GET request and return CLOB data and response header.    | <a href="#">“HTTPGETBLOBVERBOSE and HTTPGETCLOBVERBOSE table functions” on page 900</a>       |

Table 231. Foundational HTTP functions (continued)

| make HTTP requests<br>Function | Description                                                          | Documentation link                                                                        |
|--------------------------------|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| HTTPPOSTBLOB                   | Makes an HTTP POST request and return BLOB data.                     | <a href="#">“HTTPPOSTBLOB and HTTPPOSTCLOB scalar functions” on page 901</a>              |
| HTTPPOSTBLOBVERBOSE            | Makes an HTTP POST request and return BLOB data and response header. | <a href="#">“HTTPPOSTBLOBVERBOSE and HTTPPOSTCLOBVERBOSE table functions” on page 902</a> |
| HTTPPOSTCLOB                   | Makes an HTTP POST request and return CLOB data.                     | <a href="#">“HTTPPOSTBLOB and HTTPPOSTCLOB scalar functions” on page 901</a>              |
| HTTPPOSTCLOBVERBOSE            | Makes an HTTP POST request and return CLOB data and response header. | <a href="#">“HTTPPOSTBLOBVERBOSE and HTTPPOSTCLOBVERBOSE table functions” on page 902</a> |
| HTTPPUTBLOB                    | Makes an HTTP PUT request and return BLOB data.                      | <a href="#">“HTTPPUTBLOB and HTTPPUTCLOB scalar functions” on page 902</a>                |
| HTTPPUTBLOBVERBOSE             | Makes an HTTP PUT request and return BLOB data and response header.  | <a href="#">“HTTPPUTBLOBVERBOSE and HTTPPUTCLOBVERBOSE table functions” on page 903</a>   |
| HTTPPUTCLOB                    | Makes an HTTP PUT request and return CLOB data.                      | <a href="#">“HTTPPUTBLOB and HTTPPUTCLOB scalar functions” on page 902</a>                |
| HTTPPUTCLOBVERBOSE             | Makes an HTTP PUT request and return CLOB data and response header.  | <a href="#">“HTTPPUTBLOBVERBOSE and HTTPPUTCLOBVERBOSE table functions” on page 903</a>   |

These HTTP functions are passed parameters that indicate the HTTP server to access, the settings of the HTTP headers, and any data to be sent to the server. The first two parameters are the same for each HTTP function.

- The first parameter is the URL used to access the server.
- The second parameter is an XML document which indicates the HTTP header that is to be presented to the server by the request. This XML document can be provided as either a CLOB or XML value. The HTTP header XML document has the following format:

```
<httpHeader headerAttribute="headerAttributeValue">
 <header name="name" value="value" />
 <header name="Accept" value="text/plain,application/xml,*/*" />
</httpHeader>
```

The XML header document consists of an *httpHeader* element with various attributes, described below. The *httpHeader* element then contains *header* elements that are used to set additional HTTP headers. This is done by specifying the name of the header and the value that corresponds to the name.

The following attributes can be set in the *httpHeader* element.

Attribute name	Type	Default	Description
connectionTimeout	Integer	SystemDefault	Timeout in milliseconds
readTimeout	Integer	SystemDefault	Timeout in milliseconds

Attribute name	Type	Default	Description
followRedirects	Boolean	true	If true, then redirects are followed. Redirects will only be followed if the protocol does not change. For example, a redirect from "http:" to "https:" will not be followed.
useCaches	Boolean	true	If true, then caches are used.
includeErrorMsg	Boolean	false	If true, then HTTP verbose functions will return error information in the RESPONSEMSG column.

The *header* elements allow various HTTP headers to be set. The following header names are case insensitive: "Content-Type", "Content-Length", "Content-Encoding", "Accept-Encoding", "Authorization", "User-Agent". The remaining headers are sent as specified. To correctly use these headers, consult the appropriate Web server and RFC documentation on the use of HTTP request headers. Examples of RFCs are <https://tools.ietf.org/html/rfc2616> and <https://tools.ietf.org/html/rfc7231>.

For example, using the header

```
<httpHeader><header name="User-Agent" value="IBM i HTTP function"/></httpHeader>
```

will cause "User-Agent": "IBM i HTTP function" to be included in the HTTP header sent to the server.

## Generic HTTP functions

The HTTPBLOB and HTTPCLOB functions can be used to perform any of the non-verbose HTTP functions by passing the operation to perform as an argument: GET, PUT, POST, or DELETE.

Function	Description	Documentation link
HTTPBLOB	Makes an HTTP request and return BLOB data.	<a href="#">“HTTPBLOB and HTTPCLOB scalar functions” on page 896</a>
HTTPCLOB	Makes an HTTP request and return CLOB data.	<a href="#">“HTTPBLOB and HTTPCLOB scalar functions” on page 896</a>

## Utility HTTP functions

Some additional HTTP functions are utility functions for working with HTTP requests.

Function	Description	Documentation link
BASE64DECODE	Returns a bit data string that has been Base64 decoded.	<a href="#">“BASE64DECODE scalar function” on page 895</a>



Table 234. HTTP utility functions (continued)

Function	Description	Documentation link
BASE64ENCODE	Returns the Base64 encoded version of a character string.	<a href="#">“BASE64ENCODE scalar function” on page 896</a>
HTTPHEAD	Verifies the HTTP header using an HTTP HEAD request.	<a href="#">“HTTPHEAD scalar function” on page 901</a>
URLDECODE	Decodes a URL encoded string.	<a href="#">“URLDECODE scalar function” on page 904</a>
URLENCODE	Encodes a string using URL encoding.	<a href="#">“URLENCODE scalar function” on page 904</a>

## Common errors

When an error occurs, to get additional information about the error consider using the VERBOSE form of the function with `includeErrorMsg="true"`. This will return error information in the RESPONSEMSG column. For example:

```
select * from table(systools.httpgetclobverbose('http://www.w3.org/notfound.html',
 '<httpHeader includeErrorMsg="true"/>'))
```

The following are some common errors that may be encountered when using the HTTP functions.

- If the current system does not have internet access, the following error may be encountered.

```
[SQL4302] Java stored procedure or user-defined function
SYSTOOLS.HTTPGETCLOB, specific name HTTPG00005 aborted with an exception
"java.net.ConnectException:A remote host refused an attempted connect
operation. (Connection refused)".
```

- If DNS is not configured on the system, the following error may be encountered.

```
[SQL4302] Java stored procedure or user-defined function
SYSTOOLS.HTTPGETCLOB, specific name HTTPG00005 aborted with an exception
"java.net.UnknownHostException"
```

- If the Job CCSID is 65535, the following error may be encountered.

```
[SQL0332] Character conversion between CCSID 65535 and CCSID 1200 not valid.
```

- If https is being used, then the connection uses SSL. In order to work properly, the server must use a certificated signed by a recognized certificate authority. Using a self signed certificate may result in the following error.

```
[SQL4302] Java stored procedure or user-defined function
SYSTOOLS.HTTPGETCLOB, specific name HTTPG00005 aborted with
an exception "com.ibm.jsse2.util.h: PKIX path building failed:
java.security.cert.CertPathBuilderException: PKIXCertPathBuilderImpl could
not build a valid CertPath.
```

If you must access a server that uses a self signed certificate, you must add a certificate to the trust store of the JVM being used. If the current JVM is the 32-bit JDK 6, then the following command may be used in QSH after downloading the certificate from the server. The trust store of the JVM may be refreshed when Java PTFs are applied, causing the loss of local changes. This step needs to be repeated after applying Java PTFs that refresh the trust store.

```
keytool -import -trustcacerts
-keystore /QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit/jre/lib/security/cacerts
-storepass changeit -noprompt -alias myhost -file myhost.crt
```

- Some websites require the use of specific SSL versions. In that case, an error like the following may be encountered.



[SQL4302] Java stored procedure or user-defined function SYSTOOLS.HTTPGETCLOBVERBOSE, specific name HTTPG00007 aborted with an exception "Received fatal alert: handshake\_failure".

To force the use of a specific SSL version, the following may be done.

1. Create a user defined function to set a JVM property. This only needs to be done once per system.

```
CREATE FUNCTION SYSTOOLS.setProperty(PROPERTY VARCHAR(80), VALUE VARCHAR(80))
 RETURNS VARCHAR(80)
 LANGUAGE JAVA
 PARAMETER STYLE JAVA
 EXTERNAL NAME 'java.lang.System.setProperty'
```

2. Before calling the HTTP function in a job, use the following to set the default SSL protocol.

```
VALUES SYSTOOLS.setProperty('com.ibm.jsse2.overrideDefaultProtocol','TLSv12')
```

## Debugging considerations

Using STRSQL to invoke HTTP functions will not provide useful problem information. Instead, you should use ACS Run SQL Scripts to execute the HTTP functions in order to see error information.

Because the HTTP functions utilize the networking functions provided by Java, any networking problem detected by the Java code may not be exposed to the SQL level with enough detail to diagnose the problem. If the HTTP functions fail with an SQL4302 error and the information is not useful, a Java stored procedure trace may be used to see more information about the exception.

A Java stored procedure trace may be executed and viewed by using the following SQL statements.

- Enable the trace.

```
CALL QSYS2.QCMDEXC('ADDENVVAR ENVVAR(QIBM_COMPONENT_TRACE_LEVEL)
 VALUE('SQJAVA,VERBOSE')')
```

- Execute the HTTP function.
- Dump the trace to a file in QTEMP.

```
CALL QSYS2.QCMDEXC('DMPUSRTRC')
```

- Copy the trace to the current library.


```
CREATE TABLE MYTRACE AS
 (SELECT TRIM(CAST(QAP0ZDMP AS VARCHAR(200) CCSID 37)) LINE FROM QTEMP.QAP0ZDMP)
 WITH DATA
```

- View the Java exception in the trace.

```
SELECT * FROM MYTRACE
 WHERE LINE LIKE '% at %' OR LINE LIKE '%Except%'
```

## BASE64DECODE scalar function

The BASE64DECODE function returns a character string that has been Base64 decoded. Base64 encoding is widely used to represent binary data as a string.

►► BASE64DECODE (  ) ►►

The schema is SYSTOOLS.

**character-string** A character string in CCSID 1208 that is currently Base64 encoded. The length cannot exceed 4096 characters.

The result of the function is a varying length character for bit data string that contains *character-string* after being Base64 decoded.

## Example

- Decode a binary string that was originally X'1122334455'. The result is the original value.

```
VALUES SYSTOOLS.BASE64DECODE('ESIZRFU=');
```

## BASE64ENCODE scalar function

The BASE64ENCODE function returns the Base64 encoded version of the binary values of a character string.

►► BASE64ENCODE ( *character-string* ) ►►  
                   IN =>

The schema is SYSTOOLS.

***character-string*** A character expression to be encoded. The maximum length is 2732 characters.

The result of the function is a varying length character string in CCSID 1208 that contains the bytes of *character-string* as a Base64-encoded string.

## Example

- Encode a string that contains the value X'1122334455'.

```
VALUES SYSTOOLS.BASE64ENCODE(X'1122334455');
```

The result is: ESIZRFU=.

- Encode a character string in CCSID 37 that contains the value 'ABC'.

```
VALUES SYSTOOLS.BASE64ENCODE('ABC');
```

The result is: wCLD

- Encode a UTF-8 string that contains the value 'ABC'.

```
VALUES SYSTOOLS.BASE64ENCODE(CAST('ABC' AS VARCHAR(10) CCSID 1208));
```

The result is: QUJD

## HTTPBLOB and HTTPCLOB scalar functions

The HTTPBLOB and HTTPCLOB REST functions complete an HTTP request with the specified HTTP verb.

►► HTTPBLOB ( *url* ) ►►  
               HTTPCLOB ( *URL* => )

*httpmethod* ►►  
                   HTTPMETHOD =>

*httpHeader* ►►  
                   HTTPHEADER =>

*requestmsg* ►►  
                   REQUESTMSG =>

The schema is SYSTOOLS.

- url** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.
- httpmethod** Specifies the HTTP verb to use. Valid values are GET, POST, PUT, and DELETE.
- httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See “Foundational HTTP functions” on page 891 for a description of the header content.
- requestmsg** Specifies the data to update at the specified URL. For the HTTPBLOB function, this argument is defined as BLOB(2G). For the HTTPCLOB function, this argument is defined as CLOB(2G) CCSID 1208. This argument is required for the POST and PUT methods.

For the HTTPBLOB function, the result response message is returned as BLOB(2G). For the HTTPCLOB function, the result response message is returned as CLOB(2G) CCSID 1208.

## Example

- Use HTTPCLOB with the HTTP GET method to retrieve the group PTF information from the IBM PSP website.


```
VALUES SYSTOOLS.HTTPCLOB(
 URL => 'https://www.ibm.com/support/pages/sites/default/files/inline-files/
xml doc.xml',
 HTTPMETHOD => 'GET',
 HTTPHEADER => NULL,
 REQUESTMSG => NULL);
```

The same operation could be performed using the HTTPGETCLOB function.

```
VALUES SYSTOOLS.HTTPGETCLOB(
 URL => 'https://www.ibm.com/support/pages/sites/default/files/inline-files/
xml doc.xml',
 HTTPHEADER => NULL);
```

## HTTPBLOBVERBOSE and HTTPCLOBVERBOSE table functions

The HTTPBLOBVERBOSE and HTTPCLOBVERBOSE REST table functions complete an HTTP request with the specified HTTP verb. They return a one row table that contains the normal HTTP response for the request and the header information returned from the HTTP request.



The schema is SYSTOOLS.

- url** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.
- httpmethod** Specifies the HTTP verb to use. Valid values are GET, POST, PUT, and DELETE.

**httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See “Foundational HTTP functions” on page 891 for a description of the header content.

**requestmsg** Specifies the data to update at the specified URL. For the HTTPBLOBVERBOSE function, this argument is defined as BLOB(2G). For the HTTPCLOBVERBOSE function, this argument is defined as CLOB(2G) CCSID 1208. This argument is required for the POST and PUT methods.

The result of the function is a table containing the following two columns.

Column name	Data type	Description
RESPONSEMSG	BLOB(2G) or CLOB(2G) CCSID 1208 depending on which function is used	The normal HTTP response for the request.
RESPONSEHTTPHEADER	CLOB(1M) CCSID 1208 or XML based on the data type of the <i>httpHeader</i> argument.	Header information returned from the HTTP request.

## Example

- This query returns the error information from using an incorrect URL to request the group PTF information from the IBM PSP website.


```
SELECT * FROM TABLE(SYSTOOLS.HTTPCLOBVERBOSE(
 URL => 'http://www.ibm.com/support/docview.wss?uid=nas4PSPbyNum&aid=1',
 HTTPMETHOD => 'GET',
 HTTPHEADER => NULL,
 REQUESTMSG => NULL));
```

The RESPONSEHTTPHEADER result column contains the following information about the error.

```
<?xml version="1.0" encoding="UTF-8" ?><httpHeader responseCode="301">
 <responseMessage>Moved Permanently....
```

## HTTPDELETEBLOB and HTTPDELETECLOB scalar functions

The HTTPDELETEBLOB and HTTPDELETECLOB REST functions delete a binary or text-based resource from the specified URL through an HTTP DELETE request.



The schema is SYSTOOLS.


**url** Specifies the URL of the resource being accessed. This parameter is a VARCHAR(2048) CCSID 1208 value.

**httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See “[Foundational HTTP functions](#)” on page 891 for a description of the header content.

For the HTTPDELETEBLOB function, the result response message is returned as BLOB(2G). For the HTTPDELETECLOB function, the result response message is returned as CLOB(2G) CCSID 1208.

## HTTPDELETEBLOBVERBOSE and HTTPDELETECLOBVERBOSE table functions

The HTTPDELETEBLOBVERBOSE and HTTPDELETECLOBVERBOSE REST table functions delete a binary or text-based resource from the specified URL through an HTTP DELETE request. They return a one row table that contains the normal HTTP response for the request and the header information returned from the HTTP request.



The schema is SYSTOOLS.

**url** Specifies the URL of the resource being accessed. This parameter is a VARCHAR(2048) CCSID 1208 value.


**httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See “[Foundational HTTP functions](#)” on page 891 for a description of the header content.

The result of the function is a table containing the following two columns.

Column name	Data type	Description
RESPONSEMSG	BLOB(2G) or CLOB(2G) CCSID 1208 depending on which function is used	The normal HTTP response for the request.
RESPONSEHTTPHEADER	CLOB(1M) CCSID 1208 or XML based on the data type of the <i>httpHeader</i> argument.	Header information returned from the HTTP request.

## HTTPGETBLOB and HTTPGETCLOB scalar functions

The HTTPGETBLOB and HTTPGETCLOB REST functions retrieve a binary or text-based resource from the specified URL through an HTTP GET request.



The schema is SYSTOOLS.

- url** Specifies the URL of the resource being accessed. This parameter is a VARCHAR(2048) CCSID 1208 value.
- httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or XML value.  
See “[Foundational HTTP functions](#)” on page 891 for a description of the header content.


HTTPGETBLOB returns the resource as BLOB(2G) data. HTTPGETCLOB returns the resource as CLOB(2G) CCSID 1208 data.

### Example

```
SELECT SYSTOOLS.HTTPGETCLOB(
 URL => CAST('http://ws.geonames.org/countryInfo?lang=' CONCAT
 SYSTOOLS.URLENCODE('en','') CONCAT
 '&country=' CONCAT
 SYSTOOLS.URLENCODE('us','') CONCAT
 '&type=XML' AS VARCHAR(255)),
 HTTPHEADER => CAST(NULL AS CLOB(1K)))
FROM SYSIBM.SYSDUMMY1;
```

### HTTPGETBLOBVERBOSE and HTTPGETCLOBVERBOSE table functions

The HTTPGETBLOBVERBOSE and HTTPGETCLOBVERBOSE REST table functions retrieve a binary or text-based resource from the specified URL through an HTTP GET request. They return a one row table that contains the normal HTTP response for the request and the header information returned from the HTTP request.



The schema is SYSTOOLS.

- url** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.
- httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value.  
See “[Foundational HTTP functions](#)” on page 891 for a description of the header content.

The result of the function is a table containing the following two columns.

Column name	Data type	Description
RESPONSEMSG	BLOB(2G) or CLOB(2G) CCSID 1208 depending on which function is used	The normal HTTP response for the request.

Table 237. HTTPGETBLOBVERBOSE and HTTPGETCLOBVERBOSE result table (continued)

Column name	Data type	Description
RESPONSEHTTPHEADER	CLOB(1M) CCSID 1208 or XML based on the data type of the <i>httpHeader</i> argument.	Header information returned from the HTTP request.

## HTTPHEAD scalar function

The HTTPHEAD REST function verifies the HTTP header for the specified resource through an HTTP HEAD request.

►► HTTPHEAD — ( — URL — => — *url* — , — HTTPHEADER — => — *httpHeader* — ) — ►►

The schema is SYSTOOLS.

***url*** Specifies the URL of the resource being accessed. This parameter is a VARCHAR(2048) CCSID 1208 value.

***httpHeader*** Specifies an optional header XML document. This parameter is a CLOB(10K) CCSID 1208 or an XML value.

See “Foundational HTTP functions” on page 891 for a description of the header content.

The HTTP header is returned as CLOB(10K) CCSID 1208 or XML data based on the data type of the *httpHeader* argument.

## HTTPPOSTBLOB and HTTPPOSTCLOB scalar functions

The HTTPPOSTBLOB and HTTPPOSTCLOB REST functions update a binary or text-based resource under the specified URL through an HTTP POST request.

►► HTTPPOSTBLOB — ( — URL — => — *url* — , — HTTPHEADER — => — *httpHeader* — , — REQUESTMSG — => — *requestmsg* — ) — ►►

The schema is SYSTOOLS.

***url*** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.

***httpHeader*** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value.


See “Foundational HTTP functions” on page 891 for a description of the header content.

***requestmsg*** Specifies the data to update at the specified URL. For the HTTPPOSTBLOB function, this argument is defined as BLOB(2G). For the HTTPPOSTCLOB function, this argument is defined as CLOB(2G) CCSID 1208.

For the HTTPPOSTBLOB function, the result response message is returned as BLOB(2G). For the HTTPPOSTCLOB function, the result response message is returned as CLOB(2G) CCSID 1208.

## HTTPPOSTBLOBVERBOSE and HTTPPOSTCLOBVERBOSE table functions

The HTTPPOSTBLOBVERBOSE and HTTPPOSTCLOBVERBOSE REST table functions update a binary or text-based resource under the specified URL through an HTTP POST request. They return a one row table that contains the normal HTTP response for the request and the header information returned from the HTTP request.



The schema is SYSTOOLS.

- url** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.
- httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See [“Foundational HTTP functions” on page 891](#) for a description of the header content.
- requestmsg** Specifies the data to update at the specified URL. For the HTTPPOSTBLOBVERBOSE function, this argument is defined as BLOB(2G). For the HTTPPOSTCLOBVERBOSE function, this argument is defined as CLOB(2G) CCSID 1208.

The result of the function is a table containing the following two columns.

Column name	Data type	Description
RESPONSEMSG	BLOB(2G) or CLOB(2G) CCSID 1208 depending on which function is used	The normal HTTP response for the request.
RESPONSEHTTPHEADER	CLOB(1M) CCSID 1208 or XML based on the data type of the <i>httpHeader</i> argument.	Header information returned from the HTTP request.

## HTTPPUTBLOB and HTTPPUTCLOB scalar functions

The HTTPPUTBLOB and HTTPPUTCLOB REST functions create or update a binary or text-based resource under the specified URL through an HTTP PUT request.



The schema is SYSTOOLS.




- url** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.
- httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See “Foundational HTTP functions” on page 891 for a description of the header content.
- requestmsg** Specifies the data to update at the specified URL. For the HTTPPUTBLOB function, this argument is defined as BLOB(2G). For the HTTPPUTCLOB function, this argument is defined as CLOB(2G) CCSID 1208.

For the HTTPPUTBLOB function, the result response message is returned as BLOB(2G). For the HTTPPUTCLOB function, the result response message is returned as CLOB(2G) CCSID 1208.

## HTTPPUTBLOBVERBOSE and HTTPPUTCLOBVERBOSE table functions

The HTTPPUTBLOBVERBOSE and HTTPPUTCLOBVERBOSE REST table functions create or update a binary or text-based resource under the specified URL through an HTTP PUT request. They return a one row table that contains the normal HTTP response for the request and the header information returned from the HTTP request.



The schema is SYSTOOLS.


- url** Specifies the URL at which to complete the request. This argument is defined as a VARCHAR(2048) CCSID 1208 value.
- httpHeader** Specifies an optional header XML document. To use the default HTTP header, specify NULL or the empty string. This parameter is a CLOB(10K) CCSID 1208 or an XML value. See “Foundational HTTP functions” on page 891 for a description of the header content.
- requestmsg** Specifies the data to update at the specified URL. For the HTTPPUTBLOBVERBOSE function, this argument is defined as BLOB(2G). For the HTTPPUTCLOBVERBOSE function, this argument is defined as CLOB(2G) CCSID 1208.

The result of the function is a table containing the following two columns.

Column name	Data type	Description
RESPONSEMSG	BLOB(2G) or CLOB(2G) CCSID 1208 depending on which function is used	The normal HTTP response for the request.
RESPONSEHTTPHEADER	CLOB(1M) CCSID 1208 or XML based on the data type of the <i>httpHeader</i> argument.	Header information returned from the HTTP request.

## URLDECODE scalar function

The URLDECODE function completes URL decoding of the provided text.



The schema is SYSTOOLS.


**character-string** A character string that is currently encoded. The length cannot exceed 4096 characters.

**encoding** A character or graphic string that specifies the encoding that was used for encoding characters that are not values in a URL. The default and recommended value is UTF-8. Otherwise, this is a charset that follows the naming conventions defined in RFC 2278: IANA Charset Registration Procedures.

The result of the function is a varying length character string that contains *character-string* after being decoded.

## URLENCODE scalar function

The URLENCODE function completes URL encoding of the provided text.



The schema is SYSTOOLS.

**character-string** A character or graphic string expression that is to be encoded.

**encoding** A character or graphic string that specifies the encoding used for encoding characters that are not values in a URL. The default and recommended value is UTF-8. Otherwise, this is a charset that follows the naming conventions defined in RFC 2278: IANA Charset Registration Procedures.

The result of the function is a varying length character string that contains *character-string* after being encoded.

## Database monitor formats

This section contains the formats used to create the database monitor SQL tables and views.

### Database monitor SQL table format

Displays the format used to create the QSYS/QAQQDBMN performance statistics table, that is shipped with the system.

```
CREATE TABLE QSYS.QAQQDBMN (
 QQRID DECIMAL(15, 0) NOT NULL DEFAULT 0 ,
 QQTIME TIMESTAMP NOT NULL DEFAULT CURRENT_TIMESTAMP ,
 QQJFLD CHAR(46) CCSID 65535 NOT NULL DEFAULT ' ' ,
 QQRDBN CHAR(18) CCSID 37 NOT NULL DEFAULT ' ' ,
 QQSYS CHAR(8) CCSID 37 NOT NULL DEFAULT ' ' ,
 QQJOB CHAR(10) CCSID 37 NOT NULL DEFAULT ' ' ,
 QQUSER CHAR(10) CCSID 37 NOT NULL DEFAULT ' ' ,
 QQJNUM CHAR(6) CCSID 37 NOT NULL DEFAULT ' ' ,
```

```

QQUCNT DECIMAL(15, 0) DEFAULT NULL ,
QQUDEF VARCHAR(100) CCSID 37 DEFAULT NULL ,
QQSTN DECIMAL(15, 0) DEFAULT NULL ,
QQQDTN DECIMAL(15, 0) DEFAULT NULL ,
QQQDTL DECIMAL(15, 0) DEFAULT NULL ,
QQMATN DECIMAL(15, 0) DEFAULT NULL ,
QQMATL DECIMAL(15, 0) DEFAULT NULL ,
QQTLN CHAR(10) CCSID 37 DEFAULT NULL ,
QQTFN CHAR(10) CCSID 37 DEFAULT NULL ,
QQTMN CHAR(10) CCSID 37 DEFAULT NULL ,
QQPTLN CHAR(10) CCSID 37 DEFAULT NULL ,
QQPTFN CHAR(10) CCSID 37 DEFAULT NULL ,
QQPTMN CHAR(10) CCSID 37 DEFAULT NULL ,
QQILNM CHAR(10) CCSID 37 DEFAULT NULL ,
QQIFNM CHAR(10) CCSID 37 DEFAULT NULL ,
QQIMNM CHAR(10) CCSID 37 DEFAULT NULL ,
QQNTNM CHAR(10) CCSID 37 DEFAULT NULL ,
QQNLNM CHAR(10) CCSID 37 DEFAULT NULL ,
QQSTIM TIMESTAMP DEFAULT NULL ,
QQETIM TIMESTAMP DEFAULT NULL ,
QQKP CHAR(1) CCSID 37 DEFAULT NULL ,
QQKS CHAR(1) CCSID 37 DEFAULT NULL ,
QQTOTR DECIMAL(15, 0) DEFAULT NULL ,
QQTMPR DECIMAL(15, 0) DEFAULT NULL ,
QQJNP DECIMAL(15, 0) DEFAULT NULL ,
QQEPT DECIMAL(15, 0) DEFAULT NULL ,
QQDSS CHAR(1) CCSID 37 DEFAULT NULL ,
QQIDXA CHAR(1) CCSID 37 DEFAULT NULL ,
QQORDG CHAR(1) CCSID 37 DEFAULT NULL ,
QQGRPG CHAR(1) CCSID 37 DEFAULT NULL ,
QQJNG CHAR(1) CCSID 37 DEFAULT NULL ,
QQUNIN CHAR(1) CCSID 37 DEFAULT NULL ,
QQSUBQ CHAR(1) CCSID 37 DEFAULT NULL ,
QQHSTV CHAR(1) CCSID 37 DEFAULT NULL ,
QQRCDS CHAR(1) CCSID 37 DEFAULT NULL ,
QQRCOD CHAR(2) CCSID 37 DEFAULT NULL ,
QQRSS DECIMAL(15, 0) DEFAULT NULL ,
QQREST DECIMAL(15, 0) DEFAULT NULL ,
QQRIDX DECIMAL(15, 0) DEFAULT NULL ,
QQFKEY DECIMAL(15, 0) DEFAULT NULL ,
QQKSEL DECIMAL(15, 0) DEFAULT NULL ,
QQAJN DECIMAL(15, 0) DEFAULT NULL ,
QQIDXD VARCHAR(1000) ALLOCATE(48) CCSID 37 DEFAULT NULL ,
QQC11 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC12 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC13 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC14 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC15 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC16 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC18 CHAR(1) CCSID 37 DEFAULT NULL ,
QQC21 CHAR(2) CCSID 37 DEFAULT NULL ,
QQC22 CHAR(2) CCSID 37 DEFAULT NULL ,
QQC23 CHAR(2) CCSID 37 DEFAULT NULL ,
QQI1 DECIMAL(15, 0) DEFAULT NULL ,
QQI2 DECIMAL(15, 0) DEFAULT NULL ,
QQI3 DECIMAL(15, 0) DEFAULT NULL ,
QQI4 DECIMAL(15, 0) DEFAULT NULL ,
QQI5 DECIMAL(15, 0) DEFAULT NULL ,
QQI6 DECIMAL(15, 0) DEFAULT NULL ,
QQI7 DECIMAL(15, 0) DEFAULT NULL ,
QQI8 DECIMAL(15, 0) DEFAULT NULL ,
QQI9 DECIMAL(15, 0) DEFAULT NULL ,
QQIA DECIMAL(15, 0) DEFAULT NULL ,
QQF1 DECIMAL(15, 0) DEFAULT NULL ,
QQF2 DECIMAL(15, 0) DEFAULT NULL ,
QQF3 DECIMAL(15, 0) DEFAULT NULL ,
QQC61 CHAR(6) CCSID 37 DEFAULT NULL ,
QQC81 CHAR(8) CCSID 37 DEFAULT NULL ,
QQC82 CHAR(8) CCSID 37 DEFAULT NULL ,
QQC83 CHAR(8) CCSID 37 DEFAULT NULL ,
QQC84 CHAR(8) CCSID 37 DEFAULT NULL ,
QQC101 CHAR(10) CCSID 37 DEFAULT NULL ,
QQC102 CHAR(10) CCSID 37 DEFAULT NULL ,
QQC103 CHAR(10) CCSID 37 DEFAULT NULL ,
QQC104 CHAR(10) CCSID 37 DEFAULT NULL ,
QQC105 CHAR(10) CCSID 37 DEFAULT NULL ,
QQC106 CHAR(10) CCSID 37 DEFAULT NULL ,
QQC181 VARCHAR(128) ALLOCATE(18) CCSID 37 DEFAULT NULL ,
QQC182 VARCHAR(128) ALLOCATE(18) CCSID 37 DEFAULT NULL ,
QQC183 VARCHAR(128) ALLOCATE(15) CCSID 37 DEFAULT NULL ,
QQC301 VARCHAR(30) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QQC302 VARCHAR(30) ALLOCATE(10) CCSID 37 DEFAULT NULL ,

```

```

QQC303 VARCHAR(30) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QQ1000 VARCHAR(1000) ALLOCATE(48) CCSID 37 DEFAULT NULL ,
QQTIM1 TIMESTAMP DEFAULT NULL ,
QQTIM2 TIMESTAMP DEFAULT NULL ,
QVQTB VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVQLIB VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVPTBL VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVPLIB VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVINAM VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVILIB VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVQTBLI CHAR(1) CCSID 37 DEFAULT NULL ,
QVPTBLI CHAR(1) CCSID 37 DEFAULT NULL ,
QVINAMI CHAR(1) CCSID 37 DEFAULT NULL ,
QVBNDY CHAR(1) CCSID 37 DEFAULT NULL ,
QVJFANO CHAR(1) CCSID 37 DEFAULT NULL ,
QVPPARPF CHAR(1) CCSID 37 DEFAULT NULL ,
QVPPARPL CHAR(1) CCSID 37 DEFAULT NULL ,
QVC11 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC12 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC13 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC14 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC15 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC16 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC17 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC18 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC19 CHAR(1) CCSID 37 DEFAULT NULL ,
QVC1A CHAR(1) CCSID 37 DEFAULT NULL ,
QVC1B CHAR(1) CCSID 37 DEFAULT NULL ,
QVC1C CHAR(1) CCSID 37 DEFAULT NULL ,
QVC1D CHAR(1) CCSID 37 DEFAULT NULL ,
QVC1E CHAR(1) CCSID 37 DEFAULT NULL ,
QVC1F CHAR(1) CCSID 37 DEFAULT NULL ,
QWC11 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC12 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC13 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC14 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC15 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC16 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC17 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC18 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC19 CHAR(1) CCSID 37 DEFAULT NULL ,
QWC1A CHAR(1) CCSID 37 DEFAULT NULL ,
QWC1B CHAR(1) CCSID 37 DEFAULT NULL ,
QWC1C CHAR(1) CCSID 37 DEFAULT NULL ,
QWC1D CHAR(1) CCSID 37 DEFAULT NULL ,
QWC1E CHAR(1) CCSID 37 DEFAULT NULL ,
QWC1F CHAR(1) CCSID 37 DEFAULT NULL ,
QVC21 CHAR(2) CCSID 37 DEFAULT NULL ,
QVC22 CHAR(2) CCSID 37 DEFAULT NULL ,
QVC23 CHAR(2) CCSID 37 DEFAULT NULL ,
QVC24 CHAR(2) CCSID 37 DEFAULT NULL ,
QVCTIM DECIMAL(15, 0) DEFAULT NULL ,
QVPARD DECIMAL(15, 0) DEFAULT NULL ,
QVPARU DECIMAL(15, 0) DEFAULT NULL ,
QVPARRC DECIMAL(15, 0) DEFAULT NULL ,
QVRCNT DECIMAL(15, 0) DEFAULT NULL ,
QVFILES DECIMAL(15, 0) DEFAULT NULL ,
QVP151 DECIMAL(15, 0) DEFAULT NULL ,
QVP152 DECIMAL(15, 0) DEFAULT NULL ,
QVP153 DECIMAL(15, 0) DEFAULT NULL ,
QVP154 DECIMAL(15, 0) DEFAULT NULL ,
QVP155 DECIMAL(15, 0) DEFAULT NULL ,
QVP156 DECIMAL(15, 0) DEFAULT NULL ,
QVP157 DECIMAL(15, 0) DEFAULT NULL ,
QVP158 DECIMAL(15, 0) DEFAULT NULL ,
QVP159 DECIMAL(15, 0) DEFAULT NULL ,
QVP15A DECIMAL(15, 0) DEFAULT NULL ,
QVP15B DECIMAL(15, 0) DEFAULT NULL ,
QVP15C DECIMAL(15, 0) DEFAULT NULL ,
QVP15D DECIMAL(15, 0) DEFAULT NULL ,
QVP15E DECIMAL(15, 0) DEFAULT NULL ,
QVP15F DECIMAL(15, 0) DEFAULT NULL ,
QVC41 CHAR(4) CCSID 37 DEFAULT NULL ,
QVC42 CHAR(4) CCSID 37 DEFAULT NULL ,
QVC43 CHAR(4) CCSID 37 DEFAULT NULL ,
QVC44 CHAR(4) CCSID 37 DEFAULT NULL ,
QVC81 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC82 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC83 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC84 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC85 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC86 CHAR(8) CCSID 37 DEFAULT NULL ,

```

```

QVC87 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC88 CHAR(8) CCSID 37 DEFAULT NULL ,
QVC101 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC102 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC103 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC104 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC105 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC106 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC107 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC108 CHAR(10) CCSID 37 DEFAULT NULL ,
QVC1281 VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVC1282 VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVC1283 VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVC1284 VARCHAR(128) ALLOCATE(10) CCSID 37 DEFAULT NULL ,
QVC3001 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3002 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3003 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3004 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3005 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3006 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3007 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC3008 VARCHAR(300) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC5001 VARCHAR(500) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC5002 VARCHAR(500) ALLOCATE(32) CCSID 37 DEFAULT NULL ,
QVC1000 VARCHAR(1000) ALLOCATE(48) CCSID 37 DEFAULT NULL ,
QWC1000 VARCHAR(1000) ALLOCATE(48) CCSID 37 DEFAULT NULL ,
QQINT01 INTEGER DEFAULT NULL ,
QQINT02 INTEGER DEFAULT NULL ,
QQINT03 INTEGER DEFAULT NULL ,
QQINT04 INTEGER DEFAULT NULL ,
QQSMINT1 SMALLINT DEFAULT NULL ,
QQSMINT2 SMALLINT DEFAULT NULL ,
QQSMINT3 SMALLINT DEFAULT NULL ,
QQSMINT4 SMALLINT DEFAULT NULL ,
QQSMINT5 SMALLINT DEFAULT NULL ,
QQSMINT6 SMALLINT DEFAULT NULL ,
QQ1000L CLOB(2147483647) ALLOCATE(48) CCSID 37 DEFAULT NULL ,
QFC11 CHAR(1) CCSID 37 DEFAULT NULL ,
QFC12 CHAR(1) CCSID 37 DEFAULT NULL ,
QFC13 CHAR(1) CCSID 37 DEFAULT NULL ,
QQCLOB2 CLOB(2147483647) ALLOCATE(48) CCSID 37 DEFAULT NULL ,
QFC14 CHAR(1) CCSID 37 DEFAULT NULL ,
QFC15 CHAR(1) CCSID 37 DEFAULT NULL ,
QFC16 CHAR(1) CCSID 37 DEFAULT NULL ,
QQCLOB3 CLOB(2147483647) CCSID 37 DEFAULT NULL ,
QFC17 CHAR(1) CCSID 37 DEFAULT NULL ,
QFC18 CHAR(1) CCSID 37 DEFAULT NULL ,
QFC19 CHAR(1) CCSID 37 DEFAULT NULL ,
QDDBCLOB1 DBCLOB(1073741823) ALLOCATE(24) CCSID 1200 DEFAULT NULL ,
QFC1A CHAR(1) CCSID 37 DEFAULT NULL ,
QFC1B CHAR(1) CCSID 37 DEFAULT NULL ,
QFC1C CHAR(1) CCSID 37 DEFAULT NULL ,
QDDBCLOB2 DBCLOB(1073741823) CCSID 1200 DEFAULT NULL ,
QFC1D CHAR(1) CCSID 37 DEFAULT NULL ,
QFC1E CHAR(1) CCSID 37 DEFAULT NULL ,
QFC1F CHAR(1) CCSID 37 DEFAULT NULL ,
QOBLOB1 BLOB(2147483647) DEFAULT NULL ,
QXC11 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC12 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC13 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC14 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC15 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC16 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC17 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC18 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC19 CHAR(1) CCSID 37 DEFAULT NULL ,
QXC1A CHAR(1) CCSID 37 DEFAULT NULL ,
QXC1B CHAR(1) CCSID 37 DEFAULT NULL ,
QXC1C CHAR(1) CCSID 37 DEFAULT NULL ,
QXC1D CHAR(1) CCSID 37 DEFAULT NULL ,
QXC1E CHAR(1) CCSID 37 DEFAULT NULL ,
QXC21 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC22 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC23 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC24 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC25 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC26 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC27 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC28 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC29 CHAR(2) CCSID 37 DEFAULT NULL ,
QXC41 CHAR(4) CCSID 37 DEFAULT NULL ,
QXC42 CHAR(4) CCSID 37 DEFAULT NULL ,

```

```

QXC43 CHAR(4) CCSID 65535 DEFAULT NULL ,
QXC44 CHAR(4) CCSID 37 DEFAULT NULL ,
QQINT05 INTEGER DEFAULT NULL ,
QQINT06 INTEGER DEFAULT NULL ,
QQINT07 INTEGER DEFAULT NULL ,
QQINT08 INTEGER DEFAULT NULL ,
QQINT09 INTEGER DEFAULT NULL ,
QQINT0A INTEGER DEFAULT NULL ,
QQINT0B INTEGER DEFAULT NULL ,
QQINT0C INTEGER DEFAULT NULL ,
QQINT0D INTEGER DEFAULT NULL ,
QQINT0E INTEGER DEFAULT NULL ,
QQINT0F INTEGER DEFAULT NULL ,
QOSMINT7 SMALLINT DEFAULT NULL ,
QOSMINT8 SMALLINT DEFAULT NULL ,
QOSMINT9 SMALLINT DEFAULT NULL ,
QOSMINTA SMALLINT DEFAULT NULL ,
QOSMINTB SMALLINT DEFAULT NULL ,
QOSMINTC SMALLINT DEFAULT NULL ,
QOSMINTD SMALLINT DEFAULT NULL ,
QOSMINTE SMALLINT DEFAULT NULL ,
QOSMINTF SMALLINT DEFAULT NULL ,
QQTIM12A TIMESTAMP(32) DEFAULT NULL,
QQTIM12B TIMESTAMP(32) DEFAULT NULL,
QVP161 DECIMAL(15,0) DEFAULT NULL,
QVP162 DECIMAL(15,0) DEFAULT NULL,
QQBGINT1 BIGINT DEFAULT NULL,
QQBGINT2 BIGINT DEFAULT NULL)

RCDFMT QQQDBMN ;
RENAME QSYS/QQQDBMN TO SYSTEM NAME QAQQDBMN;

LABEL ON TABLE QSYS/QAQQDBMN
IS 'Database Monitor Physical File' ;

LABEL ON COLUMN QSYS.QAQQDBMN
(QQRID IS 'Record ID',
 QQTIME IS 'Created Time',
 QQJFLD IS 'Join Column',
 QQRDBN IS 'Relational Database Name',
 QOSYS IS 'System Name',
 QQJOB IS 'Job Name',
 QQUSER IS 'Job User',
 QQJNUM IS 'Job Number',
 QQUCNT IS 'Unique Counter',
 QQUDEF IS 'User Defined Column',
 QQSTN IS 'Statement Number',
 QQQDTN IS 'Subselect Number',
 QQQDTL IS 'Subselect Nested Level',
 QQMATN IS 'Subselect Number of Materialized View',
 QQMATL IS 'Subselect Level of Materialized View',
 QQTLN IS 'Library of Table Queried',
 QQTFN IS 'Name of Table Queried',
 QQTMN IS 'Member of Table Queried',
 QQPTLN IS 'Library of Base Table',
 QQPTFN IS 'Name of Base Table',
 QQPTMN IS 'Member of Base Table',
 QQILNM IS 'Library of Index Used',
 QQIFNM IS 'Name of Index Used',
 QQIMNM IS 'Member of Index Used',
 QQNTNM IS 'NLSS Table',
 QQNLNM IS 'NLSS Library',
 QQSTIM IS 'Start Time',
 QQETIM IS 'End Time',
 QQKP IS 'Key Positioning',
 QQKS IS 'Key Selection',
 QQTOTR IS 'Total Rows',
 QQTMPR IS 'Number of Rows in Temporary',
 QQJNP IS 'Join Position',
 QQEPT IS 'Estimated Processing Time',
 QQDSS IS 'Data Space Selection',
 QQIDXA IS 'Index Advised',
 QQORDG IS 'Ordering',
 QQGRPG IS 'Grouping',
 QQJNG IS 'Join',
 QQUNIN IS 'Union',
 QQSUBQ IS 'Subquery',
 QQHSTV IS 'Host Variables',
 QQRCDS IS 'Row Selection',
 QQRCOD IS 'Reason Code',
 QQRSS IS 'Number of Rows Selected',
 QQREST IS 'Estimated Number of Rows Selected',

```

```

QQRIDX IS 'Number of Entries in Index Created' ,
QQFKEY IS 'Estimated Entries for Key Positioning' ,
QQKSEL IS 'Estimated Entries for Key Selection' ,
QQAJN IS 'Estimated Number of Joined Rows' ,
QQIDX IS 'Advised Key Columns' ,
QQI9 IS 'Thread Identifier' ,
QVQTBL IS 'Queried Table Long Name' ,
QVQLIB IS 'Queried Library Long Name' ,
QVPTBL IS 'Base Table Long Name' ,
QVPLIB IS 'Base Library Long Name' ,
QVINAM IS 'Index Used Long Name' ,
QVILIB IS 'Index Used Library Name' ,
QVQTBLI IS 'Table Long Required' ,
QVPTBLI IS 'Base Long Required' ,
QVINAMI IS 'Index Long Required' ,
QVBNDY IS 'I/O or CPU Bound' ,
QVJFANO IS 'Join Fan Out' ,
QVPARPF IS 'Parallel Pre-Fetch' ,
QVPARPL IS 'Parallel Pre-Load' ,
QVCTIM IS 'Estimated Cumulative Time' ,
QVPARD IS 'Parallel Degree Requested' ,
QVPARU IS 'Parallel Degree Used' ,
QVPARRC IS 'Parallel Limited Reason Code' ,
QVRCNT IS 'Refresh Count' ,
QVFILES IS 'Number of Tables Joined') ;

```

LABEL ON COLUMN QSYS.QAQQDBMN

```

(QQRID TEXT IS 'Record ID' ,
 QQTIME TEXT IS 'Time record was created' ,
 QQJFLD TEXT IS 'Join Column' ,
 QQRDBN TEXT IS 'Relational Database Name' ,
 QQSYS TEXT IS 'System Name' ,
 QQJOB TEXT IS 'Job Name' ,
 QQUSER TEXT IS 'Job User' ,
 QQJNUM TEXT IS 'Job Number' ,
 QQUCNT TEXT IS 'Unique Counter' ,
 QQUDEF TEXT IS 'User Defined Column' ,
 QQSTN TEXT IS 'Statement Number' ,
 QQQDTN TEXT IS 'Subselect Number' ,
 QQQDTL TEXT IS 'Subselect Nested Level' ,
 QQMATN TEXT IS 'Subselect Number of Materialized View' ,
 QQMATL TEXT IS 'Subselect Level of Materialized View' ,
 QQTLN TEXT IS 'Library of Table Queried' ,
 QQTFN TEXT IS 'Name of Table Queried' ,
 QQTMN TEXT IS 'Member of Table Queried' ,
 QQPTLN TEXT IS 'Base Table Library' ,
 QQPTFN TEXT IS 'Base Table' ,
 QQPTMN TEXT IS 'Base Table Member' ,
 QQILNM TEXT IS 'Library of Index Used' ,
 QQIFNM TEXT IS 'Name of Index Used' ,
 QQIMNM TEXT IS 'Member of Index Used' ,
 QQNTNM TEXT IS 'NLSS Table' ,
 QQNLNM TEXT IS 'NLSS Library' ,
 QQSTIM TEXT IS 'Start timestamp' ,
 QQETIM TEXT IS 'End timestamp' ,
 QQKP TEXT IS 'Key positioning' ,
 QQKS TEXT IS 'Key selection' ,
 QQTOTR TEXT IS 'Total row in table' ,
 QQTMPR TEXT IS 'Number of rows in temporary' ,
 QQJNP TEXT IS 'Join Position' ,
 QQEPT TEXT IS 'Estimated processing time' ,
 QQDSS TEXT IS 'Data Space Selection' ,
 QQIDXA TEXT IS 'Index advised' ,
 QQORDG TEXT IS 'Ordering' ,
 QQGRPG TEXT IS 'Grouping' ,
 QQJNG TEXT IS 'Join' ,
 QQUNIN TEXT IS 'Union' ,
 QQSUBQ TEXT IS 'Subquery' ,
 QQHSTV TEXT IS 'Host Variables' ,
 QQRCD S TEXT IS 'Row Selection' ,
 QQRCD TEXT IS 'Reason Code' ,
 QQRSS TEXT IS 'Number of rows selected or sorted' ,
 QQREST TEXT IS 'Estimated number of rows selected' ,
 QQRIDX TEXT IS 'Number of entries in index created' ,
 QQFKEY TEXT IS 'Estimated keys for key positioning' ,
 QQKSEL TEXT IS 'Estimated keys for key selection' ,
 QQAJN TEXT IS 'Estimated number of joined rows' ,
 QQIDX TEXT IS 'Key columns for the index advised' ,
 QQI9 TEXT IS 'Thread Identifier' ,
 QVQTBL TEXT IS 'Queried Table, Long Name' ,
 QVQLIB TEXT IS 'Queried Library, Long Name' ,
 QVPTBL TEXT IS 'Base Table, Long Name' ,

```

```

QVPLIB TEXT IS 'Base Library, Long Name' ,
QVINAM TEXT IS 'Index Used, Long Name' ,
QVILIB TEXT IS 'Index Used, Library Name' ,
QVQTB LI TEXT IS 'Table Long Required' ,
QVP TBLI TEXT IS 'Base Long Required' ,
QVINAMI TEXT IS 'Index Long Required' ,
QVBNDY TEXT IS 'I/O or CPU Bound' ,
QVJFANO TEXT IS 'Join Fan out' ,
QVPARPF TEXT IS 'Parallel Pre-Fetch' ,
QVPARPL TEXT IS 'Parallel Pre-Load' ,
QVCTIM TEXT IS 'Cumulative Time' ,
QVPARD TEXT IS 'Parallel Degree, Requested' ,
QVPARU TEXT IS 'Parallel Degree, Used' ,
QVPARRC TEXT IS 'Parallel Limited, Reason Code' ,
QVRCNT TEXT IS 'Refresh Count' ,
QVFILES TEXT IS 'Number of, Tables Joined') ;

```

## Optional database monitor SQL view format

These examples show the different optional SQL view format that you can create with the SQL shown. The column descriptions are explained in the tables following each example. These views are not shipped with the system, and you must create them, if you choose to do so. These views are optional and are not required for analyzing monitor data.

Any rows that have a row identification number (QQRID) of 5000 or greater are for internal database use.

### Database monitor view 1000 - SQL Information

Displays the SQL logical view format for database monitor QQQ1000.

```

Create View QQQ1000 as
(SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQI5 as Unique_Refresh_Counter,
 QQUDEF as User_Defined,
 QQSTN as Statement_Number,
 QQC11 as Statement_Function,
 QQC21 as Statement_Operation,
 QQC12 as Statement_Type,
 QQC13 as Parse_Required,
 QQC103 as Package_Name,
 QQC104 as Package_Library,
 QQC181 as Cursor_Name,
 QQC182 as Statement_Name,
 QQSTIM as Start_Timestamp,
 QQ1000 as Statement_Text,
 QQC14 as Statement_Outcome,
 QQI2 as Result_Rows,
 QQC22 as Dynamic_Replan_Reason_Code,
 QQC16 as Data_Conversion_Reason_Code,
 QQI4 as Total_Time_Milliseconds,
 QQI3 as Rows_Fetched,
 QQETIM as End_Timestamp,
 QQI6 as Total_Time_Microseconds,
 QQI7 as SQL_Statement_Length,
 QQI1 as Insert_Unique_Count,
 QQI8 as SQLCode,
 QQC81 as SQLState,
 QVC101 as Close_Cursor_Mode,
 QVC11 as Allow_Copy_Data_Value,
 QVC12 as PseudoOpen,
 QVC13 as PseudoClose,
 QVC14 as ODP_Implementation,
 QVC21 as Dynamic_Replan_SubCode,
 QVC41 as Commitment_Control_Level,
 QWC1B as Concurrent_Access_Resolution,
 QVC15 as Blocking_Type,

```



QVC16 as Delay\_Prepare,  
 QVC1C as Explainable,  
 QVC17 as Naming\_Convention,  
 QVC18 as Dynamic\_Processing\_Type,  
 QVC19 as LOB\_Data\_Optimized,  
 QVC1A as Program\_User\_Profile\_Used,  
 QVC1B as Dynamic\_User\_Profile\_Used,  
 QVC1281 as Default\_Collection,  
 QVC1282 as Procedure\_Name,  
 QVC1283 as Procedure\_Library,  
 QQCLOB2 as SQL\_Path,  
 QVC1284 as Current\_Schema,  
 QQC18 as Binding\_Type,  
 QQC61 as Cursor\_Type,  
 QVC1D as Statement\_Originator,  
 QQC15 as Hard\_Close\_Reason\_Code,  
 QQC23 as Hard\_Close\_Subcode,  
 QVC42 as Date\_Format,  
 QWC11 as Date\_Separator,  
 QVC43 as Time\_Format,  
 QWC12 as Time\_Separator,  
 QWC13 as Decimal\_Point,  
 QVC104 as Sort\_Sequence\_Table ,  
 QVC105 as Sort\_Sequence\_Library,  
 QVC44 as Language\_ID,  
 QVC23 as Country\_ID,  
 QQIA as First\_N\_Rows\_Value,  
 QQF1 as Optimize\_For\_N\_Rows\_Value,  
 QVC24 as Access\_Plan\_Not\_Saved\_Reason\_Code,  
 QVC81 as Transaction\_Context\_ID,  
 QVP152 as Activation\_Group\_Mark,  
 QVP153 as Open\_Cursor\_Threshold,  
 QVP154 as Open\_Cursor\_Close\_Count,  
 QVP155 as Commitment\_Control\_Lock\_Limit,  
 QWC15 as Allow\_SQL\_Mixed\_Constants,  
 QWC16 as Suppress\_SQL\_Warnings,  
 QWC17 as Translate\_ASCII,  
 QWC18 as System\_Wide\_Statement\_Cache,  
 QVP159 as LOB\_Locator\_Threshold,  
 QVP156 as Max\_Decimal\_Precision,  
 QVP157 as Max\_Decimal\_Scale,  
 QVP158 as Min\_Decimal\_Divide\_Scale ,  
 QWC19 as Unicode\_Normalization,  
 QQ1000L as Statement\_Text\_Long,  
 QVP15B as Old\_Access\_Plan\_Length,  
 QVP15C as New\_Access\_Plan\_Length,  
 QVP151 as Fast\_Delete\_Count,  
 QQF2 as Statement\_Max\_Compression,  
 QVC102 as Current\_User\_Profile,  
 QVC1E as Expression\_Evaluator\_Used,  
 QVP15A as Host\_Server\_Delta,  
 QQC301 as NTS\_Lock\_Space\_Id,  
 QQC183 as IP\_Address,  
 QFC11 as IP\_Type,  
 QQSMINT2 as IP\_Port\_Number,  
 QVC3004 as NTS\_Transaction\_Id,  
 QQSMINT3 as NTS\_Format\_Id\_Length,  
 QQSMINT4 as NTS\_Transaction\_ID\_SubLength,  
 QVRCNT as Unique\_Refresh\_Counter2,  
 QVP15F as Times\_Run,  
 QVP15E as FullOpens,  
 QVC1F as Proc\_In\_Cache,  
 QWC1A as Combined\_Operation,  
 QVC3001 as Client\_Applname,  
 QVC3002 as Client\_Userid,  
 QVC3003 as Client\_Wrkstnname,  
 QVC3005 as Client\_Acctng,  
 QVC3006 as Client\_Progamid,  
 QVC5001 as Interface\_Information,  
 QVC82 as Open\_Options,  
 QWC1D as Extended\_Indicators,  
 QWC1C as DECFLOAT\_Rounding\_Mode,  
 QWC1E as SQL\_DECFLOAT\_Warnings,  
 QVP15D as Worst\_Time\_Micro,  
 QQINT05 as SQ\_Unique\_Count,  
 QFC13 as Concurrent\_Access\_Res\_Used,  
 QQSMINT8 as SQL\_UDFs\_Not\_Inlined,  
 QVC3007 as Result\_Set\_Cursor,  
 QFC12 as Implicit\_XMLPARSE\_Option,  
 QQSMINT7 as SQL\_XML\_Data\_CC SID,  
 QQSMINT5 as OPTIMIZER\_USE,  
 QFC14 as XML\_Schema\_In\_Cache,

```

 QQC105 as Current_User,
 QFC15 as Row_Column_Access_Control,
 QQTIM12A as Temporal_System_Time,
 QFC16 as SYSTIME_Bind_Option,
 QFC17 as Temporal_System_Time_Query,
 QQDBCLOB1 as DBLOB_HOSTVR,
 QQSMINT6 as StmtCmpReuseMin
FROM DbMonLib/DbMonTable
WHERE QQRID=1000)

```

Table 240. QQQ1000 - SQL Information

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query) When monitor files are created from an SQL Plan Cache snapshot, this value is the Plan Identifier
Unique_Refresh_Counter	QQI5	Unique refresh counter
User_Defined	QQUDEF	User-defined column
Statement_Number	QQSTN	Statement number (unique per statement)
Statement_Function	QQC11	Statement function: <ul style="list-style-type: none"> <li>• S - Select</li> <li>• U - Update</li> <li>• I - Insert</li> <li>• D - Delete</li> <li>• L - Data definition language</li> <li>• O - Other</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Statement_Operation	QQC21	Statement operation: <ul style="list-style-type: none"> <li>• AC - Allocate cursor</li> <li>• AD - Allocate descriptor</li> <li>• AF - Alter function</li> <li>• AL - Alter table</li> <li>• AK - Alter mask</li> <li>• AP - Alter procedure</li> <li>• AQ - Alter sequence</li> <li>• AR - Alter permission</li> <li>• AS - Associate locators</li> <li>• AT - Alter trigger</li> <li>• BE - Compound (dynamic)</li> <li>• CA - Call</li> <li>• CB - Create variable</li> <li>• CC - Create collection</li> <li>• CD - Create type</li> <li>• CF - Create function</li> <li>• CG - Create trigger</li> <li>• CI - Create index</li> <li>• CK - Create mask</li> <li>• CL - Close</li> <li>• CM - Commit</li> <li>• CN - Connect</li> <li>• CO - Comment on</li> <li>• CP - Create procedure</li> <li>• CQ - Create sequence</li> <li>• CR - Create permission</li> <li>• CS - Create alias/synonym</li> <li>• CT - Create table</li> <li>• CV - Create view</li> <li>• DA - Deallocate descriptor</li> <li>• DE - Describe</li> <li>• DI - Disconnect</li> <li>• DL - Delete</li> <li>• DM - Describe parameter marker</li> <li>• DO - Describe procedure</li> <li>• DP - Declare procedure</li> <li>• DR - Drop</li> <li>• DS - Describe cursor</li> <li>• DT - Describe table</li> <li>• EI - Execute immediate</li> <li>• EX - Execute</li> <li>• FE - Fetch</li> <li>• FL - Free locator</li> <li>• GR - Grant</li> <li>• GS - Get descriptor</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Statement_Operation (continued)	QQC21	<ul style="list-style-type: none"> <li>• HC - Hard close</li> <li>• HL - Hold locator</li> <li>• IN - Insert</li> <li>• JR - Server job reused</li> <li>• LK - Lock</li> <li>• LO - Label on</li> <li>• MG - Merge</li> <li>• MT - More text (Deprecated in V5R4)</li> <li>• OP - Open</li> <li>• PD - Prepare and describe</li> <li>• PR - Prepare</li> <li>• QF - OPNQRYF command</li> <li>• QM - Query/400 STRQMQRYP command</li> <li>• QO - OPNDBF command or Native open</li> <li>• QQ - QQQQRY() API</li> <li>• QR - RUNQRY command</li> <li>• RB - Rollback to savepoint</li> <li>• RE - Release</li> <li>• RF - Refresh Table</li> <li>• RG - Resignal</li> <li>• RM - Set current DECFLOAT rounding mode</li> <li>• RO - Rollback</li> <li>• RS - Release Savepoint</li> <li>• RT - Rename table</li> <li>• RV - Revoke</li> <li>• SA - Savepoint</li> <li>• SC - Set connection</li> <li>• SD - Set descriptor</li> <li>• SE - Set encryption password</li> <li>• SN - Set session user</li> <li>• SI - Select into</li> <li>• SO - Set current degree</li> <li>• SP - Set path</li> <li>• SR - Set result set</li> <li>• SS - Set current schema</li> <li>• ST - Set transaction</li> <li>• SV - Set variable</li> <li>• SX - Set current implicit XMLPARSE option</li> <li>• SZ - Set current temporal SYSTEM_TIME</li> <li>• TO - Transfer ownership</li> <li>• TT - Truncate</li> <li>• UP - Update</li> <li>• VI - Values into</li> <li>• X0 - Unknown statement</li> <li>• X1 - Unknown statement</li> <li>• X2 - DRDA (AS) Unknown statement</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Statement_Operation (continued)	QQC21	<ul style="list-style-type: none"> <li>• X3 - Unknown statement</li> <li>• X9 - Internal error</li> <li>• XA - X/Open API</li> <li>• ZD - Host server only activity</li> </ul>
Statement_Type	QQC12	Statement type: <ul style="list-style-type: none"> <li>• D - Dynamic statement</li> <li>• S - Static statement</li> </ul>
Parse_Required	QQC13	Parse required (Y/N)
Package_Name	QQC103	Name of the package or name of the program that contains the current SQL statement
Package_Library	QQC104	Name of the library containing the package
Cursor_Name	QQC181	Name of the cursor corresponding to this SQL statement, if applicable
Statement_Name	QQC182	Name of statement for SQL statement, if applicable
Start_Timestamp	QQSTIM	Time this statement entered
Statement_Text	QQ1000	First 1000 bytes of statement text
Statement_Outcome	QQC14	Statement outcome <ul style="list-style-type: none"> <li>• S - Successful</li> <li>• U - Unsuccessful</li> </ul>
Result_Rows	QQI2	Number of result rows returned. Will only be set for the following SQL operations and is 0 for all others: <ul style="list-style-type: none"> <li>• IN - Insert</li> <li>• UP - Update</li> <li>• DL - Delete</li> <li>• For an SQL Plan Cache snapshot, this count represents the aggregate count for all runs of this query. This count can be divided by the total number of runs, COALESCE(QVP15F,1), to determine the average rows fetched for a given query run.</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Dynamic_Replan_Reason_Code	QQC22	<p>Dynamic replan (access plan rebuilt)</p> <ul style="list-style-type: none"> <li>• NA - No replan.</li> <li>• NR - SQL QDT rebuilt for new release.</li> <li>• A1 - A table or member is not the same object as the one referenced when the access plan was last built. Some reasons why they might be different are:               <ul style="list-style-type: none"> <li>– Object was deleted and recreated.</li> <li>– Object was saved and restored.</li> <li>– Library list was changed.</li> <li>– Object was renamed.</li> <li>– Object was moved.</li> <li>– Object was overridden to a different object.</li> <li>– This run is the first run of this query after the object containing the query has been restored.</li> <li>– Mask or permission attributes changed for the object.</li> </ul> </li> <li>• A2 - Access plan was built to use a reusable Open Data Path (ODP) and the optimizer chose to use a nonreusable ODP for this call.</li> <li>• A3 - Access plan was built to use a non-reusable Open Data Path (ODP) and the optimizer chose to use a reusable ODP for this call.</li> <li>• A4 - Either the number of rows in the table member has changed by more than 10% or a selectivity or cardinality statistic has change by more than 25% since the access plan was last built.</li> <li>• A5 - A new index exists over one of the tables in the query.</li> <li>• A6 - An index that was used for this access plan no longer exists or is no longer valid.</li> <li>• A7 - IBM i Query requires the access plan to be rebuilt because of system programming changes.</li> <li>• A8 - The CCSID of the current job is different from the CCSID of the job that last created the access plan.</li> <li>• A9 - The value of one or more of the following values is different for the current job than it was for the job that last created this access plan:               <ul style="list-style-type: none"> <li>– date format</li> <li>– date separator</li> <li>– time format</li> <li>– time separator</li> </ul> </li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Dynamic_Replan_Reason_Code (continued)	QQC22	<ul style="list-style-type: none"> <li>• AA - The sort sequence table specified is different from the sort sequence table that was used when this access plan was created.</li> <li>• AB - Storage pool changed or DEGREE parameter of CHGQRYA command changed.</li> <li>• AC - The system feature Db2 Symmetric Multiprocessing has been installed or removed.</li> <li>• AD - The value of the degree query attribute has changed.</li> <li>• AE - A view is either being opened by a high-level language or a view is being materialized.</li> <li>• AF - A user-defined type or user-defined function is not the same object as the one referred to in the access plan; or the SQL Path is not the same as when the access plan was built.</li> <li>• B0 - The options specified have changed as a result of the query options file.</li> <li>• B1 - The access plan was generated with a commitment control level that is different in the current job.</li> <li>• B2 - The access plan was generated with a static cursor answer set size that is different from the previous access plan.</li> <li>• B3 - The query was reoptimized because this run is the first run of the query after it was prepared. This run is the first run with actual parameter marker values.</li> <li>• B4 - The query was reoptimized because referential or check constraints have changed.</li> <li>• B5 - The query was reoptimized because Materialized query tables have changed.</li> <li>• B6 - The query was reoptimized because the value of a host variable changed and the access plan is no longer valid.</li> <li>• B7 - The query was reoptimized because AQP determined that it was beneficial.</li> <li>• B8 - The query was reoptimized because Expression Evaluator determined that the statement should be reoptimized.</li> </ul>
Data_Conversion_Reason_Code	QQC16	<p>Data conversion</p> <ul style="list-style-type: none"> <li>• N - No.</li> <li>• 0 - Not applicable.</li> <li>• 1 - Lengths do not match.</li> <li>• 2 - Numeric types do not match.</li> <li>• 3 - C host variable is NUL-terminated.</li> <li>• 4 - Host variable or column is variable length and the other is not variable length.</li> <li>• 5 - Host variable or column is not variable length and the other is variable length.</li> <li>• 6 - Host variable or column is variable length and the other is not variable length.</li> <li>• 7 - CCSID conversion.</li> <li>• 8 - DRDA and NULL capable, variable length, contained in a partial row, derived expression, or blocked fetch with not enough host variables.</li> <li>• 9 - Target table of an insert is not an SQL table.</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Data_Conversion_Reason_Code (continued)		<ul style="list-style-type: none"> <li>• 10 - Host variable is too short to hold a TIME or TIMESTAMP value being retrieved.</li> <li>• 11 - Host variable is DATE, TIME, or TIMESTAMP and value being retrieved is a character string.</li> <li>• 12 - Too many host variables specified and records are blocked.</li> <li>• 13 - DRDA used for a blocked FETCH. Also, the number of host variables specified in the INTO clause is less than the number of result values in the select list.</li> <li>• 14 - LOB locator used and the commitment control level was not *ALL.</li> </ul>
Total_Time_Milliseconds	QQI4	<p>Total time for this statement, in milliseconds. For fetches, the time includes all fetches for this OPEN of the cursor.</p> <p>Note: When monitor files are created when using an SQL Plan Cache snapshot, this time represents the aggregate time for all runs of this query. This time can be divided by the total number of runs, COALESCE(QVP15F,1), to determine an average time for a given run of the query.</p>
Rows_Fetched	QQI3	<p>Total rows fetched for cursor</p> <p>Note: When monitor files are created when using an SQL Plan Cache snapshot, this field is not set.</p>
End_Timestamp	QQETIM	Time SQL request completed
Total_Time_Microseconds	QQI6	<p>Total time for this statement, in microseconds. For fetches, this time includes all fetches for this OPEN of the cursor.</p> <p>Note: When monitor files are created when using an SQL Plan Cache snapshot, this time represents the aggregate time for all runs of this query. This time can be divided by the total number of runs, COALESCE(QVP15F,1), to determine an average time for a given run of the query.</p>
SQL_Statement_Length	QQI7	Length of SQL Statement



Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Insert_Unique_Count	QQI1	<p>If the operation (QQC21) indicates INSERT (IN), this field contains the unique query count for the QDT associated with the INSERT. QQUCNT contains the unique query count for the QDT associated with the WHERE part of the statement.</p> <p>If the operation (QQC21) indicates DELETE (DL) or TRUNCATE (TT), this field contains the fast delete reason code.</p> <p>Possible values if the operation is a DELETE or TRUNCATE are :</p> <ul style="list-style-type: none"> <li>• 0 - Fast delete results unknown or fast delete is not relevant because the delete failed.</li> <li>• 1 - Fast delete was achieved.</li> </ul> <p>All other values if the operation is a DELETE or TRUNCATE indicate the reason the database was unable to implement the request using fast delete. Fast delete attempt denied values:</p> <ul style="list-style-type: none"> <li>• 2 - File is a DDM file.</li> <li>• 3 - File is a multi member file.</li> <li>• 4 - File is distributed file.</li> <li>• 5 - File is a logical file or SQL view.</li> <li>• 6 - File is a parent file.</li> <li>• 7 - File has one or more enabled delete triggers created over it.</li> <li>• 8 - Number of rows in table is less than 1000 OR less than the QAQQINI SQL_FAST_DELETE_ROW_COUNT value. Refer to QVP151 to see the SQL_FAST_DELETE_ROW_COUNT value in effect for this statement.</li> <li>• 9 - DBMAINT failed. This reason code could appear for many reasons, including the existence of a logical open within this job, pending record changes, ragged save in progress and possibly other reasons.</li> <li>• 10- Failed to acquire an exclusive no read (LENR) lock on the file.</li> <li>• 11- Failed to acquire an exclusive allow read (LEAR) lock on the file's data space.</li> <li>• 12- The user does not have *EXECUTE authority to the library.</li> <li>• 13- File has a delete trigger that is enabled and RESTRICT WHEN DELETE TRIGGERS was specified on a TRUNCATE statement.</li> <li>• 14- File is a temporal table.</li> <li>• 51- A WHERE clause was used on the DELETE.</li> <li>• 52- QAQQINI SQL_FAST_DELETE_ROW_COUNT indicated to disallow fast delete.</li> <li>• 53- File is an alias referring to a partition table member.</li> <li>• 54- The user does not have *DELETE authority to the file.</li> <li>• 55- File is not found.</li> </ul>
SQLCode	QQI8	SQL return code
SQLState	QQC81	SQLSTATE
Close_Cursor_Mode	QVC101	<p>Close Cursor. Possible values are:</p> <ul style="list-style-type: none"> <li>• *ENDJOB - SQL cursors are closed when the job ends.</li> <li>• *ENDMOD - SQL cursors are closed when the module ends</li> <li>• *ENDPGM - SQL cursors are closed when the program ends.</li> <li>• *ENDSQL - SQL cursors are closed when the first SQL program on the call stack ends.</li> <li>• *ENDACTGRP - SQL cursors are closed when the activation group ends.</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Allow_Copy_Data_Value	QVC11	ALWCPYDTA setting (Y/N/O) <ul style="list-style-type: none"> <li>• Y - A copy of the data might be used.</li> <li>• N - Cannot use a copy of the data.</li> <li>• O - The optimizer can choose to use a copy of the data for performance.</li> </ul>
PseudoOpen	QVC12	Pseudo Open (Y/N) for SQL operations that can trigger opens. <ul style="list-style-type: none"> <li>• OP - Open</li> <li>• IN - Insert</li> <li>• UP - Update</li> <li>• DL - Delete</li> <li>• SI - Select Into</li> <li>• SV - Set</li> <li>• VI - Values into</li> </ul> For all operations, it can be blank.
PseudoClose	QVC13	Pseudo Close (Y/N) for SQL operations that can trigger a close. <ul style="list-style-type: none"> <li>• CL - Close</li> <li>• IN - Insert</li> <li>• UP - Update</li> <li>• DL - Delete</li> <li>• SI - Select Into</li> <li>• SV - Set</li> <li>• VI - Values into</li> </ul> For all operations, it can be blank.
ODP_Implementation	QVC14	ODP implementation <ul style="list-style-type: none"> <li>• R - Reusable ODP</li> <li>• N - Nonreusable ODP</li> <li>• ' ' - Column not used</li> </ul>
Dynamic_Replan_SubCode	QVC21	Dynamic replan, subtype reason code
Commitment_Control_Level	QVC41	Commitment control level. Possible values are: <ul style="list-style-type: none"> <li>• CS - Cursor stability</li> <li>• CSKL - Cursor stability. Keep exclusive locks.</li> <li>• NC - No commit</li> <li>• RR - Repeatable read</li> <li>• RREL - Repeatable read. Keep exclusive locks.</li> <li>• RS - Read stability</li> <li>• RSEL - Read stability. Keep exclusive locks.</li> <li>• UR - Uncommitted read</li> </ul>
Concurrent_Access_Resolution	QWC1B	Indicates what method of concurrent access resolution was specified. <ul style="list-style-type: none"> <li>• N - Concurrent access resolution was not specified.</li> <li>• S - SKIP LOCKED DATA clause was specified.</li> <li>• U - USE CURRENTLY COMMITTED clause was specified.</li> <li>• W- WAIT FOR OUTCOME clause was specified.</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Blocking_Type	QVC15	Type of blocking. Possible values are: <ul style="list-style-type: none"> <li>• S - Single row, ALWBLK(*READ)</li> <li>• F - Force one row, ALWBLK(*NONE)</li> <li>• L - Limited block, ALWBLK(*ALLREAD)</li> </ul>
Delay_Prepare	QVC16	Delay prepare of statement (Y/N).
Explainable	QVC1C	The SQL statement is explainable (Y/N).
Naming_Convention	QVC17	Naming convention. Possible values: <ul style="list-style-type: none"> <li>• N - System naming convention</li> <li>• S - SQL naming convention</li> </ul>
Dynamic_Processing_Type	QVC18	Type of dynamic processing. <ul style="list-style-type: none"> <li>• E - Extended dynamic</li> <li>• S - System wide cache</li> <li>• L - Local prepared statement</li> </ul>
LOB_Data_Optimized	QVC19	Optimize LOB data types (Y/N)
Program_User_Profile_Used	QVC1A	User profile used when compiled programs are executed. Possible values are: <ul style="list-style-type: none"> <li>• N = User Profile is determined by naming conventions. For *SQL, USRPRF(*OWNER) is used. For *SYS, USRPRF(*USER) is used.</li> <li>• U = USRPRF(*USER) is used.</li> <li>• O = USRPRF(*OWNER) is used.</li> </ul>
Dynamic_User_Profile_Used	QVC1B	User profile used for dynamic SQL statements. <ul style="list-style-type: none"> <li>• U = USRPRF(*USER) is used.</li> <li>• O = USRPRF(*OWNER) is used.</li> </ul>
Default_Collection	QVC1281	Name of the default collection.
Procedure_Name	QVC1282	Procedure name on CALL to SQL.
Procedure_Library	QVC1283	Procedure library on CALL to SQL.
SQL_Path	QQCLOB2	Path used to find procedures, functions, and user-defined types for static SQL statements.
Current_Schema	QVC1284	SQL current schema.
Binding_Type	QQC18	Binding type: <ul style="list-style-type: none"> <li>• C - Column-wise binding</li> <li>• R - Row-wise binding</li> </ul>
Cursor_Type	QQC61	Cursor Type: <ul style="list-style-type: none"> <li>• NSA - Non-scrollable, asensitive, forward only</li> <li>• NSI - Non-scrollable, insensitive, forward only</li> <li>• NSS - Non-scrollable, sensitive, forward only</li> <li>• SCA - scrollable, asensitive</li> <li>• SCI - scrollable, insensitive</li> <li>• SCS - scrollable, sensitive</li> </ul>
Statement_Originator	QVC1D	SQL statement originator: <ul style="list-style-type: none"> <li>• U - User</li> <li>• S - System</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Hard_Close_Reason_Code	QQC15	<p>SQL cursor hard close reason. Possible reasons are:</p> <ul style="list-style-type: none"> <li>• 1 - Internal Error</li> <li>• 2 - Exclusive Lock</li> <li>• 3 - Interactive SQL Reuse Restriction</li> <li>• 4 - Host variable Reuse Restriction</li> <li>• 5 - Temporary Result Restriction</li> <li>• 6 - Cursor Restriction</li> <li>• 7 - Cursor Hard Close Requested</li> <li>• 8 - Internal Error</li> <li>• 9 - Cursor Threshold</li> <li>• A - Optimizer decided to Hard-Close</li> <li>• B - Reuse Cursor Error</li> <li>• C - DRDA AS Cursor Closed</li> <li>• D - DRDA AR Not WITH HOLD</li> <li>• E - Repeatable Read</li> <li>• F - Lock Conflict Or QSQPRCED Threshold - Library</li> <li>• G - Lock Conflict Or QSQPRCED Threshold - File</li> <li>• H - Execute Immediate Access Plan Space</li> <li>• I - QSQCSRTH Dummy Cursor Threshold</li> <li>• J - File Override Change</li> <li>• K - Program Invocation Change</li> <li>• L - File Open Options Change</li> <li>• M - Statement Reuse Restriction</li> <li>• N - Internal Error</li> <li>• O - Library List Changed</li> <li>• P - Exit Processing</li> <li>• Q - SET SESSION USER statement</li> </ul>
Hard_Close_Subcode	QQC23	<p>SQL cursor hard close reason subcode.</p> <p>For QQC15 Reason code 'A' the following subcodes apply:</p> <ul style="list-style-type: none"> <li>• Z7 - New Index found</li> <li>• Z8 - Data Space Size changed out side of range</li> <li>• Z9 - MQT refresh age expired</li> <li>• ZA - Host variable values are no longer compatible with current plan</li> <li>• ZB - new statistic was found</li> <li>• ZC - commit level changed</li> <li>• ZD - Reoptimze for Warm IO</li> <li>• ZE - Reoptimze and change from FIRSTIO to ALLIO</li> <li>• ZF - Host variable selectivity changes require Reoptimization</li> <li>• ZG - AQP decided to hard-close</li> </ul>

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Date_Format	QVC42	Date Format. Possible values are: <ul style="list-style-type: none"> <li>• ISO</li> <li>• USA</li> <li>• EUR</li> <li>• JIS</li> <li>• JUL</li> <li>• MDY</li> <li>• DMY</li> <li>• YMD</li> </ul>
Date_Separator	QWC11	Date Separator. Possible values are: <ul style="list-style-type: none"> <li>• "/"</li> <li>• "."</li> <li>• ","</li> <li>• "-"</li> <li>• " "</li> </ul>
Time_Format	QVC43	Time Format. Possible values are: <ul style="list-style-type: none"> <li>• ISO</li> <li>• USA</li> <li>• EUR</li> <li>• JIS</li> <li>• HMS</li> </ul>
Time_Separator	QWC12	Time Separator. Possible values are: <ul style="list-style-type: none"> <li>• ":"</li> <li>• "."</li> <li>• ","</li> <li>• " "</li> </ul>
Decimal_Point	QWC13	Decimal Point. Possible values are: <ul style="list-style-type: none"> <li>• "."</li> <li>• ","</li> </ul>
Sort_Sequence_Table	QVC104	Sort Sequence Table
Sort_Sequence_Library	QVC105	Sort Sequence Library
Language_ID	QVC44	Language ID
Country_ID	QVC23	Country ID
First_N_Rows_Value	QQIA	Value specified on the FIRST n ROWS clause.
Optimize_For_N_Rows_Value	QQF1	Value specified on the OPTIMIZE FOR n ROWS clause.

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Access_Plan_Not_Saved_Reason_Code	QVC24	Access plan not saved reason code. Possible reasons are: <ul style="list-style-type: none"> <li>• A1 - Failed to get an LSUP lock on associated space of program or package.</li> <li>• A2 - Failed to get an immediate LEAR space location lock on first byte of associated space of program.</li> <li>• A3 - Failed to get an immediate LENR space location lock on first byte of associated space of program.</li> <li>• A5 - Failed to get an immediate LEAR space location lock on first byte of ILE associated space of a program.</li> <li>• A6 - Error trying to extend space of an ILE program.</li> <li>• A7 - No room in program.</li> <li>• A8 - No room in program associated space.</li> <li>• A9 - No room in program associated space.</li> <li>• AA - No need to save. Save already done in another job.</li> <li>• AB - Query optimizer cannot lock the QDT.</li> <li>• B1 - Saved at the end of the program associated space.</li> <li>• B2 - Saved at the end of the program associated space.</li> <li>• B3 - Saved in place.</li> <li>• B4 - Saved in place.</li> <li>• B5 - Saved at the end of the program associated space.</li> <li>• B6 - Saved in place.</li> <li>• B7 - Saved at the end of the program associated space.</li> <li>• B8 - Saved at the end of the program associated space.</li> </ul>
Transaction_Context_ID	QVC81	Transaction context ID.
Activation_Group_Mark	QVP152	Activation Group Mark
Open_Cursor_Threshold	QVP153	Open cursor threshold
Open_Cursor_Close_Count	QVP154	Open cursor close count
Commitment_Control_Lock_Limit	QVP155	Commitment control lock limit
Allow_SQL_Mixed_Constants	QWC15	Using SQL mixed constants (Y/N)
Suppress_SQL_Warnings	QWC16	Suppress SQL warning messages (Y/N)
Translate_ASCII	QWC17	Translate ASCII to job (Y/N)
System_Wide_Statement_Cache	QWC18	Using system-wide SQL statement cache (Y/N)
LOB_Locator_Threshold	QVP159	LOB locator threshold
Max_Decimal_Precision	QVP156	Maximum decimal precision (63/31)
Max_Decimal_Scale	QVP157	Maximum decimal scale
Min_Decimal_Divide_Scale	QVP158	Minimum decimal divide scale
Unicode_Normalization	QWC19	Unicode data normalization requested (Y/N)
Statement_Text_Long	QQ1000L	Complete statement text
Old_Access_Plan_Length	QVP15B	Length of old access plan
New_Access_Plan_Length	QVP15C	Length of new access plan

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Fast_Delete_Count	QVP151	SQL fast delete count. Possible values are: <ul style="list-style-type: none"> <li>• 0 = *OPTIMIZE or *DEFAULT</li> <li>• 1-999,999,999,999 = User specified value</li> <li>• 'FFFFFFFFFFFFFFFF'x = *NONE</li> </ul>
Statement_Max_Compression	QQF2	SQL statement maximum compression. Possible values are: <ul style="list-style-type: none"> <li>• 1 - *DEFAULT</li> <li>• 1 - User specified queries</li> <li>• 2 - All queries, user, and system</li> <li>• 3 - System generated internal queries</li> </ul>
Current_User_Profile	QVC102	Current user profile name
Expression_Evaluator_Used	QVC1E	<ul style="list-style-type: none"> <li>• N - Not applicable</li> <li>• S - SQL mapping</li> <li>• Y - QQ expression evaluator</li> <li>• O - Expression handled by an Open</li> </ul>
Host_Server_Delta	QVP15A	Time not spent within Host Server
NTS_Lock_Space_Id	QQC301	NTS Lock Space Identifier
IP_Address	QQC183	IP Address
IP_Type	QFC11	IP address type <ul style="list-style-type: none"> <li>• '0' = No client IP address</li> <li>• '1' = IPV4 format</li> <li>• '2' = IPV6 format</li> </ul> Only applicable for database server jobs.
IP_Port_Number	QQSMINT2	IP Port Number
NTS_Transaction_Id	QVC3004	NTS Transaction Identifier
NTS_Format_Id_Length	QQSMINT3	NTS Format Identified length
NTS_Transaction_ID_SubLength	QQSMINT4	NTS Transaction Identifier sublength.
Unique_Refresh_Counter2	QVRCNT	Unique refresh counter
Times_Run	QVP15F	<p>Number of times this Statement was run. If Null, then the statement was run once.</p> <p>Note: While using an SQL Plan Cache snapshot, this value can be set by the database monitor. This value might be null if the query never completed, or was running when the snapshot was created. If there is not a plan cache snapshot, the value is null.</p>
Full_Opens	QVP15E	<p>Number of runs that were processed as full opens. If Null, then the refresh count (qvrnt) is used to determine if the open was a full open (0) or a pseudo open (&gt;0).</p> <p>Note: While using an SQL Plan Cache snapshot, this value can be set by the database monitor. This value might be null if the query never completed, or was running when the snapshot was created. If there is not a plan cache snapshot, the value is null.</p>
Proc_In_Cache	QVC1F	Procedure definition was found in an internal cache. (Y/N) Only applicable for CALL statements.
Combined_Operation	QWC1A	Statement was performed with the processing for another statement. (Y/N) Only applicable for OPEN, FETCH, and CLOSE statements.
Client_Applname	QVC3001	Client Special Register - application name

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Client_Userid	QVC3002	Client Special Register - userid
Client_Wrkstnname	QVC3003	Client Special Register - work station name
Client_Acctng	QVC3005	Client Special Register - accounting string
Client_Programid	QVC3006	Client Special Register - program name
Interface_Information	QVC5001	Part of the CLIENT special register information. Three types of info are stored in this char500 column, separated by colons. <ul style="list-style-type: none"> <li>• First part, Interface Name, varchar(127);</li> <li>• Second part, Interface Level, varchar(63);</li> <li>• Third part, Interface Type, varchar(63)</li> </ul>
Open_Options	QVC82	Open options appear as a combination of the following characters, representing the actual capability for the cursor. The character values are left-aligned and padded on the right with blanks. Example 'RU ' indicate that the cursor is both read and update capable. <ul style="list-style-type: none"> <li>• R - Read capable</li> <li>• W - Write capable</li> <li>• U - Update capable</li> <li>• D - Delete capable</li> </ul>
Extended_Indicators	QWC1D	An Update or Insert statement was enabled to use extended indicators (Y/N).
DECFLOAT_Rounding_Mode	QWC1C	Rounding mode to use for DECFLOAT computations and conversions. <ul style="list-style-type: none"> <li>• 'E' = ROUND_HALF_EVEN</li> <li>• 'C' = ROUND_CEILING</li> <li>• 'D' = ROUND_DOWN</li> <li>• 'F' = ROUND_FLOOR</li> <li>• 'G' = ROUND_HALF_DOWN</li> <li>• 'H' = ROUND_HALF_UP</li> <li>• 'U' = ROUND_UP</li> </ul>
SQL_DECFLOAT_Warnings	QWC1E	DECFLOAT computations and conversions involving division by 0, overflow, underflow, an invalid operand, an inexact result, or a subnormal number results in a warning (Y/N).
Worst_Time_Micro	QVP15D	If not null, this time is the time for the slowest single run of this query. Note: When monitor files are created when using an SQL Plan Cache snapshot, this time represents the run time for the longest single run of the query. If the value is null, then the longest run information is not available. In that case, QQI6 might be the next best answer. See documentation for QQI6 for the proper use of that field
SQ_Unique_Count	QQINT05	A unique count used to uniquely identify statements which do not have an ODP but do pass in host variables. If QQUCNT is 0 and the statement passes in host variables, this value is non-zero. An example would be a CALL statement.



Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
Concurrent_Access_Res_Used	QFC13	Specifies what method of concurrent access resolution was used. <ul style="list-style-type: none"> <li>'N' = Concurrent access resolution is not applicable. This method applies to read queries with no commit or uncommitted read.</li> <li>'S' = SKIP LOCKED DATA clause was specified and rows with incompatible locks held by other transactions are skipped.</li> <li>'U' = USE CURRENTLY COMMITTED clause was specified and the currently committed version of data being updated or deleted is used. Data being inserted is skipped.</li> <li>'W' = Wait for commit or rollback when data is in the process of being inserted, updated, or deleted. This is the default method when the isolation level does not apply, the query is processed by CQE, or when not specified by the user.</li> </ul>
SQL_UDFs_Not_Inlined	QQSMINT8	Specifies the number of SQL user-defined scalar functions (UDFs) or user-defined table functions (UDTFs) that were not inlined in a SQL query or expression.
Result_Set_Cursor	QVC3007	Result Set Cursor name. Set by Allocate Cursor, Fetch, and Close.
Implicit_XMLPARSE_Option	QFC12	CURRENT IMPLICIT XMLPARSE OPTION special register. This option is used to specify white-space handling for an implicit parse of serialized XML data. <ul style="list-style-type: none"> <li>'S' = STRIP WHITESPACE</li> <li>'P' = PRESERVE WHITESPACE</li> </ul>
SQL_XML_Data_CCSID	QQSMINT7	The CCSID used for XML columns, host variables, parameter markers, and expressions if not explicitly specified.
OPTIMIZER_USE	QQSMINT5	Which optimizer was used for the query. Set to null if the monitor predates this option. <ul style="list-style-type: none"> <li>0 = Does not apply for this statement</li> <li>1 = SQE was used (SQL Query Engine)</li> <li>2 = CQE was used (Classic Query Engine)</li> <li>3 = CQE direct was used (statements like INSERT W/VALUES)</li> </ul>
XML_Schema_In_Cache	QFC14	The XML schema binary used during XMLVALIDATE or decomposition was found in the XML cache. <ul style="list-style-type: none"> <li>'Y' = Yes</li> <li>'N' = No</li> </ul>
Current_User	QQC105	The value of the CURRENT USER special register. The value only appears in the QQC105 column if the SQL statement used CURRENT USER.
Row_Column_Access_Control	QFC15	Type of row or column access applied. <ul style="list-style-type: none"> <li>' ' = Not applicable</li> <li>'C' = Column Access Control</li> <li>'R' = Row Access Control</li> <li>'B' = Both Row and Column Access Control</li> </ul>
Temporal_System_Time	QQTIM12A	The value of the CURRENT TEMPORAL SYSTEM_TIME special register. This TIMESTAMP(12) value is used as a default FOR SYSTEM_TIME period specification when a system-period temporal table is referenced in either static or dynamic SQL statements. This value will only be set when the SQL SYSTIME bind option is set to YES and the CURRENT TEMPORAL SYSTEM_TIME was applied to a system-period temporal table or when CURRENT TEMPORAL SYSTEM_TIME is explicitly used.

Table 240. QQQ1000 - SQL Information (continued)

View Column Name	Table Column Name	Description
SYSTIME_Bind_Option	QFC16	References to system-period temporal tables in both static or dynamic SQL statements are impacted by the value of the CURRENT TEMPORAL SYSTEM_TIME special register. <ul style="list-style-type: none"> <li>'Y' = Yes</li> <li>'N' = No</li> </ul>
Temporal_System_Time_Query	QFC17	A system time period specification was included on a system-period temporal table reference causing both current and historical rows to be queried. <ul style="list-style-type: none"> <li>'Y' = Yes, this is a system time temporal query.</li> <li>'N' = No, this is not a system time temporal query.</li> </ul>
DBCLOB_HOSTVR	QQDBCLOB1	Host variables values in a DBCLOB CCSID 1200 field (max 1MB). Only set if HOSTVAR(*CONDENSED) is used on STRDBMON, otherwise null.
StmtCmpReuseMin	QQSMINT6	SQL statement compression reuse minimum. This only applies to process extended dynamic *SQLPKGs. The minimum number of reuses that a statement must have in order for its access plan information to remain in the package across a compress of the package. <p>Possible values are:</p> <ul style="list-style-type: none"> <li>0 = *DEFAULT</li> <li>1-32767 = User specified value.</li> </ul>

## Database monitor view 3000 - Table Scan

Displays the SQL logical view format for database monitor QQQ3000

```

Create View QQQ3000 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QQTOTR as Table_Total_Rows,
 QQREST as Estimated_Rows_Selected,
 QQAJN as Estimated_Join_Rows,
 QQEPT as Estimated_Processing_Time,
 QQJNP as Join_Position,
 QQI1 as DataSpace_Number,
 QQC21 as Join_Method,
 QQC22 as Join_Type,
 QQC23 as Join_Operator,
 QQI2 as Index_Advised_Columns_Count,
 QQDSS as DataSpace_Selection,
)

```

```

QQIDXA as Index_Advised,
QQRCOD as Reason_Code,
QQIDXD as Index_Advised_Columns,
QVQTBL as Table_Name,
QVQLIB as Table_Schema,
QVPTBL as Base_Table_Name,
QVPLIB as Base_Table_Schema,
QVBNDY as Bound,
QVRCNT as Unique_Refresh_Counter,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QVPPRF as Parallel_Prefetch,
QVPPRPL as Parallel_PreLoad,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QVCTIM as Estimated_Cumulative_Time,
QQC11 as Skip_Sequential_Table_Scan,
QQI3 as Table_Size,
QVC3001 as DataSpace_Selection_Columns,
QQC14 as Derived_Column_Selection,
QVC3002 as Derived_Column_Selection_Columns,
QQC18 as Read_Trigger,
QVP157 as UDTF_Cardinality,
QVC1281 as UDTF_Specific_Name,
QVC1282 as UDTF_Specific_Schema,
QVP154 as Pool_Size,
QVP155 as Pool_Id,
QQC13 as MQT_Replacement,
QQC15 as InsertTable,
QQSMINTF as Plan_Iteration_Number,
QQF1 as Average_Read_Time
FROM UserLib/DBMONTABLE
WHERE (QQRID=3000)

```

Table 241. QQQ3000 - Table Scan

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level

Table 241. QQQ3000 - Table Scan (continued)

View Column Name	Table Column Name	Description
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member name of base table
Table_Total_Rows	QQTOTR	Total rows in table
Estimated_Rows_Selected	QQREST	Estimated number of rows selected
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI1	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>

Table 241. QQQ3000 - Table Scan (continued)

View Column Name	Table Column Name	Description
Index_Advised_Columns_Count	QQI2	Number of advised columns that use index scan-key positioning
DataSpace_Selection	QQDSS	Dataspace selection <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Index_Advised	QQIDXA	Index advised <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Reason_Code	QQRCD	Reason code <ul style="list-style-type: none"> <li>• T1 - No indexes exist.</li> <li>• T2 - Indexes exist, but none can be used.</li> <li>• T3 - Optimizer chose table scan over available indexes.</li> </ul>
Index_Advised_Columns	QQIDXD	Columns for the index advised
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Schema of base table, long name
Bound	QVBNDY	I/O or CPU bound. Possible values are: <ul style="list-style-type: none"> <li>• I - I/O bound</li> <li>• C - CPU bound</li> </ul>
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (Y/N)
Parallel_Degree_Requested	QVWARD	Parallel degree requested
Parallel_Degree_Used	QVWARD	Parallel degree used
Parallel_Degree_Reason_Code	QVWARD	Reason parallel processing was limited
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds

Table 241. QQQ3000 - Table Scan (continued)

View Column Name	Table Column Name	Description
Skip_Sequential_Table_Scan	QQC11	Skip sequential table scan (Y/N)
Table_Size	QQI3	Size of table being queried
DataSpace_Selection_Columns	QVC3001	Columns used for dataspace selection
Derived_Column_Selection	QQC14	Derived column selection (Y/N)
Derived_Column_Selection_Columns	QVC3002	Columns used for derived column selection
Read_Trigger	QQC18	Read Trigger (Y/N)
UDTF_Cardinality	QVP157	User-defined table function Cardinality
UDTF_Specific_Name	QVC1281	User-defined table function specific name
UDTF_Specific_Schema	QVC1282	User-defined table function specific schema
Pool_Size	QVP154	Memory pool size
Pool_Id	QVP155	Memory pool ID
MQT_Replacement	QQC13	Materialized Query Table replaced queried table (Y/N)
Insert_Table	QQC15	This is a target table of an insert (Y/N)
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1
Average_Read_Time	QQF1	Average disk I/O time for this object

## Database monitor view 3001 - Index Used

Displays the SQL logical view format for database monitor QQQ3001

```

Create View QQQ3001 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QQILNM as System_Index_Schema,
 QQIFNM as System_Index_Name,
 QQIMNM as Index_Member_Name,
 QQTOTR as Table_Total_Rows,
 QQREST as Estimated_Rows_Selected,
 QQFKEY as Index_Probe_Keys,

```

```

QQKSEL as Index_Scan_Keys,
QQAJN as Estimated_Join_Rows,
QQEPT as Estimated_Processing_Time,
QQJNP as Join_Position,
QQI1 as DataSpace_Number,
QQC21 as Join_Method,
QQC22 as Join_Type,
QQC23 as Join_Operator,
QQI2 as Index_Advised_Probe_Count,
QQKP as Index_Probe_Used,
QQI3 as Index_Probe_Column_Count,
QQKS as Index_Scan_Used,
QQDSS as DataSpace_Selection,
QQIDXA as Index_Advised,
QQRCOD as Reason_Code,
QQIDXD as Index_Advised_Columns,
QQC11 as Constraint,
QQ1000 as Constraint_Name,
QVQTBL as Table_Name,
QVQLIB as Table_Schema,
QVPTBL as Base_Table_Name,
QVPLIB as Base_Table_Schema,
QVINAM as Index_Name,
QVILIB as Index_Schema,
QVBNDY as Bound,
QVRCNT as Unique_Refresh_Counter,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QVPARPF as Parallel_Prefetch,
QVPARPL as Parallel_Preload,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QVCTIM as Estimated_Cumulative_Time,
QVc14 as Index_Only_Access,
QQc12 as Index_Fits_In_Memory,
QQC15 as Index_Type,
QVC12 as Index_Usage,
QQI4 as Index_Entries,
QQI5 as Unique_Keys,
QQI6 as Percent_Overflow,
QQI7 as Vector_Size,
QQI8 as Index_Size,
QQIA as Index_Page_Size,
QVP154 as Pool_Size,
QVP155 as Pool_Id,
QVP156 as Table_Size,
QQC16 as Skip_Sequential_Table_Scan,
QVC13 as Tertiary_Indexes_Exist,
QVC3001 as DataSpace_Selection_Columns,
QQC14 as Derived_Column_Selection,
QVC3002 as Derived_Column_Selection_Columns,
QVC3003 as Table_Columns_For_Index_Probe,
QVC3004 as Table_Columns_For_Index_Scan,
QVC3005 as Join_Selection_Columns,
QVC3006 as Ordering_Columns,
QVC3007 as Grouping_Columns,
QQC18 as Read_Trigger,
QVP157 as UDTF_Cardinality,
QVC1281 as UDTF_Specific_Name,
QVC1282 as UDTF_Specific_Schema,
QQC13 as MQT_Replacement,
QQSMINTF as Plan_Iteration_Number,
QVC3008 as Include_Values,
QVC15 as Sparse_Index,
QQF1 as Average_Read_Time,
QVC11 as Distinct_Indicator
FROM UserLib/DBMONTTable
WHERE QQRID=3001)

```

Table 242. QQQ3001 - Index Used

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created

Table 242. QQQ3001 - Index Used (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User-defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member name of base table
System_Index_Schema	QQILNM	Schema name of index used for access
System_Index_Name	QQIFNM	Name of index used for access
Index_Member_Name	QQIMNM	Member name of index used for access
Table_Total_Rows	QQTOTR	Total rows in base table
Estimated_Rows_Selected	QQREST	Estimated number of rows selected
Index_Probe_Keys	QQFKEY	Rows selected through index scan-key positioning



Table 242. QQQ3001 - Index Used (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Index_Scan_Keys	QQKSEL	Rows selected through index scan-key selection
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI1	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Index_Advised_Probe_Count	QQI2	Number of advised key columns that use index scan-key positioning
Index_Probe_Used	QQKP	Index scan-key positioning <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Index_Probe_Column_Count	QQI3	Number of columns that use index scan-key positioning for the index used
Index_Scan_Used	QQKS	Index scan-key selection <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
DataSpace_Selection	QQDSS	Dataspace selection <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>

Table 242. QQQ3001 - Index Used (continued)

View Column Name	Table Column Name	Description
Index_Advised	QQIDXA	Index advised <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Reason_Code	QQRCOD	Reason code <ul style="list-style-type: none"> <li>• I1 - Row selection</li> <li>• I2 - Ordering/Grouping</li> <li>• I3 - Row selection and Ordering/Grouping</li> <li>• I4 - Nested loop join</li> <li>• I5 - Row selection using bitmap processing</li> <li>• I7 - Index Merge Ordering</li> <li>• I8 - Encoded Vector Index Only Column Projection</li> </ul>
Index_Advised_Columns	QQIDXD	Columns for index advised
Constraint	QQC11	Index is a constraint (Y/N)
Constraint_Name	QQ1000	Constraint name
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Schema of base table, long name
Index_Name	QVINAM	Name of index (or constraint) used, long name
Index_Schema	QVILIB	Library of index used, long name
Bound	QVBNDY	I/O or CPU bound. Possible values are: <ul style="list-style-type: none"> <li>• I - I/O bound</li> <li>• C - CPU bound</li> </ul>
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Join_Fanout	QVJFANO	Join fanout. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_Preload	QVPARPL	Parallel Preload (Y/N)
Parallel_Degree_Requested	QVPARD	Parallel degree requested
Parallel_Degree_Used	QVPARU	Parallel degree used

Table 242. QQQ3001 - Index Used (continued)

View Column Name	Table Column Name	Description
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Index_Only_Access	QVC14	Index only access (Y/N)
Index_Fits_In_Memory	QQC12	Index fits in memory (Y/N)
Index_Type	QQC15	Type of Index. Possible values are: <ul style="list-style-type: none"> <li>• B - Binary Radix Index</li> <li>• C - Constraint (Binary Radix)</li> <li>• E - Encoded Vector Index (EVI)</li> <li>• X - Query created temporary index</li> </ul>
Index_Usage	QVC12	Index Usage. Possible values are: <ul style="list-style-type: none"> <li>• P - Primary Index</li> <li>• T - Tertiary (AND or OR) Index</li> </ul>
Index_Entries	QQI4	Number of index entries
Unique_Keys	QQI5	Number of unique key values
Percent_Overflow	QQI6	Percent overflow
Vector_Size	QQI7	Vector size
Index_Size	QQI8	Index size
Index_Page_Size	QQIA	Index page size
Pool_Size	QVP154	Pool size
Pool_Id	QVP155	Pool ID
Table_Size	QVP156	Table size
Skip_Sequential_Table_Scan	QQC16	Skip sequential table scan (Y/N)
Tertiary_Indexes_Exist	QVC13	Tertiary indexes exist (Y/N)
DataSpace_Selection_Columns	QVC3001	Columns used for dataspace selection
Derived_Column_Selection	QQC14	Derived column selection (Y/N)
Derived_Column_Selection_Columns	QVC3002	Columns used for derived column selection
Table_Column_For_Index_Probe	QVC3003	Columns used for index scan-key positioning
Table_Column_For_Index_Scan	QVC3004	Columns used for index scan-key selection
Join_Selection_Columns	QVC3005	Columns used for Join selection
Ordering_Columns	QVC3006	Columns used for Ordering
Grouping_Columns	QVC3007	Columns used for Grouping
Read_Trigger	QQC18	Read Trigger (Y/N)
UDTF_Cardinality	QVP157	Cardinality for user-defined table function.
UDTF_Specific_Name	QVC1281	Specific name for user-defined table function.

Table 242. QQQ3001 - Index Used (continued)

View Column Name	Table Column Name	Description
UDTF_Specific_Schema	QVC1282	Specific schema for user-defined table function.
MQT_Replacement	QQC13	Materialized Query Table replaced queried table (Y/N)
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1
Include_Values	QVC3008	Encoded Vector indexes only. Aggregates included as part of index creation and predetermined for Grouping query request.
Sparse_Index	QVC15	Index contains sparse selection or Select/Omit selection criteria (Y/N).
Average_Read_Time	QQF1	Average disk I/O time for this object
Distinct_Indicator	QVC11	Distinct Scan or Probe was used (Y/N)

## Database monitor view 3002 - Index Created

Displays the SQL logical view format for database monitor QQQ3002.

```

Create View QQQ3002 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QQILNM as System_Index_Schema,
 QQIFNM as System_Index_Name,
 QQIMNM as Index_Member_Name,
 QQNTNM as NLSS_Table,
 QQNLNM as NLSS_Library,
 QQSTIM as Start_Timestamp,
 QQETIM as End_Timestamp,
 QQTOTR as Table_Total_Rows,
 QQRIDX as Created_Index_Entries,
 QQREST as Estimated_Rows_Selected,
 QQFKEY as Index_Probe_Keys,
 QQKSEL as Index_Scan_Keys,
 QQAJN as Estimated_Join_Rows,
 QQEPT as Estimated_Processing_Time,
 QQJNP as Join_Position,
 QQI1 as DataSpace_Number,
 QQC21 as Join_Method,
 QQC22 as Join_Type,
 QQC23 as Join_Operator,
 QQI2 as Index_Advised_Probe_Count,

```

```

QQKP as Index_Probe_Used,
QQI3 as Index_Probe_Column_Count,
QQKS as Index_Scan_Used,
QQDSS as DataSpace_Selection,
QQIDXA as Index_Advised,
QQRCOD as Reason_Code,
QQIDXD as Index_Advised_Columns,
QQ1000 as Created_Index_Columns,
QVQTBL as Table_Name,
QVQLIB as Table_Schema,
QVPTBL as Base_Table_Name,
QVPLIB as Base_Table_Schema,
QVINAM as Index_Name,
QVILIB as Index_Schema,
QVBNDY as Bound,
QVRCNT as Unique_Refresh_Counter,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QVPPFP as Parallel_Prefetch,
QVPPPL as Parallel_Preload,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QVCTIM as Estimated_Cumulative_Time,
QQC101 as Created_Index_Name,
QQC102 as Created_Index_Schema,
QQI4 as Created_Index_Page_Size,
QQI5 as Created_Index_Row_Size,
QQC14 as Created_Index_Used_ACS_Table,
QQC103 as Created_Index_ACS_Table,
QQC104 as Created_Index_ACS_Library,
QVC13 as Created_Index_Reusable,
QVC14 as Created_Index_Sparse,
QVC1F as Created_Index_Type,
QVP15F as Created_Index_Unique_EVI_Count,
QVC15 as Permanent_Index_Created,
QVC16 as Index_From_Index,
QVP151 as Created_Index_Parallel_Degree_Requested,
QVP152 as Created_Index_Parallel_Degree_Used,
QVP153 as Created_Index_Parallel_Degree_Reason_Code,
QVC17 as Index_Only_Access,
QVC18 as Index_Fits_In_Memory,
QVC1B as Index_Type,
QQI6 as Index_Entries,
QQI7 as Unique_Keys,
QVP158 as Percent_Overflow,
QVP159 as Vector_Size,
QQI8 as Index_Size,
QVP156 as Index_Page_Size,
QVP154 as Pool_Size,
QVP155 as Pool_ID,
QVP157 as Table_Size,
QVC1C as Skip_Sequential_Table_Scan,
QVC3001 as DataSpace_Selection_Columns,
QVC1E as Derived_Column_Selection,
QVC3002 as Derived_Column_Selection_Columns,
QVC3003 as Table_Column_For_Index_Probe,
QVC3004 as Table_Column_For_Index_Scan,
QQC18 as Read_Trigger,
QQC13 as MQT_Replacement,
QQC16 as Reused_Temporary_Index,
QQINT03 as Estimated_Storage,
QQSMINTF as Plan_Iteration_Number,
QQF1 as Average_Read_Time
FROM UserLib/DBMONTTable
WHERE QQRID=3002)

```

Table 243. QQQ3002 - Index Created

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created

Table 243. QQQ3002 - Index Created (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member name of base table
System_Index_Schema	QQILNM	Schema name of index used for access
System_Index_Name	QQIFNM	Name of index used for access
Index_Member_Name	QQIMNM	Member name of index used for access
NLSS_Table	QQNTNM	NLSS table
NLSS_Library	QQNLNM	NLSS library
Start_Timestamp	QQSTIM	Start timestamp, when available.
End_Timestamp	QQETIM	End timestamp, when available
Table_Total_Rows	QQTOTR	Total rows in table

Table 243. QQQ3002 - Index Created (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Created_Index_Entries	QQRIDX	Number of entries in index created
Estimated_Rows_Selected	QQREST	Estimated number of rows selected
Index_Probe_Keys	QQFKEY	Keys selected thru index scan-key positioning
Index_Scan_Keys	QQKSEL	Keys selected thru index scan-key selection
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI1	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Index_Advised_Probe_Count	QQI2	Number of advised key columns that use index scan-key positioning
Index_Probe_Used	QQKP	Index scan-key positioning <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Index_Probe_Column_Count	QQI3	Number of columns that use index scan-key positioning for the index used
Index_Scan_Used	QQKS	Index scan-key selection <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>

Table 243. QQQ3002 - Index Created (continued)

View Column Name	Table Column Name	Description
DataSpace_Selection	QQDSS	Dataspace selection <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Index_Advised	QQIDXA	Index advised <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Reason_Code	QQRCOD	Reason code <ul style="list-style-type: none"> <li>• I1 - Row selection</li> <li>• I2 - Ordering/Grouping</li> <li>• I3 - Row selection and Ordering/Grouping</li> <li>• I4 - Nested loop join</li> </ul>
Index_Advised_Columns	QQIDXD	Key columns for index advised
Created_Index_Columns	QQ1000	Key columns for index created
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Schema of base table, long name
Index_Name	QVINAM	Name of index (or constraint) used, long name
Index_Schema	QVILIB	Schema of index used, long name
Bound	QVBNDY	I/O or CPU bound. Possible values are: <ul style="list-style-type: none"> <li>• I - I/O bound</li> <li>• C - CPU bound</li> </ul>
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_Preload	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)



Table 243. QQQ3002 - Index Created (continued)

View Column Name	Table Column Name	Description
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Created_Index_Name	QQC101	Name of index created - when available
Created_Index_Schema	QQC102	Schema of index created - when available
Created_Index_Page_Size	QQI4	Page size of index created
Created_Index_Row_Size	QQI5	Row size of index created
Created_Index_Used_ACS_Table	QQC14	Index Created used Alternate Collating Sequence Table (Y/N)
Created_Index_ACS_Table	QQC103	Alternate Collating Sequence table of index created.
Created_Index_ACS_Library	QQC104	Alternate Collating Sequence library of index created.
Created_Index_Reusable	QVC13	Index created is reusable (Y/N)
Created_Index_Sparse	QVC14	Index created is sparse index (Y/N)
Created_Index_Type	QVC1F	Type of index created. Possible values: <ul style="list-style-type: none"> <li>• B - Binary Radix Index</li> <li>• E - Encoded Vector Index (EVI)</li> </ul>
Created_Index_Unique_EVI_Count	QVP15F	Number of unique values of index created if index created is an EVI index.
Permanent_Index_Created	QVC15	Permanent index created (Y/N)
Index_From_Index	QVC16	Index from index (Y/N)
Created_Index_Parallel_Degree_Requested	QVP151	Parallel degree requested (index created)
Created_Index_Parallel_Degree_Used	QVP152	Parallel degree used (index created)
Created_Index_Parallel_Degree_Reason_Code	QVP153	Reason parallel processing was limited (index created)
Index_Only_Access	QVC17	Index only access (Y/N)
Index_Fits_In_Memory	QVC18	Index fits in memory (Y/N)
Index_Type	QVC1B	Type of Index. Possible values are: <ul style="list-style-type: none"> <li>• B - Binary Radix Index</li> <li>• C - Constraint (Binary Radix)</li> <li>• E - Encoded Vector Index (EVI)</li> <li>• T - Tertiary (AND/OR) Index</li> </ul>
Index_Entries	QQI6	Number of index entries, index used
Unique_Keys	QQI7	Number of unique key values, index used
Percent_Overflow	QVP158	Percent overflow, index used

Table 243. QQQ3002 - Index Created (continued)

View Column Name	Table Column Name	Description
Vector_Size	QVP159	Vector size, index used
Index_Size	QQI8	Size of index used.
Index_Page_Size	QVP156	Index page size
Pool_Size	QVP154	Pool size
Pool_ID	QVP155	Pool id
Table_Size	QVP157	Table size
Skip_Sequential_Table_Scan	QVC1C	Skip sequential table scan (Y/N)
DataSpace_Selection_Columns	QVC3001	Columns used for dataspace selection
Derived_Column_Selection	QVC1E	Derived column selection (Y/N)
Derived_Column_Selection_Columns	QVC3002	Columns used for derived column selection
Table_Columns_For_Index_Probe	QVC3003	Columns used for index scan-key positioning
Table_Columns_For_Index_Scan	QVC3004	Columns used for index scan-key selection
Read_Trigger	QQC18	Read Trigger (Y/N)
MQT_Replacement	QQC13	Materialized Query Table replaced queried table (Y/N)
Reused_Temporary_Index	QQC16	Temporary index reused (Y/N)
Estimated_Storage	QQINT03	Estimated amount of temporary storage used, in megabytes, to create the temporary index.
Plan_Iteration_Number	QQSMINT F	AQP Plan iteration number, original optimization = 1
Average_Read_Time	QQF1	Average disk I/O time for this object

### Database monitor view 3003 - Query Sort

Displays the SQL logical view format for database monitor QQQ3003.

```

Create View QQQ3003 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQSTIM as Start_Timestamp,
 QQETIM as End_Timestamp,
)

```

```

QQRSS as Sorted_Rows,
QQI1 as Sort_Space_Size,
QQI2 as Pool_Size,
QQI3 as Pool_Id,
QQI4 as Internal_Sort_Buffer_Length,
QQI5 as External_Sort_Buffer_Length,
QQRCD as Reason_Code,
QQI7 as Union_Reason_Subcode,
QVBNDY as Bound,
QVRCNT as Unique_Refresh_Counter,
QVPPRF as Parallel_Prefetch,
QVPPRL as Parallel_PreLoad,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QQEPT as Estimated_Processing_Time,
QVCTIM as Estimated_Cumulative_Time,
QQAJN as Estimated_Join_Rows,
QQJNP as Join_Position,
QQI6 as DataSpace_Number,
QQC21 as Join_Method,
QQC22 as Join_Type,
QQC23 as Join_Operator,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QQINT03 as Estimated_Storage,
QQSMINTF as Plan_Iteration_Number,
QQF1 as Average_Read_Time
FROM UserLib/DBMONTTable
WHERE QQRID=3003)

```

Table 244. QQQ3003 - Query Sort

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects

Table 244. QQQ3003 - Query Sort (continued)

View Column Name	Table Column Name	Description
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Start_Timestamp	QQSTIM	Start timestamp, when available
End_Timestamp	QQETIM	End timestamp, when available
Sorted_Rows	QQRSS	Estimated number of rows selected or sorted.
Sort_Space_Size	QQI1	Estimated size of sort space.
Pool_Size	QQI2	Pool size
Pool_Id	QQI3	Pool id
Internal_Sort_Buffer_Length	QQI4	Internal sort buffer length
External_Sort_Buffer_Length	QQI5	External sort buffer length
Reason_Code	QQRCOD	Reason code <ul style="list-style-type: none"> <li>• F1 - Query contains grouping columns (GROUP BY) from more than one table, or contains grouping columns from a secondary table of a join query that cannot be reordered.</li> <li>• F2 - Query contains ordering columns (ORDER BY) from more than one table, or contains ordering columns from a secondary table of a join query that cannot be reordered.</li> <li>• F3 - The grouping and ordering columns are not compatible.</li> <li>• F4 - DISTINCT was specified for the query.</li> <li>• F5 - UNION was specified for the query.</li> <li>• F6 - Query had to be implemented using a sort. Key length of more than 2000 bytes or more than 120 key columns specified for ordering.</li> </ul>
Reason_Code (continued)		<ul style="list-style-type: none"> <li>• F7 - Query optimizer chose to use a sort rather than an index to order the results of the query.</li> <li>• F8 - Perform specified row selection to minimize I/O wait time.</li> <li>• FC - The query contains grouping fields and there is a read trigger on at least one of the physical files in the query.</li> </ul>
Union_Reason_Subcode	QQI7	Reason subcode for Union: <ul style="list-style-type: none"> <li>• 51 - Query contains UNION and ORDER BY</li> <li>• 52 - Query contains UNION ALL</li> </ul>

Table 244. QQQ3003 - Query Sort (continued)

View Column Name	Table Column Name	Description
Bound	QVBNDY	I/O or CPU bound. Possible values are: <ul style="list-style-type: none"> <li>• I - I/O bound</li> <li>• C - CPU bound</li> </ul>
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>

Table 244. QQQ3003 - Query Sort (continued)

View Column Name	Table Column Name	Description
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Estimated_Storage	QQINT03	Estimated amount of temporary storage used, in megabytes, to create the temporary index.
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1
Average_Read_Time	QQF1	Average disk I/O time for this object

## Database monitor view 3004 - Temp Table

Displays the SQL logical view format for database monitor QQQ3004.

```

Create View QQQ3004 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QQSTIM as Start_Timestamp,
 QQETIM as End_Timestamp,
 QQC11 as Has_Default_Values,
 QQTMPR as Table_Rows,
 QQRCOD as Reason_Code,
 VQTBL as Table_Name,
 VQLIB as Table_Schema,
 VPTBL as Base_Table_Name,
 VPLIB as Base_Table_Schema,
 QQC101 as Temporary_Table_Name,
 QQC102 as Temporary_Table_Schema,
 VBN DY as Bound,
 QVRCNT as Unique_Refresh_Counter,
 QVJFANO as Join_Fanout,
 VFILES as Join_Table_Count,
 VPARPF as Parallel_Prefetch,
 VPARPL as Parallel_PreLoad,
 VPARD as Parallel_Degree_Requested,
)

```

```

QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QQEPT as Estimated_Processing_Time,
QVCTIM as Estimated_Cumulative_Time,
QQAJN as Estimated_Join_Rows,
QQJNP as Join_Position,
QQI6 as DataSpace_Number,
QQC21 as Join_Method,
QQC22 as Join_Type,
QQC23 as Join_Operator,
QQI2 as Temporary_Table_Row_Size,
QQI3 as Temporary_Table_Size,
QQC12 as Temporary_Query_Result,
QQC13 as Distributed_Temporary_Table,
QVC3001 as Distributed_Temporary_Data_Nodes,
QQI7 as Materialized_Subquery_QDT_Level,
QQI8 as Materialized_Union_QDT_Level,
QQC14 as View_Contains_Union,
QQINT03 as Estimated_Storage,
QQSMINTF as Plan_Iteration_Number,
QQF1 as Average_Read_Time
FROM UserLib/DBMONTTable
WHERE QQRID=3004)

```

Table 245. QQQ3004 - Temp Table

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect

Table 245. QQQ3004 - Temp Table (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member name of base table
Start_Timestamp	QQSTIM	Start timestamp, when available
End_Timestamp	QQETIM	End timestamp, when available
Has_Default_Values	QQC11	Default values may be present in temporary <ul style="list-style-type: none"> <li>• Y - Yes</li> <li>• N - No</li> </ul>
Table_Rows	QQTMPR	Estimated number of rows in the temporary



Table 245. QQQ3004 - Temp Table (continued)

View Column Name	Table Column Name	Description
Reason_Code	QQRCD	<p>Reason code. Possible values are:</p> <ul style="list-style-type: none"> <li>• F1 - Query contains grouping columns (GROUP BY) from more than one table, or contains grouping columns from a secondary table of a join query that cannot be reordered.</li> <li>• F2 - Query contains ordering columns (ORDER BY) from more than one table, or contains ordering columns from a secondary table of a join query that cannot be reordered.</li> <li>• F3 - The grouping and ordering columns are not compatible.</li> <li>• F4 - DISTINCT was specified for the query.</li> <li>• F5 - UNION was specified for the query.</li> <li>• F6 - Query had to be implemented using a sort. Key length of more than 2000 bytes or more than 120 key columns specified for ordering.</li> <li>• F7 - Query optimizer chose to use a sort rather than an index to order the results of the query.</li> <li>• F8 - Perform specified row selection to minimize I/O wait time.</li> <li>• F9 - The query optimizer chose to use a hashing algorithm rather than an index to perform the grouping.</li> <li>• FA - The query contains a join condition that requires a temporary table</li> <li>• FB - The query optimizer creates a run-time temporary file in order to implement certain correlated group by queries.</li> <li>• FC - The query contains grouping fields and there is a read trigger on at least one of the physical files in the query.</li> <li>• FD - The query optimizer creates a runtime temporary file for a static-cursor request.</li> <li>• H1 - Table is a join logical file and its join type does not match the join type specified in the query.</li> <li>• H2 - Format specified for the logical table references more than one base table.</li> <li>• H3 - Table is a complex SQL view requiring a temporary table to contain the results of the SQL view.</li> <li>• H4 - For an update-capable query, a subselect references a column in this table which matches one of the columns being updated.</li> </ul>

Table 245. QQQ3004 - Temp Table (continued)

View Column Name	Table Column Name	Description
Reason_Code (continued)	QQRCOD	Reason code. Possible values are: <ul style="list-style-type: none"> <li>• H5 - For an update-capable query, a subselect references an SQL view which is based on the table being updated.</li> <li>• H6 - For a delete-capable query, a subselect references either the table from which rows are to be deleted, an SQL view, or an index based on the table from which rows are to be deleted</li> <li>• H7 - A user-defined table function was materialized.</li> <li>• H8 - The query optimizer created an OLAP window temporary.</li> </ul>
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Library of base table, long name
Temporary_Table_Name	QQC101	Temporary table name
Temporary_Table_Schema	QQC102	Temporary table schema
Bound	QVBNDY	I/O or CPU bound. Possible values are: <ul style="list-style-type: none"> <li>• I - I/O bound</li> <li>• C - CPU bound</li> </ul>
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (Y/N)
Parallel_Degree_Requested	QVPARD	Parallel degree requested
Parallel_Degree_Used	QVPARU	Parallel degree used
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds

Table 245. QQQ3004 - Temp Table (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Temporary_Table_Row_Size	QQI2	Row size of temporary table, in bytes
Temporary_Table_Size	QQI3	Estimated size of temporary table, in bytes.
Temporary_Query_Result	QQC12	Temporary result table that contains the results of the query. (Y/N)
Distributed_Temporary_Table	QQC13	Distributed Table (Y/N)
Distributed_Temporary_Data_Nodes	QVC3001	Data nodes of temporary table
Materialized_Subquery_QDT_Level	QQI7	Materialized subquery QDT level
Materialized_Union_QDT_Level	QQI8	Materialized Union QDT level
View_Contains_Union	QQC14	Union in a view (Y/N)
Estimated_Storage	QQINT03	Estimated amount of temporary storage used, in megabytes, to create the temporary index.
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1
Average_Read_Time	QQF1	Average disk I/O time for this object

## Database monitor view 3005 - Table Locked

Displays the SQL logical view format for database monitor QQQ3005.

```

Create View QQQ3005 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QQC11 as Lock_Success,
 QQC12 as Unlock_Request,
 QQRCOD as Reason_Code,
 QVQTBL as Table_Name,
 VQLIB as Table_Schema,
 QVPTBL as Base_Table_Name,
 VPLIB as Base_Table_Schema,
 QQJNP as Join_Position,
 QQI6 as DataSpace_Number,
 QQC21 as Join_Method,
 QQC22 as Join_Type,
 QQC23 as Join_Operator,
 QVJFANO as Join_Fanout,
 QVFILES as Join_Table_Count,
 QVRCNT as Unique_Refresh_Counter
 FROM UserLib/DBMONTTable
 WHERE QQRID=3005)

```

Table 246. QQQ3005 - Table Locked

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)

Table 246. QQQ3005 - Table Locked (continued)

View Column Name	Table Column Name	Description
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Numbe r	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSel ects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Co de	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member name of base table
Lock_Success	QQC11	Successful lock indicator (Y/N)
Unlock_Request	QQC12	Unlock request (Y/N)
Reason_Code	QQRCOD	Reason code <ul style="list-style-type: none"> <li>• L1 - UNION with *ALL or *CS with Keep Locks</li> <li>• L2 - DISTINCT with *ALL or *CS with Keep Locks</li> <li>• L3 - No duplicate keys with *ALL or *CS with Keep Locks</li> <li>• L4 - Temporary needed with *ALL or *CS with Keep Locks</li> <li>• L5 - System Table with *ALL or *CS with Keep Locks</li> <li>• L6 - Orderby &gt; 2000 bytes with *ALL or *CS with Keep Locks</li> <li>• L9 - Unknown</li> <li>• LA - User-defined table function with *ALL or *CS with Keep Locks</li> </ul>
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name

Table 246. QQQ3005 - Table Locked (continued)

View Column Name	Table Column Name	Description
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Schema of base table, long name
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Unique_Refresh_Counter	QVRCNT	Unique refresh counter

### Database monitor view 3006 - Access Plan Rebuilt

Displays the SQL logical view format for database monitor QQQ3006.

```

Create View QQQ3006 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,

```

```

QQUSER as Job_User,
QQJNUM as Job_Number,
QQI9 as Thread_ID,
QQUCNT as Unique_Count,
QQUDEF as User_Defined,
QQQDTN as Unique_SubSelect_Number,
QQQDTL as SubSelect_Nested_Level,
QQMATN as Materialized_View_Subselect_Number,
QQMATL as Materialized_View_Nested_Level,
QVP15E as Materialized_View_Union_Level,
QVP15A as Decomposed_Subselect_Number,
QVP15B as Total_Number_Decomposed_SubSelects,
QVP15C as Decomposed_SubSelect_Reason_Code,
QVP15D as Starting_Decomposed_SubSelect,
QQRCD as Reason_Code,
QQC21 as SubCode,
QVRCNT as Unique_Refresh_Counter,
QQTIM1 as Last_Access_Plan_Rebuild_Timestamp,
QQC11 as Reoptimization_Done,
QVC22 as Previous_Reason_Code,
QVC23 as Previous_SubCode,
QQSMINTF as Plan_Iteration_Number
FROM UserLib/DBMONTable
WHERE QQRID=3006)

```

Table 247. QQQ3006 - Access Plan Rebuilt

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Numbe r	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSele cts	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Co de	QVP15C	Decomposed query subselect reason code

Table 247. QQQ3006 - Access Plan Rebuilt (continued)

View Column Name	Table Column Name	Description
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Reason_Code	QQRCD	<p data-bbox="802 384 1292 411">Reason code why access plan was rebuilt</p> <ul style="list-style-type: none"> <li data-bbox="802 432 1451 527">• A1 - A table or member is not the same object as the one referenced when the access plan was last built. Some reasons they might be different are:               <ul style="list-style-type: none"> <li data-bbox="824 541 1260 569">– Object was deleted and recreated.</li> <li data-bbox="824 583 1224 611">– Object was saved and restored.</li> <li data-bbox="824 625 1146 653">– Library list was changed.</li> <li data-bbox="824 667 1101 695">– Object was renamed.</li> <li data-bbox="824 709 1078 737">– Object was moved.</li> <li data-bbox="824 751 1365 779">– Object was overridden to a different object.</li> <li data-bbox="824 793 1419 852">– This is the first run of this query after the object containing the query has been restored.</li> <li data-bbox="824 867 1398 926">– Mask or permission attributes changed for the object.</li> </ul> </li> <li data-bbox="802 947 1414 1041">• A2 - Access plan was built to use a reusable Open Data Path (ODP) and the optimizer chose to use a non-reusable ODP for this call.</li> <li data-bbox="802 1056 1472 1150">• A3 - Access plan was built to use a non-reusable Open Data Path (ODP) and the optimizer chose to use a reusable ODP for this call.</li> <li data-bbox="802 1165 1451 1224">• A4 - The number of rows in the table has changed by more than 10% since the access plan was last built.</li> <li data-bbox="802 1239 1438 1297">• A5 - A new index exists over one of the tables in the query</li> <li data-bbox="802 1312 1425 1371">• A6 - An index that was used for this access plan no longer exists or is no longer valid.</li> <li data-bbox="802 1386 1472 1444">• A7 - IBM i Query requires the access plan to be rebuilt because of system programming changes.</li> <li data-bbox="802 1459 1459 1518">• A8 - The CCSID of the current job is different than the CCSID of the job that last created the access plan.</li> <li data-bbox="802 1533 1472 1793">• A9 - The value of one or more of the following is different for the current job than it was for the job that last created this access plan:               <ul style="list-style-type: none"> <li data-bbox="824 1640 992 1667">– date format</li> <li data-bbox="824 1682 1027 1709">– date separator</li> <li data-bbox="824 1724 992 1751">– time format</li> <li data-bbox="824 1766 1027 1793">– time separator.</li> </ul> </li> </ul>



Table 247. QQQ3006 - Access Plan Rebuilt (continued)

View Column Name	Table Column Name	Description
Reason_Code (continued)	QQRCOD	<ul style="list-style-type: none"> <li>• AA - The sort sequence table specified is different than the sort sequence table that was used when this access plan was created.</li> <li>• AB - Storage pool changed.</li> <li>• AC - The system feature Db2 multisystem has been installed or removed.</li> <li>• AD - The value of the degree query attribute has changed.</li> <li>• AE - A view is either being opened by a high level language or a view is being materialized.</li> <li>• AF - A sequence object or user-defined type or function is not the same object as the one referred to in the access plan; or, the SQL path used to generate the access plan is different than the current SQL path.</li> <li>• B0 - The options specified have changed as a result of the query options file.</li> <li>• B1 - The access plan was generated with a commitment control level that is different in the current job.</li> <li>• B2 - The access plan was generated with a static cursor answer set size that is different than the previous access plan.</li> <li>• B3 - The query was reoptimized because this is the first run of the query after a prepare. That is, it is the first run with real actual parameter marker values.</li> <li>• B4 - The query was reoptimized because referential or check constraints have changed.</li> <li>• B5 - The query was reoptimized because MQTs have changed.</li> <li>• B6 - The query was reoptimized because the value of a host variable changed and the access plan is no longer valid.</li> <li>• B7 - The query was reoptimized because AQP determined that the query should be reoptimized.</li> <li>• B8 - The query was reoptimized because Expression Evaluator determined that the statement should be reoptimized</li> </ul>
SubCode	QQC21	If the access plan rebuild reason code was A7 this two-byte hex value identifies which specific reason for A7 forced a rebuild.
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Last_Access_Plan_Rebuild_Timestamp	QQTIM1	Timestamp of last access plan rebuild

Table 247. QQQ3006 - Access Plan Rebuilt (continued)

View Column Name	Table Column Name	Description
Reoptimization_Done	QQC11	Required optimization for this plan. <ul style="list-style-type: none"> <li>• Y - Yes, plan was really optimized.</li> <li>• N - No, the plan was not reoptimized because of the QAAQINI option for the REOPTIMIZE_ACCESS_PLAN parameter value</li> </ul>
Previous_Reason_Code	QVC22	Previous reason code
Previous_SubCode	QVC23	Previous reason subcode
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

## Database monitor view 3007 - Optimizer Timed Out

Displays the SQL logical view format for database monitor QQQ3007.

```

Create View QQQ3007 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QQ1000 as Index_Names,
 QQC11 as Optimizer_Timed_Out,
 QQC301 as Reason_Codes,
 QVQTBL as Table_Name,
 QVQLIB as Table_Schema,
 QVPTBL as Base_Table_Name,
 QVPLIB as Base_Table_Schema,
 QQJNP as Join_Position,
 QQI6 as DataSpace_Number,
 QQC21 as Join_Method,
 QQC22 as Join_Type,
 QQC23 as Join_Operator,
 QVJFANO as Join_Fanout,
 QVFILES as Join_Table_Count,
 QVRCNT as Unique_Refresh_Counter,
 QQIDXNL as Index_Names_2,
 QQSMINTF as Plan_iteration_number
 FROM UserLib/DBMONTTable
 WHERE QQRID=3007)

```

Table 248. QQQ3007 - Optimizer Timed Out

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member name of base table

---

Table 248. QQQ3007 - Optimizer Timed Out (continued)

---

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Index_Names	QQ1000	<p>Names of indexes not used and reason code.</p> <ol style="list-style-type: none"><li>1. Access path was not in a valid state. The system invalidated the access path.</li><li>2. Access path was not in a valid state. The user requested that the access path be rebuilt.</li><li>3. Access path is a temporary access path (resides in library QTEMP) and was not specified as the file to be queried.</li><li>4. The cost to use this access path, as determined by the optimizer, was higher than the cost associated with the chosen access method.</li><li>5. The keys of the access path did not match the fields specified for the ordering/grouping criteria. For distributed file queries, the access path keys must exactly match the ordering fields if the access path is to be used when ALWCOPYDTA(*YES or *NO) is specified.</li><li>6. The keys of the access path did not match the fields specified for the join criteria.</li><li>7. Use of this access path will not minimize delays when reading records from the file. The user requested to minimize delays when reading records from the file.</li><li>8. The access path cannot be used for a secondary file of the join query because it contains static select/omit selection criteria. The join-type of the query does not allow the use of select/omit access paths for secondary files.</li><li>9. File contains record ID selection. The join-type of the query forces a temporary access path to be built to process the record ID selection.</li><li>10. The user specified ignore decimal data errors on the query. This disallows the use of permanent access paths.</li></ol>

---

Table 248. QQQ3007 - Optimizer Timed Out (continued)

View Column Name	Table Column Name	Description
Index_Names (continued)	QQ1000	<ul style="list-style-type: none"> <li>• 11. The access path contains static select/omit selection criteria which is not compatible with the selection in the query.</li> <li>• 12. The access path contains static select/omit selection criteria whose compatibility with the selection in the query cannot be determined. Either the select/omit criteria or the query selection became too complex during compatibility processing.</li> <li>• 13. The access path contains one or more keys which may be changed by the query during an insert or update.</li> <li>• 14. The access path is being deleted or is being created in an uncommitted unit of work in another process.</li> <li>• 15. The keys of the access path matched the fields specified for the ordering/grouping criteria. However, the sequence table associated with the access path did not match the sequence table associated with the query.</li> <li>• 16. The keys of the access path matched the fields specified for the join criteria. However, the sequence table associated with the access path did not match the sequence table associated with the query.</li> <li>• 17. The left-most key of the access path did not match any fields specified for the selection criteria. Therefore, key row positioning cannot be performed, making the cost to use this access path higher than the cost associated with the chosen access method.</li> <li>• 18. The left-most key of the access path matched a field specified for the selection criteria. However, the sequence table associated with the access path did not match the sequence table associated with the query. Therefore, key row positioning cannot be performed, making the cost to use this access path higher than the cost associated with the chosen access method.</li> <li>• 19. The access path cannot be used because the secondary file of the join query is a select/omit logical file. The join-type requires that the select/omit access path associated with the secondary file be used or, if dynamic, that an access path be created by the system.</li> <li>• 20. The access path cannot be used because Encoded Vector Index does not fit in memory.</li> <li>• 21. The access path cannot be used because not all referenced columns had an Encoded Vector Index.</li> </ul>

Table 248. QQQ3007 - Optimizer Timed Out (continued)

View Column Name	Table Column Name	Description
Optimizer_Timed_Out	QQC11	Optimizer timed out (Y/N)
Reason_Codes	QQC301	List of unique reason codes used by the indexes that timed out (each index has a corresponding reason code associated with it)
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Schema of base table, long name
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Index_Names_2	QQ1000L	Index names when the list will not fit into QQ1000. Set to null otherwise

Table 248. QQQ3007 - Optimizer Timed Out (continued)

View Column Name	Table Column Name	Description
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

### Database monitor view 3008 - Subquery Processing

Displays the SQL logical view format for database monitor QQQ3008.

```

Create View QQQ3008 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QQI1 as Original_QDT_Count,
 QQI2 as Merged_QDT_Count,
 QQI3 as Final_QDT_Count,
 QVRCNT as Unique_Refresh_Counter,
 QQSMINTF as PlanIterNum
 FROM UserLib/DBMONTable
 WHERE QQRID=3008)

```

Table 249. QQQ3008 - Subquery Processing

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level

Table 249. QQQ3008 - Subquery Processing (continued)

View Column Name	Table Column Name	Description
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Original_QDT_Count	QQI1	Original number of QDTs
Merged_QDT_Count	QQI2	Number of QDTs merged
Final_QDT_Count	QQI3	Final number of QDTs
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
PlanIterNum	QQSMINTF	AQP Plan iteration number, original optimization = 1

### Database monitor view 3010 - Host Variable & ODP Implementation

Displays the SQL logical view format for database monitor QQQ3010.

```

Create View QQQ3010 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQI5 as Unqie_Refresh_Counter2,
 QQUDEF as User_Defined,
 QQC11 as ODP_Implementation,
 QQC12 as Host_Variable_Implementation,
 QQ1000 as Host_Variable_Values,
 QVRCNT as Unique_Refresh_Counter,
 QQDBCLOB1 as DBCLOB_CCSID,
 QQI7 as DBCLOB_Length,
 QQINT05 as SQ_Unique_Count,
 QVC11 as HV_Truncated
 FROM UserLib/DBMONTTable
 WHERE QQRID=3010)

```

Table 250. QQQ3010 - HostVar & ODP Implementation

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name



Table 250. QQQ3010 - HostVar & ODP Implementation (continued)

View Column Name	Table Column Name	Description
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
Unqiue_Refresh_Counter2	QQI5	Unique refresh counter
User_Defined	QQUDEF	User defined column
ODP_Implementation	QQC11	ODP implementation <ul style="list-style-type: none"> <li>• R - Reusable ODP</li> <li>• N - Nonreusable ODP</li> <li>• ' ' - Column not used</li> </ul>
Host_Variable_Implementati on	QQC12	Host variable implementation <ul style="list-style-type: none"> <li>• I - Interface supplied values (ISV)</li> <li>• V - Host variables treated as literals (V2)</li> <li>• U - Table management row positioning (UP)</li> <li>• S - SQL Insert/Update host variable value</li> </ul>
Host_Variable_Values	QQ1000	Host variable values
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
DBCLOB_CCSID	QQDBCLOB1	Host variables values in a DBCLOB CCSID 1200 field
DBCLOB_Length	QQI7	Length of host variables in the DBCLOB column.
SQ_Unique_Count	QQINT05	A unique count used to uniquely identify statements which do not have an ODP but do pass in host variables. If QQUCNT is 0 and the statement passes in host variables, this value will be non-zero. An example would be a CALL statement.
HV_Truncated	QVC11	Host variable has been truncated (Y/N).

## Database monitor view 3011 - Array Host Variables

Displays the SQL logical view format for database monitor QQQ3011.

```

Create View QQQ3011 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQC11 as ODP_Implementation,
 QQC12 as Array_Variable_Implementation,
 QQC101 as Array_Name,
 QVRCNT as Unique_Refresh_Counter,

```

```

 QQDBCLOB1 as Array_Values,
 QQINT05 as SQ_Unique_Count,
 QVC11 AS HV_Truncated,
 QVC1281 as Array_Name,
 QVC1282 as Array_Library,
 QQI1 as Max_Cardinality,
 QQI2 as Cur_Cardinality,
 QQI3 as Index_Position
FROM
WHERE
 UserLib/DBMONTable
 QQRID=3011)

```

Table 251. QQQ3011 - Array Host Variables

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
ODP_Implementation	QQC11	ODP implementation: <ul style="list-style-type: none"> <li>• R - Reusable ODP</li> <li>• N - Nonreusable ODP</li> <li>• ' ' - Column not used</li> </ul>
Array_Variable_Implementati on	QQC12	Array variable implementation: <ul style="list-style-type: none"> <li>• I - Interface supplied values (ISV)</li> <li>• S- SQL Insert/Update array variable value</li> </ul>
Array_Name	QQC101	Array name generated by the optimizer. Matches the array value in the QQ1000 QQHVAR field in the 3010 record.
Unique_Refresh_Counter	QVRCNT	Unique refresh counter.
Array_Values	QQDBCLOB1	Array variables values in a DBCLOB CCSID 1200 field (max 1 MB).
SQ_Unique_Count	QQINT05	A unique count used to uniquely identify statements which do not have an ODP but do pass in Arrays. If QQUCNT is 0 and the statement passes in Arrays, this value will be non-zero. An example would be a CALL statement.
HV_Truncated	QVC11	Host variable has been truncated (Y/N).
Array_Name	QVC1281	Name of Array UDT.
Array_Library	QVC1282	Library of Array UDT.

Table 251. QQQ3011 - Array Host Variables (continued)

View Column Name	Table Column Name	Description
Max_Cardinality	QQI1	Maximum cardinality of Array.
Cur_Cardinality	QQI2	Current cardinality of Array.
Index_Position	QQI3	Index position in the Array designated in the QQ1000 QQHVAR field in the 3010 record.

## Database monitor view 3012 - Global Variables

Displays the SQL logical view format for database monitor QQQ3012.

```

Create View QQQ3012 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQI5 as Unique_Refresh_Counter2,
 QQUDEF as User_Defined,
 QVRCNT as Unique_Refresh_Counter,
 QQDBCLOB1 as DBCLOB_Global_Variable,
 QQINT05 as SQ_Unique_Count,
 QVC11 as GV_Truncated
 FROM UserLib/DBMONTable
 WHERE QQRID=3012)

```

Table 252. QQQ3012 - Global Variables

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
Unique_Refresh_Counter2	QQI5	Unique refresh counter
User_Defined	QQUDEF	User-defined column
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
DBCLOB_Global_Variable	QQDBCLOB1	Global session variable values in a DBCLOB CCSID 1200 field.

Table 252. QQQ3012 - Global Variables (continued)

View Column Name	Table Column Name	Description
SQ_Unique_Count	QQINT05	A unique count used to uniquely identify statements which do not have an ODP but do pass in global variables. If QQUCNT is 0 and the statement passes in global variables, this value is non-zero. An example would be a CALL statement.
GV_Truncated	QVC11	Host variable has been truncated (Y/N).

## Database monitor view 3014 - Generic QQ Information

Displays the SQL logical view format for database monitor QQQ3014.

```

Create View QQQ3014 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQREST as Estimated_Rows_Selected,
 QQEPT as Estimated_Processing_Time,
 QQI1 as Open_Time,
 QQORDG as Has_Ordering,
 QQGRPG as Has_Grouping,
 QQJNG as Has_Join,
 QQC22 as Join_Type,
 QQUNIN as Has_Union,
 QQSUBQ as Has_Subquery,
 QWC1F as Has_Scalar_Subselect,
 QQHSTV as Has_Host_Variables,
 QQRCD5 as Has_Row_Selection,
 QQC11 as Query_Governor_Enabled,
 QQC12 as Stopped_By_Query_Governor,
 QQC101 as Open_Id,
 QQC102 as Query_Options_Library,
 QQC103 as Query_Options_Table_Name,
 QQC13 as Early_Exit,
 QVRCNT as Unique_Refresh_Counter,
 QQI5 as Optimizer_Time,
 QQTIM1 as Access_Plan_Timestamp,
 QVC11 as Ordering_Implementation,
 QVC12 as Grouping_Implementation,
 QVC13 as Join_Implementation,
 QVC14 as Has_Distinct,
 QVC15 as Is_Distributed,
 QVC3001 as Distributed_Nodes,
 QVC105 as NLSS_Table,
 QVC106 as NLSS_Library,
 QVC16 as ALWCPYDATA,
 QVC21 as Access_Plan_Reason_Code,
 QVC22 as Access_Plan_Reason_SubCode,
 QVC3002 as Summary,
 QWC16 as Last_Union_Subselect,
 QVP154 as Query_PoolSize,
 QVP155 as Query_PoolID,
 QQI2 as Query_Time_Limit,
 QVC81 as Parallel_Degree,

```

```

QQI3 as Max_Number_of_Tasks,
QVC17 as Apply_CHGQRYA_Remote,
QVC82 as Async_Job_Usage,
QVC18 as Force_Join_Order_Indicator,
QVC19 as Print_Debug_Messages,
QVC1A as Parameter_Marker_Conversion,
QQI4 as UDF_Time_Limit,
QVC1283 as Optimizer_Limitations,
QVC1E as Reoptimize_Requested,
QVC87 as Optimize_All_Indexes,
QQC14 as Has_Final_Decomposed_QDT,
QQC15 as Is_Final_Decomposed_QDT,
QQC18 as Read_Trigger,
QQC81 as Star_Join,
SUBSTR(QVC23,1,1) as Optimization_Goal,
SUBSTR(QVC24,1,1) as VE_Diagram_Type,
SUBSTR(QVC24,2,1) as Ignore_Like_Redunant_Shifts,
QQC23 as Union_QDT,
QQC21 as Unicode_Normalization,
QVP153 as Pool_Fair_Share,
QQC82 as Force_Join_Order_Requested,
QVP152 as Force_Join_Order_Dataspace1,
QQI6 as No_Parameter_Marker_Reason_Code,
QVP151 as Hash_Join_Reason_Code,
QQI7 as MQT_Refresh_Age,
SUBSTR(QVC42,1,1) as MQT_Usage,
QVC43 as SQE_NotUsed_Reason_Code,
QVP156 as Estimated_IO_Count,
QVP157 as Estimated_Processing_Cost,
QVP158 as Estimated_CPU_Cost,
QVP159 as Estimated_IO_Cost,
SUBSTR(QVC44,1,1) as Has_Implicit_Numeric_Conversion,
QVCTIM as Accumulated_Est_Process_Time,
QQINT01 as Query_Gov_Storage_Limit,
QQINT02 as Estimated_Storage,
QQINT03 as Adjusted_Temp_Storage,
QQINT04 as Original_Cost_Estimate,
QQI8 as Parallel_Degree_Percentage,
QFC12 as FieldProc_Encoded_Comparison,
QFC13 as Allow_Array_Changes_INI_Opt,
QFC11 as SQL_Concurrent_Access_Resolution,
QQSMINTF as Plan_Iteration_Number,
QXC11 as Warm_IO_Requested,
QXC12 as Warm_IO_Used,
QXC13 as Optimization_Goal_Override,
QXC1E as Plan_Signature_Match,
QXC14 as Check_HostVars,
QXC15 as FullOptimization,
QXC16 as Pseudo_Open_Replace_Reason,
QXC17 as WorkloadGroup,
QXC18 as Concurrent_Access_Behavior,
QQINT05 as Memory_Pool_Preference,
QQINT06 as Longest_Key_Range_Estimate,
QXC19 as KeyRangeEst_Timeout,
QQC16 as SQE_Used_Indicator,
QQC83 as QRO_Hash,
QQSMINT1 as Shared_CTE_Usage
FROM UserLib/DBMONTTable
WHERE QQRID=3014)

```

Table 253. QQQ3014 - Generic QQ Information

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User-defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Estimated_Rows_Selected	QQREST	Estimated number of rows selected
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Open_Time	QQI1	Time spent to open cursor, in milliseconds
Has_Ordering	QQORDG	Ordering (Y/N)
Has_Grouping	QQGRPG	Grouping (Y/N)
Has_Join	QQJNG	Join Query (Y/N)
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Has_Union	QQUNIN	Union Query (Y/N)
Has_Subquery	QQSUBQ	Subquery (Y/N)
Has_Scalar_Subselect	QWC1F	Scalar Subselects (Y/N)
Has_Host_Variables	QQHSTV	Host variables (Y/N)
Has_Row_Selection	QQRCDS	Row selection (Y/N)
Query_Governor_Enabled	QQC11	Query governor enabled (Y/N)

Table 253. QQQ3014 - Generic QQ Information (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Stopped_By_Query_Governor	QQC12	Query governor stopped the query (Y/N)
Open_Id	QQC101	Query open ID
Query_Options_Library	QQC102	Query Options library name
Query_Options_Table_Name	QQC103	Query Options file name
Early_Exit	QQC13	Query early exit value
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Optimizer_Time	QQI5	Time spent in optimizer, in milliseconds
Access_Plan_Timestamp	QQTIM1	Access Plan rebuilt timestamp, last time access plan was rebuilt.
Ordering_Implementation	QVC11	Ordering implementation. Possible values are: <ul style="list-style-type: none"> <li>• I - Index</li> <li>• S - Sort</li> </ul>
Grouping_Implementation	QVC12	Grouping implementation. Possible values are: <ul style="list-style-type: none"> <li>• I - Index</li> <li>• H - Hash grouping</li> </ul>
Join_Implementation	QVC13	Join Implementation. Possible values are: <ul style="list-style-type: none"> <li>• N - Nested Loop join</li> <li>• H - Hash join</li> <li>• C - Combination of Nested Loop and Hash</li> </ul>
Has_Distinct	QVC14	Distinct query (Y/N)
Is_Distributed	QVC15	Distributed query (Y/N)
Distributed_Nodes	QVC3001	Distributed nodes
NLSS_Table	QVC105	Sort Sequence Table
NLSS_Library	QVC106	Sort Sequence Library
ALWCPYDATA	QVC16	ALWCPYDTA setting
Access_Plan_Reason_Code	QVC21	Reason code why access plan was rebuilt
Access_Plan_Reason_SubCode	QVC22	Subcode why access plan was rebuilt
Summary	QVC3002	Summary of query implementation. Shows dataspace number and name of index used for each table being queried.
Last_Union_Subselect	QWC16	Last part (last QDT) of Union (Y/N)
Query_PoolSize	QVP154	Pool size
Query_PoolID	QVP155	Pool id
Query_Time_Limit	QQI2	Query time limit

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
Parallel_Degree	QVC81	Parallel Degree <ul style="list-style-type: none"> <li>• *SAME - Do not change current setting</li> <li>• *NONE - No parallel processing is allowed</li> <li>• *I/O - Any number of tasks might be used for I/O processing. SMP parallel processing is not allowed.</li> <li>• *OPTIMIZE - The optimizer chooses the number of tasks to use for either I/O or SMP parallel processing.</li> <li>• *MAX - The optimizer chooses to use either I/O or SMP parallel processing.</li> <li>• *SYSVAL - Use the current system value to process the query.</li> <li>• *ANY - Has the same meaning as *I/O.</li> <li>• *NBRTASKS - The number of tasks for SMP parallel processing is specified in column QVTASKN.</li> </ul>
Max_Number_of_Tasks	QQI3	Max number of tasks
Apply_CHGQRYA_Remote	QVC17	Apply CHGQRYA remotely (Y/N)
Async_Job_Usage	QVC82	Asynchronous job usage <ul style="list-style-type: none"> <li>• *SAME - Do not change current setting</li> <li>• *DIST - Asynchronous jobs might be used for queries with distributed tables</li> <li>• *LOCAL - Asynchronous jobs might be used for queries with local tables only</li> <li>• *ANY - Asynchronous jobs might be used for any database query</li> <li>• *NONE - No asynchronous jobs are allowed</li> </ul>
Force_Join_Order_Indicator	QVC18	Force join order (Y/N)
Print_Debug_Messages	QVC19	Print debug messages (Y/N)
Parameter_Marker_Conversion	QVC1A	Parameter marker conversion (Y/N)
UDF_Time_Limit	QQI4	User Defined Function time limit
Optimizer_Limitations	QVC1283	Optimizer limitations. Possible values: <ul style="list-style-type: none"> <li>• *PERCENT followed by 2 byte integer containing the percent value</li> <li>• *MAX_NUMBER_OF_RECORDS followed by an integer value that represents the maximum number of rows</li> </ul>



Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
Reoptimize_Requested		<p>Reoptimize access plan requested. Possible values are:</p> <ul style="list-style-type: none"> <li>• O - Only reoptimize the access plan when required. Do not reoptimize for subjective reasons.</li> <li>• Y - Yes, force the access plan to be reoptimized.</li> <li>• N - No, do not reoptimize the access plan, unless optimizer determines that it is necessary. May reoptimize for subjective reasons.</li> </ul>
Optimize_All_Indexes		<p>Optimize all indexes requested</p> <ul style="list-style-type: none"> <li>• *SAME - Do not change current setting</li> <li>• *YES - Examine all indexes</li> <li>• *NO - Allow optimizer to time out</li> <li>• *TIMEOUT - Force optimizer to time out</li> </ul>
Has_Final_Decomposed_QDT	QQC14	Final decomposed QDT built indicator (Y/N)
Is_Final_Decomposed_QDT	QQC15	The final decomposed QDT indicator (Y/N)
Read_Trigger	QQC18	One of the files contains a read trigger (Y/N)
Star_Join	QQC81	<p>Star join optimization requested.</p> <ul style="list-style-type: none"> <li>• *NO - Star join optimization is not performed.</li> <li>• *COST - The optimizer determines if any EVIs can be used for star join optimization.</li> <li>• *FORCE - The optimizer adds any EVIs that can be used for star join optimization.</li> </ul>
Optimization_Goal	QVC23	<p>Byte 1 = Optimization goal. Possible values are:</p> <ul style="list-style-type: none"> <li>• F - First I/O, optimize the query to return the first screen full of rows as quickly as possible.</li> <li>• A - All I/O, optimize the query to return all rows as quickly as possible.</li> </ul>
VE_Diagram_Type	QVC24	<p>Byte 1 = Type of Visual Explain diagram. Possible values are:</p> <ul style="list-style-type: none"> <li>• D - Detail</li> <li>• B - Basic</li> </ul>
Ignore_Like_Redunant_Shifts	QVC24	<p>Byte 2 - Ignore LIKE redundant shifts. Possible values are:</p> <ul style="list-style-type: none"> <li>• O - Optimize, the query optimizer determines which redundant shifts to ignore.</li> <li>• A - All redundant shifts are ignored.</li> </ul>

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
Union_QDT	QQC23	Byte 1 = This QDT is part of a UNION that is contained within a view (Y/N).  Byte 2 = This QDT is the last subselect of the UNION that is contained within a view (Y/N).
Unicode_Normalization	QQC21	Unicode data normalization requested (Y/N)
Pool_Fair_Share	QVP153	Fair share of the pool size as determined by the optimizer
Force_Join_Order_Requested	QQC82	Force Join Order requested. Possible values are: <ul style="list-style-type: none"> <li>• *NO - The optimizer was allowed to reorder join files</li> <li>• *YES - The optimizer was not allowed to reorder join files as part of its optimization process</li> <li>• *SQL - The optimizer only forced the join order for those queries that used the SQL JOIN syntax</li> <li>• *PRIMARY - The optimizer was only required to force the primary dial for the join.</li> </ul>
Force_Join_Order_Dataspace1	QVP152	Primary dial to force if Force_Join_Order_Indicator is *PRIMARY.
No_Parameter_Marker_Reason_Code	QQI6	Reason code for why Parameter Marker Conversion was not performed: <ol style="list-style-type: none"> <li>1. Argument of function must be a literal</li> <li>2. LOCALTIME or LOCALTIMESTAMP</li> <li>3. Duration literal in arithmetic expression</li> <li>4. UPDATE query with no WHERE clause</li> <li>5. BLOB literal</li> <li>6. Special register in UPDATE or INSERT with values</li> <li>7. Result expression for CASE</li> <li>8. GROUP BY expression</li> <li>9. ESCAPE character</li> <li>10. Double Negative value -(-1)</li> <li>11. INSERT or UPDATE with a mix of literals, parameter markers, and NULLs</li> <li>12. UPDATE with a mix of literals and parameter markers</li> <li>13. INSERT with VALUES containing NULL value and expressions</li> <li>14. UPDATE with list of expressions</li> </ol> <p>99. Parameter marker conversion disabled by QAQQINI</p>
Hash_Join_Reason_Code	QVP151	Reason code why hash join was not used.

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
MQT_Refresh_Age	QQI7	Value of the MATERIALIZED_QUERY_TABLE_REFRESH_AGE duration. If the QAQQINI parameter value is set to *ANY, the timestamp duration is 9999999999999999.
MQT_Usage	QVC42,1,1	Byte 1 - Contains the MATERIALIZED_QUERY_TABLE_USAGE. Possible values are: <ul style="list-style-type: none"> <li>• N - *NONE - no materialized query tables used in query optimization and implementation</li> <li>• A - *ALL - User-maintained. Refresh-deferred query tables can be used.</li> <li>• U - *USER - Only user-maintained materialized query tables can be used.</li> </ul>
SQE_NotUsed_Reason_Code	QVC43	SQE Not Used Reason Code. Possible values: <ul style="list-style-type: none"> <li>• LF - DDS logical file specified in query definition</li> <li>• DK - An index with derived key or select/omit was found over a queried table</li> <li>• NF - Too many tables in query</li> <li>• NS - Not an SQL query or query not run through an SQL interface</li> <li>• DF - Distributed table in query</li> <li>• RT - Read Trigger defined on queried table</li> <li>• PD - Program described file in query</li> <li>• WC - WHERE CURRENT OF a partition table</li> <li>• IO - Simple INSERT query</li> <li>• CV - Create view statement</li> </ul>
Estimated_IO_Count	QVP156	Estimated I/O count
Estimated_Processing_Cost	QVP157	Estimated processing cost in milliseconds
Estimated_CPU_Cost	QVP158	Estimated CPU cost in milliseconds
Estimated_IO_Cost	QVP159	Estimated I/O cost in milliseconds
Has_Implicit_Numeric_Conversion	QVC44	Byte 1: Implicit numeric conversion (Y/N)
Accumulated_Est_Process_Time	QVCTIM	Accumulated estimated processing time across all subselects, in seconds.
Query_Gov_Storage_Limit	QQINT01	Specified query governor storage limit, in megabytes
Estimated_Storage	QQINT02	Original estimated temporary storage used, in megabytes.
Adjusted_Temp_Storage	QQINT03	Adjusted temporary storage used, in Adjusted megabytes. This value accumulates the actual time and storage it took to create any temporary indexes and temporary tables. Set by CQE only.

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
Original_Cost_Estimate	QQINT04	Original cost estimate as determined by the CQE query optimizer. Set by CQE only.
Parallel_Degree_Percentage	QQI8	Percentage specified on Parallel_Degree *OPTIMIZE and *MAX.
FieldProc_Encoded_Comparison	QFC12	FIELDPROC_ENCODED_COMPARISON option active for this query. Specifies the amount of optimization that the optimizer might use when queried columns have attached field procedures. <ul style="list-style-type: none"> <li>• 'N' - NONE</li> <li>• 'E' - ALLOW_EQUAL</li> <li>• 'R' - ALLOW_RANGE</li> <li>• 'A' - ALL</li> </ul>
Allow_Array_Change_INI_Opt	QFC13	ALLOW_ARRAY_VALUE_CHANGES QAQQINI option active for this query. <ul style="list-style-type: none"> <li>• 'N' - Do not allow changes to values in arrays referenced in the query to be visible after the query is opened.</li> <li>• 'Y' - Allow changes to values in arrays to be visible to the query while the query is running.</li> </ul>
SQL_Concurrent_Access_Resolution	QFC11	SQL_CONCURRENT_ACCESS_RESOLUTION QAQQINI option active for this query. <ul style="list-style-type: none"> <li>• 'U' - USE CURRENTLY COMMITTED</li> <li>• 'W' - WAIT FOR OUTCOME</li> </ul>
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number; original optimization = 1
Warm_IO_Requested	QXC11	Warm I/O value that was requested. <ul style="list-style-type: none"> <li>• 'Y' - Yes, use Warm I/O</li> <li>• 'N' - No, do not use Warm I/O</li> <li>• 'D' - Default</li> </ul>
Warm_IO_Used	QXC12	Warm I/O values used to implement the query. <ul style="list-style-type: none"> <li>• 'Y' - Yes, use Warm I/O</li> <li>• 'N' - No, do not use Warm I/O</li> <li>• 'D' - Default</li> </ul>
Optimization_Goal_Override	QXC13	Optimization Goal Override. <ul style="list-style-type: none"> <li>• 'O' - Override the specified Optimize For N Rows value and use Optimize For All Rows.</li> <li>• 'D' - Default, use the specified Optimize For N Rows value.</li> </ul>

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
Plan_Signature_Match	QXC1E	Plan signature match. <ul style="list-style-type: none"> <li>• 'Y' - New plan matched old plan it replaced; same plan signature.</li> <li>• 'N' - New plan different from old plan it replaced; different plan signature.</li> </ul>
Check_HostVars	QXC14	Indicates if this query is enabled for host variable selectivity checking at pseudo-open time. <ul style="list-style-type: none"> <li>• 'N' – No pseudo-open host variable checking enabled</li> <li>• 'O' – This query is a candidate for pseudo-open host variable checking and the QAAQINI option was *OPTIMIZE</li> <li>• 'Y' - This query is a candidate for pseudo-open host variable checking and the QAAQINI option was *YES</li> </ul>
FullOptimization	QXC15	Plan was rebuilt (Y/N)
Pseudo_Open_Replace_Reason	QXC16	Indicates if the plan was replaced due to QAAQINI PSEUDO_OPEN_CHECK_HOST_VARS option <ul style="list-style-type: none"> <li>• '0' - Plan was not replaced</li> <li>• '1' - Plan was replaced</li> </ul>
WorkloadGroup	QXC17	Workload Group is in effect (Y/N)
Concurrent_Access_Behavior	QXC18	Controls how queries with an isolation level of Cursor Stability (CS) or Read Stability (RS) interact with uncommitted table changes. <ul style="list-style-type: none"> <li>• 'N' - Not Applicable</li> <li>• 'O' - OPTIMIZE - Uncommitted changes that delete or update records so that they are no longer selected by the query will not be considered as candidates for query synchronization.</li> <li>• 'S' - STRICTSCAN - All records referenced by a table scan query access plan will synchronize with any changes that are not yet committed.</li> </ul>
Memory_Pool_Preference	QQINT05	The pool identifier where paging will occur for database operation that supports targeted paging. The pool identifier will be 0 if paging will occur in the same pool as the job is running in.
Longest_Key_Range_Estimate	QQINT06	Number of seconds required to run the longest key range estimate (or EKR) used to optimize this query.

Table 253. QQQ3014 - Generic QQ Information (continued)

View Column Name	Table Column Name	Description
KeyRangeEst_Timeout	QXC19	<p>Indicates whether a key range estimate exceeded the time limit. This limit may be specified in the QAQQINI file.</p> <ul style="list-style-type: none"> <li>'N' - No</li> <li>'Y' - One or more estimates exceeded the time limit. The estimates have been queued for background processing.</li> </ul>
SQE_Used_Indicator	QQC16	<p>Indicator of which optimizer was used</p> <ul style="list-style-type: none"> <li>'Y' - Query ran using SQE</li> <li>'N' - Query ran using CQE</li> </ul>
QRO_HASH	QQC83	<p>For SQE queries, the QRO hash is an internally generated identifier for an SQE query. In general, this identifier will be unique for each SQE query and uses implicit schema qualification among other data to generate the QRO hash. If the SQE optimizer generates multiple plans for the same query, then multiple plans will have the same QRO hash.</p> <p>The QRO hash for a statement may change on release boundaries or after loading PTFs.</p>
Shared_CTE_Usage	QQSMINT1	<p>For SQE queries, this field describes whether the query uses a common table expression (CTE) that is referenced more than once in the query. This is referred to as a shared CTE and requires special processing by the optimizer.</p> <ul style="list-style-type: none"> <li><b>0</b> Query does not contain CTEs with multiple references.</li> <li><b>1</b> Query contains at least one CTE with multiple references.</li> <li><b>2</b> Query contains at least one CTE with multiple references and at least one of these CTEs is considered complex by the optimizer.</li> </ul>

## Database monitor view 3015 - Statistics Information

Displays the SQL logical view format for database monitor QQQ3015.

```

Create View QQQ3015 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,

```

```

QQQDTL as SubSelect_Nested_Level,
QQMATN as Materialized_View_Subselect_Number,
QQMATL as Materialized_View_Nested_Level,
QVP15E as Materialized_View_Union_Level,
QVP15A as Decomposed_Subselect_Number,
QVP15B as Total_Number_Decomposed_SubSelects,
QVP15C as Decomposed_SubSelect_Reason_Code,
QVP15D as Starting_Decomposed_SubSelect,
QQTLN as System_Table_Schema,
QQTFN as System_Table_Name,
QQTMN as Member_Name,
QQPTLN as System_Base_Table_Schema,
QQPTFN as System_Base_Table_Name,
QQPTMN as Base_Member_Name,
QVQTBL as Table_Name,
QVQLIB as Table_Schema,
QVPTBL as Base_Table_Name,
QVPLIB as Base_Table_Schema,
QQNTNM as NLSS_Table,
QQNLNM as NLSS_Library,
QQC11 as Statistic_Status,
QQI2 as Statistic_Importance,
QQ1000 as Statistic_Columns,
QVC1000 as Statistic_ID,
QQSMINTF as Plan_Iteration_Number
FROM UserLib/DBMONTAbLe
WHERE QQRID=3015)

```

Table 254. QQQ3015 - Statistic Information

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code

Table 254. QQQ3015 - Statistic Information (continued)

View Column Name	Table Column Name	Description
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of the base table queried
Base_Member_Name	QQPTMN	Member name of base table
Table_Name	QVQTBL	Queried table, long name
Table_Schema	QVQLIB	Schema of queried table, long name
Base_Table_Name	QVPTBL	Base table, long name
Base_Table_Schema	QVPLIB	Schema of base table, long name
NLSS_Table	QQNTNM	NLSS table
NLSS_Library	QQNLNM	NLSS library
Statistic_Status	QQC11	Statistic Status. Possible values are: <ul style="list-style-type: none"> <li>• 'N' - No statistic</li> <li>• 'S' - Stale statistic</li> <li>• '' - Unknown</li> </ul>
Statistic_Importance	QQI2	Importance of this statistic
Statistic_Columns	QQ1000	Columns for the statistic advised
Statistic_ID	QVC1000	Statistic identifier
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

### Database monitor view 3018 - STRDBMON/ENDDBMON

Displays the SQL logical view format for database monitor QQQ3018.

```

Create View QQQ3018 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQC11 as Monitored_Job_type,
 QQC12 as Monitor_Command,
 QQ301 as Monitor_Job_Information,
 QQ1000L as STRDBMON_Command_Text,
 QQC101 as Monitor_ID,
 QQC102 as Version_Release_Mod,
 QQC103 as Group_PTF,
 QVC11 as Initial_AQP_Processing,
 QQSMINT1 as Current_Plan_Cache_Threshold,

```



```

QQSMINT2 as Current_Plan_Cache_Subcaches,
QQINT01 as Number_of_Currently_Active_Queries,
QQINT02 as Current_Plan_Cache_Size,
QQINT03 as Current_Plan_Cache_Size_Threshold,
QQINT04 as Number_of_SMP_Threads,
QQINT05 as Current_Number_of_MTIs,
QQINT06 as Number_of_Pruning_Listeners,
QQINT07 as Number_of_Plan_Cache_Awakenings,
QQINT08 as Number_of_Plan_Cache_Naps,
QQINT09 as Number_Pseudo_Open_queries_Hard_Closed
QQINT0A as Number_of_MTIs_Created,
QQINT0B as Number_of_MTIs_Deleted,
QQINT0C as Number_AQP_Wakeups,
QQINT0D as Number_AQP_Plans_Replaced,
QQINT0E as Number_of_Active_Queries,
QVP151 as Number_of_Plans_in_Cache,
QVP152 as Number_of_ROQs_in_Cache,
QVP153 as Number_of_Temp_ROQs_in_Cache,
QVP154 as Number_of_Reuseable_ROQs_in_Cache,
QVP155 as Size_of_Temporary_Objects_stored_in_Cache,
QVP156 as Number_of_Plans_Built_Since_Start,
QVP157 as Number_of_Plans_Used_ROQ,
QVP158 as Number_of_Plans_Used_nonROQ,
QVP159 as Number_of_Plans_Used_No_ROQ,
QVP15A as Number_of_Plan_Cache_Probes,
QVP15B as Number_of_Plans_Used_from_Cache,
QVP15C as Number_of_Plan_Cache_No_Matches,
QVP15D as Number_of_Plans_Pruned,
QVP15E as Number_of_Plans_Removed,
QVP15F as Number_of_Queries_Run_Since_Start,
QQI1 as Number_of_Query_Full_Opens_Since_Start,
QQI2 as Number_of_Full_Opens_Which_Reused_ROQ,
QQI3 as Number_Full_Optimizations,
QQI4 as Number_Reopts_with_Existing_Valid_Plans,
QQDBCLOB1 as Plan_Cache_Properties)
FROM UserLib/DBMONTable
WHERE QQRID=3018)

```

Table 255. QQQ3018 - STRDBMON/ENDDBMON

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Monitored_Job_type	QQC11	Type of job monitored <ul style="list-style-type: none"> <li>• C - Current</li> <li>• J - Job name</li> <li>• A - All</li> </ul>
Monitor_Command	QQC12	Command type <ul style="list-style-type: none"> <li>• S - STRDBMON</li> <li>• E - ENDDBMON</li> </ul>
Monitor_Job_Information	QQC301	Monitored job information <ul style="list-style-type: none"> <li>• * - Current job</li> <li>• Job number/User/Job name</li> <li>• *ALL - All jobs</li> </ul>

Table 255. QQQ3018 - STRDBMON/ENDDBMON (continued)

View Column Name	Table Column Name	Description
STRDBMON_Command_Text	QQ1000L	STRDBMON command text.
Monitor_ID	QQC101	Monitor ID
Version_Release_Mod	QQC102	Version Release and modification level
Group_PTF	QQC103	Installed Group PTF number and level. For example, 'SF99601 3' indicates that Database Group release V6R1M0, version 3 is installed.
Initial_AQP_Processing	QVC11	Initial AQP processing.  Y indicates that the monitor file has already been processed to handle the AQP plan iteration number. All iterations other than the last one have been changed to a negative number.  All other values indicate that the monitor file has not been processed.
Current_PC_Threshold	QQSMINT1	Database Plan Cache threshold, a value between 1-100 to represent a percent
Current_PC_Subcaches	QQSMINT2	Number of sub-caches within the database Plan Cache
Num_Active_Queries	QQINT01	Number of queries currently active
Current_PC_Size	QQINT02	Current size of database Plan Cache, in MB
Current_PC_Size_Threshold	QQINT03	Size threshold of database Plan Cache, in MB
Number_of_SMP_Threads	QQINT04	Number of SMP Threads
Current_Number_of_MTIs	QQINT05	Number of temporary indexes
Num_of_Pruning_Monitors	QQINT06	Number of event monitors used to prune the database Plan Cache
Num_of_PC_Prunnings	QQINT07	Number of times plans were pruned from the database Plan Cache
Num_of_Plan_Cache_Naps	QQINT08	Number of times the database Plan Cache became inactive
Num_POpen_Hard_Closed	QQINT09	Number of pseudo-opened queries that were hard closed
Num_of_MTIs_Created	QQINT0A	Number of temporary indexes that were created
Num_of_MTIs_Deleted	QQINT0B	Number of temporary indexes that were deleted
Num_AQP_Active	QQINT0C	Number of times AQP became active, whether a new plan was created or not
Num_AQP_Plans_Replaced	QQINT0D	Number of plans that were rebuilt due to AQP
Num_of_Plans_in_Cache	QVP151	Number of plans in the database Plan Cache
Num_of_ROQs_in_Cache	QVP152	Number of plans with a Read Only Query (ROQ) in the database Plan Cache
Num_of_Temp_ROQs_in_Cache	QVP153	Number of plans with temporary ROQs in the database Plan Cache
Num_of_Reuseable_ROQs	QVP154	Number of plans with reusable ROQs in the database Plan Cache
Size_Temp_Objects_in_Cache	QVP155	Size of temporary objects in the database Plan Cache, in MB
Num_Plans_Built_Since_Start	QVP156	Number of new plans built
Num_Plans_Used_ROQ	QVP157	Number of times a plan with a reusable ROQ was run
Num_Plans_Used_nonROQ	QVP158	Number of times a plan with a non-reusable ROQ was run
Num_Plans_Used_No_ROQ	QVP159	Number of times a plan without a ROQ was run
Num_Plan_Cache_Probes	QVP15A	Number of times the database Plan Cache was probed in hopes of finding a matching plan
Num_Plans_Used_from_Cache	QVP15B	Number of times a matching plan was found in the database Plan Cache
Num_Plan_Cache_No_Matches	QVP15C	Number of times a matching plan was not found in the database Plan Cache
Num_Plans_Pruned	QVP15D	Number of plans that were removed from the database Plan Cache due to size restrictions
Num_Plans_Removed	QVP15E	Number of obsolete plans that were removed from the database Plan Cache

Table 255. QQQ3018 - STRDBMON/ENDDBMON (continued)

View Column Name	Table Column Name	Description
Num_Run_Since_Start	QVP15F	Number of queries that were run
Num_FullOpens_Since_Start	QQI1	Number of queries that performed a full open when they were run
Num_FullOpens_Reused_ROQ	QQI2	Number of queries that performed a full open and reused an existing ROQ when run
Num_Full_Optimizations	QQI3	Number of queries that required a full optimization when run
Num_Valid_Plan_Reopts	QQI4	Number of queries that performed a full optimization even when a valid plan existed
Plan_Cache_Properties	QQDBCLOB1	Plan cache properties in XML format.

## Database monitor view 3019 - Rows retrieved

Displays the SQL logical view format for database monitor QQQ3019.

```

Create View QQQ3019 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQI1 as CPU_Time_to_Return_All_Rows,
 QQI2 as Clock_Time_to_Return_All_Rows,
 QQI3 as Number_Synchronous_Database_Reads,
 QQI4 as Number_Synchronous_Database_Writes,
 QQI5 as Number_Asynchronous_Database_Reads,
 QQI6 as Number_Asynchronous_Database_Writes,
 QVP151 as Number_Page_Faults,
 QQI7 as Number_Rows_Returned,
 QQI8 as Number_of_Calls_for_Returned_Rows,
 QVP15F as Number_of_Times_Statement_was_Run,
 QQINT03 as Temporary_Storage,
 QQC11 as DBMON_Temp_Result_Reused,
 QQC21 as DBMON_Temp_Reused_RC,
 QQINT01 as DBMON_Temp_Reuse_Count,
 QQIA as Skip_Lock_Row_Count,
 QQINT05 as Skip_Lock_Row_Runs,
 QQINT06 as Skip_Lock_On_Runs,
 QQF1 as Adjusted_Average_Run_Time,
 QVRCNT as Unique_Refresh_Counter,
 QVP152 as Committed_Journal_Search_Requests,
 QVP153 as Committed_Journal_Search_Failures,
 QVP154 as Committed_Journal_Search_Time_Limit
 FROM UserLib/DBMONTable
 WHERE QQRID=3019)

```

Table 256. QQQ3019 - Rows retrieved

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
CPU_Time_to_Return_All_Rows	QQI1	CPU time to return all rows, in milliseconds
Clock_Time_to_Return_All_Rows	QQI2	Clock time to return all rows, in milliseconds
Number_Synchronous_Database_Reads	QQI3	Number of synchronous database reads
Number_Synchronous_Database_Writes	QQI4	Number of synchronous database writes
Number_Asynchronous_Database_Reads	QQI5	Number of asynchronous database reads
Number_Asynchronous_Database_Writes	QQI6	Number of asynchronous database writes
Number_Page_Faults	QVP151	Number of page faults
Number_Rows_Returned	QQI7	Number of rows returned
Number_of_Calls_for_Returned_Rows	QQI8	Number of calls to retrieve rows returned
Number_of_Times_Statement_was_Run	QVP15F	Number of times this Statement was run. If Null, then the statement was run once.
Temporary_Storage	QQINT03	Amount of temporary storage used.

Table 256. QQQ3019 - Rows retrieved (continued)

View Column Name	Table Column Name	Description
DBMON_Temp_Result_Reused	QQC11	Indicates if the query temporary result was reused (Y/N).
DBMON_Temp_Reused_RC	QQC21	Reason code why the query temporary result was reused.
DBMON_Temp_Reuse_Count	QQINT01	Number of times the query temporary result was reused.
Skip_Lock_Row_Count	QQIA	Number of locked rows that were skipped.
Skip_Lock_Row_Runs	QQINT05	Number of runs where some rows were skipped.
Skip_Lock_On_Runs	QQINT06	Number of runs where Skip Lock was active.
Adjusted_Average_Run_Time	QQF1	Average runtime for the query, adjusted to not include individual runs that are well outside the norm. Units in microseconds.
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Committed_Journal_Search_Requests	QVP152	Number of times the database manager searched the journal for the currently committed version of a record.
Committed_Journal_Search_Failures	QVP153	Number of times the database manager failed to find the currently committed version of a record in the journal.
Committed_Journal_Search_Time_Limit	QVP154	Maximum amount of time allowed to search the journal for the currently committed version of a record.

## Database monitor view 3020 - Index advised (SQE)

Displays the SQL logical view format for database monitor QQQ3020.

```

Create View QQQ3020 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QQTLN as System_Table_Schema,
 QQTFN as System_Table_Name,
 QQTMN as Member_Name,
 QQPTLN as System_Base_Table_Schema,
 QQPTFN as System_Base_Table_Name,
 QQPTMN as Base_Member_Name,
 QVPLIB as Base_Table_Schema,
 QVPTBL as Base_Table_Name,

```

```

 QQTOTR as Table_Total_Rows,
 QQEPT as Estimated_Processing_Time,
 QQIDXA as Index_is_Advised,
 QQIDXD as Index_Advised_Columns_Short_List,
 QQ1000L as Index_Advised_Columns_Long_List,
 QQI1 as Number_of_Advised_Columns,
 QQI2 as Number_of_Advised_Primary_Columns,
 QQRCD as Reason_Code,
 QVRCNT as Unique_Refresh_Counter,
 QVC1F as Type_of_Index_Advised,
 QQNTNM as NLSS_Table,
 QQNLNM as NLSS_Library,
 QQSMINTF as Plan_Iteration_Number,
 QQ13 as DEPENDENT_ADVICE_ID
FROM UserLib/DBMONTTable
WHERE QQRID=3020)

```

Table 257. QQQ3020 - Index advised (SQE)

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Numbe r	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSele cts	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Cod e	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
System_Table_Schema	QQTLN	Schema of table queried
System_Table_Name	QQTFN	Name of table queried
Member_Name	QQTMN	Member name of table queried

Table 257. QQQ3020 - Index advised (SQE) (continued)

View Column Name	Table Column Name	Description
System_Base_Table_Schema	QQPTLN	Schema name of base table
System_Base_Table_Name	QQPTFN	Name of base table for table queried
Base_Member_Name	QQPTMN	Member of base table
Base_Table_Schema	QVPLIB	Schema of base table, long name
Base_Table_Name	QVPTBL	Base table, long name
Table_Total_Rows	QQTOTR	Number of rows in the table
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Index_is_Advised	QQIDXA	Index advised (Y/N)
Index_Advised_Columns_Short_List	QQIDXD	Columns for the index advised, first 1000 bytes
Index_Advised_Columns_Long_List	QQ1000L	Column for the index advised
Number_of_Advised_Columns	QQI1	Number of indexes advised
Number_of_Advised_Primary_Columns	QQI2	Number of advised columns that use index scan-key positioning
Reason_Code	QQRCOD	Reason code <ul style="list-style-type: none"> <li>• I1 - Row selection</li> <li>• I2 - Ordering/Grouping</li> <li>• I3 - Row selection and Ordering/Grouping</li> <li>• I5 - Row selection using bitmap processing</li> <li>• I6 - Source of statistics</li> <li>• I8 - Encoded Vector Index Only Column Projection</li> </ul>
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Type_of_Index_Advised	QVC1F	Type of index advised. Possible values are: <ul style="list-style-type: none"> <li>• B - Radix index</li> <li>• E - Encoded vector index</li> </ul>
NLSS_Table	QQNTNM	Sort Sequence Table
NLSS_Library	QQNLNM	Sort Sequence Library
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1
Dependent_Advice_ID	QQ13	Unique identifier used to link together OR predicate index advice

**Related reference**

[Index advisor](#)

The query optimizer analyzes the row selection in the query and determines, based on default values, if creation of a permanent index improves performance. If the optimizer determines that a permanent index might be beneficial, it returns the key columns necessary to create the suggested index.

## Database monitor view 3021 - Bitmap Created

Displays the SQL logical view format for database monitor QQQ3021.

```

Create View QQQ3021 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QVPARPF as Parallel_Prefetch,
 QVPARPL as Parallel_PreLoad,
 QVPARD as Parallel_Degree_Requested,
 QVPARU as Parallel_Degree_Used,
 QVPARRC as Parallel_Degree_Reason_Code,
 QQEPT as Estimated_Processing_Time,
 QVCTIM as Estimated_Cumulative_Time,
 QQREST as Estimated_Rows_Selected,
 QQAJN as Estimated_Join_Rows,
 QQJNP as Join_Position,
 QQI6 as DataSpace_Number,
 QQC21 as Join_Method,
 QQC22 as Join_Type,
 QQC23 as Join_Operator,
 QVJFANO as Join_Fanout,
 QVFILES as Join_Table_Count,
 QQI2 as Bitmap_Size,
 QVP151 as Bitmap_Count,
 QVC3001 as Bitmap_IDs,
 QQINT03 as Storage_Estimate,
 QQSMINTF as Plan_Iteration_Number,
 QQC11 as Bitmap_Type
 FROM UserLib/DBMONTable
 WHERE QQRID=3021)

```

Table 258. QQQ3021 - Bitmap Created

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name



Table 258. QQQ3021 - Bitmap Created (continued)

View Column Name	Table Column Name	Description
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated rows selected
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>

Table 258. QQQ3021 - Bitmap Created (continued)

View Column Name	Table Column Name	Description
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Bitmap_Size	QQI2	Bitmap size
Bitmap_Count	QVP151	Number of bitmaps created
Bitmap_IDs	QVC3001	Internal bitmap IDs
Storage_Estimate	QQINT03	Estimated amount of temporary storage used, in megabytes, to create the temporary index.
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1.
Bitmap_Type	QQC11	Type of bitmap temporary <ul style="list-style-type: none"> <li>• 'E' - EVI Candidate Bitmap</li> <li>• 'L' - RRN List</li> <li>• 'B' - RRN Bitmap</li> </ul>

## Database monitor view 3022 - Bitmap Merge

Displays the SQL logical view format for database monitor QQQ3022

```

Create View QQQ3022 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,

```

```

QQJOB as Job_Name,
QQUSER as Job_User,
QQJNUM as Job_Number,
QQI9 as Thread_ID,
QQUCNT as Unique_Count,
QQUDEF as User_Defined,
QQQDTN as Unique_SubSelect_Number,
QQQDTL as SubSelect_Nested_Level,
QQMATN as Materialized_View_Subselect_Number,
QQMATL as Materialized_View_Nested_Level,
QVP15E as Materialized_View_Union_Level,
QVP15A as Decomposed_Subselect_Number,
QVP15B as Total_Number_Decomposed_SubSelects,
QVP15C as Decomposed_SubSelect_Reason_Code,
QVP15D as Starting_Decomposed_SubSelect,
QVRCNT as Unique_Refresh_Counter,
QVPARPF as Parallel_Prefetch,
QVPARPL as Parallel_PreLoad,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QQEPT as Estimated_Processing_Time,
QVCTIM as Estimated_Cumulative_Time,
QQREST as Estimated_Rows_Selected,
QQAJN as Estimated_Join_Rows,
QQJNP as Join_Position,
QQI6 as DataSpace_Number,
QQC21 as Join_Method,
QQC22 as Join_Type,
QQC23 as Join_Operator,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QQI2 as Bitmap_Size,
QVC101 as Bitmap_ID,
QVC3001 as Bitmaps_Merged,
QQINT03 as Storage_Estimate,
 QQSMINTF as Plan_Iteration_Number
FROM UserLib/DBMONTTable
WHERE QQRID=3022)

```

Table 259. QQQ3022 - Bitmap Merge

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number

Table 259. QQQ3022 - Bitmap Merge (continued)

View Column Name	Table Column Name	Description
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated rows selected
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>

Table 259. QQQ3022 - Bitmap Merge (continued)

View Column Name	Table Column Name	Description
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Bitmap_Size	QQI2	Bitmap size
Bitmap_ID	QVC101	Internal bitmap ID
Bitmaps_Merged	QVC3001	IDs of bitmaps merged together
Storage_Estimate	QQINT03	Estimated amount of temporary storage used, in megabytes, to create the final bitmap. Only set by CQE
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

## Database monitor view 3023 - Temp Hash Table Created

Displays the SQL logical view format for database monitor QQQ3023.

```

Create View QQQ3023 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QVPARPF as Parallel_Prefetch,
)

```

```

QVPARPL as Parallel_PreLoad,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QQEPT as Estimated_Processing_Time,
QVCTIM as Estimated_Cumulative_Time,
QQREST as Estimated_Rows_Selected,
QQAJN as Estimated_Join_Rows,
QQJNP as Join_Position,
QQI6 as DataSpace_Number,
QQC21 as Join_Method,
QQC22 as Join_Type,
QQC23 as Join_Operator,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QVC1F as HashTable_ReasonCode,
QQI2 as HashTable_Entries,
QQI3 as HashTable_Size,
QQI4 as HashTable_Row_Size,
QQI5 as HashTable_Key_Size,
QQIA as HashTable_Element_Size,
QQI7 as HashTable_PoolSize,
QQI8 as HashTable_PoolID,
QVC101 as HashTable_Name,
QVC102 as HashTable_Library,
QVC3001 as HashTable_Columns,
QQINT03 as Storage_Estimate,
QQSMINTF as Plan_Iteration_Number
FROM UserLib/DBMONTable
WHERE QQRID=3023)

```

Table 260. QQQ3023 - Temp Hash Table Created

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects

Table 260. QQQ3023 - Temp Hash Table Created (continued)

View Column Name	Table Column Name	Description
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated rows selected
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>

Table 260. QQQ3023 - Temp Hash Table Created (continued)

View Column Name	Table Column Name	Description
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
HashTable_ReasonCode	QVC1F	Hash table reason code <ul style="list-style-type: none"> <li>• J - Created for hash join</li> <li>• G - Created for hash grouping</li> </ul>
HashTable_Entries	QQI2	Hash table entries
HashTable_Size	QQI3	Hash table size
HashTable_Row_Size	QQI4	Hash table row size
HashTable_Key_Size	QQI5	Hash table key size
HashTable_Element_Size	QQIA	Hash table element size
HashTable_PoolSize	QQI7	Hash table pool size
HashTable_PoolID	QQI8	Hash table pool ID
HashTable_Name	QVC101	Hash table internal name
HashTable_Library	QVC102	Hash table library
HashTable_Columns	QVC3001	Columns used to create hash table
Storage_Estimate	QQINT03	Estimated amount of temporary storage used, in megabytes, to create the hash table. Only set by CQE.
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

## Database monitor view 3025 - Distinct Processing

Displays the SQL logical view format for database monitor QQQ3025.

```

Create View QQQ3025 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,

```



```

QVP15A as Decomposed_Subselect_Number,
QVP15B as Total_Number_Decomposed_SubSelects,
QVP15C as Decomposed_SubSelect_Reason_Code,
QVP15D as Starting_Decomposed_SubSelect,
QVRCNT as Unique_Refresh_Counter,
QVPARPF as Parallel_Prefetch,
QVPARPL as Parallel_PreLoad,
QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QQEPT as Estimated_Processing_Time,
QVCTIM as Estimated_Cumulative_Time,
QQREST as Estimated_Rows_Selected,
QQSMINTF as Plan_Iteration_Number
FROM UserLib/DBMONTTable
WHERE QQRID=3025)

```

Table 261. QQQ3025 - Distinct Processing

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Numbe r	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSele cts	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Cod e	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)

Table 261. QQQ3025 - Distinct Processing (continued)

View Column Name	Table Column Name	Description
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated rows selected
Plan_Iteration_Number	QQSMINT F	AQP Plan iteration number, original optimization = 1

## Database monitor view 3026 - Set operation

Displays the SQL logical view format for database monitor QQQ3026.

```

Create View QQQ3026 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QVPARPF as Parallel_Prefetch,
 QVPARPL as Parallel_PreLoad,
 QVPARD as Parallel_Degree_Requested,
 QVPARU as Parallel_Degree_Used,
 QVPARRC as Parallel_Degree_Reason_Code,
 QQEPT as Estimated_Processing_Time,
 QVCTIM as Estimated_Cumulative_Time,
 QQREST as Estimated_Rows_Selected,
 QQC11 as Union_Type,
 QVFILES as Join_Table_Count,
 QQUNIN as Has_Union,
 QWC16 as Last_Union_Subselect,
 QQC23 as Set_in_a_View,
 QQC22 as Set_Operator,
 QQSMINTF as Plan_Iteration_Number
 FROM UserLib/DBMONTable
 WHERE QQRID=3026)

```

Table 262. QQQ3026 - Set operatoin

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Numbe r	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSel ects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Co de	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (Y/N)
Parallel_Degree_Requested	QVPARD	Parallel degree requested
Parallel_Degree_Used	QVPARU	Parallel degree used
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated number of rows selected

Table 262. QQQ3026 - Set operatoin (continued)

View Column Name	Table Column Name	Description
Union_Type	QQC11	Type of union. Possible values are: <ul style="list-style-type: none"> <li>• A - Union All</li> <li>• U - Union</li> </ul>
Join_Table_Count	QVFILES	Number of tables queried
Has_Union	QQUNIN	Union subselect (Y/N)
Last_Union_Subselect	QWC16	This is the last subselect, or only subselect, for the query. (Y/N)
Set_in_a_View	QQC23	Set operation within a view. <ul style="list-style-type: none"> <li>• Byte 1 of 2 (Y/N): This subselect is part of a query that is contained within a view and it contains a set operation (for example, Union).</li> <li>• Byte 2 of 2 (Y/N): This is the last subselect of the query that is contained within a view.</li> </ul>
Set_Operator	QQC22	Type of set operation. Possible values are: <ul style="list-style-type: none"> <li>• UU - Union</li> <li>• UA - Union All</li> <li>• UR - Union Recursive</li> <li>• EE - Except</li> <li>• EA - Except All</li> <li>• II - Intersect</li> <li>• IA - Intersect All</li> </ul>
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

## Database monitor view 3027 - Subquery Merge

Displays the SQL logical view format for database monitor QQQ3027.

```

Create View QQQ3027 as
(SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QVPARPF as Parallel_Prefetch,
 QVPARPL as Parallel_PreLoad,

```

```

QVPARD as Parallel_Degree_Requested,
QVPARU as Parallel_Degree_Used,
QVPARRC as Parallel_Degree_Reason_Code,
QQEPT as Estimated_Processing_Time,
QVCTIM as Estimated_Cumulative_Time,
QQREST as Estimated_Rows_Selected,
QQAJN as Estimated_Join_Rows,
QQJNP as Join_Position,
QQI1 as DataSpace_Number,
QQC21 as Join_Method,
QQC22 as Join_Type,
QQC23 as Join_Operator,
QVJFANO as Join_Fanout,
QVFILES as Join_Table_Count,
QVP151 as Subselect_Number_of_Inner_Subquery,
QVP152 as Subselect_Level_of_Inner_Subquery,
QVP153 as Materialized_View_Subselect_Number_of_Inner,
QVP154 as Materialized_View_Nested_Level_of_Inner,
QVP155 as Materialized_View_Union_Level_of_Inner,
QQC101 as Subquery_Operator,
QVC21 as Subquery_Type,
QQC11 as Has_Correlated_Columns,
QVC3001 as Correlated_Columns,
QQSMINTF as Plan_Iteration_Number
FROM UserLib/DBMONTable
WHERE QQRID=3027)

```

Table 263. QQQ3027 - Subquery Merge

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Subselect number for outer subquery
SubSelect_Nested_Level	QQQDTL	Subselect level for outer subquery
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number for outer subquery
Materialized_View_Nested_Level	QQMATL	Materialized view subselect level for outer subquery
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects

Table 263. QQQ3027 - Subquery Merge (continued)

View Column Name	Table Column Name	Description
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated rows selected
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position - when available
DataSpace_Number	QQI6	Dataspace number
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>

Table 263. QQQ3027 - Subquery Merge (continued)

View Column Name	Table Column Name	Description
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
Subselect_Number_of_Inner_Subquery	QVP151	Subselect number for inner subquery
Subselect_Level_of_Inner_Subquery	QVP152	Subselect level for inner subquery
Materialized_View_Subselect_Number_of_Inner	QVP153	Materialized view subselect number for inner subquery
Materialized_View_Nested_Level_of_Inner	QVP154	Materialized view subselect level for inner subquery
Materialized_View_Union_Level_of_Inner	QVP155	Materialized view union level for inner subquery
Subquery_Operator	QQC101	Subquery operator. Possible values are: <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not Equal</li> <li>• LT - Less Than or Equal</li> <li>• LT - Less Than</li> <li>• GE - Greater Than or Equal</li> <li>• GT - Greater Than</li> <li>• IN</li> <li>• LIKE</li> <li>• EXISTS</li> <li>• NOT - Can precede IN, LIKE or EXISTS</li> </ul>
Subquery_Type	QVC21	Subquery type. Possible values are: <ul style="list-style-type: none"> <li>• SQ - Subquery</li> <li>• SS - Scalar subselect</li> <li>• SU - Set Update</li> </ul>
Has_Correlated_Columns	QQC11	Correlated columns exist (Y/N)
Correlated_Columns	QVC3001	List of correlated columns with corresponding QDT number
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

## Database monitor view 3028 - Grouping

Displays the SQL logical view format for database monitor QQQ3028.

```

Create View QQQ3028 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QVPARPF as Parallel_Prefetch,
 QVPARPL as Parallel_PreLoad,
 QVPARD as Parallel_Degree_Requested,
 QVPARU as Parallel_Degree_Used,
 QVPARRC as Parallel_Degree_Reason_Code,
 QQEPT as Estimated_Processing_Time,
 QVCTIM as Estimated_Cumulative_Time,
 QQREST as Estimated_Rows_Selected,
 QQAJN as Estimated_Join_Rows,
 QQJNP as Join_Position,
 QQI1 as DataSpace_Number,
 QQC21 as Join_Method,
 QQC22 as Join_Type,
 QQC23 as Join_Operator,
 QVJFANO as Join_Fanout,
 QVFILES as Join_Table_Count,
 QQC11 as GroupBy_Implementation,
 QQC101 as GroupBy_Index_Name,
 QQC102 as GroupBy_Index_Library,
 QVINAM as GroupBy_Index_Long_Name,
 QVILIB as GroupBy_Index_Long_Library,
 QQC12 as Has_Having_Selection,
 QQC13 as Having_to_Where_Selection_Conversion,
 QQI2 as Estimated_Number_of_Groups,
 QQI3 as Average_Number_Rows_per_Group,
 QVC3001 as GroupBy_Columns,
 QVC3002 as MIN_Columns,
 QVC3003 as MAX_Columns,
 QVC3004 as SUM_Columns,
 QVC3005 as COUNT_Columns,
 QVC3006 as AVG_Columns,
 QVC3007 as STDDEV_Columns,
 QVC3008 as VAR_Columns,
 QQSMINTF as Plan_Iteration_Number
 FROM UserLib/DBMONTTable
 WHERE QQRID=3028)

```

Table 264. QQQ3028 - Grouping

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name



Table 264. QQQ3028 - Grouping (continued)

<b>View Column Name</b>	<b>Table Column Name</b>	<b>Description</b>
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job user
Job_Number	QQJNUM	Job number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Numbe r	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSele cts	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Cod e	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (index used)
Parallel_Degree_Requested	QVPARD	Parallel degree requested (index used)
Parallel_Degree_Used	QVPARU	Parallel degree used (index used)
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited (index used)
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated rows selected
Estimated_Join_Rows	QQAJN	Estimated number of joined rows
Join_Position	QQJNP	Join position
DataSpace_Number	QQI1	Dataspace number

Table 264. QQQ3028 - Grouping (continued)

View Column Name	Table Column Name	Description
Join_Method	QQC21	Join method - when available <ul style="list-style-type: none"> <li>• NL - Nested loop</li> <li>• MF - Nested loop with selection</li> <li>• HJ - Hash join</li> </ul>
Join_Type	QQC22	Join type - when available <ul style="list-style-type: none"> <li>• IN - Inner join</li> <li>• PO - Left partial outer join</li> <li>• EX - Exception join</li> </ul>
Join_Operator	QQC23	Join operator - when available <ul style="list-style-type: none"> <li>• EQ - Equal</li> <li>• NE - Not equal</li> <li>• GT - Greater than</li> <li>• GE - Greater than or equal</li> <li>• LT - Less than</li> <li>• LE - Less than or equal</li> <li>• CP - Cartesian product</li> </ul>
Join_Fanout	QVJFANO	Join fan out. Possible values are: <ul style="list-style-type: none"> <li>• N - Normal join situation where fanout is allowed and each matching row of the join fanout is returned.</li> <li>• D - Distinct fanout. Join fanout is allowed however none of the join fanout rows are returned.</li> <li>• U - Unique fanout. Join fanout is not allowed. Error situation if join fanout occurs.</li> </ul>
Join_Table_Count	QVFILES	Number of tables joined
GroupBy_Implementation	QQC11	Group by implementation <ul style="list-style-type: none"> <li>• ' ' - No grouping</li> <li>• I - Index</li> <li>• H - Hash</li> </ul>
GroupBy_Index_Name	QQC101	Index, or constraint, used for grouping
GroupBy_Index_Library	QQC102	Library of index used for grouping
GroupBy_Index_Long_Name	QVINAM	Long name of index, or constraint, used for grouping
GroupBy_Index_Long_Library	QVILIB	Long name of index, or constraint, library used for grouping
Has_Having_Selection	QQC12	Having selection exists (Y/N)
Having_to_Where_Selection_Conversion	QQC13	Having to Where conversion (Y/N)

Table 264. QQQ3028 - Grouping (continued)

View Column Name	Table Column Name	Description
Estimated_Number_of_Groups	QQI2	Estimated number of groups
Average_Number_Rows_per_Group	QQI3	Average number of rows in each group
GroupBy_Columns	QVC3001	Grouping columns
MIN_Columns	QVC3002	MIN columns
MAX_Columns	QVC3003	MAX columns
SUM_Columns	QVC3004	SUM columns
COUNT_Columns	QVC3005	COUNT columns
AVG_Columns	QVC3006	AVG columns
STDDEV_Columns	QVC3007	STDDEV columns
VAR_Columns	QVC3008	VAR columns
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

### Database monitor view 3030 - Materialized query tables

Displays the SQL logical view format for database monitor QQQ3030.

```

Create View QQQ3030 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QQ1000 as Materialized_Query_Tables,
 QQC301 as MQT_Reason_Codes,
 QQSMINTF as Plan_Iteration_Number
 FROM UserLib/DBMONTable
 WHERE QQRID=3030)

```

Table 265. QQQ3030 - Materialized query tables

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created

Table 265. QQQ3030 - Materialized query tables (continued)

View Column Name	Table Column Name	Description
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job User
Job_Number	QQJNUM	Job Number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Materialized_Query_Tables	QQ1000	Materialized query tables examined and reason why used or not used: <ul style="list-style-type: none"> <li>• 0 - The materialized query table was used</li> <li>• 1 - The cost to use the materialized query table, as determined by the optimizer, was higher than the cost associated with the chosen implementation.</li> <li>• 2 - The join specified in the materialized query was not compatible with the query.</li> <li>• 3 - The materialized query table had predicates that were not matched in the query.</li> <li>• 4 - The grouping specified in the materialized query table is not compatible with the grouping specified in the query.</li> </ul>

Table 265. QQQ3030 - Materialized query tables (continued)

View Column Name	Table Column Name	Description
Materialized_Query_Tables (continued)		<ul style="list-style-type: none"> <li>• 5 - The query specified columns that were not in the select-list of the materialized query table.</li> <li>• 6 - The materialized query table query contains functionality that is not supported by the query optimizer.</li> <li>• 7 - The materialized query table specified the DISABLE QUERY OPTIMIZATION clause.</li> <li>• 8 - The ordering specified in the materialized query table is not compatible with the ordering specified in the query.</li> <li>• 9 - The query contains functionality that is not supported by the materialized query table matching algorithm.</li> <li>• 10 - Materialized query tables may not be used for this query.</li> <li>• 11 - The refresh age of this materialized query table exceeds the duration specified by the MATERIALIZED_QUERY_TABLE_REFRESH_AGE QAQQINI option.</li> <li>• 12 - The commit level of the materialized query table is lower than the commit level specified for the query.</li> <li>• 13 - The distinct specified in the materialized query table is not compatible with the distinct specified in the query.</li> <li>• 14 - The FETCH FOR FIRST n ROWS clause of the materialized query table is not compatible with the query.</li> <li>• 15 - The QAQQINI options used to create the materialized query table are not compatible with the QAQQINI options used to run this query.</li> <li>• 16 - The materialized query table is not usable.</li> <li>• 17 - The union specified in the materialized query table is not compatible with the query.</li> <li>• 18 - The constants specified in the materialized query table are not compatible with host variable values specified in the query.</li> <li>• 19 - The Materialized query table is in check pending status.</li> <li>• 20 - The UDTF specified in the materialized query table was not compatible with the query.</li> <li>• 21 - The VALUES clause specified in the materialized query table was not compatible with the query.</li> <li>• 22 - The UNNEST clause specified in the materialized query table was not compatible with the query.</li> </ul>

Table 265. QQQ3030 - Materialized query tables (continued)

View Column Name	Table Column Name	Description
MQT_Reason_Codes	QQC301	List of unique reason codes used by the materialized query tables (each materialized query table has a corresponding reason code associated with it)
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

### Database monitor view 3031 - Recursive common table expressions

Displays the SQL logical view format for database monitor QQQ3031.

```

Create View QQQ3031 as
 (SELECT QQRID as Row_ID,
 QQTIME as Time_Created,
 QQJFLD as Join_Column,
 QQRDBN as Relational_Database_Name,
 QQSYS as System_Name,
 QQJOB as Job_Name,
 QQUSER as Job_User,
 QQJNUM as Job_Number,
 QQI9 as Thread_ID,
 QQUCNT as Unique_Count,
 QQUDEF as User_Defined,
 QQQDTN as Unique_SubSelect_Number,
 QQQDTL as SubSelect_Nested_Level,
 QQMATN as Materialized_View_Subselect_Number,
 QQMATL as Materialized_View_Nested_Level,
 QVP15E as Materialized_View_Union_Level,
 QVP15A as Decomposed_Subselect_Number,
 QVP15B as Total_Number_Decomposed_SubSelects,
 QVP15C as Decomposed_SubSelect_Reason_Code,
 QVP15D as Starting_Decomposed_SubSelect,
 QVRCNT as Unique_Refresh_Counter,
 QVPARPF as Parallel_Prefetch,
 QVPARPL as Parallel_PreLoad,
 QVPARD as Parallel_Degree_Requested,
 QVPARU as Parallel_Degree_Used,
 QVPARRC as Parallel_Degree_Reason_Code,
 QQEPT as Estimated_Processing_Time,
 QVCTIM as Estimated_Cumulative_Time,
 QQREST as Estimated_Rows_Selected,
 QQC11 as Recursive_Query_Cycle_Check,
 QQC15 as Recursive_Query_Search_Option,
 QQI2 as Number_of_Recursive_Values,
 QQSMINTF as Plan_Iteration_Number
 FROM UserLib/DBMONTable
 WHERE QQRID=3031)

```

Table 266. QQQ3031 - Recursive common table expressions

View Column Name	Table Column Name	Description
Row_ID	QQRID	Row identification
Time_Created	QQTIME	Time row was created
Join_Column	QQJFLD	Join column (unique per job)
Relational_Database_Name	QQRDBN	Relational database name
System_Name	QQSYS	System name
Job_Name	QQJOB	Job name
Job_User	QQUSER	Job User

Table 266. QQQ3031 - Recursive common table expressions (continued)

View Column Name	Table Column Name	Description
Job_Number	QQJNUM	Job Number
Thread_ID	QQI9	Thread identifier
Unique_Count	QQUCNT	Unique count (unique per query)
User_Defined	QQUDEF	User defined column
Unique_SubSelect_Number	QQQDTN	Unique subselect number
SubSelect_Nested_Level	QQQDTL	Subselect nested level
Materialized_View_Subselect_Number	QQMATN	Materialized view subselect number
Materialized_View_Nested_Level	QQMATL	Materialized view nested level
Materialized_View_Union_Level	QVP15E	Materialized view union level
Decomposed_Subselect_Number	QVP15A	Decomposed query subselect number, unique across all decomposed subselects
Total_Number_Decomposed_SubSelects	QVP15B	Total number of decomposed subselects
Decomposed_SubSelect_Reason_Code	QVP15C	Decomposed query subselect reason code
Starting_Decomposed_SubSelect	QVP15D	Decomposed query subselect number for the first decomposed subselect
Unique_Refresh_Counter	QVRCNT	Unique refresh counter
Parallel_Prefetch	QVPARPF	Parallel Prefetch (Y/N)
Parallel_PreLoad	QVPARPL	Parallel Preload (Y/N)
Parallel_Degree_Requested	QVPARD	Parallel degree requested
Parallel_Degree_Used	QVPARU	Parallel degree used
Parallel_Degree_Reason_Code	QVPARRC	Reason parallel processing was limited
Estimated_Processing_Time	QQEPT	Estimated processing time, in seconds
Estimated_Cumulative_Time	QVCTIM	Estimated cumulative time, in seconds
Estimated_Rows_Selected	QQREST	Estimated number of rows selected
Recursive_Query_Cycle_Check	QQC11	CYCLE option: <ul style="list-style-type: none"> <li>• Y - checking for cyclic data</li> <li>• N - not checking for cyclic data</li> </ul>
Recursive_Query_Search_Option	QQC15	SEARCH option: <ul style="list-style-type: none"> <li>• N - None specified</li> <li>• D - Depth first</li> <li>• B - Breadth first</li> </ul>

Table 266. QQQ3031 - Recursive common table expressions (continued)

View Column Name	Table Column Name	Description
Number_of_Recursive_Values	QQI2	Number of values put on queue to implement recursion. Includes values necessary for CYCLE and SEARCH options.
Plan_Iteration_Number	QQSMINTF	AQP Plan iteration number, original optimization = 1

## Query optimizer messages reference

See the following for query optimizer message reference:

### Query optimization performance information messages

You can evaluate the structure and performance of the SQL statements in a program using informational messages. These messages are put in the job log by the database manager.

The messages are issued for an SQL program or interactive SQL when running in the debug mode. The database manager could send any of the following messages when appropriate. The ampersand variables (&1, &X) are replacement variables that contain either an object name or other substitution value when you see the message in the job log. These messages provide feedback on how a query was run. In some cases, the messages indicate improvements you can make to help the query run faster.

The messages contain message help that provides information about the cause for the message, object name references, and possible user responses.

The time at which the message is sent does not necessarily indicate when the associated function was performed. Some messages are sent altogether at the start of a query run.

<b>CPI4321 - Access path built for &amp;18 &amp;19</b>	
<b>Message Text:</b>	Access path built for &18 &19.



**CPI4321 - Access path built for &18 &19****Cause Text:**

A temporary access path was built to access records from member &6 of &18 &19 in library &5 for reason code &10. This process took &11 minutes and &12 seconds. The access path built contains &15 entries. The access path was built using &16 parallel tasks. A zero for the number of parallel tasks indicates that parallelism was not used. The reason codes and their meanings follow:

- 1 - Perform specified ordering/grouping criteria.
- 2 - Perform specified join criteria.
- 3 - Perform specified record selection to minimize I/O wait time.

The access path was built using the following key fields. The key fields and their corresponding sequence (ASCEND or DESCEND) will be shown:

&17.

A key field of \*MAP indicates the key field is an expression (derived field).

The access path was built using sequence table &13 in library &14.

A sequence table of \*N indicates the access path was built without a sequence table. A sequence table of \*I indicates the table was an internally derived table that is not available to the user.

If &18 &19 in library &5 is a logical file then the access path is built over member &9 of physical file &7 in library &8.

A file name starting with \*QUERY or \*N indicates the access path was built over a temporary file.

**Recovery Text:**

If this query is run frequently, you may want to create an access path (index) similar to this definition for performance reasons. Create the access path using sequence table &13 in library &14, unless the sequence table is \*N. If an access path is created, it is possible the query optimizer may still choose to create a temporary access path to process the query.

If \*MAP is returned for one of the key fields or \*I is returned for the sequence table, then a permanent access path cannot be created. A permanent access path cannot be built with these specifications.

**CPI4322 - Access path built from keyed file &1****Message Text:**

Access path built from keyed file &1.

**CPI4322 - Access path built from keyed file &1****Cause Text:**

A temporary access path was built using the access path from member &3 of keyed file &1 in library &2 to access records from member &6 of file &4 in library &5 for reason code &10. This process took &11 minutes and &12 seconds. The access path built contains &15 entries. The reason codes and their meanings follow:

- 1 - Perform specified ordering/grouping criteria.
- 2 - Perform specified join criteria.
- 3 - Perform specified record selection to minimize I/O wait time.

The access path was built using the following key fields. The key fields and their corresponding sequence (ASCEND or DESCEND) will be shown:

&17.

A key field of \*MAP indicates the key field is an expression (derived field).

The temporary access path was built using sequence table &13 in library &14.

A sequence table of \*N indicates the access path was built without a sequence table. A sequence table of \*I indicates the table was an internally derived table that is not available to the user.

If file &4 in library &5 is a logical file then the temporary access path is built over member &9 of physical file &7 in library &8. Creating an access path from a keyed file generally results in improved performance.

**Recovery Text:**

If this query is run frequently, you may want to create an access path (index) similar to this definition for performance reasons. Create the access path using sequence table &13 in library &14, unless the sequence table is \*N. If an access path is created, it is possible the query optimizer may still choose to create a temporary access path to process the query.

If \*MAP is returned for one of the key fields or \*I is returned for the sequence table, then a permanent access path cannot be created. A permanent access path cannot be built with these specifications.

A temporary access path can only be created using index only access if all of the fields that were used by this temporary access path are also key fields for the access path from the keyed file.

**CPI4323 - The query access plan has been rebuilt****Message Text:**

The query access plan has been rebuilt.

**CPI4323 - The query access plan has been rebuilt**

<b>Cause Text:</b>	<p>The access plan was rebuilt for reason code &amp;13. The reason codes and their meanings follow:</p> <ul style="list-style-type: none"><li>0 - A new access plan was created.</li><li>1 - A file or member is not the same object as the one referred to in the access plan. Some reasons include the object being recreated, restored, or overridden to a new object.</li><li>2 - Access plan was using a reusable Open Data Path (ODP), and the optimizer chose to use a non-reusable ODP.</li><li>3 - Access plan was using a non-reusable Open Data Path (ODP) and the optimizer chose to use a reusable ODP.</li><li>4 - The number of records in member &amp;3 of file &amp;1 in library &amp;2 has changed by more than 10%.</li><li>5 - A new access path exists over member &amp;6 of file &amp;4 in library &amp;5.</li><li>6 - An access path over member &amp;9 of file &amp;7 in library &amp;8 that was used for this access plan no longer exists or is no longer valid.</li><li>7 - The query access plan had to be rebuilt because of system programming changes.</li><li>8 - The CCSID (Coded Character Set Identifier) of the current job is different than the CCSID used in the access plan.</li><li>9 - The value of one of the following is different in the current job: date format, date separator, time format, or time separator.</li><li>10 - The sort sequence table specified has changed.</li><li>11 - The number of active processors or the size or paging option of the storage pool has changed.</li><li>12 - The system feature DB2 Symmetric Multiprocessing has either been installed or removed.</li><li>13 - The value of the degree query attribute has changed either by the <b>CHGSYSVAL</b> or <b>CHGQRYA</b> CL commands or with the query options file &amp;15 in library &amp;16.</li><li>14 - A view is either being opened by a high level language open, or is being materialized.</li><li>15 - A sequence object or user-defined type or function is not the same object as the one referred to in the access plan; or, the SQL path used to generate the access plan is different than the current SQL path.</li><li>16 - Query attributes have been specified from the query options file &amp;15 in library &amp;16.</li><li>17 - The access plan was generated with a commitment control level that is different in the current job.</li><li>18 - The access plan was generated with a different static cursor answer set size.</li><li>19 - This is the first run of the query since a prepare or compile.</li><li>20 and greater -- View the second level message text of the next message issued (CPI4351) for an explanation of these reason codes.</li></ul> <p>If the reason code is 4, 5, 6, 20, or 21 and the file specified in the reason code explanation is a logical file, then member &amp;12 of physical file &amp;10 in library &amp;11 is the file with the specified change.</p>
<b>Recovery Text:</b>	Excessive rebuilds should be avoided and may indicate an application design problem.

**CPI4324 - Temporary file built for file &1**

<b>Message Text:</b>	Temporary file built for file &1.
----------------------	-----------------------------------

**CPI4324 - Temporary file built for file &1**

<b>Cause Text:</b>	<p>A temporary file was built for member &amp;3 of file &amp;1 in library &amp;2 for reason code &amp;4. This process took &amp;5 minutes and &amp;6 seconds. The temporary file was required in order for the query to be processed. The reason codes and their meanings follow:</p> <p>1 - The file is a join logical file and its join-type (JDFTVAL) does not match the join-type specified in the query.</p> <p>2 - The format specified for the logical file references more than one physical file.</p> <p>3 - The file is a complex SQL view, or nested table expression, or common table expression, or is a data change table reference that requires a temporary file.</p> <p>4 - For an update-capable query, a subselect references a field in this file which matches one of the fields being updated.</p> <p>5 - For an update-capable query, a subselect references SQL view &amp;1, which is based on the file being updated.</p> <p>6 - For a delete-capable query, a subselect references either the file from which records are to be deleted or an SQL view or logical file based on the file from which records are to be deleted.</p> <p>7 - The file is user-defined table function &amp;8 in &amp;2, and all the records were retrieved from the function. The processing time is not returned for this reason code.</p> <p>8 - The file is a partition file requiring a temporary file for processing the grouping or join.</p>
<b>Recovery Text:</b>	You may want to change the query to refer to a file that does not require a temporary file to be built.

**CPI4325 - Temporary result file built for query**

<b>Message Text:</b>	Temporary result file built for query.
----------------------	----------------------------------------

**CPI4325 - Temporary result file built for query**

<b>Cause Text:</b>	<p>A temporary result file was created to contain the results of the query for reason code &amp;4. This process took &amp;5 minutes and &amp;6 seconds. The temporary file created contains &amp;7 records. The reason codes and their meanings follow:</p> <ol style="list-style-type: none"><li>1 - The query contains grouping fields (GROUP BY) from more than one file, or contains grouping fields from a secondary file of a join query that cannot be reordered.</li><li>2 - The query contains ordering fields (ORDER BY) from more than one file, or contains ordering fields from a secondary file of a join query that cannot be reordered.</li><li>3 - The grouping and ordering fields are not compatible.</li><li>4 - DISTINCT was specified for the query.</li><li>5 - Set operator (UNION, EXCEPT, or INTERSECT) was specified for the query.</li><li>6 - The query had to be implemented using a sort. More than 120 key fields specified for ordering.</li><li>7 - The query optimizer chose to use a sort rather than an access path to order the results of the query.</li><li>8 - Perform specified record selection to minimize I/O wait time.</li><li>9 - The query optimizer chose to use a hashing algorithm rather than an access path to perform the grouping for the query.</li><li>10 - The query contains a join condition that requires a temporary file.</li><li>11 - The query optimizer creates a run-time temporary file in order to implement certain correlated group by queries.</li><li>12 - The query contains grouping fields (GROUP BY, MIN/MAX, COUNT, etc.) and there is a read trigger on one or more of the underlying physical files in the query.</li><li>13 - The query involves a static cursor or the SQL FETCH FIRST clause.</li></ol>
<b>Recovery Text:</b>	For more information on why a temporary result was used, refer to <a href="#">“Data access methods”</a> on page 16.

**CPI4325 - Temporary result file built for query**

<b>Message Text:</b>	&12 &13 processed in join position &10.
----------------------	-----------------------------------------

**CPI4325 - Temporary result file built for query**

**Cause Text:** Access path for member &5 of file &3 in library &4 was used to access records in member &2 of file &13 in library &1 for reason code &9. The reason codes and their meanings follow:

- 1 - Perform specified record selection.
- 2 - Perform specified ordering/grouping criteria.
- 3 - Record selection and ordering/grouping criteria.
- 4 - Perform specified join criteria.

If file &13 in library &1 is a logical file then member &8 of physical file &6 in library &7 is the actual file in join position &10.

A file name starting with \*TEMPX for the access path indicates it is a temporary access path built over file &6.

A file name starting with \*N or \*QUERY for the file indicates it is a temporary file.

Index only access was used for this file within the query: &11.

A value of \*YES for index only access processing indicates that all of the fields used from this file for this query can be found within the access path of file &3. A value of \*NO indicates that index only access could not be performed for this access path.

Index only access is generally a performance advantage since all of the data can be extracted from the access path and the data space does not have to be paged into active memory.

**Recovery Text:** Generally, to force a file to be processed in join position 1, specify an order by field from that file only.

If ordering is desired, specifying ORDER BY fields over more than one file forces the creation of a temporary file and allows the optimizer to optimize the join order of all the files. No file is forced to be first.

An access path can only be considered for index only access if all of the fields used within the query for this file are also key fields for that access path.

Refer to the “Data access methods” on page 16 for additional tips on optimizing a query's join order and index only access.

In some cases, creating a temporary result table provides the fastest way to run a query. Other queries that have many rows to be copied into the temporary result table can take a significant amount of time. However, if the query is taking more time and resources than can be allowed, consider changing the query so that a temporary result table is not required.

**CPI4326 - &12 &13 processed in join position &10**

**Message Text:** &12 &13 processed in join position &10.

**CPI4326 - &12 &13 processed in join position &10**

**Cause Text:** Access path for member &5 of file &3 in library &4 was used to access records in member &2 of file &13 in library &1 for reason code &9. The reason codes and their meanings follow:

- 1 - Perform specified record selection.
- 2 - Perform specified ordering/grouping criteria.
- 3 - Record selection and ordering/grouping criteria.
- 4 - Perform specified join criteria.

If file &13 in library &1 is a logical file then member &8 of physical file &6 in library &7 is the actual file in join position &10.

A file name starting with \*TEMPX for the access path indicates it is a temporary access path built over file &6.

A file name starting with \*N or \*QUERY for the file indicates it is a temporary file.

Index only access was used for this file within the query: &11.

A value of \*YES for index only access processing indicates that all of the fields used from this file for this query can be found within the access path of file &3. A value of \*NO indicates that index only access could not be performed for this access path.

Index only access is generally a performance advantage since all of the data can be extracted from the access path and the data space does not have to be paged into active memory.

**Recovery Text:** Generally, to force a file to be processed in join position 1, specify an order by field from that file only.

If ordering is desired, specifying ORDER BY fields over more than one file forces the creation of a temporary file and allows the optimizer to optimize the join order of all the files. No file is forced to be first.

An access path can only be considered for index only access if all of the fields used within the query for this file are also key fields for that access path.

Refer to the “Data access methods” on page 16 for additional tips on optimizing a query's join order and index only access.

In some cases, creating a temporary result table provides the fastest way to run a query. Other queries that have many rows to be copied into the temporary result table can take a significant amount of time. However, if the query is taking more time and resources than can be allowed, consider changing the query so that a temporary result table is not required.

This message provides the join position of the specified table when an index is used to access the table data. **Join position** pertains to the order in which the tables are joined.

**CPI4327 - File &12 &13 processed in join position &10**

**Message Text:** &12 &13 processed in join position &10.

**Cause Text:** Arrival sequence access was used to select records from member &2 of file &13 in library &1.

If file &13 in library &1 is a logical file then member &8 of physical file &6 in library &7 is the actual file in join position &10.

A file name that starts with \*QUERY for the file indicates it is a temporary file.

**CPI4327 - File &12 &13 processed in join position &10****Recovery Text:**

Generally, to force a file to be processed in join position 1, specify an order by field from that file only.

Refer to the [“Data access methods”](#) on page 16 for additional tips on optimizing a query's join order.

**CPI4328 - Access path of file &3 was used by query****Message Text:**

Access path of file &3 was used by query.

**Cause Text:**

Access path for member &5 of file &3 in library &4 was used to access records from member &2 of &12 &13 in library &1 for reason code &9. The reason codes and their meanings follow:

- 1 - Record selection.
- 2 - Ordering/grouping criteria.
- 3 - Record selection and ordering/grouping criteria.

If file &13 in library &1 is a logical file then member &8 of physical file &6 in library &7 is the actual file being accessed.

Index only access was used for this query: &11.

A value of \*YES for index only access processing indicates that all of the fields used for this query can be found within the access path of file &3. A value of \*NO indicates that index only access could not be performed for this access path.

Index only access is generally a performance advantage since all of the data can be extracted from the access path and the data space does not have to be paged into active memory.

**Recovery Text:**

An access path can only be considered for index only access if all of the fields used within the query for this file are also key fields for that access path.

Refer to the [“Data access methods”](#) on page 16. for additional tips on index only access.

**CPI4329 - Arrival sequence access was used for &12 &13****Message Text:**

Arrival sequence access was used for &12 &13.

**Cause Text:**

Arrival sequence access was used to select records from member &2 of file &13 in library &1.

If file &13 in library &1 is a logical file then member &8 of physical file &6 in library &7 is the actual file from which records are being selected.

A file name starting with \*N or \*QUERY for the file indicates it is a temporary file.

**Recovery Text:**

The use of an access path may improve the performance of the query if record selection is specified.

If an access path does not exist, you may want to create one whose left-most key fields match fields in the record selection. Matching more key fields in the access path with fields in the record selection will result in improved performance.

Generally, to force the use of an existing access path, specify order by fields that match the left-most key fields of that access path.

For more information refer to [“Data access methods”](#) on page 16.



<b>CPI432A - Query optimizer timed out for file &amp;1</b>	
<b>Message Text:</b>	Query optimizer timed out for file &1.
<b>Cause Text:</b>	<p>The query optimizer timed out before it could consider all access paths built over member &amp;3 of file &amp;1 in library &amp;2.</p> <p>The list below shows the access paths considered before the optimizer timed out. If file &amp;1 in library &amp;2 is a logical file then the access paths specified are actually built over member &amp;9 of physical file &amp;7 in library &amp;8. Following each access path name in the list is a reason code which explains how the optimizer considered the access path.</p> <p>&amp;11.</p> <p>The reason codes and their meanings follow:</p> <p>0 - The access path was used to implement the query.</p> <p>1 - Access path was not in a valid state. The system invalidated the access path.</p> <p>2 - Access path was not in a valid state. The user requested that the access path be rebuilt.</p> <p>3 - Access path is a temporary access path (resides in library QTEMP) and was not specified as the file to be queried.</p> <p>4 - The cost to use this access path, as determined by the optimizer, was higher than the cost associated with the chosen access method.</p> <p>5 - The keys of the access path did not match the fields specified for the ordering/grouping criteria.</p> <p>6 - The keys of the access path did not match the fields specified for the join criteria.</p> <p>7 - Use of this access path would not minimize delays when reading records from the file as the user requested.</p> <p>8 - The access path cannot be used for a secondary file of the join query because it contains static select/omit selection criteria. The join-type of the query does not allow the use of select/omit access paths for secondary files.</p> <p>9 - File &amp;1 contains record ID selection. The join-type of the query forces a temporary access path to be built to process the record ID selection.</p> <p>10 and greater - View the second level message text of the next message issued (CPI432D) for an explanation of these reason codes.</p>
<b>Recovery Text:</b>	<p>To ensure an access path is considered for optimization specify that access path to be the queried file. The optimizer will first consider the access path of the file specified on the query. SQL-created indexes cannot be queried but can be deleted and recreated to increase the chance they will be considered during query optimization.</p> <p>The user may want to delete any access paths no longer needed.</p>

<b>CPI432B - Subselects processed as join query</b>	
<b>Message Text:</b>	Subselects processed as join query.
<b>Cause Text:</b>	Two or more SQL subselects were combined together by the query optimizer and processed as a join query. Processing subselects as a join query generally results in improved performance.
<b>Recovery Text:</b>	None — Generally, this method of processing is a good performing option.

<b>CPI432C - All access paths were considered for file &amp;1</b>	
<b>Message Text:</b>	All access paths were considered for file &1.
<b>Cause Text:</b>	<p>The query optimizer considered all access paths built over member &amp;3 of file &amp;1 in library &amp;2.</p> <p>The list below shows the access paths considered. If file &amp;1 in library &amp;2 is a logical file then the access paths specified are actually built over member &amp;9 of physical file &amp;7 in library &amp;8. Following each access path name in the list is a reason code which explains how the optimizer considered the access path.</p> <p>&amp;11.</p> <p>The reason codes and their meanings follow:</p> <p>0 - The access path was used to implement the query.</p> <p>1 - Access path was not in a valid state. The system invalidated the access path.</p> <p>2 - Access path was not in a valid state. The user requested that the access path be rebuilt.</p> <p>3 - Access path is a temporary access path (resides in library QTEMP) and was not specified as the file to be queried.</p> <p>4 - The cost to use this access path, as determined by the optimizer, was higher than the cost associated with the chosen access method.</p> <p>5 - The keys of the access path did not match the fields specified for the ordering/grouping criteria. For distributed file queries, the access path keys must exactly match the ordering fields if the access path is to be used when ALWCPYDTA(*YES or *NO) is specified.</p> <p>6 - The keys of the access path did not match the fields specified for the join criteria.</p> <p>7 - Use of this access path would not minimize delays when reading records from the file. The user requested to minimize delays when reading records from the file.</p> <p>8 - The access path cannot be used for a secondary file of the join query because it contains static select/omit selection criteria. The join-type of the query does not allow the use of select/omit access paths for secondary files.</p> <p>9 - File &amp;1 contains record ID selection. The join-type of the query forces a temporary access path to be built to process the record ID selection.</p> <p>10 and greater - View the second level message text of the next message issued (CPI432D) for an explanation of these reason codes.</p>
<b>Recovery Text:</b>	The user may want to delete any access paths no longer needed.
<b>CPI432D - Additional access path reason codes were used</b>	
<b>Message Text:</b>	Additional access path reason codes were used.

**CPI432D - Additional access path reason codes were used**

<b>Cause Text:</b>	<p>Message CPI432A or CPI432C was issued immediately before this message. Because of message length restrictions, some of the reason codes used by messages CPI432A and CPI432C are explained below rather than in those messages.</p> <p>The reason codes and their meanings follow:</p> <p>10 - The user specified ignore decimal data errors on the query. This disallows the use of permanent access paths.</p> <p>11 - The access path contains static select/omit selection criteria which is not compatible with the selection in the query.</p> <p>12 - The access path contains static select/omit selection criteria whose compatibility with the selection in the query could not be determined. Either the select/omit criteria or the query selection became too complex during compatibility processing.</p> <p>13 - The access path cannot be used because it contains one or more keys which may be changed by the query during an insert or update.</p> <p>14 - The access path is being deleted or is being created in an uncommitted unit of work in another process.</p> <p>15 - The keys of the access path matched the fields specified for the ordering/grouping criteria. However, the sequence table associated with the access path did not match the sequence table associated with the query.</p> <p>16 - The keys of the access path matched the fields specified for the join criteria. However, the sequence table associated with the access path did not match the sequence table associated with the query.</p> <p>17 - The left-most key of the access path did not match any fields specified for the selection criteria. Therefore, key row positioning could not be performed, making the cost to use this access path higher than the cost associated with the chosen access method.</p> <p>18 - The left-most key of the access path matched a field specified for the selection criteria. However, the sequence table associated with the access path did not match the sequence table associated with the query. Therefore, key row positioning could not be performed, making the cost to use this access path higher than the cost associated with the chosen access method.</p> <p>19 - The access path cannot be used because the secondary file of the join query is a select/omit logical file. The join-type requires that the select/omit access path associated with the secondary file be used or, if dynamic, that an access path be created by the system.</p> <p>99 - The access path was used to gather statistics information for the query optimizer.</p>
<b>Recovery Text:</b>	See prior message CPI432A or CPI432C for more information.

Because of message length restrictions, some of the reason codes used by messages CPI432A and CPI432C are explained in the message help of CPI432D. Use the message help from this message to interpret the information returned from message CPI432A or CPI432C.

**CPI432E - Selection fields mapped to different attributes**

<b>Message Text:</b>	Selection fields mapped to different attributes.
----------------------	--------------------------------------------------

<b>CPI432E - Selection fields mapped to different attributes</b>	
<b>Cause Text:</b>	<p>The data type, digits, decimal position, or length of each of the following selection fields was changed so that the field could be properly compared to the literal, host variable, or field operand associated with it. Therefore, an access path cannot be used to process that selection, since no key field has attributes that match the new attributes of the field. &amp;1.</p> <p>The data type of the field may have been changed to match the comparison operand. For a numeric field, the number of total digits or fractional digits of the comparison operand may have exceeded that of the field.</p>
<b>Recovery Text:</b>	<p>You may want to change each comparison operand as follows:</p> <ol style="list-style-type: none"> <li>1 - For a literal, change the literal value so that its attributes match the field's attributes. Normally, an attributes mismatch is caused by a numeric literal that has non-significant leading or trailing zeroes.</li> <li>2 - For a host variable, either change the host variable's definition to match the field's definition or define a new host variable that matches the field's definition.</li> <li>3 - For a field, change the attributes of one of the fields to match the other's attributes.</li> </ol>

<b>CPI432F - Access path suggestion for file &amp;1</b>	
<b>Message Text:</b>	Access path suggestion for file &1.
<b>Cause Text:</b>	<p>To improve performance the query optimizer is suggesting a permanent access path be built with the key fields it is recommending. The access path will access records from member &amp;3 of file &amp;1 in library &amp;2.</p> <p>In the list of key fields that follow, the query optimizer is recommending the first &amp;10 key fields as primary key fields. The remaining key fields are considered secondary key fields and are listed in order of expected selectivity based on this query. Primary key fields are fields that significantly reduce the number of keys selected based on the corresponding selection predicate. Secondary key fields are fields that may or may not significantly reduce the number of keys selected. It is up to the user to determine the true selectivity of secondary key fields and to determine whether those key fields should be used when creating the access path.</p> <p>The query optimizer is able to perform key positioning over any combination of the primary key fields, plus one additional secondary key field. Therefore it is important that the first secondary key field be the most selective secondary key field. The query optimizer will use key selection with any remaining secondary key fields. While key selection is not as fast as key positioning it can still reduce the number of keys selected. Hence, secondary key fields that are fairly selective should be included. When building the access path all primary key fields should be specified first followed by the secondary key fields which are prioritized by selectivity. The following list contains the suggested primary and secondary key fields:</p> <p>&amp;11.</p> <p>If file &amp;1 in library &amp;2 is a logical file then the access path should be built over member &amp;9 of physical file &amp;7 in library &amp;8.</p>
<b>Recovery Text:</b>	<p>If this query is run frequently, you may want to create the suggested access path for performance reasons. It is possible that the query optimizer will choose not to use the access path just created.</p> <p>For more information, refer to <a href="#">“Data access methods”</a> on page 16.</p>

<b>CPI4330 - &amp;6 tasks used for parallel &amp;10 scan of file &amp;1</b>	
<b>Message Text:</b>	&6 tasks used for parallel &10 scan of file &1.
<b>Cause Text:</b>	<p>&amp;6 is the average numbers of tasks used for a &amp;10 scan of member &amp;3 of file &amp;1 in library &amp;2.</p> <p>If file &amp;1 in library &amp;2 is a logical file, then member &amp;9 of physical file &amp;7 in library &amp;8 is the actual file from which records are being selected.</p> <p>A file name starting with *QUERY or *N for the file indicates a temporary result file is being used.</p> <p>The query optimizer has calculated that the optimal number of tasks is &amp;5 which was limited for reason code &amp;4. The reason code definitions are:</p> <ol style="list-style-type: none"> <li>1 - The *NBRTASKS parameter value was specified for the DEGREE parameter of the CHGQRYA CL command.</li> <li>2 - The optimizer calculated the number of tasks which would use all of the central processing units (CPU).</li> <li>3 - The optimizer calculated the number of tasks which can efficiently run in this job's share of the memory pool.</li> <li>4 - The optimizer calculated the number of tasks which can efficiently run using the entire memory pool.</li> <li>5 - The optimizer limited the number of tasks to equal the number of disk units which contain the file's data.</li> </ol> <p>The database manager may further limit the number of tasks used if the allocation of the file's data is not evenly distributed across disk units.</p>
<b>Recovery Text:</b>	<p>To disallow usage of parallel &amp;10 scan, specify *NONE on the query attribute degree.</p> <p>A larger number of tasks might further improve performance. The following actions based on the optimizer reason code might allow the optimizer to calculate a larger number:</p> <ol style="list-style-type: none"> <li>1 - Specify a larger number of tasks value for the DEGREE parameter of the CHGQRYA CL command. Start with a value for number of tasks which is a slightly larger than &amp;5.</li> <li>2 - Simplify the query by reducing the number of fields being mapped to the result buffer or by removing expressions. Also, try specifying a number of tasks as described by reason code 1.</li> <li>3 - Specify *MAX for the query attribute DEGREE.</li> <li>4 - Increase the size of the memory pool.</li> <li>5 - Use the CHGPF CL command or the SQL ALTER statement to redistribute the file's data across more disk units.</li> </ol>

<b>CPI4331 - &amp;6 tasks used for parallel index created over file</b>	
<b>Message Text:</b>	&6 tasks used for parallel index created over file &1.

**CPI4331 - &6 tasks used for parallel index created over file**

**Cause Text:** &6 is the average numbers of tasks used for an index created over member &3 of file &1 in library &2.

If file &1 in library &2 is a logical file, then member &9 of physical file &7 in library &8 is the actual file over which the index is being built.

A file name starting with \*QUERY or \*N for the file indicates a temporary result file is being used.

The query optimizer has calculated that the optimal number of tasks is &5 which was limited for reason code &4. The definition of reason codes are:

- 1 - The \*NBRTASKS parameter value was specified for the DEGREE parameter of the CHGQRYA CL command.
- 2 - The optimizer calculated the number of tasks which would use all of the central processing units (CPU).
- 3 - The optimizer calculated the number of tasks which can efficiently run in this job's share of the memory pool.
- 4 - The optimizer calculated the number of tasks which can efficiently run using the entire memory pool.

The database manager may further limit the number of tasks used for the parallel index build if either the allocation of the file's data is not evenly distributed across disk units or the system has too few disk units.

**Recovery Text:** To disallow usage of parallel index build, specify \*NONE on the query attribute degree.

A larger number of tasks might further improve performance. The following actions based on the reason code might allow the optimizer to calculate a larger number:

- 1 - Specify a larger number of tasks value for the DEGREE parameter of the CHGQRYA CL command. Start with a value for number of tasks which is a slightly larger than &5 to see if a performance improvement is achieved.
- 2 - Simplify the query by reducing the number of fields being mapped to the result buffer or by removing expressions. Also, try specifying a number of tasks for the DEGREE parameter of the CHGQRYA CL command as described by reason code 1.
- 3 - Specify \*MAX for the query attribute degree.
- 4 - Increase the size of the memory pool.

**CPI4332 - &1 host variables used in query**

**Message Text:** &1 host variables used in query.

**Cause Text:** There were &1 host variables defined for use in the query. The values used for the host variables for this open of the query follow: &2.

The host variables values displayed above may have been special values. An explanation of the special values follow:

- DBCS data is displayed in hex format.
- \*N denotes a value of NULL.
- \*Z denotes a zero length string.
- \*L denotes a value too long to display in the replacement text.
- \*U denotes a value that could not be displayed.

<b>CPI4332 - &amp;1 host variables used in query</b>	
<b>Recovery Text:</b>	None

<b>CPI4333 - Hashing algorithm used to process join</b>	
<b>Message Text:</b>	Hashing algorithm used to process join.
<b>Cause Text:</b>	<p>The hash join method is typically used for longer running join queries. The original query will be subdivided into hash join steps.</p> <p>Each hash join step will be optimized and processed separately. Debug messages which explain the implementation of each hash join step follow this message in the joblog.</p> <p>The list below shows the names of the files or the table functions used in this query. If the entry is for a file, the format of the entry in this list is the number of the hash join step, the filename as specified in the query, the member name as specified in the query, the filename actually used in the hash join step, and the member name actually used in the hash join step. If the entry is for a table function, the format of the entry in this list is the number of the hash join step and the function name as specified in the query.</p> <p>If there are two or more files or functions listed for the same hash step, then that hash step is implemented with nested loop join.</p>
<b>Recovery Text:</b>	The hash join method is usually a good implementation choice, however, if you want to disallow the use of this method specify ALWCOPYDTA(*YES).

<b>CPI4334 - Query implemented as reusable ODP</b>	
<b>Message Text:</b>	Query implemented as reusable ODP.
<b>Cause Text:</b>	The query optimizer built the access plan for this query such that a reusable open data path (ODP) will be created. This plan will allow the query to be run repeatedly for this job without having to rebuild the ODP each time. This normally improves performance because the ODP is created only once for the job.
<b>Recovery Text:</b>	Generally, reusable ODPs perform better than non-reusable ODPs.

<b>CPI4335 - Optimizer debug messages for hash join step &amp;1 follow</b>	
<b>Message Text:</b>	Optimizer debug messages for hash join step &1 follow:
<b>Cause Text:</b>	This join query is implemented using the hash join algorithm. The optimizer debug messages that follow provide the query optimization information about hash join step &1.
<b>Recovery Text:</b>	Refer to “Data access methods” on <a href="#">page 16</a> for more information about hashing algorithm for join processing.

<b>CPI4336 - Group processing generated</b>	
<b>Message Text:</b>	Group processing generated.
<b>Cause Text:</b>	Group processing (GROUP BY) was added to the query step. Adding the group processing reduced the number of result records which should, in turn, improve the performance of subsequent steps.

<b>CPI4336 - Group processing generated</b>	
<b>Recovery Text:</b>	For more information refer to <a href="#">“Data access methods” on page 16</a>

<b>CPI4337 - Temporary hash table build for hash join step &amp;1</b>	
<b>Message Text:</b>	Temporary hash table built for hash join step &1.
<b>Cause Text:</b>	A temporary hash table was created to contain the results of hash join step &1. This process took &2 minutes and &3 seconds. The temporary hash table created contains &4 records. The total size of the temporary hash table in units of 1024 bytes is &5. A list of the fields which define the hash keys follow:
<b>Recovery Text:</b>	Refer to <a href="#">“Data access methods” on page 16</a> for more information about hashing algorithm for join processing.

<b>CPI4338 - &amp;1 Access path(s) used for bitmap processing of file &amp;2</b>	
<b>Message Text:</b>	&1 Access path(s) used for bitmap processing of file &2.
<b>Cause Text:</b>	<p>Bitmap processing was used to access records from member &amp;4 of file &amp;2 in library &amp;3.</p> <p>Bitmap processing is a method of allowing one or more access path(s) to be used to access the selected records from a file. Using bitmap processing, record selection is applied against each access path, similar to key row positioning, to create a bitmap. The bitmap has marked in it only the records of the file that are to be selected. If more than one access path is used, the resulting bitmaps are merged together using boolean logic. The resulting bitmap is then used to reduce access to just those records actually selected from the file.</p> <p>Bitmap processing is used in conjunction with the two primary access methods: arrival sequence (CPI4327 or CPI4329) or keyed access (CPI4326 or CPI4328). The message that describes the primary access method immediately precedes this message.</p> <p>When the bitmap is used with the keyed access method then it is used to further reduce the number of records selected by the primary access path before retrieving the selected records from the file.</p> <p>When the bitmap is used with arrival sequence then it allows the sequential scan of the file to skip records which are not selected by the bitmap. This is called skip sequential processing.</p> <p>The list below shows the names of the access paths used in the bitmap processing:</p> <p>&amp;8</p> <p>If file &amp;2 in library &amp;3 is a logical file then member &amp;7 of physical file &amp;5 in library &amp;6 is the actual file being accessed.</p>
<b>Recovery Text:</b>	Refer to <a href="#">“Data access methods” on page 16</a> for more information about bitmap processing.

<b>CPI433A - Unable to retrieve query options file</b>	
<b>Message Text:</b>	Unable to retrieve query options file.



<b>CPI433A - Unable to retrieve query options file</b>	
<b>Cause Text:</b>	<p>Unable to retrieve the query options from member &amp;3 in file &amp;2 in library &amp;1 for reason code &amp;4. The reason codes and their meanings follow:</p> <ul style="list-style-type: none"> <li>1 - Library &amp;1 was not found.</li> <li>2 - File &amp;2 in library &amp;1 was not found.</li> <li>3 - The file was damaged.</li> <li>4 - The file was locked by another process which prevented successful retrieval of the query options.</li> <li>5 - File &amp;2 and the internal query options structures are out of sync.</li> <li>6 - An unexpected error occurred while trying to retrieve the options file.</li> </ul> <p>The query options file is used by the Query Optimizer to determine how a query will be implemented.</p>
<b>Recovery Text:</b>	<p>Default query options will be used, unless one of the following actions are taken, based on the reason code above.</p> <ul style="list-style-type: none"> <li>1 - Either create the library (CRTLIB command) or correct the library name and then try the request again.</li> <li>2 - Either specify the library name that contains the query options file or create a duplicate object (CRTDUPOBJ command) of file &amp;2 from library QSYS into the specified library.</li> <li>4 - Wait for lock on file &amp;2 in library &amp;1 to be released and try the request again.</li> <li>3, 5, or 6 - Delete query options file &amp;2 in library &amp;1 and then duplicate it from QSYS. If the problem still persists, report the problem (ANZPRB command).</li> </ul>

<b>CPI433B - Unable to update query options file</b>	
<b>Message Text:</b>	Unable to update query options file.
<b>Cause Text:</b>	<p>An error occurred while trying to update the query options from member &amp;3, file &amp;2, library &amp;1 for reason code &amp;4. The reason codes and their meanings follow:</p> <ul style="list-style-type: none"> <li>1 - The library &amp;1 was not found.</li> <li>2 - The file &amp;2 in library &amp;1 was not found.</li> <li>3 - The parameter &amp;5 was not found.</li> <li>4 - The value &amp;6 for parameter &amp;5 was not valid.</li> <li>5 - An unexpected error occurred while trying to update the options file.</li> </ul>
<b>Recovery Text:</b>	<p>Do one of the following actions based on the reason code above.</p> <ul style="list-style-type: none"> <li>1 - Either create the library (CRTLIB) command or correct the library name and then try the request again.</li> <li>2 - Either specify the library name that contains the query options file or create duplicate object (CRTDUPOBJ) command of QAQQINI from library QSYS into the specified library.</li> <li>3 - Either specify a valid parameter or correct the parameter name and then try the request again.</li> <li>4 - Either specify a valid parameter value or correct the parameter value and then try the request again. (WRKJOB) command.</li> </ul>

<b>CPI433C - Library &amp;1 not found</b>	
<b>Message Text:</b>	Library &1 not found.
<b>Cause Text:</b>	The specified library does not exist, or the name of the library is not spelled correctly.
<b>Recovery Text:</b>	Correct the spelling of the library name, or specify the name of an existing library. Then try the request again.

<b>CPI433D - Query options used to build the query access plan</b>	
<b>Message Text:</b>	Query options used to build the query access plan.
<b>Cause Text:</b>	The access plan that was saved was created with query options retrieved from file &2 in library &1.
<b>Recovery Text:</b>	None

<b>CPI433E - User-defined function &amp;4 found in library &amp;1</b>	
<b>Message Text:</b>	User-defined function &4 found in library &1.
<b>Cause Text:</b>	Function &4 was resolved to library &1. The specific name of the function is &5.  If the function is defined to use an external program, the associated program or service program is &3 in library &2.
<b>Recovery Text:</b>	Refer to the <a href="#">SQL programming</a> topic collection, for more information on user-defined functions.

<b>CPI433F - Multiple join classes used to process join</b>	
<b>Message Text:</b>	Multiple join classes used to process join.
<b>Cause Text:</b>	Multiple join classes are used when join queries are written that have conflicting operations or cannot be implemented as a single query.  Each join class step will be optimized and processed separately. Debug messages detailing the implementation of each join class follow this message in the joblog.  The list below shows the file names of the files used in this query. The format of each entry in this list is the number of the join class step, the number of the join position in the join class step, the file name as specified in the query, the member name as specified in the query, the file name actually used in the join class step, and the member name actually used in the join class step.
<b>Recovery Text:</b>	Refer to <a href="#">“Join optimization”</a> on page 66 for more information about join classes.

<b>CPI4340 - Optimizer debug messages for join class step &amp;1 follow</b>	
<b>Message Text:</b>	Optimizer debug messages for join class step &1 follow:
<b>Cause Text:</b>	This join query is implemented using multiple join classes. The optimizer debug messages that follow provide the query optimization information about join class step &1.

**CPI4340 - Optimizer debug messages for join class step &1 follow**

<b>Recovery Text:</b>	Refer to <a href="#">“Join optimization”</a> on page 66 for more information about join classes.
-----------------------	--------------------------------------------------------------------------------------------------

**CPI4341 - Performing distributed query**

<b>Message Text:</b>	Performing distributed query.
----------------------	-------------------------------

<b>Cause Text:</b>	Query contains a distributed file. The query was processed in parallel on the following nodes: &1.
--------------------	----------------------------------------------------------------------------------------------------

<b>Recovery Text:</b>	For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> topic collection.
-----------------------	---------------------------------------------------------------------------------------------------------------------------------------------

**CPI4342 - Performing distributed join for query**

<b>Message Text:</b>	Performing distributed join for query.
----------------------	----------------------------------------

<b>Cause Text:</b>	<p>Query contains join criteria over a distributed file and a distributed join was performed, in parallel, on the following nodes: &amp;1.</p> <p>The library, file and member names of each file involved in the join follow: &amp;2.</p> <p>A file name beginning with *QQTDF indicates it is a temporary distributed result file created by the query optimizer and it will not contain an associated library or member name.</p>
--------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<b>Recovery Text:</b>	For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> .
-----------------------	-----------------------------------------------------------------------------------------------------------------------------

**CPI4343 - Optimizer debug messages for distributed query step &1 of &2 follow**

<b>Message Text:</b>	Optimizer debug messages for distributed query step &1 of &2 follow:
----------------------	----------------------------------------------------------------------

<b>Cause Text:</b>	A distributed file was specified in the query which caused the query to be processed in multiple steps. The optimizer debug messages that follow provide the query optimization information about distributed step &1 of &2 total steps.
--------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

<b>Recovery Text:</b>	For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> .
-----------------------	-----------------------------------------------------------------------------------------------------------------------------

**CPI4345 - Temporary distributed result file &3 built for query**

<b>Message Text:</b>	Temporary distributed result file &3 built for query.
----------------------	-------------------------------------------------------

**CPI4345 - Temporary distributed result file &3 built for query**

<b>Cause Text:</b>	<p>Temporary distributed result file &amp;3 was created to contain the intermediate results of the query for reason code &amp;6. The reason codes and their meanings follow:</p> <ol style="list-style-type: none"><li>1 - Data from member &amp;2 of &amp;7 &amp;8 in library &amp;1 was directed to other nodes.</li><li>2 - Data from member &amp;2 of &amp;7 &amp;8 in library &amp;1 was broadcast to all nodes.</li><li>3 - Either the query contains grouping fields (GROUP BY) that do not match the partitioning keys of the distributed file or the query contains grouping criteria but no grouping fields were specified or the query contains a subquery.</li><li>4 - Query contains join criteria over a distributed file and the query was processed in multiple steps.</li></ol> <p>A library and member name of *N indicates the data comes from a query temporary distributed file.</p> <p>File &amp;3 was built on nodes: &amp;9.</p> <p>It was built using partitioning keys: &amp;10.</p> <p>A partitioning key of *N indicates no partitioning keys were used when building the temporary distributed result file.</p>
<b>Recovery Text:</b>	<p>If the reason code is:</p> <ol style="list-style-type: none"><li>1 - Generally, a file is directed when the join fields do not match the partitioning keys of the distributed file. When a file is directed, the query is processed in multiple steps and processed in parallel. A temporary distributed result file is required to contain the intermediate results for each step.</li><li>2 - Generally, a file is broadcast when join fields do not match the partitioning keys of either file being joined or the join operator is not an equal operator. When a file is broadcast the query is processed in multiple steps and processed in parallel. A temporary distributed result file is required to contain the intermediate results for each step.</li><li>3 - Better performance may be achieved if grouping fields are specified that match the partitioning keys.</li><li>4 - Because the query is processed in multiple steps, a temporary distributed result file is required to contain the intermediate results for each step. See preceding message CPI4342 to determine which files were joined together.</li></ol> <p>For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a></p>

**CPI4346 - Optimizer debug messages for query join step &1 of &2 follow**

<b>Message Text:</b>	Optimizer debug messages for query join step &1 of &2 follow:
<b>Cause Text:</b>	Query processed in multiple steps. The optimizer debug messages that follow provide the query optimization information about join step &1 of &2 total steps.
<b>Recovery Text:</b>	No recovery necessary.

**CPI4347 - Query being processed in multiple steps**

<b>Message Text:</b>	Query being processed in multiple steps.
----------------------	------------------------------------------

<b>CPI4347 - Query being processed in multiple steps</b>	
<b>Cause Text:</b>	<p>The original query will be subdivided into multiple steps.</p> <p>Each step will be optimized and processed separately. Debug messages which explain the implementation of each step follow this message in the joblog.</p> <p>The list below shows the file names of the files used in this query. The format of each entry in this list is the number of the join step, the filename as specified in the query, the member name as specified in the query, the filename actually used in the step, and the member name actually used in the step.</p>
<b>Recovery Text:</b>	No recovery necessary.

<b>CPI4348 - The ODP associated with the cursor was hard closed</b>	
<b>Message Text:</b>	The ODP associated with the cursor was hard closed.
<b>Cause Text:</b>	<p>The Open Data Path (ODP) for this statement or cursor has been hard closed for reason code &amp;1. The reason codes and their meanings follow:</p> <ol style="list-style-type: none"> <li>1 - Either the length of the new LIKE pattern is zero and the length of the old LIKE pattern is nonzero or the length of the new LIKE pattern is nonzero and the length of the old LIKE pattern is zero.</li> <li>2 - An additional wildcard was specified in the LIKE pattern on this invocation of the cursor.</li> <li>3 - SQL indicated to the query optimizer that the cursor cannot be refreshed.</li> <li>4 - The system code could not obtain a lock on the file being queried.</li> <li>5 - The length of the host variable value is too large for the the host variable as determined by the query optimizer.</li> <li>6 - The size of the ODP to be refreshed is too large.</li> <li>7 - Refresh of the local ODP of a distributed query failed.</li> <li>8 - SQL hard closed the cursor prior to the fast path refresh code.</li> </ol>
<b>Recovery Text:</b>	In order for the cursor to be used in a reusable mode, the cursor cannot be hard closed. Look at the reason why the cursor was hard closed and take the appropriate actions to prevent a hard close from occurring.

<b>CPI4349 - Fast past refresh of the host variables values is not possible</b>	
<b>Message Text:</b>	Fast past refresh of the host variable values is not possible.

<b>CPI4349 - Fast past refresh of the host variables values is not possible</b>	
<b>Cause Text:</b>	<p>The Open Data Path (ODP) for this statement or cursor could not invoke the fast past refresh code for reason code &amp;1. The reason codes and their meanings follow:</p> <p>1 - The new host variable value is not null and old host variable value is null or the new host variable value is zero length and the old host variable value is not zero length.</p> <p>2 - The attributes of the new host variable value are not the same as the attributes of the old host variable value.</p> <p>3 - The length of the host variable value is either too long or too short. The length difference cannot be handled in the fast path refresh code.</p> <p>4 - The host variable has a data type of IGC ONLY and the the length is not even or is less than 2 bytes.</p> <p>5 - The host variable has a data type of IGC ONLY and the new host variable value does not contain an even number of bytes.</p> <p>6 - A translate table with substitution characters was used.</p> <p>7 - The host variable contains DBCS data and a CCSID translate table with substitution characters is required.</p> <p>8 - The host variable contains DBCS that is not well formed. That is, a shift-in without a shift-out or visa versa.</p> <p>9 - The host variable must be translated with a sort sequence table and the sort sequence table contains substitution characters.</p> <p>10 - The host variable contains DBCS data and must be translated with a sort sequence table that contains substitution characters.</p> <p>11 - The host variable is a Date, Time or Timestamp data type and the length of the host variable value is either too long or too short.</p>
<b>Recovery Text:</b>	Look at the reason why fast path refresh could not be used and take the appropriate actions so that fast path refresh can be used on the next invocation of this statement or cursor.

<b>CPI434 - Member &amp;3 was opened with fewer open options than were specified</b>	
<b>Message Text:</b>	Member &3 was opened with fewer open options than were specified.
<b>Cause Text:</b>	An INSTEAD OF trigger is being used for some of the open options. However there is an additional INSTEAD OF trigger on an underlying SQL view file whose trigger actions cannot be used. An open request can support INSTEAD OF triggers from only one SQL view file. The member could not be opened with the following open options: &4.
<b>Recovery Text:</b>	When adding an INSTEAD OF trigger, specify trigger actions for all of the requested open options.

<b>CPI434E - Query could not be run using SQE</b>	
<b>Message Text:</b>	Query could not be run using SQE.
<b>Cause Text:</b>	<p>The query was run using CQE (Current Query Engine). The query could not be run using SQE (SQL Query Engine) for reason code &amp;1. The reason codes and their meanings follow:</p> <p>1 -- Sort sequence table &amp;2 in library &amp;3 is an ICU (International Components of Unicode) sort sequence table that is not supported by SQE.</p>

<b>CPI434E - Query could not be run using SQE</b>	
---------------------------------------------------	--

<b>Recovery Text:</b>	Recovery for reason code 1: To run the query using SQE, specify a version of the ICU sort sequence table that is &4 or later.
-----------------------	-------------------------------------------------------------------------------------------------------------------------------

<b>CPI4350 - Materialized query tables were considered for optimization</b>	
-----------------------------------------------------------------------------	--

<b>Message Text:</b>	Materialized query tables were considered for optimization.
----------------------	-------------------------------------------------------------

**CPI4350 - Materialized query tables were considered for optimization****Cause Text:**

The query optimizer considered usage of materialized query tables for this query.

Following each materialized query table name in the list is a reason code which explains why the materialized query table was not used. A reason code of 0 indicates that the materialized query table was used to implement the query.

The reason codes and their meanings follow:

- 1 - The cost to use the materialized query table, as determined by the optimizer, was higher than the cost associated with the chosen implementation.
- 2 - The join specified in the materialized query was not compatible with the query.
- 3 - The materialized query table had predicates that were not matched in the query.
- 4 - The grouping or distinct specified in the materialized query table is not compatible with the grouping or distinct specified in the query.
- 5 - The query specified columns that were not in the select-list of the materialized query table.
- 6 - The materialized query table query contains functionality that is not supported by the query optimizer.
- 7 - The materialized query table specified the DISABLE QUERY OPTIMIZATION clause.
- 8 - The ordering specified in the materialized query table is not compatible with the ordering specified in the query.
- 9 - The query contains functionality that is not supported by the materialized query table matching algorithm.
- 10 - Materialized query tables may not be used for this query.
- 11 - The refresh age of this materialized query table exceeds the duration specified by the MATERIALIZED\_QUERY\_TABLE\_REFRESH\_AGE QAQQINI option.
- 12 - The commit level of the materialized query table is lower than the commit level specified for the query.
- 14 - The FETCH FOR FIRST n ROWS clause of the materialized query table is not compatible with the query.
- 15 - The QAQQINI options used to create the materialized query table are not compatible with the QAQQINI options used to run this query.
- 16 - The materialized query table is not usable.
- 17 - The UNION specified in the materialized query table is not compatible with the query.
- 18 - The constants specified in the materialized query table are not compatible with host variable values specified in the query.
- 19 - The materialized query table is in Check Pending status and cannot be used.
- 20 - The UDTF specified in the materialized query table is not compatible with UDTF in the query.
- 21 - The Values clause specified in the materialized query table is not compatible with Values specified in the query.

**Recovery Text:**

The user may want to delete any materialized query tables that are no longer needed.



<b>CPI4351 - Additional reason codes for query access plan has been rebuilt</b>	
<b>Message Text:</b>	Additional reason codes for query access plan has been rebuilt.
<b>Cause Text:</b>	<p>Message CPI4323 was issued immediately before this message. Because of message length restrictions, some of the reason codes used by message CPI4323 are explained below rather than in that message. The CPI4323 message was issued for reason code &amp;13. The additional reason codes and their meaning follow:</p> <p>20 - Referential or check constraints for member &amp;19 of file &amp;17 in library &amp;18 have changed since the access plan was generated.</p> <p>21 - Materialized query tables for member &amp;22 of file &amp;20 in library &amp;21 have changed since the access plan was generated. If the file is *N then the file name is not available.</p> <p>22 - The value of a host variable changed and the access plan is no longer valid.</p> <p>23 - Adaptive Query Processing (AQP) determined that a new access plan is needed.</p>
<b>Recovery Text:</b>	See the prior message CPI4323 for more information.

<b>CPI436A - Database monitor started for job &amp;1, monitor ID &amp;2</b>	
<b>Message Text:</b>	Database monitor started for job &1, monitor ID &2.
<b>Cause Text:</b>	<p>The database monitor was started for job &amp;1. The system generated monitor ID for this database monitor is &amp;2.</p> <p>If multiple monitors have been started using the same generic job name, the monitor ID is needed to uniquely identify which monitor is to be ended with the ENDDBMON command.</p>
<b>Recovery Text:</b>	If multiple monitors have been started using the same generic job name, remember the monitor ID. The monitor ID will be required when using the ENDDBMON command to end this specific monitor.

## Query optimization performance information messages and open data paths

Several of the following SQL runtime messages refer to open data paths.

An open data path (ODP) definition is an internal object that is created when a cursor is opened or when other SQL statements are run. It provides a direct link to the data so that I/O operations can occur. ODPs are used on OPEN, INSERT, UPDATE, DELETE, and SELECT INTO statements to perform their respective operations on the data.

Even though SQL cursors are closed and SQL statements have run, in many cases, the database manager saves the associated ODPs of the SQL operations. These ODPs are then reused the next time the statement is run. For example, an SQL CLOSE statement could close the SQL cursor, but leave the ODP available to use again the next time the cursor is opened. This technique can significantly reduce the processing and response time in running SQL statements.

The ability to reuse ODPs when SQL statements are run repeatedly is an important consideration in achieving faster performance.

<b>SQL7910 - All SQL cursors closed</b>	
<b>Message Text:</b>	SQL cursors closed.

<b>SQL7910 - All SQL cursors closed</b>	
<b>Cause Text:</b>	SQL cursors have been closed and all Open Data Paths (ODPs) have been deleted, except those that were opened by programs with the CLOSQLCSR(*ENDJOB) option or were opened by modules with the CLOSQLCSR(*ENDACTGRP) option. All SQL programs on the call stack have completed, and the SQL environment has been exited. This process includes the closing of cursors, the deletion of ODPs, the removal of prepared statements, and the release of locks.
<b>Recovery Text:</b>	To keep cursors, ODPs, prepared statements, and locks available after the completion of a program, use the CLOSQLCSR precompile parameter. -- The *ENDJOB option will allow the user to keep the SQL resources active for the duration of the job. -- The *ENDSQL option will allow the user to keep SQL resources active across program calls, provided the SQL environment stays resident. Running an SQL statement in the first program of an application will keep the SQL environment active for the duration of that application. -- The *ENDPGM option, which is the default for non-Integrated Language Environment® (ILE) programs, causes all SQL resources to only be accessible by the same invocation of a program. Once an *ENDPGM program has completed, if it is called again, the SQL resources are no longer active. -- The *ENDMOD option causes all SQL resources to only be accessible by the same invocation of the module. -- The *ENDACTGRP option, which is the default for ILE modules, will allow the user to keep the SQL resources active for the duration of the activation group.

<b>SQL7911 - ODP reused</b>	
<b>Message Text:</b>	ODP reused.
<b>Cause Text:</b>	An ODP that was previously created has been reused. There was a reusable Open Data Path (ODP) found for this SQL statement, and it has been used. The reusable ODP may have been from the same call to a program or a previous call to the program. A reuse of an ODP will not generate an OPEN entry in the journal.
<b>Recovery Text:</b>	None

<b>SQL7912 - ODP created</b>	
<b>Message Text:</b>	ODP created.
<b>Cause Text:</b>	An Open Data Path (ODP) has been created. No reusable ODP could be found. This occurs in the following cases: -- This is the first time the statement has been run. -- A RCLRSC has been issued since the last run of this statement. -- The last run of the statement caused the ODP to be deleted. -- If this is an OPEN statement, the last CLOSE of this cursor caused the ODP to be deleted. -- The Application Server (AS) has been changed by a CONNECT statement.

<b>SQL7912 - ODP created</b>	
<b>Recovery Text:</b>	If a cursor is being opened many times in an application, it is more efficient to use a reusable ODP, and not create an ODP every time. This also applies to repeated runs of INSERT, UPDATE, DELETE, and SELECT INTO statements. If ODPs are being created on every open, see the close message to determine why the ODP is being deleted.

The first time that the statement is run or the cursor is opened for a process, an ODP must always be created. However, if this message appears on every statement run or cursor open, use the tips recommended in “Retaining cursor positions for non-ILE program calls” on page 297 in your application.

<b>SQL7913 - ODP deleted</b>	
<b>Message Text:</b>	ODP deleted.
<b>Cause Text:</b>	The Open Data Path (ODP) for this statement or cursor has been deleted. The ODP was not reusable. This could be caused by using a host variable in a LIKE clause, ordering on a host variable, or because the query optimizer chose to accomplish the query with an ODP that was not reusable.
<b>Recovery Text:</b>	See previous query optimizer messages to determine how the cursor was opened.

<b>SQL7914 - ODP not deleted</b>	
<b>Message Text:</b>	ODP not deleted.
<b>Cause Text:</b>	The Open Data Path (ODP) for this statement or cursor has not been deleted. This ODP can be reused on a subsequent run of the statement. This will not generate an entry in the journal.
<b>Recovery Text:</b>	None

<b>SQL7915 - Access plan for SQL statement has been built</b>	
<b>Message Text:</b>	Access plan for SQL statement has been built.
<b>Cause Text:</b>	SQL had to build the access plan for this statement at run time. This occurs in the following cases: -- The program has been restored from a different release and this is the first time this statement has been run. -- All the files required for the statement did not exist at precompile time, and this is the first time this statement has been run. -- The program was precompiled using SQL naming mode, and the program owner has changed since the last time the program was called.
<b>Recovery Text:</b>	This is normal processing for SQL. Once the access plan is built, it will be used on subsequent runs of the statement.

<b>SQL7916 - Blocking used for query</b>	
<b>Message Text:</b>	Blocking used for query.

**SQL7916 - Blocking used for query**

<b>Cause Text:</b>	Blocking has been used in the implementation of this query. SQL will retrieve a block of records from the database manager on the first FETCH statement. Additional FETCH statements have to be issued by the calling program, but they do not require SQL to request more records, and therefore will run faster.
<b>Recovery Text:</b>	SQL attempts to utilize blocking whenever possible. In cases where the cursor is not update capable, and commitment control is not active, there is a possibility that blocking will be used.

**SQL7917 - Access plan not updated**

<b>Message Text:</b>	Access plan not updated.
<b>Cause Text:</b>	The query optimizer rebuilt the access plan for this statement, but the program could not be updated. Another job may be running the program. The program cannot be updated with the new access plan until a job can obtain an exclusive lock on the program. The exclusive lock cannot be obtained if another job is running the program, if the job does not have proper authority to the program, or if the program is currently being saved. The query will still run, but access plan rebuilds will continue to occur until the program is updated.
<b>Recovery Text:</b>	See previous messages from the query optimizer to determine why the access plan has been rebuilt. To ensure that the program gets updated with the new access plan, run the program when no other active jobs are using it.

**SQL7918 - Reusable ODP deleted**

<b>Message Text:</b>	Reusable ODP deleted. Reason code &1.
----------------------	---------------------------------------

**SQL7918 - Reusable ODP deleted**

<b>Cause Text:</b>	<p>An existing Open Data Path (ODP) was found for this statement, but it could not be reused for reason &amp;1. The statement now refers to different files or uses different override options than are in the ODP. Reason codes and their meanings are:</p> <ul style="list-style-type: none"><li>1 -- Commitment control isolation level is not compatible.</li><li>2 -- The statement contains SQL special register USER, CURRENT DEBUG MODE, CURRENT DECFLOAT ROUNDING MODE, or CURRENT TIMEZONE, and the value for one of these registers has changed.</li><li>3 -- The PATH used to locate an SQL function has changed.</li><li>4 -- The job default CCSID has changed.</li><li>5 -- The library list has changed, such that a file is found in a different library. This only affects statements with unqualified table names, when the table exists in multiple libraries.</li><li>6 -- The file, library, or member for the original ODP was changed with an override.</li><li>7 -- An OVRDBF or DLTOVR command has been issued. A file referred to in the statement now refers to a different file, library, or member.</li><li>8 -- An OVRDBF or DLTOVR command has been issued, causing different override options, such as different SEQONLY or WAITRCD values.</li><li>9 -- An error occurred when attempting to verify the statement override information is compatible with the reusable ODP information.</li><li>10 -- The query optimizer has determined the ODP cannot be reused.</li><li>11 -- The client application requested not to reuse ODPs.</li></ul>
<b>Recovery Text:</b>	Do not change the library list, the override environment, or the values of the special registers if reusable ODPs are to be used.

**SQL7919 - Data conversion required on FETCH or embedded SELECT**

<b>Message Text:</b>	Data conversion required on FETCH or embedded SELECT.
----------------------	-------------------------------------------------------

**SQL7919 - Data conversion required on FETCH or embedded SELECT****Cause Text:**

Host variable &2 requires conversion. The data retrieved for the FETCH or embedded SELECT statement can not be directly moved to the host variables. The statement ran correctly. Performance, however, would be improved if no data conversion was required. The host variable requires conversion for reason &1.

-- Reason 1 - host variable &2 is a character or graphic string of a different length than the value being retrieved.

-- Reason 2 - host variable &2 is a numeric type that is different than the type of the value being retrieved.

-- Reason 3 - host variable &2 is a C character or C graphic string that is NUL-terminated, the program was compiled with option \*CNULRQD specified, and the statement is a multiple-row FETCH.

-- Reason 4 - host variable &2 is a variable length string and the value being retrieved is not.

-- Reason 5 - host variable &2 is not a variable length string and the value being retrieved is.

-- Reason 6 - host variable &2 is a variable length string whose maximum length is different than the maximum length of the variable length value being retrieved.

-- Reason 7 - a data conversion was required on the mapping of the value being retrieved to host variable &2, such as a CCSID conversion.

-- Reason 8 - a DRDA connection was used to get the value being retrieved into host variable &2. The value being retrieved is either null capable or varying-length, is contained in a partial row, or is a derived expression.

-- Reason 10 - the length of host variable &2 is too short to hold a TIME or TIMESTAMP value being retrieved.

-- Reason 11 - host variable &2 is of type DATE, TIME or TIMESTAMP, and the value being retrieved is a character string.

-- Reason 12 - too many host variables were specified and records are blocked. Host variable &2 does not have a corresponding column returned from the query.

-- Reason 13 - a DRDA connection was used for a blocked FETCH and the number of host variables specified in the INTO clause is less than the number of result values in the select list.

-- Reason 14 - a LOB Locator was used and the commitment control level of the process was not \*ALL.

**Recovery Text:**

To get better performance, attempt to use host variables of the same type and length as their corresponding result columns.

**SQL7939 - Data conversion required on INSERT or UPDATE****Message Text:**

Data conversion required on INSERT or UPDATE.

<b>SQL7939 - Data conversion required on INSERT or UPDATE</b>	
<b>Cause Text:</b>	<p>The INSERT or UPDATE values can not be directly moved to the columns because the data type or length of a value is different than one of the columns. The INSERT or UPDATE statement ran correctly. Performance, however, would be improved if no data conversion was required. The reason data conversion is required is &amp;1.</p> <p>-- Reason 1 is that the INSERT or UPDATE value is a character or graphic string of a different length than column &amp;2.</p> <p>-- Reason 2 is that the INSERT or UPDATE value is a numeric type that is different than the type of column &amp;2.</p> <p>-- Reason 3 is that the INSERT or UPDATE value is a variable length string and column &amp;2 is not.</p> <p>-- Reason 4 is that the INSERT or UPDATE value is not a variable length string and column &amp;2 is.</p> <p>-- Reason 5 is that the INSERT or UPDATE value is a variable length string whose maximum length is different than the maximum length of column &amp;2.</p> <p>-- Reason 6 is that a data conversion was required on the mapping of the INSERT or UPDATE value to column &amp;2, such as a CCSID conversion.</p> <p>-- Reason 7 is that the INSERT or UPDATE value is a character string and column &amp;2 is of type DATE, TIME, or TIMESTAMP.</p> <p>-- Reason 8 is that the target table of the INSERT is not a SQL table.</p>
<b>Recovery Text:</b>	To get better performance, try to use values of the same type and length as their corresponding columns.

## PRTSQLINF message reference

The following messages are returned from **PRTSQLINF**.

<b>SQL400A - Temporary distributed result file &amp;1 was created to contain join result</b>	
<b>Message Text:</b>	Temporary distributed result file &1 was created to contain join result. Result file was directed.
<b>Cause Text:</b>	Query contains join criteria over a distributed file and a distributed join was performed in parallel. A temporary distributed result file was created to contain the results of the distributed join.
<b>Recovery Text:</b>	For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL400B - Temporary distributed result file &amp;1 was created to contain join result</b>	
<b>Message Text:</b>	Temporary distributed result file &1 was created to contain join result. Result file was broadcast.
<b>Cause Text:</b>	Query contains join criteria over a distributed file and a distributed join was performed in parallel. A temporary distributed result file was created to contain the results of the distributed join.
<b>Recovery Text:</b>	For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL400C - Optimizer debug messages for distributed query step &amp;1 and &amp;2 follow</b>	
<b>Message Text:</b>	Optimizer debug messages for distributed query step &1 of &2 follow:
<b>Cause Text:</b>	A distributed file was specified in the query which caused the query to be processed in multiple steps. The optimizer debug messages that follow provide the query optimization information about the current step.
<b>Recovery Text:</b>	For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL400D - GROUP BY processing generated</b>	
<b>Message Text:</b>	GROUP BY processing generated.
<b>Cause Text:</b>	GROUP BY processing was added to the query step. Adding the GROUP BY reduced the number of result rows which should, in turn, improve the performance of subsequent steps.
<b>Recovery Text:</b>	For more information refer to the <a href="#">SQL programming</a> topic collection.

<b>SQL400E - Temporary distributed result file &amp;1 was created while processing distributed subquery</b>	
<b>Message Text:</b>	Temporary distributed result file &1 was created while processing distributed subquery.
<b>Cause Text:</b>	A temporary distributed result file was created to contain the intermediate results of the query. The query contains a subquery which requires an intermediate result.
<b>Recovery Text:</b>	Generally, if the fields correlated between the query and subquery do not match the partition keys of the respective files, the query must be processed in multiple steps and a temporary distributed file will be built to contain the intermediate results. For more information about processing of distributed files, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL4001 - Temporary result created</b>	
<b>Message Text:</b>	Temporary result created.
<b>Cause Text:</b>	<p>Conditions exist in the query which cause a temporary result to be created. One of the following reasons may be the cause for the temporary result:</p> <ul style="list-style-type: none"> <li>-- The table is a join logical file and its join type (JDFTVAL) does not match the join-type specified in the query.</li> <li>-- The format specified for the logical file refers to more than one physical table.</li> <li>-- The table is a complex SQL view requiring a temporary table to contain the results of the SQL view.</li> <li>-- The query contains grouping columns (GROUP BY) from more than one table, or contains grouping columns from a secondary table of a join query that cannot be reordered.</li> </ul>
<b>Recovery Text:</b>	Performance may be improved if the query can be changed to avoid temporary results.



<b>SQL4002 - Reusable ODP sort used</b>	
<b>Message Text:</b>	Reusable ODP sort used.
<b>Cause Text:</b>	<p>Conditions exist in the query which cause a sort to be used. This allowed the open data path (ODP) to be reusable. One of the following reasons may be the cause for the sort:</p> <ul style="list-style-type: none"> <li>-- The query contains ordering columns (ORDER BY) from more than one table, or contains ordering columns from a secondary table of a join query that cannot be reordered.</li> <li>-- The grouping and ordering columns are not compatible.</li> <li>-- DISTINCT was specified for the query.</li> <li>-- UNION was specified for the query.</li> <li>-- The query had to be implemented using a sort. Key length of more than 2000 bytes, more than 120 ordering columns, or an ordering column containing a reference to an external user-defined function was specified for ordering.</li> <li>-- The query optimizer chose to use a sort rather than an index to order the results of the query.</li> </ul>
<b>Recovery Text:</b>	A reusable ODP generally results in improved performance when compared to a non-reusable ODP.

<b>SQL4003 - UNION</b>	
<b>Message Text:</b>	UNION, EXCEPT, or INTERSECT.
<b>Cause Text:</b>	A UNION, EXCEPT, or INTERSECT operator was specified in the query. The messages preceding this keyword delimiter correspond to the subselect preceding the UNION, EXCEPT, or INTERSECT operator. The messages following this keyword delimiter correspond to the subselect following the UNION, EXCEPT, or INTERSECT operator.
<b>Recovery Text:</b>	None

<b>SQL4004 - SUBQUERY</b>	
<b>Message Text:</b>	SUBQUERY.
<b>Cause Text:</b>	The SQL statement contains a subquery. The messages preceding the SUBQUERY delimiter correspond to the subselect containing the subquery. The messages following the SUBQUERY delimiter correspond to the subquery.
<b>Recovery Text:</b>	None

<b>SQL4005 - Query optimizer timed out for table &amp;1</b>	
<b>Message Text:</b>	Query optimizer timed out for table &1.
<b>Cause Text:</b>	The query optimizer timed out before it could consider all indexes built over the table. This is not an error condition. The query optimizer may time out in order to minimize optimization time. The query can be run in debug mode (STRDBG) to see the list of indexes which were considered during optimization. The table number refers to the relative position of this table in the query.

<b>SQL4005 - Query optimizer timed out for table &amp;1</b>	
<b>Recovery Text:</b>	To ensure an index is considered for optimization, specify the logical file of the index as the table to be queried. The optimizer will first consider the index of the logical file specified on the SQL select statement. Note that SQL created indexes cannot be queried. An SQL index can be deleted and recreated to increase the chances it will be considered during query optimization. Consider deleting any indexes no longer needed.

<b>SQL4006 - All indexes considered for table &amp;1</b>	
<b>Message Text:</b>	All indexes considered for table &1.
<b>Cause Text:</b>	The query optimizer considered all index built over the table when optimizing the query. The query can be run in debug mode (STRDBG) to see the list of indexes which were considered during optimization. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	None

<b>SQL4007 - Query implementation for join position &amp;1 table &amp;2</b>	
<b>Message Text:</b>	Query implementation for join position &1 table &2.
<b>Cause Text:</b>	The join position identifies the order in which the tables are joined. A join position of 1 indicates this table is the first, or left-most, table in the join order. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	Join order can be influenced by adding an ORDER BY clause to the query. Refer to “Join optimization” on page 66 for more information about join optimization and tips to influence join order.

<b>SQL4008 - Index &amp;1 used for table &amp;2</b>	
<b>Message Text:</b>	Index &1 used for table &2.
<b>Cause Text:</b>	The index was used to access rows from the table for one of the following reasons: -- Row selection. -- Join criteria. -- Ordering/grouping criteria. -- Row selection and ordering/grouping criteria. The table number refers to the relative position of this table in the query. The query can be run in debug mode (STRDBG) to determine the specific reason the index was used.
<b>Recovery Text:</b>	None

<b>SQL4009 - Index created for table &amp;1</b>	
<b>Message Text:</b>	Index created for table &1.

<b>SQL4009 - Index created for table &amp;1</b>	
<b>Cause Text:</b>	A temporary index was built to access rows from the table for one of the following reasons:  -- Perform specified ordering/grouping criteria.  -- Perform specified join criteria.  The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	To improve performance, consider creating a permanent index if the query is run frequently. The query can be run in debug mode (STRDBG) to determine the specific reason the index was created and the key columns used when creating the index. NOTE: If permanent index is created, it is possible the query optimizer may still choose to create a temporary index to access the rows from the table.

<b>SQL401A - Processing grouping criteria for query containing a distributed table</b>	
<b>Message Text:</b>	Processing grouping criteria for query containing a distributed table.
<b>Cause Text:</b>	Grouping for queries that contain distributed tables can be implemented using either a one or two step method. If the one step method is used, the grouping columns (GROUP BY) match the partitioning keys of the distributed table. If the two step method is used, the grouping columns do not match the partitioning keys of the distributed table or the query contains grouping criteria but no grouping columns were specified. If the two step method is used, message SQL401B will appear followed by another SQL401A message.
<b>Recovery Text:</b>	For more information about processing of distributed tables, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL401B - Temporary distributed result table &amp;1 was created while processing grouping criteria</b>	
<b>Message Text:</b>	Temporary distributed result table &1 was created while processing grouping criteria.
<b>Cause Text:</b>	A temporary distributed result table was created to contain the intermediate results of the query. Either the query contains grouping columns (GROUP BY) that do not match the partitioning keys of the distributed table or the query contains grouping criteria but no grouping columns were specified.
<b>Recovery Text:</b>	For more information about processing of distributed tables, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL401C - Performing distributed join for query</b>	
<b>Message Text:</b>	Performing distributed join for query.
<b>Cause Text:</b>	Query contains join criteria over a distributed table and a distributed join was performed in parallel. See the following SQL401F messages to determine which tables were joined together.
<b>Recovery Text:</b>	For more information about processing of distributed tables, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL401D - Temporary distributed result table &amp;1 was created because table &amp;2 was directed</b>	
<b>Message Text:</b>	Temporary distributed result table &1 was created because table &2 was directed.

<b>SQL401D - Temporary distributed result table &amp;1 was created because table &amp;2 was directed</b>	
<b>Cause Text:</b>	Temporary distributed result table was created to contain the intermediate results of the query. Data from a distributed table in the query was directed to other nodes.
<b>Recovery Text:</b>	Generally, a table is directed when the join columns do not match the partitioning keys of the distributed table. When a table is directed, the query is processed in multiple steps and processed in parallel. A temporary distributed result file is required to contain the intermediate results for each step. For more information about processing of distributed tables, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL401E - Temporary distributed result table &amp;1 was created because table &amp;2 was broadcast</b>	
<b>Message Text:</b>	Temporary distributed result table &1 was created because table &2 was broadcast.
<b>Cause Text:</b>	Temporary distributed result table was created to contain the intermediate results of the query. Data from a distributed table in the query was broadcast to all nodes.
<b>Recovery Text:</b>	Generally, a table is broadcast when join columns do not match the partitioning keys of either table being joined or the join operator is not an equal operator. When a table is broadcast the query is processed in multiple steps and processed in parallel. A temporary distributed result table is required to contain the intermediate results for each step. For more information about processing of distributed tables, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL401F - Table &amp;1 used in distributed join</b>	
<b>Message Text:</b>	Table &1 used in distributed join.
<b>Cause Text:</b>	Query contains join criteria over a distributed table and a distributed join was performed in parallel.
<b>Recovery Text:</b>	For more information about processing of distributed tables, refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL4010 - Table scan access for table &amp;1</b>	
<b>Message Text:</b>	Table scan access for table &1.
<b>Cause Text:</b>	Table scan access was used to select rows from the table. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	Table scan is generally a good performing option when selecting a high percentage of rows from the table. The use of an index, however, may improve the performance of the query when selecting a low percentage of rows from the table.

<b>SQL4011 - Index scan-key row positioning used on table &amp;1</b>	
<b>Message Text:</b>	Index scan-key row positioning used on table &1.
<b>Cause Text:</b>	Index scan-key row positioning is defined as applying selection against the index to position directly to ranges of keys that match some or all of the selection criteria. Index scan-key row positioning only processes a subset of the keys in the index and is a good performing option when selecting a small percentage of rows from the table.  The table number refers to the relative position of this table in the query.

<b>SQL4011 - Index scan-key row positioning used on table &amp;1</b>	
<b>Recovery Text:</b>	Refer to “Data access methods” on page 16 for more information about index scan-key row positioning.

<b>SQL4012 - Index created from index &amp;1 for table &amp;2</b>	
<b>Message Text:</b>	Index created from index &1 for table &2.
<b>Cause Text:</b>	A temporary index was created using the specified index to access rows from the queried table for one of the following reasons: -- Perform specified ordering/grouping criteria. -- Perform specified join criteria. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	Creating an index from an index is generally a good performing option. Consider creating a permanent index for frequently run queries. The query can be run in debug mode (STRDBG) to determine the key columns used when creating the index. NOTE: If a permanent index is created, it is possible the query optimizer may still choose to create a temporary index to access the rows from the table.

<b>SQL4013 - Access plan has not been built</b>	
<b>Message Text:</b>	Access plan has not been built.
<b>Cause Text:</b>	An access plan was not created for this query. Possible reasons may include: -- Tables were not found when the program was created. -- The query was complex and required a temporary result table. -- Dynamic SQL was specified.
<b>Recovery Text:</b>	If an access plan was not created, review the possible causes. Attempt to correct the problem if possible.

<b>SQL4014 - &amp;1 join column pair(s) are used for this join position</b>	
<b>Message Text:</b>	&1 join column pair(s) are used for this join position.
<b>Cause Text:</b>	The query optimizer may choose to process join predicates as either join selection or row selection. The join predicates used in join selection are determined by the final join order and the index used. This message indicates how many join column pairs were processed as join selection at this join position. Message SQL4015 provides detail on which columns comprise the join column pairs. If 0 join column pairs were specified then index scan-key row positioning with row selection was used instead of join selection.
<b>Recovery Text:</b>	If fewer join pairs are used at a join position than expected, it is possible no index exists which has keys matching the desired join columns. Try creating an index whose keys match the join predicates. If 0 join column pairs were specified then index scan-key row positioning was used. Index scan-key row positioning is normally a good performing option. Message SQL4011 provides more information on index scan-key row positioning.

<b>SQL4015 - From-column &amp;1.&amp;2, to-column &amp;3.&amp;4, join operator &amp;5, join predicate &amp;6</b>	
<b>Message Text:</b>	From-column &1.&2, to-column &3.&4, join operator &5, join predicate &6.
<b>Cause Text:</b>	<p>Identifies which join predicate was implemented at the current join position. The replacement text parameters are:</p> <p>-- &amp;1: The join 'from table' number. The table number refers to the relative position of this table in the query.</p> <p>-- &amp;2: The join 'from column' name. The column within the join from table which comprises the left half of the join column pair. If the column name is *MAP, the column is an expression (derived field).</p> <p>-- &amp;3: The join 'to table' number. The table number refers to the relative position of this table in the query.</p> <p>-- &amp;4: The join 'to column' name. The column within the join to column which comprises the right half of the join column pair. If the column name is *MAP, the column is an expression (derived field).</p> <p>-- &amp;5: The join operator. Possible values are EQ (equal), NE (not equal), GT (greater than), LT (less than), GE (greater than or equal), LE (less than or equal), and CP (cross join or cartesian product).</p> <p>-- &amp;6: The join predicate number. Identifies the join predicate within this set of join pairs.</p>
<b>Recovery Text:</b>	Refer to <a href="#">"Join optimization"</a> on page 66 for more information about joins.

<b>SQL4016 - Subselects processed as join query</b>	
<b>Message Text:</b>	Subselects processed as join query.
<b>Cause Text:</b>	The query optimizer chose to implement some or all of the subselects with a join query. Implementing subqueries with a join generally improves performance over implementing alternative methods.
<b>Recovery Text:</b>	None

<b>SQL4017 - Host variables implemented as reusable ODP</b>	
<b>Message Text:</b>	Host variables implemented as reusable ODP.
<b>Cause Text:</b>	The query optimizer has built the access plan allowing for the values of the host variables to be supplied when the query is opened. This query can be run with different values being provided for the host variables without requiring the access plan to be rebuilt. This is the normal method of handling host variables in access plans. The open data path (ODP) that will be created from this access plan will be a reusable ODP.
<b>Recovery Text:</b>	Generally, reusable open data paths perform better than non-reusable open data paths.

<b>SQL4018 - Host variables implemented as non-reusable ODP</b>	
<b>Message Text:</b>	Host variables implemented as non-reusable ODP.

<b>SQL4018 - Host variables implemented as non-reusable ODP</b>	
<b>Cause Text:</b>	The query optimizer has implemented the host variables with a non-reusable open data path (ODP).
<b>Recovery Text:</b>	This can be a good performing option in special circumstances, but generally a reusable ODP gives the best performance.

<b>SQL4019 - Host variables implemented as file management row positioning reusable ODP</b>	
<b>Message Text:</b>	Host variables implemented as file management row positioning reusable ODP.
<b>Cause Text:</b>	The query optimizer has implemented the host variables with a reusable open data path (ODP) using file management row positioning.
<b>Recovery Text:</b>	Generally, a reusable ODP performs better than a non-reusable ODP.

<b>SQL402A - Hashing algorithm used to process join</b>	
<b>Message Text:</b>	Hashing algorithm used to process join.
<b>Cause Text:</b>	The hash join algorithm is typically used for longer running join queries. The original query will be subdivided into hash join steps. Each hash join step will be optimized and processed separately. Access plan implementation information for each of the hash join steps is not available because access plans are not saved for the individual hash join dials. Debug messages detailing the implementation of each hash dial can be found in the joblog if the query is run in debug mode using the STRDBG CL command.
<b>Recovery Text:</b>	The hash join method is usually a good implementation choice, however, if you want to disallow the use of this method specify ALWCPYDTA(*YES). Refer to the &qryopt. for more information on hashing algorithm for join processing.

<b>SQL402B - Table &amp;1 used in hash join step &amp;2</b>	
<b>Message Text:</b>	Table &1 used in hash join step &2.
<b>Cause Text:</b>	This message lists the table number used by the hash join steps. The table number refers to the relative position of this table in the query. If there are two or more of these messages for the same hash join step, then that step is a nested loop join. Access plan implementation information for each of the hash join step are not available because access plans are not saved for the individual hash steps. Debug messages detailing the implementation of each hash step can be found in the joblog if the query is run in debug mode using the STRDBG CL command.
<b>Recovery Text:</b>	Refer to <a href="#">“Data access methods”</a> on page 16 for more information about hashing.

<b>SQL402C - Temporary table created for hash join results</b>	
<b>Message Text:</b>	Temporary table created for hash join results.
<b>Cause Text:</b>	The results of the hash join were written to a temporary table so that query processing could be completed. The temporary table was required because the query contained one or more of the following: GROUP BY or summary functions ORDER BY DISTINCT Expression containing columns from more than one table Complex row selection involving columns from more than one table

**SQL402C - Temporary table created for hash join results**

<b>Recovery Text:</b>	Refer to <a href="#">“Data access methods”</a> on page 16 for more information about the hashing algorithm for join processing.
-----------------------	---------------------------------------------------------------------------------------------------------------------------------

**SQL402D - Query attributes overridden from query options file &2 in library &1**

<b>Message Text:</b>	Query attributes overridden from query options file &2 in library &1.
----------------------	-----------------------------------------------------------------------

<b>Cause Text:</b>	None
--------------------	------

<b>Recovery Text:</b>	None
-----------------------	------

**SQL4020 - Estimated query run time is &1 seconds**

<b>Message Text:</b>	Estimated query run time is &1 seconds.
----------------------	-----------------------------------------

<b>Cause Text:</b>	The total estimated time, in seconds, of executing this query.
--------------------	----------------------------------------------------------------

<b>Recovery Text:</b>	None
-----------------------	------

**SQL4021 - Access plan last saved on &1 at &2**

<b>Message Text:</b>	Access plan last saved on &1 at &2.
----------------------	-------------------------------------

<b>Cause Text:</b>	The date and time reflect the last time the access plan was successfully updated in the program object.
--------------------	---------------------------------------------------------------------------------------------------------

<b>Recovery Text:</b>	None
-----------------------	------

**SQL4022 - Access plan was saved with SRVQRY attributes active**

<b>Message Text:</b>	Access plan was saved with SRVQRY attributes active.
----------------------	------------------------------------------------------

<b>Cause Text:</b>	The access plan that was saved was created while SRVQRY was active. Attributes saved in the access plan may be the result of SRVQRY.
--------------------	--------------------------------------------------------------------------------------------------------------------------------------

<b>Recovery Text:</b>	The query will be re-optimized the next time it is run so that SRVQRY attributes will not be permanently saved.
-----------------------	-----------------------------------------------------------------------------------------------------------------

**SQL4023 - Parallel table prefetch used**

<b>Message Text:</b>	Parallel table prefetch used.
----------------------	-------------------------------

<b>Cause Text:</b>	The query optimizer chose to use a parallel prefetch access method to reduce the processing time required for the table scan.
--------------------	-------------------------------------------------------------------------------------------------------------------------------



### SQL4023 - Parallel table prefetch used

<b>Recovery Text:</b>	<p>Parallel prefetch can improve the performance of queries. Even though the access plan was created to use parallel prefetch, the system will actually run the query only if the following are true:</p> <ul style="list-style-type: none"><li>-- The query attribute degree was specified with an option of *IO or *ANY for the application process.</li><li>-- There is enough main storage available to cache the data being retrieved by multiple I/O streams. Normally, 5 megabytes would be a minimum. Increasing the size of the shared pool may improve performance.</li></ul> <p>For more information about parallel table prefetch, refer to <a href="#">“Data access methods” on page 16.</a></p>
-----------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

### SQL4024 - Parallel index preload access method used

<b>Message Text:</b>	Parallel index preload access method used.
<b>Cause Text:</b>	The query optimizer chose to use a parallel index preload access method to reduce the processing time required for this query. This means that the indexes used by this query will be loaded into active memory when the query is opened.
<b>Recovery Text:</b>	<p>Parallel index preload can improve the performance of queries. Even though the access plan was created to use parallel preload, the system will actually use parallel preload only if the following are true:</p> <ul style="list-style-type: none"><li>-- The query attribute degree was specified with an option of *IO or *ANY for the application process.</li><li>-- There is enough main storage to load all of the index objects used by this query into active memory. Normally, a minimum of 5 megabytes would be a minimum. Increasing the size of the shared pool may improve performance.</li></ul> <p>For more information about parallel table prefetch, refer to <a href="#">“Data access methods” on page 16.</a></p>

### SQL4025 - Parallel table preload access method used

<b>Message Text:</b>	Parallel table preload access method used.
<b>Cause Text:</b>	The query optimizer chose to use a parallel table preload access method to reduce the processing time required for this query. This means that the data accessed by this query will be loaded into active memory when the query is opened.
<b>Recovery Text:</b>	<p>Parallel table preload can improve the performance of queries. Even though the access plan was created to use parallel preload, the system will actually use parallel preload only if the following are true:</p> <ul style="list-style-type: none"><li>-- The query attribute degree must have been specified with an option of *IO or *ANY for the application process.</li><li>-- There is enough main storage available to load all of the data in the file into active memory. Normally, 5 megabytes would be a minimum. Increasing the size of the shared pool may improve performance.</li></ul> <p>For more information about parallel table prefetch, refer to <a href="#">“Data access methods” on page 16.</a></p>

<b>SQL4026 - Index only access used on table number &amp;1</b>	
<b>Message Text:</b>	Index only access used on table number &1.
<b>Cause Text:</b>	Index only access is primarily used in conjunction with either index scan-key row positioning index scan-key selection. This access method will extract all of the data from the index rather than performing random I/O to the data space. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	Refer to <a href="#">“Data access methods”</a> on page 16 for more information about index only access.

<b>SQL4027 - Access plan was saved with DB2 Symmetric Multiprocessing installed on the system</b>	
<b>Message Text:</b>	Access plan was saved with DB2 Symmetric Multiprocessing installed on the system.
<b>Cause Text:</b>	The access plan saved was created while the system feature DB2 Symmetric Multiprocessing was installed on the system. The access plan may have been influenced by the presence of this system feature. Having this system feature installed may cause the implementation of the query to change.
<b>Recovery Text:</b>	For more information about how the system feature DB2 Symmetric Multiprocessing can influence a query, refer to the <a href="#">“Controlling parallel processing for queries”</a> on page 241.

<b>SQL4028 - The query contains a distributed table</b>	
<b>Message Text:</b>	The query contains a distributed table.
<b>Cause Text:</b>	A distributed table was specified in the query which may cause the query to be processed in multiple steps. If the query is processed in multiple steps, additional messages will detail the implementation for each step. Access plan implementation information for each step is not available because access plans are not saved for the individual steps. Debug messages detailing the implementation of each step can be found in the joblog if the query is run in debug mode using the STRDBG CL command.
<b>Recovery Text:</b>	For more information about how a distributed table can influence the query implementation refer to the <a href="#">Distributed database programming</a> topic collection.

<b>SQL4029 - Hashing algorithm used to process the grouping</b>	
<b>Message Text:</b>	Hashing algorithm used to process the grouping.
<b>Cause Text:</b>	The grouping specified within the query was implemented with a hashing algorithm.
<b>Recovery Text:</b>	Implementing the grouping with the hashing algorithm is generally a performance advantage since an index does not have to be created. However, if you want to disallow the use of this method simply specify ALWCPYDTA(*YES). Refer to <a href="#">“Data access methods”</a> on page 16 for more information about the hashing algorithm.

<b>SQL4030 - &amp;1 tasks specified for parallel scan on table &amp;2</b>	
<b>Message Text:</b>	&1 tasks specified for parallel scan on table &2.
<b>Cause Text:</b>	The query optimizer has calculated the optimal number of tasks for this query based on the query attribute degree. The table number refers to the relative position of this table in the query.

**SQL4030 - &1 tasks specified for parallel scan on table &2**

<b>Recovery Text:</b>	Parallel table or index scan can improve the performance of queries. Even though the access plan was created to use the specified number of tasks for the parallel scan, the system may alter that number based on the availability of the pool in which this job is running or the allocation of the table's data across the disk units. Refer to <a href="#">“Data access methods”</a> on page 16 for more information about parallel scan.
-----------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**SQL4031 - &1 tasks specified for parallel index create over table &2**

<b>Message Text:</b>	&1 tasks specified for parallel index create over table &2.
<b>Cause Text:</b>	The query optimizer has calculated the optimal number of tasks for this query based on the query attribute degree. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	Parallel index create can improve the performance of queries. Even though the access plan was created to use the specified number of tasks for the parallel index build, the system may alter that number based on the availability of the pool in which this job is running or the allocation of the table's data across the disk units. Refer to <a href="#">“Data access methods”</a> on page 16 for more information about parallel index create.

**SQL4032 - Index &1 used for bitmap processing of table &2**

<b>Message Text:</b>	Index &1 used for bitmap processing of table &2.
<b>Cause Text:</b>	The index was used, in conjunction with query selection, to create a bitmap. The bitmap, in turn, was used to access rows from the table. This message may appear more than once per table. If this occurs, then a bitmap was created from each index of each message. The bitmaps were then combined into one bitmap using boolean logic and the resulting bitmap was used to access rows from the table. The table number refers to the relative position of this table in the query.
<b>Recovery Text:</b>	The query can be run in debug mode (STRDBG) to determine more specific information. Also, refer to <a href="#">“Data access methods”</a> on page 16 for more information about bitmap processing.

**SQL4033 - &1 tasks specified for parallel bitmap create using &2**

<b>Message Text:</b>	&1 tasks specified for parallel bitmap create using &2.
<b>Cause Text:</b>	The query optimizer has calculated the optimal number of tasks to use to create the bitmap based on the query attribute degree.
<b>Recovery Text:</b>	Using parallel index scan to create the bitmap can improve the performance of queries. Even though the access plan was created to use the specified number of tasks, the system may alter that number based on the availability of the pool in which this job is running or the allocation of the file's data across the disk units. Refer to <a href="#">“Data access methods”</a> on page 16 for more information about parallel scan.

**SQL4034 - Multiple join classes used to process join**

<b>Message Text:</b>	Multiple join classes used to process join.
----------------------	---------------------------------------------

<b>SQL4034 - Multiple join classes used to process join</b>	
<b>Cause Text:</b>	Multiple join classes are used when join queries are written that have conflicting operations or cannot be implemented as a single query. Each join class will be optimized and processed as a separate step of the query with the results written out to a temporary table. Access plan implementation information for each of the join classes is not available because access plans are not saved for the individual join class dials. Debug messages detailing the implementation of each join dial can be found in the joblog if the query is run in debug mode using the STRDBG CL command.
<b>Recovery Text:</b>	Refer to <a href="#">“Join optimization” on page 66</a> for more information about join classes.

<b>SQL4035 - Table &amp;1 used in join class &amp;2</b>	
<b>Message Text:</b>	Table &1 used in join class &2.
<b>Cause Text:</b>	This message lists the table numbers used by each of the join classes. The table number refers to the relative position of this table in the query. All of the tables listed for the same join class will be processed during the same step of the query. The results from all of the join classes will then be joined together to return the final results for the query. Access plan implementation information for each of the join classes are not available because access plans are not saved for the individual classes. Debug messages detailing the implementation of each join class can be found in the joblog if the query is run in debug mode using the STRDBG CL command.
<b>Recovery Text:</b>	Refer to <a href="#">“Join optimization” on page 66</a> for more information about join classes.

## Code license and disclaimer information

IBM grants you a nonexclusive copyright license to use all programming code examples from which you can generate similar function tailored to your own specific needs.

SUBJECT TO ANY STATUTORY WARRANTIES WHICH CANNOT BE EXCLUDED, IBM, ITS PROGRAM DEVELOPERS AND SUPPLIERS MAKE NO WARRANTIES OR CONDITIONS EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, REGARDING THE PROGRAM OR TECHNICAL SUPPORT, IF ANY.

UNDER NO CIRCUMSTANCES IS IBM, ITS PROGRAM DEVELOPERS OR SUPPLIERS LIABLE FOR ANY OF THE FOLLOWING, EVEN IF INFORMED OF THEIR POSSIBILITY:

1. LOSS OF, OR DAMAGE TO, DATA;
2. DIRECT, SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES, OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES; OR
3. LOST PROFITS, BUSINESS, REVENUE, GOODWILL, OR ANTICIPATED SAVINGS.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF DIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, SO SOME OR ALL OF THE ABOVE LIMITATIONS OR EXCLUSIONS MAY NOT APPLY TO YOU.

## Notices

---

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Software Interoperability Coordinator, Department YBWA  
3605 Highway 52 N  
Rochester, MN 55901  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. \_enter the year or years\_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

## Programming interface information

---

This Database performance and query optimization publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM i.

## Trademarks

---

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be

trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Other product and service names might be trademarks of IBM or other companies.

## Terms and conditions

---

Permissions for the use of these publications is granted subject to the following terms and conditions.

**Personal Use:** You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

**Commercial Use:** You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.









Product Number: 5770-SS1