

IBM solidDB
Version 7.0

SQL Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 421.

First edition, fifth revision

This edition applies to V7.0 Fix Pack 8 of IBM solidDB (product number 5724-V17) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Oy IBM Finland Ab 1993, 2013

Contents

Figures **ix**

Tables **xi**

Summary of changes **xiii**

About this manual **xv**

Typographic conventions xv

Syntax notation conventions xvi

1 Database concepts **1**

1.1 Relational databases. 1

1.1.1 Tables, rows, and columns 1

1.1.2 Relating data in different tables 2

1.2 Client-Server architecture 4

1.3 Multi-user capability 4

1.4 Transactions 5

1.5 Transaction logging and recovery 5

1.5.1 Background 5

1.6 Summary 6

2 Getting started with SQL **7**

2.1 Tables, rows, and columns 7

2.2 SQL 7

2.3 The mathematical origins of SQL 9

2.4 Creating tables with related data 10

2.4.1 Table aliases 12

2.4.2 Subqueries 12

2.5 Which formats are used for each data type. 13

2.5.1 BLOBs (or binary data types). 14

2.5.2 NULL IS NOT NULL (or "how to say 'none of the above' in SQL"). 15

2.5.3 NOT NULL 17

2.5.4 Expressions and casts 17

2.5.5 Row value constructors 19

2.6 More about transactions 21

2.7 Summary 21

2.8 Where to find additional information about SQL 21

3 SQL extensions **23**

3.1 Stored procedures 23

3.1.1 Stored procedures – SQL 24

3.1.2 Stored procedures – external 62

3.2 Functions 62

3.3 Triggers 63

3.3.1 Triggers – principles of operation 63

3.3.2 Creating and modifying triggers. 64

3.3.3 Triggers and procedures 64

3.3.4 Triggers and transactions 67

3.3.5 Trigger privileges and security 74

3.3.6 Raising errors from inside triggers 74

3.3.7 Obtaining trigger information 75

3.3.8 Trigger parameter settings. 75

3.3.9 Trigger example 76

3.4 Sequences. 79

3.5 Events 80

3.5.1 Using events - Example 1 82

3.5.2 Using events - Example 2 84

3.5.3 Using events - Example 3 86

4 Using SQL for database administration **89**

4.1 Using solidDB SQL syntax 89

4.1.1 solidDB SQL data types 89

4.1.2 solidDB ADMIN COMMANDs 89

4.1.3 Scalar functions 89

4.2 Managing user privileges and roles 90

4.2.1 User privileges 90

4.2.2 User roles 90

4.2.3 Creating users. 91

4.2.4 Deleting users. 91

4.2.5 Changing a password 91

4.2.6 Creating roles 92

4.2.7 Deleting roles 92

4.2.8 Granting privileges to a user or a role. 92

4.2.9 Granting privileges to a user by giving the user a role. 92

4.2.10 Revoking privileges from a user or a role 92

4.2.11 Revoking privileges by revoking the role of a user 93

4.2.12 Granting administrator privileges to a user 93

4.2.13 Changing administrator's user name and password 93

4.3 Managing tables 94

4.3.1 Accessing system tables 94

4.3.2 Creating tables 95

4.3.3 Removing tables 95

4.3.4 Adding columns to a table 96

4.3.5 Deleting columns from a table 96

4.3.6 Viewing table names and definitions 97

4.4 Managing indexes 97

4.4.1 Primary key indexes 97

4.4.2 Secondary key indexes 98

4.4.3 Protection against duplicate indexes 99

4.4.4 Creating an index on a table 99

4.4.5 Creating a unique index on a table 100

4.4.6 Deleting an index 100

4.5 Referential integrity 100

4.5.1 Primary keys and candidate keys 100

4.5.2 Foreign keys 101

4.5.3 Referential actions 103

4.5.4 Dynamic constraint management 104

4.6 Managing database objects. 105

4.6.1 Catalogs 106

4.6.2 Schemas 106

4.6.3 Uniquely identifying objects within catalogs and schemas. 107

4.6.4 Creating a catalog 108

4.6.5 Setting a catalog and schema context. 108

4.6.6 Deleting a catalog	108
4.6.7 Creating a schema	108
4.6.8 Deleting a schema	109

5 Managing transactions 111

5.1 Defining read-only and read-write transactions	111
5.2 Concurrency control and locking	111
5.2.1 PESSIMISTIC vs. OPTIMISTIC concurrency control	112
5.2.2 Locks and lock modes	114
5.2.3 Setting concurrency control	118
5.3 Choosing transaction durability level	121
5.3.1 Setting transaction durability level	121

6 Diagnostics and troubleshooting for SQL 123

6.1 Observing performance	123
6.1.1 SQL Info facility	123
6.1.2 EXPLAIN PLAN FOR statement	124
6.2 Tracing facilities for stored procedures and triggers	129
6.2.1 User-definable trace output from procedure code	129
6.2.2 Procedure execution trace	130
6.3 Measuring and improving performance of START AFTER COMMIT statements	131
6.3.1 Tuning performance of START AFTER COMMIT statements	131
6.3.2 Analyzing failures in START AFTER COMMIT statements	131

7 Tuning performance with SQL 133

7.1 Tuning SQL statements and applications	133
7.1.1 Evaluating application performance	133
7.1.2 Using stored procedure language	134
7.1.3 Executing prepared statements (prepare once, execute many)	134
7.2 Optimizing single-table SQL queries	135
7.3 Using indexes to improve query performance	135
7.3.1 Full table scan	136
7.3.2 Concatenated indexes	137
7.4 Waiting on events	137
7.5 Optimizing batch inserts and updates	138
7.5.1 Increasing speed of batch inserts and updates	138
7.6 Using optimizer hints	139
7.7 Diagnosing poor performance	140

Appendix A. IBM solidDB SQL statements 141

A.1 ADMIN COMMAND	141
A.2 ADMIN EVENT	156
A.3 ALTER REMOTE SERVER	157
A.4 ALTER TABLE	157
A.4.1 ALTER TABLE ... SET HISTORY COLUMNS	159
A.4.2 ALTER TABLE ... SET SYNCHISTORY	160
A.5 ALTER TRIGGER	162
A.6 ALTER USER	162
A.6.1 ALTER USER (replica)	163

A.7 CALL	165
A.7.1 Access rights for remote stored procedure calls	168
A.8 COMMIT WORK	169
A.9 CREATE CACHE SEGMENT	169
A.10 CREATE CATALOG	170
A.11 CREATE EVENT	172
A.12 CREATE FUNCTION	173
A.13 CREATE FUNCTION (external)	175
A.14 CREATE INDEX	179
A.15 CREATE PROCEDURE	180
A.15.1 Usage	181
A.15.2 parameter_modes	182
A.15.3 prepare_statement	183
A.15.4 execute_statement	184
A.15.5 fetch_statement	185
A.15.6 post_statement	185
A.15.7 wait_register_statement	185
A.15.8 wait_event_statement	186
A.15.9 control_statement	186
A.15.10 writetrace_statement	188
A.15.11 exec_direct_statement	188
A.15.12 Procedure stack functions	188
A.15.13 Dynamic cursor names	188
A.15.14 Examples of CREATE PROCEDURE statement	190
A.16 CREATE PROCEDURE (external)	193
A.17 CREATE [OR REPLACE] PUBLICATION	197
A.18 CREATE [OR REPLACE] REMOTE SERVER	201
A.19 CREATE ROLE	202
A.20 CREATE SCHEMA	203
A.21 CREATE SEQUENCE	204
A.22 CREATE SYNC BOOKMARK	206
A.23 CREATE TABLE	207
A.24 CREATE TRIGGER	210
A.24.1 trigger_name	211
A.24.2 BEFORE AFTER clause	211
A.24.3 INSERT UPDATE DELETE clause	213
A.24.4 table_name	214
A.24.5 trigger_body	214
A.24.6 REFERENCING clause	214
A.24.7 Triggers usage notes and restrictions	215
A.25 CREATE USER	216
A.26 CREATE VIEW	217
A.27 DELETE	217
A.28 DESCRIBE	218
A.29 DROP CACHE SEGMENT	219
A.30 DROP CATALOG	219
A.31 DROP EVENT	220
A.32 DROP FUNCTION	220
A.33 DROP INDEX	221
A.34 DROP MASTER	221
A.35 DROP PROCEDURE	222
A.36 DROP PUBLICATION	222
A.37 DROP PUBLICATION REGISTRATION	222
A.38 DROP REMOTE SERVER	223
A.39 DROP REPLICA	223
A.40 DROP ROLE	224
A.41 DROP SCHEMA	224
A.42 DROP SEQUENCE	225
A.43 DROP SUBSCRIPTION	225

A.44 DROP SYNC BOOKMARK	227
A.45 DROP TABLE	228
A.46 DROP TRIGGER	228
A.47 DROP USER	229
A.48 DROP VIEW	229
A.49 EXPLAIN PLAN FOR	229
A.50 EXPORT SUBSCRIPTION	231
A.51 EXPORT SUBSCRIPTION TO REPLICA	233
A.52 GRANT	234
A.53 GRANT PASSTHROUGH	235
A.54 GRANT REFRESH	235
A.55 HINT	236
A.55.1 Hints supported by solidDB	240
A.56 IMPORT	242
A.57 INSERT	244
A.58 LIST	245
A.59 LOCK TABLE	247
A.60 MESSAGE	249
A.60.1 MESSAGE APPEND	249
A.60.2 MESSAGE BEGIN	252
A.60.3 MESSAGE DELETE	253
A.60.4 MESSAGE DELETE CURRENT TRANSACTION	254
A.60.5 MESSAGE END	256
A.60.6 MESSAGE EXECUTE	257
A.60.7 MESSAGE FORWARD	259
A.60.8 MESSAGE FROM REPLICA DELETE	264
A.60.9 MESSAGE FROM REPLICA EXECUTE	264
A.60.10 MESSAGE FROM REPLICA RESTART	265
A.60.11 MESSAGE GET REPLY	266
A.61 POST EVENT	270
A.62 REFRESH	270
A.63 REGISTER EVENT	273
A.64 REVOKE	273
A.65 REVOKE PASSTHROUGH	274
A.66 REVOKE REFRESH	274
A.67 ROLLBACK WORK	275
A.68 SAVE	275
A.69 SAVE PROPERTY	277
A.70 SELECT	278
A.71 SET	280
A.71.1 Differences between SET and SET TRANSACTION	281
A.71.2 SET (read/write level)	282
A.71.3 SET CATALOG	282
A.71.4 SET DELETE CAPTURE	282
A.71.5 SET DURABILITY	283
A.71.6 SET ISOLATION LEVEL	283
A.71.7 SET PASSTHROUGH	284
A.71.8 SET SAFENESS	284
A.71.9 SET SCHEMA	284
A.71.10 SET SEQUENCE	285
A.71.11 SET SQL	286
A.71.12 SET STATEMENT MAXTIME	287
A.71.13 SET SYNC	288
A.71.14 SET TIMEOUT	295
A.72 SET TRANSACTION	295
A.72.1 Differences between SET and SET TRANSACTION	296
A.72.2 SET TRANSACTION (read/write level)	297
A.72.3 SET TRANSACTION DELETE CAPTURE	297

A.72.4 SET TRANSACTION DURABILITY	298
A.72.5 SET TRANSACTION ISOLATION LEVEL	298
A.72.6 SET TRANSACTION PASSTHROUGH	298
A.72.7 SET TRANSACTION SAFENESS	299
A.73 START AFTER COMMIT	299
A.74 TRUNCATE TABLE	301
A.75 UNLOCK TABLE	302
A.76 UNREGISTER EVENT	303
A.77 UPDATE	303
A.78 WAIT EVENT	304
A.79 Common clauses	304
A.79.1 check_condition	304
A.79.2 data_type	305
A.79.3 expression	306
A.79.4 query_specification	307
A.79.5 search_condition	308
A.79.6 table_reference	309
A.79.7 Pseudo columns in SELECT statement	309
A.80 Date and time literals	310
A.81 Wildcard characters	311
A.81.1 Using SQL wildcards	311
A.81.2 Wildcard characters as literals	312

Appendix B. Functions 313

B.1 String functions	313
B.2 Numeric functions	314
B.3 Date time functions	316
B.4 System functions	318
B.5 Miscellaneous functions	319
B.6 Advanced replication functions	319
B.6.1 GET_PARAM() function	319
B.6.2 PUT_PARAM() function	321
B.7 Trigger functions	322

Appendix C. Data types 323

C.1 Character data types	324
C.2 Numeric data types	325
C.3 Binary data types	326
C.4 Date data type	327
C.5 Time data type	327
C.6 Timestamp data type	328
C.7 Smallest possible non-zero numbers	328
C.8 BLOBs and CLOBs	328

Appendix D. Reserved words 331

Appendix E. Database system tables and system views 345

E.1 System tables	345
E.1.1 SQL_LANGUAGES	345
E.1.2 SYS_ATTAUTH	345
E.1.3 SYS_AUDIT_TRAIL	346
E.1.4 SYS_BACKGROUNDJOB_INFO	347
E.1.5 SYS_BLOBS	347
E.1.6 SYS_CACHESEGMENT_CONDITIONS	348
E.1.7 SYS_CACHESEGMENTS	349
E.1.8 SYS_CARDINAL	349
E.1.9 SYS_CATALOGS	350
E.1.10 SYS_CHECKSTRINGS	350

E.1.11	SYS_COLUMNS	350
E.1.12	SYS_COLUMNS_AUX	351
E.1.13	SYS_DL_REPLICA_CONFIG	351
E.1.14	SYS_DL_REPLICA_DEFAULT	352
E.1.15	SYS_EVENTS	352
E.1.16	SYS_FEDT_DB_PARTITION	353
E.1.17	SYS_FEDT_TABLE_PARTITION	353
E.1.18	SYS_FORKEYPARTS	354
E.1.19	SYS_FORKEYS	354
E.1.20	SYS_HOTSTANDBY	355
E.1.21	SYS_INFO	355
E.1.22	SYS_KEYPARTS	355
E.1.23	SYS_KEYS	356
E.1.24	SYS_LOGPOS	356
E.1.25	SYS_PROCEDURES	357
E.1.26	SYS_PROCEDURE_COLUMNS	357
E.1.27	SYS_PROPERTIES	358
E.1.28	SYS_RELAUTH	359
E.1.29	SYS_SCHEMAS	359
E.1.30	SYS_SEQUENCES	359
E.1.31	SYS_SERVER	360
E.1.32	SYS_SYNC_REPLICA_PROPERTIES	360
E.1.33	SYS_SYNONYM	360
E.1.34	SYS_TABLEMODES	361
E.1.35	SYS_TABLES	361
E.1.36	SYS_TRIGGERS	362
E.1.37	SYS_TYPES	363
E.1.38	SYS_URole	363
E.1.39	SYS_USERS	364
E.1.40	SYS_VIEWS	364
E.2	System tables for data synchronization	364
E.2.1	SYS_BULLETIN_BOARD	365
E.2.2	SYS_PUBLICATION_ARGS	365
E.2.3	SYS_PUBLICATION_REPLICA_ARGS	365
E.2.4	SYS_PUBLICATION_REPLICA_STMTARGS	366
E.2.5	SYS_PUBLICATION_REPLICA_STMTS	366
E.2.6	SYS_PUBLICATION_STMTARGS	366
E.2.7	SYS_PUBLICATION_STMTS	367
E.2.8	SYS_PUBLICATIONS	367
E.2.9	SYS_PUBLICATIONS_REPLICA	368
E.2.10	SYS_SYNC_BOOKMARKS	368
E.2.11	SYS_SYNC_HISTORY_COLUMNS	369
E.2.12	SYS_SYNC_INFO	369
E.2.13	SYS_SYNC_MASTER_MSGINFO	369
E.2.14	SYS_SYNC_MASTER_RECEIVED_BLOB_REFS	371
E.2.15	SYS_SYNC_MASTER_RECEIVED_MSGPARTS	371
E.2.16	SYS_SYNC_MASTER_RECEIVED_MSGS	371
E.2.17	SYS_SYNC_MASTER_STORED_BLOB_REFS	372
E.2.18	SYS_SYNC_MASTER_STORED_MSGPARTS	372
E.2.19	SYS_SYNC_MASTER_STORED_MSGS	372
E.2.20	SYS_SYNC_MASTER_SUBSC_REQ	373
E.2.21	SYS_SYNC_MASTER_VERSIONS	373
E.2.22	SYS_SYNC_MASTERS	374
E.2.23	SYS_SYNC_RECEIVED_BLOB_ARGS	374
E.2.24	SYS_SYNC_RECEIVED_STMTS	375
E.2.25	SYS_SYNC_REPLICA_MSGINFO	375

E.2.26	SYS_SYNC_REPLICA_RECEIVED_BLOB_REFS	376
E.2.27	SYS_SYNC_REPLICA_RECEIVED_MSGPARTS	377
E.2.28	SYS_SYNC_REPLICA_RECEIVED_MSGS	377
E.2.29	SYS_SYNC_REPLICA_STORED_BLOB_REFS	378
E.2.30	SYS_SYNC_REPLICA_STORED_MSGS	378
E.2.31	SYS_SYNC_REPLICA_STORED_MSGPARTS	378
E.2.32	SYS_SYNC_REPLICA_VERSIONS	379
E.2.33	SYS_SYNC_REPLICAS	379
E.2.34	SYS_SYNC_SAVED_BLOB_ARGS	379
E.2.35	SYS_SYNC_SAVED_STMTS	380
E.2.36	SYS_SYNC_TRX_PROPERTIES	380
E.2.37	SYS_SYNC_USERMAPS	381
E.2.38	SYS_SYNC_USERS	381
E.3	System views	381
E.3.1	COLUMNS	382
E.3.2	SERVER_INFO	382
E.3.3	TABLES	383
E.3.4	USERS	383
E.4	Synchronization-related views	383
E.4.1	SYNC_FAILED_MESSAGES	384
E.4.2	SYNC_FAILED_MASTER_MESSAGES	384
E.4.3	SYNC_ACTIVE_MESSAGES	385
E.4.4	SYNC_ACTIVE_MASTER_MESSAGES	385

Appendix F. Database virtual tables 387

F.1	SYS_LOG	387
-----	---------	-----

Appendix G. System stored procedures 391

G.1	System stored procedures for advanced replication	391
G.1.1	SYNC_SETUP_CATALOG	391
G.1.2	SYNC_REGISTER_REPLICA	392
G.1.3	SYNC_UNREGISTER_REPLICA	393
G.1.4	SYNC_REGISTER_PUBLICATION	394
G.1.5	SYNC_UNREGISTER_PUBLICATION	395
G.1.6	SYNC_SHOW_SUBSCRIPTIONS	396
G.1.7	SYNC_SHOW_REPLICA_SUBSCRIPTIONS	397
G.1.8	SYNC_DELETE_MESSAGES	398
G.1.9	SYNC_DELETE_REPLICA_MESSAGES	399
G.1.10	SYNC_GET_CATALOG_SYNC_NODE_NAME_AND_MODE	400
G.2	Miscellaneous stored procedures	400
G.2.1	SYS_GETBACKGROUNDJOB_INFO	400
G.2.2	SYS_HSBSTATE_WAIT	400
G.2.3	SYS_TF1_WAIT	400

Appendix H. System events 401

H.1	Miscellaneous events	401
H.2	Errors that cause SYS_EVENT_ERROR	408
H.3	Conditions or warnings that cause SYS_EVENT_MESSAGES	409

Index	411
Notices	421

Figures

1. Sync pull notify	59	4. Execution graph 1	128
2. Example: Tables with referential constraints	101	5. Execution graph 2	129
3. Self-referential constraint.	102		

Tables

1. Typographic conventions	xv	52. MESSAGE DELETE return values from replica	254
2. Syntax notation conventions.	xvi	53. MESSAGE DELETE return values from master	254
3. Example database table	1	54. MESSAGE DELETE CURRENT TRANSACTION return values	255
4. Example database table	7	55. MESSAGE END return values from replica	256
5. Example database table	7	56. MESSAGE END return values from master	256
6. Comparison operators	30	57. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set	257
7. Logical operators: NOT	31	58. MESSAGE EXECUTE return values	259
8. Logical operators: AND	31	59. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set	261
9. Logical operators: OR	31	60. MESSAGE FORWARD return values from replica	262
10. Determining data type from parameters	40	61. MESSAGE FORWARD return values from master	264
11. INSERT/UPDATE/DELETE operations for BEFORE/AFTER triggers	69	62. MESSAGE FROM REPLICA EXECUTE return values	265
12. Example Entry 1	72	63. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set	267
13. Example entry 2	72	64. MESSAGE GET REPLY return values from replica	268
14. Sequences in stored procedures	80	65. MESSAGE GET REPLY return values from master	269
15. System roles	91	66. REFRESH return values	271
16. Viewing tables and granting access.	94	67. REVOKE REFRESH return values.	275
17. Expressions and operators	104	68. SAVE return values	276
18. SQL Info levels	123	69. SAVE PROPERTY return values	277
19. EXPLAIN PLAN FOR units.	124	70. SQL Info levels	286
20. EXPLAIN PLAN FOR table columns	125	71. SET SYNC return values.	289
21. Texts in the unit INFO column.	125	72. SET SYNC CONNECT return values	289
22. EXPLAIN PLAN FOR, Example 1.	127	73. How different operations apply to synchronization history tables	291
23. EXPLAIN PLAN FOR, Example 2.	128	74. SET SYNC MODE return values	291
24. Diagnosing poor performance	140	75. SET SYNC NODE return values	293
25. ADMIN COMMAND syntax and options	142	76. SET SYNC PARAMETER Return Values	293
26. ALTER TABLE SET HISTORY COLUMNS return values.	160	77. UNLOCK TABLE return values	303
27. ALTER TABLE SET SYNCHISTORY return values	162	78. check_condition.	304
28. ALTER USER return values	165	79. data_type	305
29. Comparison of the parameter modes	183	80. Expression	306
30. Control statements.	187	81. query_specification	307
31. CREATE PUBLICATION return values	200	82. search_condition	308
32. System roles	202	83. table_reference	309
33. Sequences in stored procedures	205	84. Pseudo columns	309
34. CREATE SYNC BOOKMARK return values	207	85. Date and time literals.	310
35. Statement atomicity in a trigger	215	86. Wildcard characters	311
36. DROP MASTER return values	221	87. String functions.	313
37. DROP PUBLICATION return values	222	88. Numeric functions	314
38. DROP PUBLICATION REGISTRATION Return Values	223	89. Date time functions	316
39. DROP REPLICA return values	224	90. System functions	318
40. DROP SUBSCRIPTION return values	226	91. Miscellaneous functions	319
41. DROP SYNC BOOKMARK return values	227	92. GET_PARAM return values.	320
42. EXPLAIN PLAN FOR units.	230	93. PUT_PARAM() return values	321
43. EXPLAIN PLAN FOR table columns	230	94. Abbreviations in data type tables	323
44. EXPORT SUBSCRIPTION return values	233	95. Character Data Types	324
45. EXPORT SUBSCRIPTION TO REPLICA return values.	234		
46. GRANT REFRESH return values	236		
47. Supported hints.	240		
48. IMPORT return values	243		
49. MESSAGE APPEND return values	251		
50. MESSAGE BEGIN return values from replica	252		
51. MESSAGE BEGIN return values from master	253		

96. Numeric data types	325	154. SYS_SYNC_MASTER_MSGINFO	370
97. Binary Data Types	326	155. SYS_SYNC_MASTER_	
98. Date Data Type	327	RECEIVED_BLOB_REFS	371
99. Time data type	327	156. SYS_SYNC_MASTER_RECEIVED	
100. Timestamp data type	328	_MSGPARTS	371
101. Smallest possible non-zero numbers	328	157. SYS_SYNC_MASTER_RECEIVED_MSGS	371
102. Reserved words list	331	158. SYS_SYNC_MASTER_STORED_BLOB_REFS	372
103. SQL_LANGUAGES system table	345	159. SYS_SYNC_MASTER_STORED_MSGPARTS	372
104. SYS_ATTAUTH	345	160. SYS_SYNC_MASTER_STORED_MSGS	372
105. SYS_AUDIT_TRAIL table definition	346	161. SYS_SYNC_MASTER_SUBSC_REQ	373
106. SYS_BACKGROUNDJOB_INFO	347	162. SYS_SYNC_MASTER_VERSIONS	373
107. SYS_BLOBS	347	163. SYS_SYNC_MASTERS	374
108. SYS_CACHESEGMENT_CONDITIONS table		164. SYS_SYNC_RECEIVED_BLOB_ARGS	374
definition	348	165. SYS_SYNC_RECEIVED_STMTS	375
109. SYS_CACHESEGMENTS table definition	349	166. SYS_SYNC_REPLICA_MSGINFO	376
110. SYS_CARDINAL	349	167. SYS_SYNC_REPLICA_RECEIVED_	
111. SYS_CATALOGS	350	BLOB_REFS	377
112. SYS_CHECKSTRINGS	350	168. SYS_SYNC_REPLICA_RECEIVED_	
113. SYS_COLUMNS	350	MSGPARTS	377
114. SYS_COLUMNS_AUX	351	169. SYS_SYNC_REPLICA_RECEIVED_MSGS	377
115. SYS_DL_REPLICA_CONFIG	352	170. SYS_SYNC_REPLICA_STORED_BLOB_REFS	378
116. SYS_DL_REPLICA_DEFAULT	352	171. SYS_SYNC_REPLICA_STORED_MSGS	378
117. SYS_EVENTS	352	172. SYS_SYNC_REPLICA_STORED_MSGPARTS	378
118. SYS_FEDT_DB_PARTITION table definition	353	173. SYS_SYNC_REPLICA_VERSIONS	379
119. SYS_FEDT_TABLE_PARTITION table		174. SYS_SYNC_REPLICAS	379
definition	353	175. SYS_SYNC_SAVED_BLOB_ARGS	380
120. SYS_FORKEYPARTS	354	176. SYS_SYNC_SAVED_STMTS	380
121. SYS_FORKEYS	354	177. SYS_SYNC_TRX_PROPERTIES	381
122. SYS_INFO	355	178. SYS_SYNC_USERMAPS	381
123. SYS_KEYPARTS	355	179. SYS_SYNC_USERS	381
124. SYS_KEYS	356	180. COLUMNS	382
125. SYS_LOGPOS table definition	357	181. SERVER_INFO	383
126. SYS_PROCEDURES	357	182. TABLES	383
127. SYS_PROCEDURE_COLUMNS	358	183. USERS	383
128. SYS_PROPERTIES	358	184. SYNC_FAILED_MESSAGES	384
129. SYS_RELAUTH	359	185. SYNC_FAILED_MASTER_MESSAGES	384
130. SYS_SCHEMAS	359	186. SYNC_ACTIVE_MESSAGES	385
131. SYS_SEQUENCES	359	187. SYNC_ACTIVE_MASTER_MESSAGES	385
132. SYS_SERVER table definition	360	188. SYS_LOG table definition	387
133. SYS_SYNC_REPLICA_PROPERTIES	360	189. Record identifier descriptions	388
134. SYS_SYNONYM	360	190. DATA type descriptions for row data	389
135. SYS_TABLEMODES	361	191. DATA type descriptions for DDL data	390
136. SYS_TABLES	361	192. SYNC_SETUP_CATALOG error codes	391
137. SYS_TRIGGERS	362	193. SYNC_REGISTER_REPLICA error codes	392
138. SYS_TYPES	363	194. SYNC_UNREGISTER_REPLICA error codes	393
139. SYS_UROLE	363	195. SYNC_REGISTER_PUBLICATION error codes	394
140. SYS_USERS	364	196. SYNC_UNREGISTER_PUBLICATION error	
141. SYS_VIEWS	364	codes	395
142. SYS_BULLETIN_BOARD	365	197. SYNC_SHOW_SUBSCRIPTIONS result set	396
143. SYS_PUBLICATION_ARGS	365	198. SYNC_SHOW_SUBSCRIPTIONS error codes	397
144. SYS_PUBLICATION_REPLICA_ARGS	365	199. SYNC_SHOW_REPLICA_SUBSCRIPTIONS	
145. SYS_PUBLICATION_REPLICA_STMTARGS	366	result set	397
146. SYS_PUBLICATION_REPLICA_STMTS	366	200. SYNC_SHOW_REPLICA_SUBSCRIPTIONS	
147. SYS_PUBLICATION_STMTARGS	366	error codes	398
148. SYS_PUBLICATION_STMTS	367	201. SYNC_DELETE_MESSAGES error codes	398
149. SYS_PUBLICATIONS	367	202. SYNC_DELETE_REPLICA_MESSAGES error	
150. SYS_PUBLICATIONS_REPLICA	368	codes	399
151. SYS_SYNC_BOOKMARKS	368	203. Miscellaneous events	402
152. SYS_SYNC_HISTORY_COLUMNS	369	204. Errors that cause SYS_EVENT_ERROR	408
153. SYS_SYNC_INFO	369	205. Warnings that cause SYS_EVENT_MESSAGES	409

Summary of changes

Changes for revision 05

- Added new SQL statements for creating and modifying cache segments:
 - CREATE CACHE SEGMENT
 - DROP CACHE SEGMENT
 - ALTER TABLE ... ADD CACHE SEGMENT
 - ALTER TABLE ... DROP CACHE SEGMENT
- Modified SQL statement ALTER USER by adding the keyword CONNECTIONS in section ALTER USER. The CONNECTIONS keyword defines the maximum number of connections for a user.
- Clarified usage of named and positional parameters in stored procedures in sections CALL and Using parameters.
- New section Executing prepared statements (prepare once, execute many) added with information about optimizing SQL statement performance.
- Detailed information about **ADMIN COMMAND 'status'** output added in section ADMIN COMMAND.
- A previously undocumented command **ADMIN COMMAND 'getreadonlyflag'** added in section ADMIN COMMAND.

Changes for revision 04

- Editorial corrections.

Changes for revision 03

- Editorial corrections.

Changes for revision 02

- The use of quotation marks in ADMIN COMMAND clarified in section ADMIN COMMAND: the *command_name* in ADMIN COMMAND '*command_name*' must be given with single quotation marks.

Changes for revision 01

- The default scale of DECIMAL and NUMERIC data types clarified in section Numeric data types.
- Information on supported ROWNUM constructions added in section Pseudo columns in SELECT statement.
- Examples of C programming conventions for CREATE PROCEDURE (external) and CREATE FUNCTION (external) updated.

About this manual

This guide introduces you to the relational database server theory and the SQL programming language. It also includes appendices that show the syntax of all SQL statement supported by IBM® solidDB®, and describes the data types that can be used in tables and SQL statements.

This guide is for users who want to learn about SQL in general or who want to learn about solidDB specific SQL.

Typographic conventions

solidDB documentation uses the following typographic conventions:

Table 1. Typographic conventions

Format	Used for
Database table	This font is used for all ordinary text.
NOT NULL	Uppercase letters on this font indicate SQL keywords and macro names.
solid.ini	These fonts indicate file names and path expressions.
SET SYNC MASTER YES; COMMIT WORK;	This font is used for program code and program output. Example SQL statements also use this font.
run.sh	This font is used for sample command lines.
TRIG_COUNT()	This font is used for function names.
java.sql.Connection	This font is used for interface names.
LockHashSize	This font is used for parameter names, function arguments, and Windows registry entries.
<i>argument</i>	Words emphasized like this indicate information that the user or the application must provide.
<i>Administrator Guide</i>	This style is used for references to other documents, or chapters in the same document. New terms and emphasized issues are also written like this.
File path presentation	Unless otherwise indicated, file paths are presented in the UNIX format. The slash (/) character represents the installation root directory.

Table 1. *Typographic conventions (continued)*

Format	Used for
Operating systems	If documentation contains differences between operating systems, the UNIX format is mentioned first. The Microsoft Windows format is mentioned in parentheses after the UNIX format. Other operating systems are separately mentioned. There may also be different chapters for different operating systems.

Syntax notation conventions

solidDB documentation uses the following syntax notation conventions:

Table 2. *Syntax notation conventions*

Format	Used for
INSERT INTO <i>table_name</i>	Syntax descriptions are on this font. Replaceable sections are on <i>this</i> font.
solid.ini	This font indicates file names and path expressions.
[]	Square brackets indicate optional items; if in bold text, brackets must be included in the syntax.
	A vertical bar separates two mutually exclusive choices in a syntax line.
{ }	Curly brackets delimit a set of mutually exclusive choices in a syntax line; if in bold text, braces must be included in the syntax.
...	An ellipsis indicates that arguments can be repeated several times.
· · ·	A column of three dots indicates continuation of previous lines of code.

1 Database concepts

If you are not already familiar with relational database servers such as solidDB, you may want to read this chapter.

This chapter explains the following concepts:

- Relational databases
 - Tables, rows, and columns
 - Relating data in different tables
- Multi-user capability / Concurrency control and locking
- Client-Server architecture
- Transactions
- Transaction logging and Recovery

1.1 Relational databases

1.1.1 Tables, rows, and columns

Most relational database servers, including the solidDB family, use a programming language known as the Structured Query Language (SQL). SQL is a set-oriented programming language that is designed to allow people to query and update tables of information. This chapter discusses tables, and how data is represented within tables. Later in the manual, we will discuss the syntax of the SQL language in more detail.

All information is stored in tables. A table is divided into rows and columns. (SQL theorists refer to columns as "attributes" and rows as "tuples", but we will use the more familiar terms "columns" and "rows". We will also use the terms "record" and "row" interchangeably.) Each database contains 0 or more tables. Most databases contain many tables. An example of a table is shown below.

Table 3. Example database table

ID	NAME	ADDRESS
1	Beethoven	23 Ludwig Lane
2	Dylan	46 Robert Road
3	Nelson	79 Willie Way

This table contains 3 rows of data. (The top "row", which has the labels "ID", "NAME", and "ADDRESS" is shown here for the convenience of the reader. The actual table in the database does not have such a row.) The table contains 3 columns (ID, NAME, and ADDRESS).

SQL provides commands to create tables, insert rows into tables, update data in tables, delete rows from tables, and query the rows in tables.

Tables in SQL, unlike arrays in programming languages like C, are not homogeneous. In SQL one column may have one data type (such as INTEGER), while an adjacent column may have a very different data type (such as CHAR(20), which means an array of 20 characters).

A table may have varying numbers of rows. Rows may be inserted and deleted at any time; you do not need to preallocate space for a maximum number of rows. (All database servers have some maximum number of rows that they can handle. For example, most database servers that run on 32-bit operating systems have a limit of approximately two billion rows. In most applications, the maximum is far more than you are likely to need.)

Each row ("record") must have at least one value, or combination of values, that is unique. If we have two composers named David Jones to our table, and we need to update the address of only one of them, then we need some way to tell them apart. In some cases, you can find a combination of columns that is unique, even if you can't find any single column that contains unique values. For example, if the name column is not sufficient, then perhaps the combination of name and address will be unique. However, without knowing all the data ahead of time, it is difficult to absolutely guarantee that each value will be unique. Most database designers add an "extra" column that has no purpose other than to uniquely and easily identify each record. In our table above, for example, the ID numbers are unique. As you may have noticed, when we actually try to update or delete a record, we identify it by its unique ID (for example, "... WHERE id = 1") rather than by using another value, such as name, that might not be unique.

1.1.2 Relating data in different tables

If SQL could only handle one table at a time, it would be convenient, but not very powerful. The true power of SQL and relational databases lies in the fact that tables can be related to each other in useful ways, and SQL queries can gather data from multiple tables and display that data in a logical fashion.

We will show how multiple tables are useful by using a bank as an example.

Each customer of a bank may have more than 1 account. There is no real limit to the number of accounts a person might have. One customer might have a checking account, savings account, certificate of deposit, mortgage, credit card, and so on. Furthermore, a person may have multiple accounts of the same type. For example, a customer might have one savings account with retirement money and another savings account (of the same type) that has money for her daughter's college fund. We describe the "relationship" between a person and her accounts as a "one to many" relationship -- one person may have many accounts.

Because there is no limit to the number of accounts a person may have, there is no way to design a record structure ahead of time that can handle all possible combinations of accounts. And if you created a record structure that held the maximum number of accounts that anyone actually owned, you'd have to waste a lot of space. Let's suppose that we tried to build a single table that held all the information about one bank customer and her accounts. Our first draft might look like the following:

```
Customer ID Number
Customer Name
Customer Address
Checking Account #1 ID
Checking Account #1 Balance
CD #1 ID
```

CD #1 Balance
CD #2 ID
CD #2 Balance
...

As you can see, we just do not know when to stop because there is no obvious limit to the number of accounts that each person might own.

Another solution is to create multiple records, one for each account, and duplicate the customer information for each account. So we have a table that looks like:

Customer Name
Customer Address
Account ID
Account Balance

If a customer has more than one account, we merely create a complete record for each account. This works reasonably well, but it means that every single account record holds all the information about the customer. This wastes storage space and also makes it harder to update the customer's address if the customer moves (you may have to update the address in several places).

Relational databases, such as solidDB's, are designed to solve this problem. We will create one table for customers, and another table for accounts. (In a real bank, we would probably divide the accounts into multiple tables, too, with one table for checking accounts, another table for savings accounts, and so on) Then we create a "link" between the customer and each of her accounts. This allows us to waste very little space and yet still have complete information available to us.

As we mentioned earlier, in our example of composers, every record should have a unique value that allows us to identify that record. The unique value is usually just an integer. We'll use that unique integer to help us "relate" a customer to her accounts. This is discussed in more detail in 2, "Getting started with SQL," on page 7.

When we create an account for a customer, we store that customer's ID number as part of the account information. Specifically, each row in the accounts table has a `customer_id` value, and that `customer_id` value matches the id of the customer who owns that account. Smith has customer id 1, and each of Smith's accounts has a 1 in the `customer_id` field. That means that we can find all of Smith's account records by doing the following:

1. Look up Smith's record in the customers table.
2. When we find Smith's record, look at the id number in that record. (In Smith's case, the id is 1.)
3. Now look up all accounts in the accounts table that have a value of 1 in the `customer_id` field.

It's as though you taped a copy of your home telephone number onto the forehead of each of your children when they went to school. If there is an emergency and you need to send a taxi driver to find and pick up your children at school, you can simply tell the taxi driver your phone number and he can check every child in the school to see if the child has your phone number. (This isn't very efficient, but it works.) By knowing the parent's id number, you can identify all the children. Conversely, by knowing each child, you can identify the parent. If, for example, one of your children is lost on a field trip away from the school, any helpful person can simply read the telephone number off the child's forehead and call you.

As you can see, the parent and child are linked to each other without any sort of physical contact. Simply having the id number (or phone number) is enough to determine which children belong to a parent and which parent belongs to each child. The technique works regardless of how many children you have.

Relational databases use the same technique. Note that join operations are not limited to two tables. It's possible to create joins with an almost arbitrary number of tables. As a realistic extension of our banking example, we might have another table, "checks", which holds information about each check written. Thus we would have not only a 1-to-many relationship from each customer to her accounts, but also a 1-to-many relationship from each checking account to all of the checks written on that account. It's quite possible to write a query that will list all the checks that a customer has written, even if that customer has multiple checking accounts.

1.2 Client-Server architecture

solidDB uses the client-server model. In a client-server model, a single "server" may process requests from 1 or more "clients". This is quite similar to the way that a restaurant works — a single waiter and cook may handle requests from many customers.

In a client-server database model, the server is a specialized computer program that knows how to store and retrieve data efficiently. The server typically accepts four basic types of requests:

- Insert a new piece of information
- Update an existing piece of information
- Retrieve an existing piece of information
- Delete an existing piece of information

The server can store almost any type of data, but generally doesn't know the "meaning" of the data. The server typically knows little or nothing about "business issues", such as accounting, inventory, and so on. It doesn't know whether a particular piece of information is an inventory record, a description of a bank deposit, or a digitized copy of the song "American Pie".

The "clients" are responsible for knowing something about the particular business issues and about the "meaning" of the data. For example, we might write a client program that knows something about accounting. The client program might know how to calculate interest on late payments, for example. Or, the client might recognize that a particular piece of data is a song, and might convert the digital data to analog audio output.

It is possible to write a single program that does both the "client" and the "server" part of the work. A program that reads digitized music and plays it could also store that data to disk and look it up on request. However, it's not very efficient for every company to write its own data storage and retrieval routines. It is usually more efficient to buy an off-the-shelf data storage solution that is general enough to meet your needs, yet has relatively high performance.

1.3 Multi-user capability

An important advantage of client-server architecture is that it usually makes it easier to have more than one client. solidDB, like most relational database servers, allows multiple users to access the data in a table.

When two users try to update the same data, there is potential danger. If the updates are not the same, then one user's updates could write over the other user's updates. solidDB uses concurrency control mechanisms to prevent this. For more information, see 5.2, "Concurrency control and locking," on page 111.

1.4 Transactions

SQL allows you to group multiple statements into a single "atomic" (indivisible) piece of work called a transaction. For example, if you write a check to a grocery store, then the grocery store's bank account should be given the money at the same instant that the money is withdrawn from your account. It wouldn't make sense for you to pay the money without the grocery store receiving it, and it wouldn't make sense for the grocery store to be paid without your account having the money subtracted. If either of these operations (adding to the grocery store's account or subtracting from yours) fails, then the other one ought to fail, too. If both statements are in the same transaction, and either statement fails, you can use the ROLLBACK command to restore things as they were before the transaction started — this prevents half-successful transactions from occurring. If both halves of our financial transaction are successful, we would like our database transaction to be successful, too. Successful transactions are preserved with the command COMMIT WORK. Below is a simplistic example.

```
COMMIT WORK; -- Finish the previous transaction.
UPDATE stores SET balance = balance + 199.95
  WHERE store_name = 'Big Tyke Bikes';
UPDATE checking_accounts SET balance = balance - 199.95
  WHERE name = 'Jay Smith';
COMMIT WORK;
```

1.5 Transaction logging and recovery

One of the major advantages of buying a commercial database server is that most such servers have been designed to protect data if the database server shuts down unexpectedly for any reason, such as a power failure, a hardware failure, or a failure in the database software itself.

There are a number of different ways to help protect data. We will focus on one such way, called Transaction Logging.

1.5.1 Background

Suppose that you are writing data to a disk drive (or other permanent storage medium) and suddenly the power fails. The data that you write might not be written completely. For example, you might try to write the account balance "122.73", but because of the power failure you just write "12". The person whose account is missing some money will be quite displeased. How do we ensure that we always write complete data? Part of the solution is to use what is called a "transaction log".

Note:

In the world of computers, many different things are called "logs". For example, the solidDB writes multiple log files, including a transaction log file and an error message log file. For the moment, we are discussing only the transaction log file.

As we mentioned previously, work is usually done in "transactions". An entire transaction is either committed or rolled back. No partial transactions are allowed. In the situation described here, where we started to write a person's new account

balance to disk but lost power before we could finish, we would like to roll back this transaction. Any transactions that were already completed and were correctly written to disk should be preserved.

To help us track what data has been written successfully and what data has not been written successfully, we actually write data to a "transaction log" as well as to the database tables. The transaction log is essentially a linear sequence of the operations that have been performed — that is, the transactions that have been committed. There are markers in the file to indicate the end of each transaction. If the last transaction in the file does not have an "end-of-transaction" marker, then we know that fractional transaction was not completed, and it should be rolled back rather than committed.

When the server restarts after a failure, it reads the transaction log and applies the completed transactions one by one. In other words, it updates the tables in the database, using the information in the transaction log file. This is called "recovery". When done properly, recovery can even protect against power failures during the recovery process itself.

This is not a complete description of how transaction logging protects against data corruption. We have explained how the server makes sure that it doesn't lose transactions. But we haven't really explained how the server prevents the database file from becoming corrupted if a write failure occurs while the server is in the middle of writing a record to a table in the disk drive. That topic is more advanced and is not discussed here.

1.6 Summary

This brief introduction to relational databases has explained the concepts that you need to start using a relational database. You should now be able to answer the following questions:

What are tables, rows, and columns?

Can you work with data in more than one table at a time?

How do transactions help keep data consistent?

Why do we write ("log") transaction data to the disk drive?

2 Getting started with SQL

This chapter gives you a quick overview (or refresher) in SQL.

2.1 Tables, rows, and columns

SQL is a set-oriented programming language that is designed to allow people to query and update tables of information.

All information is stored in tables. A table is divided into rows and columns. (SQL theorists refer to columns as "attributes" and rows as "tuples", but we will use the more familiar terms "columns" and "rows". We will also use the terms "record" and "row" interchangeably.) Each database contains 0 or more tables. Most databases contain many tables. An example of a table is shown below.

Table 4. Example database table

ID	NAME	ADDRESS
1	Beethoven	23 Ludwig Lane
2	Dylan	46 Robert Road
3	Nelson	79 Willie Way

This table contains three rows of data. (The top "row", which has the labels "ID", "NAME", and "ADDRESS" is shown here for the convenience of the reader. The actual table in the database does not have such a row.) The table contains three columns (ID, NAME, and ADDRESS). SQL provides commands to create tables, insert rows into tables, update data in tables, delete rows from tables, and query the rows in tables.

2.2 SQL

The following SQL "program" creates the table shown in the example below:

```
CREATE TABLE composers (id INTEGER PRIMARY KEY, name CHAR(20),  
address CHAR(50));  
INSERT INTO composers (id, name, address) VALUES (1, 'Beethoven',  
'23 Ludwig Lane');  
INSERT INTO composers (id, name, address) VALUES (2, 'Dylan',  
'46 Robert Road');  
INSERT INTO composers (id, name, address) VALUES (3, 'Nelson',  
'79 Willie Way');
```

Table 5. Example database table

ID	NAME	ADDRESS
1	Beethoven	23 Ludwig Lane
2	Dylan	46 Robert Road
3	Nelson	79 Willie Way

the column "id" is designated to be the "primary key" of the table. This means that each row may be uniquely identified by using this column. From now on, the system will guarantee that the value of "id" is unique and it always exists (that is, it has the NOT NULL property).

If Mr. Dylan moves to 61 Bob Street, you can update his data with the command:

```
UPDATE composers SET ADDRESS = '61 Bob Street' WHERE ID = 2;
```

Because the ID field is unique for each composer, and because the WHERE clause in this command specifies only one ID, this update will be performed on only one composer.

If Mr. Beethoven dies and you need to delete his record, you can do so with the command:

```
DELETE FROM composers WHERE ID = 1;
```

Finally, if you would like to list all the composers in your table, you can use the command:

```
SELECT id, name, address FROM composers;
```

Note that the SELECT statement, unlike the UPDATE and DELETE statements listed above, did not include a WHERE clause. Therefore, the command applied to ALL records in the specified table. Thus the result of this SQL statement is to select (and list) all of the composers listed in the table.

ID	NAME	ADDRESS
1	Beethoven	23 Ludwig Lane
2	Dylan	46 Robert Road
3	Nelson	79 Willie Way

Note that although you entered the strings with quotation marks, they are displayed without quotation marks.

The above simple commands help show some important points about SQL.

- SQL is a relatively "high level" language. A single command can create a table with as many columns as you wish. Similarly, a single command can execute an UPDATE of almost any complexity. Although not shown here, you can update multiple columns at a time, and you can even update more than one row at a time. Operations that might take dozens, or hundreds, of lines of code in languages like C or Java™ can be executed in a single SQL command.
- Unlike some other computer languages, SQL uses single quotation marks to delimit strings. For example, 'Beethoven' is a string. "Beethoven" is something different. (Technically, it is a delimited identifier, which is not discussed in this chapter.) If you are used to programming languages like C, which use double quotation marks to delimit strings (character arrays) and single quotation marks to delimit individual characters, you will have to adjust to SQL's way of doing things.

Although the example above does not clearly show it, there are several additional points you need to know about basic SQL:

- Although SQL is a very powerful high-level language, it is also a very limited one. SQL is designed for table-oriented and record-oriented operations. It has very few low-level operations. For example, there is no direct way to open a file, or to shift bits leftward or rightward. It is also hardware-independent, which is both an advantage and disadvantage. You have very little control over the format of the output from SQL queries; you may choose the order of the

columns, and by using the ORDER BY clause you may control the order of the rows, but you cannot do things such as control the size of the font on the screen, or print page numbers at the bottom of each printed page of output. SQL simply is not a complete programming language such as C, Java, PASCAL, and so on.

- Each SQL implementation has a fixed set of data types. The data types in solidDB (and most other implementations of SQL) include INTEGER, CHARACTER array, FLOATing point, DATE, and TIME.
- SQL is generally an "interpreted" language rather than a "compiled" language. To execute one or more SQL statements, you typically execute a separate program that reads your script and then executes it. No "compiled program" or "executable" is generated and stored for later use. Each time you run the program, it is interpreted again. (Stored Procedures can be reused without necessarily reinterpreting them. Stored Procedures are discussed briefly in Appendix A, "IBM solidDB SQL statements," on page 141 and extensively in 3, "SQL extensions," on page 23.
- Table and column names are case-insensitive in SQL. In our examples, keywords (such as CREATE, INSERT, SELECT) are capitalized, and table and column names are shown in lower case. However, this is only a convention, not a requirement.
- SQL is also not very picky about whether commands are written on a single line or are split across multiple lines. There are examples of multiline statements later in this chapter.
- SQL commands can get extremely complicated, with multiple nested "layers" of queries within queries. Figuring out how to write a complex query can be quite difficult - and figuring out how to understand a query that someone else wrote can be equally difficult. As in any programming language, it is a good idea to document your code!
- To help you document your code, SQL allows "comments". Comments are only for the human reader; they are skipped over by the SQL interpreter. To create a comment, you have two options:
 - Line comment: put two dashes (--) at the start of the line and end the comment with a line break. The comment cannot extend to a new line.
 - Block (multiline) comment: Begin the comment with a slash and an asterisk (/*) and end the comment with an asterisk and a slash (*/).All the subsequent characters up to the end of the line will be ignored. (There is an exception for "optimizer hints", another advanced topic that we will not discuss in this chapter.)

2.3 The mathematical origins of SQL

Relational databases and SQL were originally based in part upon the mathematical concept of set theory. If you are familiar with set theory, it will help you understand how relational databases work. If you are not familiar with set theory, then do not worry about it; this is merely one way of looking at relational databases and SQL.

A table can be thought of as a mathematical set, where each element of the set is a row. (In our example above, each person, or composer, is an element of a set. The table contains all of the elements of the set 'composers'.) In mathematics, sets are unordered. Similarly, in SQL, tables are largely treated as unordered, even though if you could look at the bits and bytes on the disk you would find that at any given time the records are stored in a particular order.

This lack of ordering is important, because it means that the results of a query may be shown in a different order each time that you run the query. With small data sets stored on a single disk drive, you will usually see the same rows in the same order each time, but this is not necessarily the case when data is spread across multiple files or disk drives.

Because SQL is a set-oriented language, you can use it to perform some set-oriented operations, such as UNIONS (that is, combining two sets of input into one set of output). However, operations such as UNION require that the sets match each other - i.e. that they have the same number of columns, and that they have the same data type (or compatible data type) in corresponding columns. You can't perform a UNION operation if the first column in set1 is of type DATETIME and the first column in set2 is INTEGER, for example.

Again, if you are not comfortable with set theory, do not worry about it. This is just another way of looking at relational databases.

2.4 Creating tables with related data

As described in the previous chapter, each customer of a bank may have more than one account. We describe the "relationship" between a person and her accounts as a "one to many" relationship — one person may have many accounts.

Because there is no limit to the number of accounts a person may have, there is no way to design a record structure ahead of time that can handle all possible combinations of accounts.

Relational databases, are designed to solve this problem. We will create one table for customers, and another table for accounts. (In a real bank, we would probably divide the accounts into multiple tables, too, with one table for checking accounts, another table for savings accounts, and so on.) Then we create a "link" between the customer and each of her accounts. This allows us to waste very little space and yet still have complete information available to us.

As we mentioned earlier, in our example of composers, every record should have a primary key that allows us to identify that record. It is usually just an integer. We will now use that unique integer to help us "relate" a customer to her accounts. Below are the commands to create and populate the customer table:

```
CREATE TABLE customers (id INTEGER PRIMARY KEY, name CHAR(20),
address CHAR(40));
INSERT INTO customers (id, name, address) VALUES (1, 'Smith',
'123 Main Street');
INSERT INTO customers (id, name, address) VALUES (2, 'Jones',
'456 Fifth Avenue');
```

We have inserted two customers, named Smith and Jones. Let us create the account table:

```
CREATE TABLE accounts (id INTEGER PRIMARY KEY, balance FLOAT,
customer_id INT REFERENCES customers);
```

Here, we have designated the column *customer_id* to be a "foreign key" pointing to the customer table (this is indicated by the REFERENCES keyword). The value of this column is supposed to be exactly the same as the "id" value (the primary key) in the corresponding customer row in the "customers" table. This way we will associate account rows with customer rows. The feature of a database allowing to maintain such relationships in a reliable way is called "referential integrity", and

the corresponding SQL syntax elements used to define such relationships are called "referential integrity constraints". For more on referential integrity, see 4.5, "Referential integrity," on page 100.

Customer Smith has two accounts, and customer Jones has 1 account.

```
INSERT INTO accounts (id, balance, customer_id)
VALUES (1001, 200.00, 1);
INSERT INTO accounts (id, balance, customer_id)
VALUES (1002, 5000.00, 1);
INSERT INTO accounts (id, balance, customer_id)
VALUES (1003, 222.00, 2);
```

As Smith has two accounts, each of Smith's accounts has a 1 in the *customer_id* field. That means that a user can find all of Smith's account records by doing the following:

1. Look up Smith's record in the customers table.
2. When we find Smith's record, look at the id number in that record. (In Smith's case, the id is 1.)
3. Now look up all accounts in the accounts table that have a value of 1 in the *customer_id* field.

It's as though you taped a copy of your home telephone number onto the forehead of each of your children when they went to school. If there is an emergency and you need to send a taxi driver to find and pick up your children at school, you can simply tell the taxi driver your phone number and he can check every child in the school to see if the child has your phone number. (This isn't very efficient, but it works.) By knowing the parent's id number, you can identify all the children. Conversely, by knowing each child, you can identify the parent. If, for example, one of your children is lost on a field trip away from the school, any helpful person can simply read the telephone number off the child's forehead and call you.

As you can see, the parent and child are linked to each other without any sort of physical contact. Simply having the id number (or phone number) is enough to determine which children belong to a parent and which parent belongs to each child. The technique works regardless of how many children you have.

Relational databases use the same technique. Now that we've created our customer table and our accounts table, we can show each customer and each of the accounts that she has. To do this, we use what SQL programmers call a "join" operation. The WHERE clause in the SELECT statement "joins" those pairs of records where the account's *customer_id* number matches the customer's id number.

```
SELECT name, balance
FROM customers, accounts
WHERE accounts.customer_id = customers.id;
```

The output of this query is similar to the following:

```
NAME BALANCE
Smith 200.0
Smith 5000.0
Jones 222.0
```

If a person has multiple accounts, she might want to know the total amount of money that she has in all accounts. The computer can provide this information by using the following query:

```
SELECT customers.id, SUM(balance)
   FROM customers, accounts
   WHERE accounts.customer_id = customers.id
   GROUP BY customers.id;
```

The output of this query is similar to the following:

```
NAME  BALANCE
Smith 5200.0
Jones 222.0
```

Note that this time, Smith appears only once, and she appears with the total amount of money in all her accounts.

This query uses the GROUP BY clause and an aggregate function named SUM(). The topic of GROUP BY clauses is more complex than we want to go into during this simple introduction to SQL. This query is just to give you a little taste of the type of useful work that SQL can do in a single statement. Getting the same result in a language like C would take many statements.

Note that join operations are not limited to two tables. It's possible to create joins with an almost arbitrary number of tables. As a realistic extension of our banking example, we might have another table, "checks", which holds information about each check written. Thus we would have not only a 1-to-many relationship from each customer to her accounts, but also a 1-to-many relationship from each checking account to all of the checks written on that account. It's quite possible to write a query that will list all the checks that a customer has written, even if that customer has multiple checking accounts.

2.4.1 Table aliases

SQL allows you to use an *alias* in place of a table name in some queries. In the following query, alias "a" is used for the accounts table and "c" for the customers table.

```
SELECT name, balance
   FROM customers c, accounts a
   WHERE a.customer_id = c.id;
```

The alias is defined in the FROM clause and then used in the WHERE clause.

2.4.2 Subqueries

SQL allows one query to contain another query, called a "subquery".

Returning to our bank example, over time, some customers add accounts and other customers terminate accounts. In some cases, a customer might gradually terminate accounts until he has no more accounts. Our bank may want to identify all customers that do not have any accounts so that those customers' records can be deleted, for example. One way to identify the customers who do not have any accounts is to use a subquery and the EXISTS clause.

To try this out, we need to create a customer who doesn't have any accounts:

```
INSERT INTO customers (id, name, address) VALUES (3, 'Zu', 'B St');
```

Before we list all customers who do not have accounts, let's list all customers who do have accounts.

```
SELECT id, name
   FROM customers c
   WHERE EXISTS (SELECT * FROM accounts a WHERE a.customer_id = c.id);
```

The subquery (also called the "inner query") is the query inside the parentheses. The inner query is executed once for each record selected by the outer query. (This functions a lot like nested loops would function in another programming language, except that with SQL we can do nested loops in a single statement.) If there are any accounts for the particular customer that the outer loop is processing, then those account records are returned to the outer query.

The "EXISTS" clause in the outer query says, effectively, "We do not care what values are in those records; all we care about is whether there are any records or not." Thus EXISTS returns TRUE if the customer has any accounts. If the customer has no accounts, then the EXISTS returns false. The EXISTS clause doesn't care whether there are multiple accounts or single accounts. It doesn't care what values are in the accounts. All the EXISTS wants to know is "Is there at least one record?"

Thus, the entire statement lists those customers who have at least one account. No matter how many accounts the customer has (as long as it's at least 1), the customer is listed only once.

Now let's list all those customers who do not have any accounts:

```
SELECT id, name
FROM customers c
WHERE NOT EXISTS (SELECT * FROM accounts a WHERE a.customer_id = c.id);
```

Merely adding the keyword NOT reverses the sense of the query.

Subqueries may themselves have subqueries. In fact, subqueries may be nested almost arbitrarily deep.

2.5 Which formats are used for each data type

As we've already shown above, SQL requires that values be expressed in a particular way. For example, character strings must be delimited by single quote marks.

Other values also must be formatted properly. The exact format required depends upon the data type. Several data types other than CHARACTER data types also require single quotation marks to delimit the values that you enter.

Below are some examples of how to format input data for most of the data types that solidDB supports. We'll show this in the form of a simple SQL script that you can execute if you wish. Note that in this script, many commands are split across multiple lines. This is quite legal in SQL. It's one of the reasons that most SQL interpreters expect a semicolon to separate each SQL statement, even though the ANSI Standard for SQL doesn't actually require a semicolon at the end of each statement.

```
CREATE TABLE one_of_almost_everything (
  int_col INTEGER,
  float_col FLOAT,
  string_col CHAR(20),
  wide_string_col WCHAR(20), -- "wide" means wide chars, for example, unicode.
  varchar_col VARCHAR, -- Note that we did not have to specify width.
  date_col DATE,
  time_col TIME,
  timestamp_col TIMESTAMP
);

INSERT INTO one_of_almost_everything (
  int_col,
```

```

float_col,
string_col,
wide_string_col,
varchar_col,
date_col,
time_col,
timestamp_col
)
VALUES (
1,
2.0,
'three',
'four',
'five point zero zero zero zero zero zero zero zero zero zero ...',
'2002-12-31',
'11:59:00',
'1999-12-31 23:59:59.00000'
);

```

As you can see, timestamp values are entered in order from the "most significant" digit to the "least significant" digit. Similarly, date and time values are also entered from the most significant digit to the least significant digit. And all 3 of these data types (timestamp, date, time) use punctuation to separate individual fields.

The reason for requiring particular formats is that some of the other possible formats are ambiguous. For example, to someone in the U.S., '07-04-1776' is July 4, 1776, since Americans usually write dates in the 'mm-dd-yyyy' (or 'mm/dd/yyyy' format). But to a person from Europe, this date is obviously April 7, not July 4th, since most Europeans write dates in the format 'dd-mm-yyyy'. Although it may seem that the problem of having too many formats is not well solved by adding still another format, there are some advantages to SQL's approach of using a format that starts with the most significant digit and moves steadily toward the least significant digit. First, it means that all three data types (date, time, and timestamp) follow the same rule. Second, the date format and the time format are both perfect subsets of the timestamp format. Third, although it's yet another format to memorize, the rule is reasonably simple and is consistent with the way that "western" languages write numbers (most significant digit is furthest to the left). Finally, by being obviously incompatible with the existing formats, there's no chance that a person will accidentally write one date (for example, '07-04-1776') and have it interpreted by the machine as another date.

2.5.1 BLOBs (or binary data types)

So far, we have discussed data types that store data that is intended to be read by humans. Some types of data are not intended to be read directly by humans, but can still be stored in a database. For example, a picture from a digital camera, or a song from a CD, is stored as a series of numbers. These numbers are almost meaningless to a human. Digitized pictures and sounds can be stored as BINARY data, however. solidDB supports three binary data types: BINARY, VARBINARY, and LONG VARBINARY (or BLOB).

In most cases, you will read and write binary data using the ODBC (Open DataBase Connectivity) API from a C program, or the JDBC API from a Java program. However, it is possible to insert data into a binary field using a utility that executes SQL statements. To insert a value into a binary field, you must represent the value as a series of hexadecimal numbers inside single quotation marks. For example, if you wanted to insert a series of bytes with the values 1, 9, 11, 255 into a binary field, you would execute:

```
INSERT INTO table1 (binary_col) VALUES (CAST('01090BFF' AS VARBINARY));
```

Because this command instructs the server to CAST the value to type VARBINARY, the server automatically interprets the string as a series of hexadecimal numbers, not as a string literal.

You may also insert a string literal directly, for example:

```
INSERT INTO table1 (binary_col) VALUES ('Thank you');
```

When you retrieve the data via solsql (solidDB utility for executing SQL statements), the return value from a binary column is expressed in hexadecimal, whether or not you originally entered it as hexadecimal. Thus, after you insert the value 'Thank you', if you select this value from the table you will see:

```
5468616E6B20796F75
```

where 54 represents capital 'T', 68 represents lower-case 'h', 61 represents lower-case 'a', 6E represents lower-case 'n', and so on.

Note also that for long values only the first several digits are shown.

2.5.2 NULL IS NOT NULL (or "how to say 'none of the above' in SQL")

Sometimes you do not have enough information to fill out a form completely. SQL uses the keyword NULL to represent "Unknown" or "No Value". (This is different from the meaning of NULL in programming languages such as C.) For example, if you are inserting a record for Joni Mitchell into the table of composers and you do not know Joni Mitchell's address, you might execute the following:

```
INSERT INTO composers (id, name, address) VALUES (5, 'Mitchell', NULL);
```

If you do not specify the address field, it will contain NULL by default.

```
INSERT INTO composers (id, name) VALUES (5, 'Mitchell');
```

The following example is a sample program that show some unusual characteristics of the value NULL.

```
-- Data of any data type may contain NULL.
-- For example, a column of type INTEGER may contain not
-- only valid integer values, but also NULL.

-- Set up for experiments...
CREATE TABLE table1 (x INTEGER, name CHAR(30));

-- The value NULL means "there is no value".
-- NULL is not the same as zero, or an empty string.
-- (It is also not a pointer value, as it is in
-- programming languages such as C.)
-- To help show this, insert 3 rows, one of which has
-- "normal" values, one of which has a 0 and an empty string,
-- and one of which has two NULL values.
INSERT INTO table1 (x, name) VALUES (2, 'Ludwig Von Beethoven');
INSERT INTO table1 (x, name) VALUES (0, '');
INSERT INTO table1 (x, name) VALUES (NULL, NULL);
-- This returns only the row containing 0,
-- not the row containing NULL.
SELECT * FROM table1 WHERE x = 0;
-- This returns only the row containing the empty string,
-- not the row containing NULL.
SELECT * FROM table1 WHERE name = '';

-- As expected, NULL does not match other values.
-- However NULL does not match even itself.
```



```

-- (A mathematician would say that NULL violates the
-- reflexive property "a = a".)
SELECT * FROM table1 WHERE x = x;

-- Since NULL does not equal NULL, what will the following query return?
SELECT * FROM table1 WHERE x <> x;

-- Similarly, although you might think that the
-- expression below is always true, it is actually
-- always false.
SELECT * FROM table1 WHERE NULL IN (NULL, 2);

-- The result set will contain 2 (since 2 is in
-- the set (NULL, 2)), but the result set will
-- not contain NULL.
SELECT * FROM table1 WHERE x IN (NULL, 2);

-- But suppose that you *want* to find all the records that
-- have NULL values. How to do that if it is not possible to say ... = NULL?
SELECT * FROM table1 WHERE x IS NULL;
-- And the opposite query is ...
SELECT * FROM table1 WHERE x IS NOT NULL;

-- Set up for more experiments...
CREATE TABLE parent (id INTEGER, name CHAR(20));
CREATE TABLE children (id INTEGER, name CHAR(12), parent_id INT);
INSERT INTO parent (id, name) VALUES (1, 'Smith');
INSERT INTO children (id, name, parent_id) VALUES (11, 'Smith child', 1);
INSERT INTO children (id, name, parent_id) VALUES (131, 'orphan', NULL);
INSERT INTO parent (id, name) VALUES (NULL, 'Has Null');

-- Since NULL <> NULL, if a "parent" record has NULL and a "child"
-- record has NULL, the child's value will not match the parent's value.
-- This result set will contain 'Smith', but not 'Has Null'.
SELECT p.name FROM parent p, children c
WHERE c.parent_id = p.id;

-- Note that a row that contains nothing but a
-- single NULL is still a row.
-- In the following query, an EXISTS clause is used.
-- It evaluates to TRUE if the subquery returns
-- any rows. Even a row that contains nothing but a
-- single NULL value is still a row, and so if the
-- subquery returns a single NULL the EXISTS clause
-- still evaluates to TRUE.
-- Even though the subquery below returns NULL rather than a name
-- or ID, the EXISTS expression evaluates to TRUE, and Smith is printed.
SELECT name FROM parent p
WHERE EXISTS(SELECT NULL FROM children c WHERE c.parent_id = p.id);

-- Now that you recognize that NULL <> NULL,
-- the following may confuse you since it breaks the pattern.
-- Contrary to what you might expect, the UNIQUE keyword
-- DOES filter out multiple NULL values.
INSERT INTO table1 (x, name) VALUES (NULL, 'any name');
-- Now the table has more than one row in which x is NULL,
-- but a query with UNIQUE nonetheless returns only a
-- single NULL value.
SELECT DISTINCT x FROM table1;
-- You may be interested to know that a UNIQUE index

```



```
-- will allow only a single NULL value. (Note that a primary key
-- will not allow any NULL values.)
```

```
-- Clean up.
DROP TABLE parent;
DROP TABLE children;
DROP TABLE table1;
```

2.5.3 NOT NULL

As opposed to NULL, NOT NULL is one of the SQL data constraints. NOT NULL indicates that null values are not allowed in any row of the table for the specified column. For more information and examples, refer to Appendix A, "IBM solidDB SQL statements," on page 141.

2.5.4 Expressions and casts

SQL allows expressions in some parts of SQL statements. For example, the following statement multiplies the value in a column by 12:

```
SELECT monthly_average * 12 FROM table1;
```

As another example, the following statement uses the built-in SQRT function to calculate the square root of each value in the column named "variance".

```
SELECT SQRT(variance) FROM table1;
```

Our next example uses the "REPLACE" function to convert numbers from U.S. format to European format. In U.S. format, numbers use the period character ('.') as the decimal point, but in Europe the comma (',') is used. For example, in the U.S. the approximation of pi is written as "3.14", while in Europe it is written as "3,14". We can use the REPLACE function to replace the '.' character with the ',' character. The following series of statements shows an example of this.

```
CREATE TABLE number_strings (n VARCHAR);
INSERT INTO number_strings (n) VALUES ('3.14'); -- input in US format.
SELECT REPLACE(n, '.', ',') FROM number_strings; -- output in European.
```

The output looks like

```
n
-----
3,14
```

Note that one function can call another. The following expression takes the square root of a number and then takes the natural log of that square root:

```
SELECT LOG(SQRT(x)) FROM table1;
```

solidDB SQL does not accept completely general expressions in all clauses. For example, in the SELECT clause, you may use predefined functions, but you may not call stored procedures that you have created. Even if you have created a stored procedure named "foo", the following will not work:

```
SELECT foo(column1) FROM table1;
```

When you use expressions, you may want to specify a new name for a column. For example, if you use the expression

```
SELECT monthly_average * 12 FROM table1;
```

you probably do not want the output column to be called "monthly_average". solidDB server will actually use the expression itself as the name of the column. In this case, the name of the column would be "monthly_average * 12". That's

certainly descriptive, but for a long expression this can get very messy. You can use the "AS" keyword to give an output column a specific name. In the following example, the output will have the column heading "yearly_average".

```
SELECT monthly_average * 12 AS yearly_average FROM table1;
```

Note that the AS clause works for any output column, not just for expressions. If you like, you may do something like the following:

```
SELECT ssn AS SocialSecurityNumber FROM table2;
```

A CASE clause allows you to control the output based on the input. Below is a simple example, which converts a number (1-12) to the name of a month:

```
CREATE TABLE dates (m INT);
INSERT INTO dates (m) VALUES (1);
-- ...
INSERT INTO dates (m) VALUES (12);
INSERT INTO dates (m) VALUES (13);

SELECT
    CASE m
        WHEN 1 THEN 'January'
        --
        WHEN 12 THEN 'December'
        ELSE 'Invalid value for month'
    END
    AS month_name
FROM dates;
```

Note that this not only allows you to convert valid values, but also allows you to generate appropriate output if there is an error. The "ELSE" clause allows you to specify an alternative value if you get an input value that you weren't expecting.

In some situations, you may want to cast a value to a different data type. For example, when inserting BLOB data, it is convenient to create a string that contains your data, and then insert that string into a BINARY column. You may use a cast as shown below:

```
CREATE TABLE table1 (b BINARY(4));
INSERT INTO table1 VALUES ( CAST('FF00AA55' AS BINARY));
```

This cast allows you to take data that is a series of hexadecimal digits and input it as though it were a string. Each of the hexadecimal pairs in the quoted string represents a single byte of data. There are 8 hexadecimal digits, and thus 4 bytes of input.

A cast can be used to change output as well as input. In the rather complex code sample below, the expression in the CASE clause converts the output from the format '2003-01-20 15:33:40' to '2003-Jan-20 15:33:40'.

```
CREATE TABLE sample1(dt TIMESTAMP);
COMMIT WORK;

INSERT INTO sample1 VALUES ('2003-01-20 15:33:40');
COMMIT WORK;

SELECT
    CASE MONTH(dt)
        WHEN 1 THEN REPLACE(CAST(dt AS varchar), '-01-', '-Jan-')
        WHEN 2 THEN REPLACE(CAST(dt AS varchar), '-02-', '-Feb-')
        WHEN 3 THEN REPLACE(CAST(dt AS varchar), '-03-', '-Mar-')
        WHEN 4 THEN REPLACE(CAST(dt AS varchar), '-04-', '-Apr-')
        WHEN 5 THEN REPLACE(CAST(dt AS varchar), '-05-', '-May-')
        WHEN 6 THEN REPLACE(CAST(dt AS varchar), '-06-', '-Jun-')
```

```

        WHEN 7 THEN REPLACE(CAST(dt AS varchar), '-07-', '-Jul-')
        WHEN 8 THEN REPLACE(CAST(dt AS varchar), '-08-', '-Aug-')
        WHEN 9 THEN REPLACE(CAST(dt AS varchar), '-09-', '-Sep-')
        WHEN 10 THEN REPLACE(CAST(dt AS varchar), '-10-', '-Oct-')
        WHEN 11 THEN REPLACE(CAST(dt AS varchar), '-11-', '-Nov-')
        WHEN 12 THEN REPLACE(CAST(dt AS varchar), '-12-', '-Dec-')
    END
    AS formatted_date
FROM sample1;

```

This takes a value from a column named `dt`, converts that value from timestamp to `VARCHAR`, then replaces the month number with an abbreviation for the month (for example, it replaces `"-01-"` with `"-Jan-"`). By using the `CASE/WHEN/END` syntax, we can specify exactly what output we want for each possible input. Note that because this expression is so complicated, it is almost mandatory to use an `AS` clause to specify the column header in the output.

2.5.5 Row value constructors

This section explains one of the less familiar types of expressions, the *row value constructor* (RVC), and how it is used with relational operators, such as greater than, less than, and so on.

A row value constructor is an ordered sequence of values delimited by parentheses, for example:

```

(1, 4, 9)
('Smith', 'Lisa')

```

You can think of this as constructing a row based on a series of elements/values, just like a row of a table is composed of a series of fields.

Row value constructors, like individual values, may be used in comparisons. For example, just as you may have expressions like:

```

WHERE x > y;
WHERE 2 > 1;

```

you also may have expressions like:

```

WHERE (2, 3, 4) > (1, 2, 3);
WHERE (t1.last_name, t1.first_name) = (t2.last_name, t2.first_name);

```

Comparisons using row value constructors must be done carefully. Rather than give the technical definition of comparisons (which you can find in section 8.2 (comparison predicates) of the SQL-92 standard), we will give examples and an analogy to help you see the pattern.

The following expressions are true:

```

(9, 9, 9) > (1, 1, 1)
('Baker', 'Barbara') > ('Alpert', 'Andy')
(1, 1) = (1, 1)
(3, 2, 1) <> (4, 3, 2)

```

The examples above are simple, because the expression is correct for each corresponding pair of elements and is therefore true for the RVCs. For example, `'Baker' > 'Alpert'` and `'Barbara' > 'Andy'`, and therefore `('Baker', 'Barbara') > ('Alpert', 'Andy')`

However, when comparing row value constructors, it is not necessary that the expression be true for each corresponding element. In a row value constructor, the further left an element is, the more significance it has. Thus the following expressions are also true:

```
(9, 1, 1) > (1, 9, 9)
('Zoomer', 'Andy') > ('Alpert', 'Zelda')
```

In these examples, since the most significant element of the first RCV is greater than the corresponding element of the second RCV, the expression is true, regardless of the values of the remaining elements. Similarly, in the examples below, the first elements are identical, but the expressions overall are true:

```
(1, 1, 2) > (1, 1, 1)
(1, 2, 1) > (1, 1, 1)
('Baker', 'Zelda') > ('Baker', 'Allison')
```

Again, in a row value constructor, the further left an element is, the more significance it has. This is similar to the way that we compare multi-digit numbers. In a 3-digit number, such as 911, the hundreds-place digit is more significant than the tens-place digit, and the tens-place digit is more significant than the ones-place digit. Thus, the number 911 is greater than the number 199, even though not all digits of 911 are greater than the corresponding digits of 199.

This is useful when comparing multiple columns that are related. A practical application of this is when comparing people's names. For example, suppose that we have 2 tables, each of which has an *lname* (last name) and *fname* (first name) column. Suppose that we want to find all people whose names are less than Michael Morley's. In this situation, we want the last name to have more significance than the first name. The following names are shown in the correct alphabetical order (by last name):

Adams, Zelda

Morley, Michael

Young, Anna

If we want to list all persons whose names are less than Michael Morley's, then we do NOT want to use the following:

```
table1.lname < 'Morley' and table1.fname < 'Michael'
```

If we used this expression, we would reject Zelda Adams because her first name is alphabetically after Michael Morley's first name. One correct solution is to use the row value constructor approach:

```
(table1.lname, table1.fname) < ('Morley', 'Michael')
```

Note that when using equality, the expression must be true for ALL elements of the RCVs. For example:

```
(1, 2, 3) = (1, 2, 3)
```

For inequality, the expression must be true for only one element:

```
(1, 2, 1) <> (1, 1, 1)
```

2.6 More about transactions

As described in the previous chapter, SQL allows you to group multiple statements into a single "atomic" (indivisible) piece of work called a transaction. Successful transactions are preserved with the command `COMMIT WORK`. Below is a simplistic example.

```
COMMIT WORK; -- Finish the previous transaction.
UPDATE stores SET balance = balance + 199.95
  WHERE store_name = 'Big Tyke Bikes';
UPDATE checking_accounts SET balance = balance - 199.95
  WHERE name = 'Jay Smith';
COMMIT WORK;
```

If you do not want to keep a particular transaction, you can roll it back by using the command:

```
ROLLBACK WORK;
```

If you do not explicitly commit or roll back your work, then the server will roll it back for you. In other words, unless you confirm that you want to keep the data (by committing it), the data will be discarded.

2.7 Summary

This brief introduction to SQL and relational databases has explained the concepts that you need to start using SQL. You should now be able to answer the following questions:

What are tables, rows, and columns?

How do I create a table?

How do I put data into a table?

How do I update data in a table?

How do I delete data from a table?

How do I list data in a table?

How do I list related data in two different tables?

How do I ensure that multiple statements are executed together (so that all fail or all succeed as a group)?

2.8 Where to find additional information about SQL

Other chapters in this manual explain more about SQL and solidDB -specific features. However, this manual is neither a complete tutorial nor a comprehensive reference on SQL. You may wish to acquire additional documents on SQL.

There are many books on SQL. These books are not specific to solidDB's implementation of SQL; most of the material is generic and will apply to any database server, such as solidDB's, that conforms to the ANSI standards. General SQL books include:

- *Introduction to SQL: Mastering the Relational Database Language*, by Rick van der Lans, published by Addison-Wesley.

ANSI standards on SQL include:

- Database Language - SQL with Integrity Enhancement, ANSI, 1989 ANSI X3.135-1989.
- Database Language - SQL: ANSI X3H2 and ISO/IEC JTC1/SC21/WG3 9075:1992 (SQL-92).

For more information about ANSI standards, see <http://www.ansi.org/>.

ISO (International Standards Organization) also has standards for SQL. See <http://www.iso.org> for more information.

3 SQL extensions

The SQL support in solidDB is comparable to any advanced SQL-based system; solidDB offers the most commonly expected features and a set of useful extensions employing solidDB-specific (nonstandard) SQL syntax. Procedural SQL extensions such as stored procedures, stored functions, and triggers enable moving parts of the application logic into the database. These extensions help reduce network traffic, thus improving performance.

solidDB also supports both internal and external user-defined stored functions and procedures. External functions and procedures can be written in C programming language and used as migration tools; missing functionality can be added at the cost of writing a C program.

Note: External functions and procedures are supported in shared memory access (SMA) and linked library access (LLA) setups.

3.1 Stored procedures

Stored procedures are simple programs, or procedures, that are executed in solidDB databases. solidDB supports both SQL-written procedures stored in the server and external procedures written in C programming language, stored as dynamic libraries.

There are two types of stored procedures:

- *Stored procedures* written in the solidDB proprietary SQL procedure language
- *External stored procedures* written in C programming language.

The external stored procedures are loaded at runtime using standard dynamic library interfaces as provided by the operating system. To enable use of external stored procedures, you must link your application dynamically to solidDB using shared memory access (SMA) or linked library access (LLA).

Because stored procedures are stored and executed directly in the server, usage of stored procedures reduces network traffic and can thus improve performance. For example, complex, data-bound transactions can be run on the server itself. External stored procedures can also be used as migration tools; missing functionality can be added at the cost of writing a C program.

You can create a procedure that contains several SQL statements or a whole transaction and execute it with a single call statement. In addition to SQL statements, 3GL type control structures can be used enabling procedural control. You can also create nested stored procedures where one procedure is executed from within another.

Stored procedures can also be used for controlling access rights and database operations. Granting execute rights on a stored procedure automatically invokes the necessary access rights to all database objects used in the procedure. Therefore, administering database access rights may be greatly simplified by allowing access to critical data through procedures.

Stored procedures are registered and called using SQL statements.

There are three calling methods for the stored procedures:

- *Local procedures* are executed on a local database server.
- *Remote procedures* are procedures that are stored on one server and called by another. Remote stored procedures are applicable only to advanced replication setups.
- *Deferred procedures* are procedures that are called after commit has been processed.

There are also system stored procedures that you can use to simplify tasks related to administering advanced replication operations.

3.1.1 Stored procedures – SQL

This section explains in detail how to use SQL stored procedures. In the beginning of this section, the general concepts of using the procedures are explained. Later sections go more in-depth and describe the actual syntax of different statements in the procedures. The end of this section discusses transaction management, sequences and other advanced stored procedure features.

Basic procedure structure

A stored procedure is a standard solidDB database object that can be manipulated using standard DDL statements CREATE and DROP.

In its simplest form a stored procedure definition looks like:

```
"CREATE PROCEDURE procedure_name
parameter_section
BEGIN
declare_section_local_variables
procedure_body
END";
```

The following example creates a procedure called TEST:

```
"CREATE PROCEDURE test
BEGIN
END"
```

Procedures can be run by issuing a CALL statement followed by the name of the procedure to be invoked:

```
CALL test
```

Naming procedures

Procedure names have to be unique within a database schema.

All the standard naming restrictions applicable to database objects, like using reserved words and identifier lengths, apply to stored procedure names. For an overview and complete list of reserved words, see Appendix D, “Reserved words,” on page 331.

Parameter section

A stored procedure communicates with the calling program using parameters. solidDB supports two methods to return values to the calling program. The first method is the standard SQL-99 method, which uses parameters, and the other is a solidDB proprietary method, RETURNS, which uses result sets.

Using parameters: Parameters are the standard SQL-99 method to return data in stored procedures. The solidDB stored procedures accept three types of parameters:

- Input parameters (IN), which are used as input to the procedure. Parameters are input parameters by default. Thus, the keyword IN is optional in parameter definitions.
- Output parameters (OUT), which are returned values from the procedure.
- Input/output parameters (INOUT), which pass values into the procedure and return a value back to the calling procedure.

When you declare input parameters in the procedure heading, you can access their values inside the procedure by referring to the parameter name. The parameter data type must also be declared.

The parameter declaration in the stored procedure definition uses the following syntax

```
parameter_definition ::= [parameter_mode] parameter_name data_type
parameter_mode ::= IN | OUT | INOUT
```

There can be any number of parameters. When you call the procedure, you typically supply the input parameters in the same order as they are defined in the procedure.

You can specify default values to the parameters in the procedure definition. When you declare the parameter, add an equals character (=) and the default value after the parameter data type. For example:

```
"CREATE PROCEDURE participants( adults integer = 1,
children integer = '0',
pets integer = '0')
BEGIN
END"
```

When you call a procedure that has default values for parameters, you do not have to give values for all the parameters. To use default values for all parameters, you can omit the parameters. For example:

```
CALL participants()
```

To give a value to a parameter in the procedure call, you have two options:

Named parameters

Use the parameter name in the CALL statement and assign the parameter value by using the equals (=) character.

For example, to use default values for parameters adults and pets, and value 2 for children, use the following statement:

```
CALL participants(children = 2)
```

If a name is given to a parameter, all parameters following it must also have a name.

For example, the following statement returns an error because a parameter name is not defined for value 2, which is a parameter in the procedure definition that follows the parameter adults. The default value is used for parameter pets.

```
CALL participants(adults = 7,2)
```

The following statement succeeds because the unnamed parameter value 7 precedes the named parameter value children = 2. The default value is used for parameter pets.

```
CALL participants(7,children = 2)
```

Positional parameters

Supply the input parameters in the same order as they are defined in the procedure.

For example, to specify value 7 for adults, 3 for children, and 5 for pets, use the following statement:

```
CALL participants(7,3,5)
```

You can omit positional parameters only in reverse order (starting from the last positional parameter) and only if the parameters have default values.

For example, you can use the following statement to specify value 7 for adults and default values for children and *pets*:

```
CALL participants(7)
```

Related reference:

A.7, "CALL," on page 165

Related information:

Appendix C, "Data types," on page 323

Using RETURNS: You can use stored procedures to return a result set table with several rows of data in separate columns. This is a solidDB proprietary method to return data and it is performed by using the RETURNS structure.

When you use the RETURNS structure, you must separately declare result set column names for the output data rows. There can be any number of result set column names. The result set column names are declared in the RETURNS section of the procedure definition:

```
"CREATE PROCEDURE procedure_name
[ (IN input_param1 datatype [,
input_param2 datatype, ... ] ) ]
[ RETURNS
(output_column_definition1 datatype [,
output_column_definition2 datatype, ... ] ) ]
BEGIN
END";
```

By default, the procedure only returns one row of data containing the values as they were at the moment when the stored procedure was run or was forced to exit. However, it is also possible to return result sets from a procedure using the following syntax:

```
return row;
```

Every RETURN ROW call adds a new row into the returned result set where column values are the current values of the result set column names.

The following statement creates a procedure that has two input parameters and two result set column names for output rows:

```
"CREATE PROCEDURE PHONEBOOK_SEARCH
(IN FIRST_NAME VARCHAR, LAST_NAME VARCHAR)
RETURNS (PHONE_NR NUMERIC, CITY VARCHAR)
BEGIN
-- procedure_body
END";
```

This procedure should be called by using two input parameter of data type VARCHAR. The procedure returns an output table consisting of two columns named PHONE_NR of type NUMERIC and CITY of type VARCHAR.

For example:

```
call phonebook_search ('JOHN','DOE');
```

The result looks as follows (when the procedure body has been programmed):

```
PHONE_NR    CITY
3433555     NEW YORK
2345226     LOS ANGELES
```

The following statement creates a calculator procedure:

```
"create procedure calc(i1 float, op char(1),
  i2 float)
  returns (calcrest float)
begin
  declare i integer;

  if op = '+' then
    calcrest := i1 + i2;
  elseif op = '-' then
    calcrest := i1 - i2;
  elseif op = '*' then
    calcrest := i1 * i2;
  elseif op = '/' then
    calcrest := i1 / i2;
  else
    calcrest := 'Error: illegal op';
  end if
end";
```

You can test the calculator with the command:

```
call calc(1,'/',3);
```

With RETURNS, select statements can also be wrapped into database procedures. The following statement creates a procedure that uses a select statement to return backups created from the database:

```
"create procedure show_backups
  returns (backup_number varchar, date_created varchar)
begin
  -- First set action for failing statements.
  exec sql whenever sqlerror rollback, abort;

  -- Prepare and execute the select statement
  exec sql prepare sel_cursor select
    replace(property, 'backup ', ''),
    substring(value_str, 1, 19) from sys_info
    where property like 'backup %';
  exec sql execute sel_cursor into (backup_number, date_created);

  -- Fetch first row;
  exec sql fetch sel_cursor;
  -- Loop until end of table
  while sqlsuccess loop
  -- Return the fetched row
    return row;
  -- Fetch next
  exec sql fetch sel_cursor;
  end loop;
end";
```

Declare section

Local variables that are used inside the procedure for temporary storage of column and control values are defined in a separate section of the stored procedure directly following the BEGIN keyword.

The syntax of declaring a variable is:

```
DECLARE variable_name datatype;
```

Note that every declare statement should be ended with a semicolon (;).

The variable name is an alphanumeric string that identifies the variable. The data type of the variable can be any valid SQL data type supported. For supported data types, see Appendix C, “Data types,” on page 323.

For example:

```
"CREATE PROCEDURE PHONEBOOK_SEARCH
  (FIRST_NAME VARCHAR, LAST_NAME VARCHAR)
  RETURNS (PHONE_NR NUMERIC, CITY VARCHAR)
BEGIN
DECLARE i INTEGER;

DECLARE dat DATE;

END";
```

Note that input and output parameters are treated like local variables within a procedure with the exception that input parameters have a preset value and output parameter values are returned or can be appended to the returned result set.

Procedure body

The procedure body contains the actual stored procedure program based on assignments, expressions, and SQL statements.

Any type of expression, including scalar functions, can be used in a procedure body. For valid expressions, see A.79.3, “expression,” on page 306.

Assignments: To assign values to variables either of the following syntax is used:

```
SET variable_name = expression;
```

or

```
variable_name := expression;
```

Example:

```
SET i = i + 20 ;

i := 100;
```

Scalar functions with assignments: A scalar function is an operation denoted by a function name followed by a pair of parentheses enclosing zero or more specified arguments. Each scalar function returns one value. Note that scalar functions can be used with assignments, as in:

```
"CREATE PROCEDURE scalar_sample
  RETURNS (string_var VARCHAR(20))
BEGIN
-- CHAR(39) is the single quotation mark/apostrophe
string_var := 'Joe' + {fn CHAR (39)} + 's Garage';
END";
```

The result of this stored procedure is the output:

```
Joe's Garage
```

For a list of solidDB-supported scalar functions (SQL-92), see Appendix A, "IBM solidDB SQL statements," on page 141. Note that *solidDB Programmer Guide* contains an appendix that describes ODBC scalar functions, which contain some differences for SQL-92.

Variables, constants, and parameters in assignments: Variables and constants are initialized every time a procedure is executed. By default, variables are initialized to NULL. Unless a variable has been explicitly initialized, its value is NULL, as the following example shows:

```
BEGIN
DECLARE   total   INTEGER;
...
total := total + 1; -- assigns a null to total
...
```

Therefore, a variable should never be referenced before it has been assigned a value.

The expression following the assignment operator can be arbitrarily complex, but it must yield a data type that is the same as or convertible to the data type of the variable.

When possible, solidDB procedure language can provide conversion of data types implicitly. This makes it possible to use literals, variables and parameters of one type where another type is expected.

Implicit conversion is not possible if:

- information would be lost in the conversion
- a string to be converted to an integer contains non-numeric data

Examples:

```
DECLARE integer_var INTEGER;
integer_var := 'NR:123';
```

returns an error.

```
DECLARE string_var CHAR(3);
string_var := 123.45;
```

results in value '123' in variable string_var.

```
DECLARE string_var VARCHAR(2);
string_var := 123.45;
```

returns an error.

Single quotation marks and apostrophes in string assignments: Strings are delimited by single quotation marks. If you want to have a single quotation mark within a string, then you can put two single quotation marks (""), side by side, to produce one quotation mark in your output. This is commonly known as an "escape sequence." Following is a stored procedure that uses this technique:

```
"CREATE PROCEDURE q
RETURNS (string_var VARCHAR(20))
BEGIN
string_var := 'Joe''s Garage';
END";
CALL q;
```

The result is:

Joe's Garage

Here are some other examples:

'I'm writing.'

becomes:

I'm writing.

and

'Here are two single quotation marks:'''''

becomes:

Here are two single quotation marks:''

Note that in the last example there are five single quotation marks in a row at the end of the string. The last of these is the delimiter (the closing quotation mark); the preceding four are part of the data. The four quotation marks are treated as two pairs of quotation marks, and each pair of quotation marks is treated as an escape sequence representing one single quotation mark.

Expressions:

Comparison operators:

Comparison operators compare one expression to another. The result is always TRUE, FALSE, or NULL. Typically, comparisons are used in conditional control statements and allow comparisons of arbitrarily complex expressions.

Table 6. Comparison operators

Operator	Meaning
=	is equal to
<>	is not equal to
<	is less than
>	is greater than
<=	is less than or equal to
>=	is greater than or equal to
!=	is not equal to Note: != notation cannot be used inside a stored procedure, use the ANSI-SQL compliant <> instead.

Logical operators: The logical operators can be used to build more complex queries. The logical operators AND, OR, and NOT operate according to the tri-state logic illustrated by the truth tables shown below. AND and OR are binary operators; NOT is a unary operator.

Table 7. Logical operators: NOT

NOT	true	false	null
	false	true	null

Table 8. Logical operators: AND

AND	true	false	null
true	true	false	null
false	false	false	false
null	null	false	null

Table 9. Logical operators: OR

OR	true	false	null
true	true	true	true
false	true	false	null
null	true	null	null

As the truth tables show, AND returns the value TRUE only if both its operands are true. On the other hand, OR returns the value TRUE if either of its operands is true. NOT returns the opposite value (logical negation) of its operand. For example, NOT TRUE returns FALSE.

NOT NULL returns NULL because nulls are indeterminate.

When not using parentheses to specify the order of evaluation, operator precedence determines the order.

Note that 'true' and 'false' are not literals accepted by SQL parser but values. Logical expression value can be interpreted as a numeric variable:

false = 0 or NULL

true = 1 or any other numeric value

Example:

IF *expression* = TRUE THEN

can be simply written

IF *expression* THEN

IS NULL operator: The IS NULL operator returns the Boolean value TRUE if its operand is null, or FALSE if it is not null. Comparisons involving nulls always yield NULL. To test whether a value is NULL, do not use the expression,

IF *variable* = NULL THEN...

because it never evaluates to TRUE.

Instead, use the following statement:

```
IF variable IS NULL THEN...
```

Note that when using multiple logical operators in solidDB stored procedures the individual logical expressions should be enclosed in parentheses like:

```
((A >= B) AND (C = 2)) OR (A = 3)
```

Control structures: The following sections describe the statements that can be used in the procedure body, including branch and loop statements.

IF statement: Often, it is necessary to take alternative actions depending on circumstances. The IF statement executes a sequence of statements conditionally. There are three forms of IF statements: IF-THEN, IF-THEN-ELSE, and IF-THEN-ELSEIF.

IF-THEN: The simplest form of IF statement associates a condition with a statement list enclosed by the keywords THEN and END IF (not ENDIF), as follows:

```
IF condition THEN
  statement_list;
END IF
```

The sequence of statements is executed only if the condition evaluates to TRUE. If the condition evaluates to FALSE or NULL, the IF statement does nothing. In either case, control passes to the next statement. An example follows:

```
IF sales > quota THEN
  SET pay = pay + bonus;
END IF
```

IF-THEN-ELSE: The second form of IF statement adds the keyword ELSE followed by an alternative statement list, as follows:

```
IF condition THEN
  statement_list1;
ELSE
  statement_list2;
END IF
```

The statement list in the ELSE clause is executed only if the condition evaluates to FALSE or NULL. Thus, the ELSE clause ensures that a statement list is executed. In the following example, the first or second assignment statement is executed when the condition is true or false, respectively:

```
IF trans_type = 'CR' THEN
  SET balance = balance + credit;
ELSE
  SET balance = balance - debit;
END IF
```

THEN and ELSE clauses can include IF statements. That is, IF statements can be nested, as the following example shows:

```
IF trans_type = 'CR' THEN
  SET balance = balance + credit ;
ELSE
  IF balance >= minimum_balance THEN
    SET balance = balance - debit ;
  ELSE
    SET balance = minimum_balance;
  END IF
END IF
```


IF-THEN-ELSEIF: Occasionally it is necessary to select an action from several mutually exclusive alternatives. The third form of IF statement uses the keyword ELSEIF to introduce additional conditions, as follows:

```
IF condition1 THEN
  statement_list1;
ELSEIF condition2 THEN
  statement_list2;
ELSE
  statement_list3;
END IF
```

If the first condition evaluates to FALSE or NULL, the ELSEIF clause tests another condition. An IF statement can have any number of ELSEIF clauses; the final ELSE clause is optional. Conditions are evaluated one by one from top to bottom. If any condition evaluates to TRUE, its associated statement list is executed and the rest of the statements (inside the IF-THEN-ELSEIF) are skipped. If all conditions evaluate to FALSE or NULL, the sequence in the ELSE clause is executed. Consider the following example:

```
IF sales > 50000 THEN
  bonus := 1500;
ELSEIF sales > 35000 THEN
  bonus := 500;
ELSE
  bonus := 100;
END IF
```

If the value of "sales" is more than 50000, the first and second conditions are true. Nevertheless, "bonus" is assigned the proper value of 1500 since the second condition is never tested. When the first condition evaluates to TRUE, its associated statement is executed and control passes to the next statement following the IF-THEN-ELSEIF.

When possible, use the ELSEIF clause instead of nested IF statements. That way, the code will be easier to read and understand. Compare the following IF statements:

```
IF condition1 THEN
  statement_list1;
ELSE
  IF condition2 THEN
    statement_list2;
  ELSE
    IF condition3 THEN
      statement_list3;
    END IF
  END IF
END IF

IF condition1 THEN
  statement_list1;
ELSEIF condition2 THEN
  statement_list2;
ELSEIF condition3 THEN
  statement_list3;
END IF
```

These statements are logically equivalent, but the first statement obscures the flow of logic, whereas the second statement reveals it.

Use of parentheses in IF-THEN statements: Parentheses can be used in IF-THEN statements with the following rules:

- A single logical expression in an IF condition can use parentheses, but they are not required.

Example: parentheses used

```
IF (x > 0) THEN
  x := x - 1;
END IF;
```

Example: parentheses not used

```

IF x > 0 THEN
x := x - 1;
END IF;

```

- If there are multiple expressions inside a logical condition, parentheses are required around each subexpression.

```

IF (x > 0) AND (y > 0) THEN
x := x - 1;
END IF;

```

Example 1

This example shows use of valid logical conditions in IF statements.

```

"CREATE PROCEDURE sample_if_conditions
BEGIN
DECLARE x INT;
DECLARE y INT;
x := 2;
y := 2;

```

Example 2

This example shows use of additional parentheses around the entire expression.

```

IF ((x > 0) AND (y > 0)) THEN
x := x - 1;
END IF;

```

WHILE-LOOP: The WHILE-LOOP statement associates a condition with a sequence of statements enclosed by the keywords LOOP and END LOOP, as follows:

```

WHILE condition LOOP
    statement_list;
END LOOP

```

Before each iteration of the loop, the condition is evaluated. If the condition evaluates to TRUE, the statement list is executed, then control resumes at the top of the loop. If the condition evaluates to FALSE or NULL, the loop is bypassed and control passes to the next statement. An example follows:

```

WHILE total <= 25000 LOOP
    ...
    total := total + salary;
END LOOP

```

The number of iterations depends on the condition and is unknown until the loop completes. Since the condition is tested at the top of the loop, the sequence might execute zero times. In the latter example, if the initial value of "total" is greater than 25000, the condition evaluates to FALSE and the loop is bypassed altogether.

Loops can be nested. When an inner loop is finished, control is returned to the next loop. The procedure continues from the next statement after END LOOP.

Leaving loops: It may be necessary to force the procedure to leave a loop prematurely. This can be implemented using the LEAVE keyword:

```

WHILE total < 25000 LOOP
    total := total + salary;
    IF exit_condition THEN
        LEAVE;
    END IF
END LOOP
statement_list2

```

Upon successful evaluation of the *exit_condition* the loop is left, and the procedure continues at the *statement_list2*.

Note:

Although solidDB databases support the ANSI-SQL CASE syntax, the CASE construct cannot be used inside a stored procedure as a control structure.

Use of parentheses in WHILE loops: The following code illustrates the rules for using parentheses in WHILE loops. Refer also to the release notes for additional information about using parentheses in WHILE loops.

```
--- This piece of code shows examples of valid logical conditions in
--- WHILE loops.
```

```
"CREATE PROCEDURE sample_while_conditions
BEGIN
DECLARE x INT;
DECLARE y INT;
x := 2;
y := 2;
```

```
--- As shown below, a single logical expression in a WHILE condition
--- may use parentheses.
```

```
WHILE (x > 0) LOOP
x := x - 1;
END LOOP;
```

```
--- As shown below, although a single logical expression in a WHILE
--- condition may use parentheses, the parentheses are not required.
```

```
WHILE x > 0 LOOP
x := x - 1;
END LOOP;
```

```
--- As shown below, if there are multiple expressions inside a
--- logical condition, then you need parentheses around EACH
--- individual expression.
```

```
WHILE (x > 0) AND (y > 0) LOOP
x := x - 1;
y := y - 1;
END LOOP;
```

```
--- The example below is the same as the preceding example,
--- except that this has additional parentheses around the
--- entire expression.
```

```
WHILE ((x > 0) AND (y > 0)) LOOP
x := x - 1;
y := y - 1;
END LOOP;
```

NOT operator: Applying the logical operator NOT to a null yields NULL. Thus, the following two statements are not always equivalent:

IF x > y THEN	IF NOT (x > y) THEN
high := x;	high := y;
ELSE	ELSE
high := y;	high := x;
END IF	END IF

The sequence of statements in the ELSE clause is executed when the IF condition evaluates to FALSE or NULL. If either or both "x" and "y" are NULL, the first IF statement assigns the value of "y" to "high", but the second IF statement assigns the value of "x" to "high". If neither "x" nor "y" is NULL, both IF statements assign the corresponding value to "high".

Handling nulls: Nulls can cause confusing behavior. To avoid some common errors, observe the following rules:

- comparisons involving nulls always yield NULL
- applying the logical operator NOT to a null yields NULL
- in conditional control statements, if the condition evaluates to NULL, its associated sequence of statements is not executed

In the example below, you might expect the statement list to execute because "x" and "y" seem unequal. Remember though that nulls are indeterminate. Whether "x" is equal to "y" or not is unknown. Therefore, the IF condition evaluates to NULL and the statement list is bypassed.

```
x := 5;
  y := NULL;
  ...
  IF x <> y THEN -- evaluates to NULL, not TRUE
    statement_list; -- not executed
  END IF
```

In the next example, one might expect the statement list to execute because "a" and "b" seem equal. But, again, this is unknown, so the IF condition evaluates to NULL and the statement list is bypassed.

```
a := NULL;
  b := NULL;
  ...
  IF a = b THEN -- evaluates to NULL, not TRUE
    statement_list; -- not executed
  END IF
```

Zero-length strings: Zero length strings are treated by a solidDB server like they are a string of zero length, instead of a null. NULL values should be specifically assigned as in the following:

```
SET a = NULL;
```

This also means that checking for NULL values will return FALSE when applied to a zero-length string.

Exiting a procedure: A procedure may be exited prematurely by issuing the keyword

```
RETURN;
```

at any location. After this keyword, control is directly handed to the program calling the procedure returning the values bound to the result set column names as indicated in the RETURNS section of the procedure definition.

Returning data: You can return data with the OUT parameter mode, which is a standard SQL-99 method of returning data. This method allows you to pass data back to the program from the procedure. For syntax information, refer to Appendix A, "IBM solidDB SQL statements," on page 141.

The OUT parameter mode has the following characteristics:

- The OUT parameter mode allows you to pass data back to the calling program from the procedure. Inside the calling program, the OUT parameter acts like a variable. That means you can use an OUT parameter as if it were a local variable. You can change its value or reference the value in any way.
- The actual parameter that corresponds to an OUT parameter must be a variable; it cannot be a constant or an expression.

- Like variables, OUT parameters are initialized to NULL.

Before exiting a procedure, you must explicitly assign values to all OUT parameters. Otherwise, the corresponding actual parameters will be null. If you exit successfully, solidDB assigns values to the actual parameters. However, if you exit with an unhandled exception, solidDB does not assign values to the actual parameters.

For a solidDB proprietary method of returning data, see “Using RETURNS” on page 26.

Example stored procedure: Following is an example of a simple procedure that determines whether a person is an adult on the basis of a birthday as input parameter.

Note the usage of {fn ...} on scalar functions, and semicolons to end assignments.

```
"CREATE PROCEDURE grown_up
(birth_date DATE)
RETURNS (description VARCHAR)
BEGIN
DECLARE age INTEGER;
-- determine the number of years since the day of birth
age := {fn TIMESTAMPDIFF(SQL_TSI_YEAR, birth_date, now())};
IF age >= 18 THEN
-- If age is at least 18, then it's an adult
description := 'ADULT';
ELSE
-- otherwise it's still a minor
description := 'MINOR';
END IF
END";
```

SQL in a stored procedure: Using SQL statements inside a stored procedure is somewhat different from issuing SQL directly from tools like solidDB SQL Editor (**sql**) – a special syntax is required.

There are two ways to execute SQL statements inside a procedure:

- Use EXECDIRECT syntax to execute a statement.
- Treat the SQL statement as a "cursor".

EXECDIRECT: The EXECDIRECT syntax is particularly appropriate for statements where there is no result set, and where you do not have to use any variable to specify a parameter value. For example, the following statement inserts a single row of data:

```
EXEC SQL EXECDIRECT insert into table1 (id, name) values (1, 'Smith');
```

For more information, see A.15.11, “exec_direct_statement,” on page 188.

Using a cursor: Cursors are appropriate for statements where there is a result set, or where you want to repeat a single basic statement but use different values from a local variable as a parameter (for example, in a loop).

A cursor is a specific allocated part of the server process memory that keeps track of the statement being processed. Memory space is allocated for holding one row of the underlying statement, together with some status information about the current row (in SELECTS) or the number of rows affected by the statement (in UPDATES, INSERTS and DELETES).

In this way query results are processed one row at a time. The stored procedure logic should take care of the actual handling of the rows, and the positioning of the cursor on the required row(s).

There are five basic steps in handling a cursor:

1. Preparing the cursor - the definition
2. Executing the cursor - executing the statement
3. Fetching on the cursor (for select procedure calls) - getting the results row by row
4. Closing the cursor after use - still enabling it to re-execute
5. Dropping the cursor from memory - removing it

1. Preparing the cursor: A cursor is defined (prepared) using the following syntax:
`EXEC SQL PREPARE cursor_name SQL_statement;`

By preparing a cursor, memory space is allocated to accommodate one row of the result set of the statement, and the statement is parsed and optimized.

A cursor name given for the statement must be unique within the connection. This means procedures that contain cursors cannot be called recursively (at least not from a statement that is after a PREPARE CURSOR and before the corresponding DROP CURSOR). When a cursor is prepared, a solidDB server checks that no other cursor of this name is currently open. If there is one, error number 14504 is returned.

Note that statement cursors can also be opened using the ODBC API. These cursor names need to be different from the cursors opened from procedures.

Example:

```
EXEC SQL PREPARE sel_tables
SELECT table_name
FROM sys_tables
WHERE table_name LIKE 'SYS%';
```

This statement will prepare the cursor named *sel_tables*, but will not execute the statement that it contains.

2. Executing the cursor: After a statement has been successfully prepared it can be executed. An execute binds possible input and output variables to it and runs the actual statement.

Syntax of the execute statement is:

```
EXEC SQL EXECUTE cursor_name
[ INTO ( var1 [, var2...] ) ];
```

The optional section INTO binds result data of the statement to variables.

Variables listed in parentheses after the INTO keyword are used when running a SELECT or CALL statement. The resulting columns of the SELECT or CALL statement are bound to these variables when the statement is executed. The variables are bound starting from the leftmost column listed in the statement. Binding of variables continues to the following column until all variables in the list of variables have been bound. For example to extend the sequence for the cursor *sel_tables* that was prepared earlier we need to run the following statements:

```
EXEC SQL PREPARE sel_tables
  SELECT table_name
  FROM sys_tables
  WHERE table_name LIKE 'SYS%'
```

```
EXEC SQL EXECUTE sel_tables INTO (tab);
```

The statement is now executed and the resulting table names will be returned into variable `tab` in the subsequent Fetch statements.

3. Fetching on the cursor: When a SELECT or CALL statement has been prepared and executed, it is ready for fetching data from it. Other statements (UPDATE, INSERT, DELETE, DDL) do not require fetching as there will be no result set. Fetching results is done using the fetch syntax:

```
EXEC SQL FETCH cursor_name;
```

This command fetches a single row from the cursor to the variables that were bound with INTO keyword when the statement was executed.

To complete the previous example to actually get result rows back, the statements will look like:

```
EXEC SQL PREPARE sel_tables
  SELECT table_name
  FROM sys_tables
  WHERE table_name LIKE 'SYS%'
EXEC SQL EXECUTE sel_tables INTO (tab);
EXEC SQL FETCH sel_tables;
```

After this the variable `tab` will contain the table name of the first table found conforming to the WHERE clause.

Subsequent calls to fetch on the cursor `sel_tables` will get the next row(s) if the select found more than one.

To fetch all table names a loop construct may be used:

```
WHILE expression LOOP
  EXEC SQL FETCH sel_tables;
END LOOP
```

Note that after the completion of the loop, the variable `tab` will contain the last fetched table name.

4. Closing the cursor: Cursors may be closed by issuing the statement

```
EXEC SQL CLOSE cursor_name;
```

This will not remove the actual cursor definition from memory; it may be re-executed when the need arises.

5. Dropping the cursor: Cursors may be dropped from memory, releasing all resources by the statement:

```
EXEC SQL DROP cursor_name;
```

Parameter markers in cursors: In order to make a cursor more dynamic, a SQL statement can contain parameter markers that indicate values that are bound to the actual parameter values at execute time. The '?' symbol is used as a parameter marker.

Syntax example:

```
EXEC SQL PREPARE sel_tabs
SELECT table_name
FROM sys_tables
WHERE table_name LIKE ?
AND table_schema LIKE ?;
```

The execution statement is adapted by including a USING keyword to accommodate the binding of a variable to the parameter marker.

```
EXEC SQL EXECUTE sel_tabs USING ( var1, var2 ) INTO (tabs);
```

In this way a single cursor can be used multiple times without having to re-prepare the cursor. As preparing a cursor involves also the parsing and optimizing of the statement, significant performance gains can be achieved by using reusable cursors.

Note that the USING list only accepts variables; data can not be directly passed in this way. So if for example an insert into a table should be made, one column value of which should always be the same (status = 'NEW') then the following syntax would be wrong:

```
EXEC SQL EXECUTE ins_tab USING (nr, desc, dat, 'NEW');
```

The correct way would be to define the constant value in the prepare section:

```
EXEC SQL PREPARE ins_tab
INSERT INTO my_tab (id, descript, in_date, status)
VALUES (?, ?, ?, 'NEW');
EXEC SQL EXECUTE ins_tab USING (nr, desc, dat);
```

Note that variables can be used multiple times in the using list.

The parameters in an SQL statement have no intrinsic data type or explicit declaration. Therefore, parameter markers can be included in an SQL statement only if their data types can be inferred from another operand in the statement.

For example, in an arithmetic expression such as ? + COLUMN1, the data type of the parameter can be inferred from the data type of the named column represented by COLUMN1. A procedure cannot use a parameter marker if the data type cannot be determined.

The following table describes how a data type is determined for several types of parameters.

Table 10. Determining data type from parameters

Location of Parameter	Assumed Data Type
One operand of a binary arithmetic or comparison operator	Same as the other operand
The first operand in a BETWEEN clause	Same as the other operand
The second or third operand in a BETWEEN clause	Same as the first operand
An expression used with IN	Same as the first value or the result column of the subquery
A value used with IN	Same as the expression
A pattern value used with LIKE	VARCHAR

Table 10. Determining data type from parameters (continued)

Location of Parameter	Assumed Data Type
An update value used with UPDATE	Same as the update column

An application cannot place parameter markers in the following locations:

- As an SQL identifier (name of a table, name of a column and so on)
- In a SELECT list.
- As both expressions in a comparison-predicate.
- As both operands of a binary operator.
- As both the first and second operands of a BETWEEN operation.
- As both the first and third operands of a BETWEEN operation.
- As both the expression and the first value of an IN operation.
- As the operand of a unary + or - operation.
- As the argument of a set-function-reference.

For more information, see the ANSI SQL-92 specification.

In the following example, a stored procedure will read rows from one table and insert parts of them in another, using multiple cursors:

```
"CREATE PROCEDURE tabs_in_schema (schema_nm VARCHAR)
RETURNS (nr_of_rows INTEGER)
BEGIN
DECLARE tab_nm VARCHAR;
EXEC SQL PREPARE sel_tab
SELECT table_name
FROM sys_tables
WHERE table_schema = ?;
EXEC SQL PREPARE ins_tab
INSERT INTO my_table (table_name, schema) VALUES (?,?);

nr_of_rows := 0;

EXEC SQL EXECUTE sel_tab USING (schema_nm) INTO (tab_nm);
EXEC SQL FETCH sel_tab;
WHILE SQLSUCCESS LOOP
nr_of_rows := nr_of_rows + 1;
EXEC SQL EXECUTE ins_tab USING (tab_nm, schema_nm);
IF SQLROWCOUNT <> 1 THEN
RETURN SQLERROR OF ins_tab;
END IF
EXEC SQL FETCH sel_tab;
END LOOP
END";
```

Error handling:

SQLSUCCESS: The return value of the latest EXEC SQL statement executed inside a procedure body is stored into variable SQLSUCCESS. This variable is automatically generated for every procedure. If the previous SQL statement was successful, the value 1 is stored into SQLSUCCESS. After a failed SQL statement, a value 0 is stored into SQLSUCCESS.

The value of SQLSUCCESS may be used, for instance, to determine when the cursor has reached the end of the result set as in the following example:

```

EXEC SQL FETCH sel_tab;
-- loop as long as last statement in loop is successful
WHILE SQLSUCCESS LOOP
    -- do something with the results, for example, return a row
    EXEC SQL FETCH sel_tab;

END LOOP

```

SQLERRNUM: This variable contains the error code of the latest SQL statement executed. It is automatically generated for every procedure. After successful execution, SQLERRNUM contains zero (0).

SQLERRSTR: This variable contains the error string from the last failed SQL statement.

SQLROWCOUNT: After the execution of UPDATE, INSERT and DELETE statements, an additional variable is available to check the result of the statement. Variable SQLROWCOUNT contains the number of rows affected by the last statement.

SQLERROR: To generate user errors from procedures, the SQLERROR variable may be used to return an actual error string that caused the statement to fail to the calling application. The syntax is:

```

RETURN SQLERROR 'error string'
RETURN SQLERROR char_variable

```

The error is returned in the following format:

User error: *error_string*

SQLERROR OF cursorname: For error checking of EXEC SQL statements, the SQLSUCCESS variable may be used as described under SQLSUCCESS in the beginning of this section. To return the actual error that caused the statement to fail to the calling application, the following syntax may be used:

```

EXEC SQL PREPARE cursorname sql_statement;
EXEC SQL EXECUTE cursorname;
IF NOT SQLSUCCESS THEN
    RETURN SQLERROR OF cursorname;
END IF
END IF

```

Processing will stop immediately when this statement is executed and the procedure return code is SQLERROR. The actual database error can be returned using the SQLError function:

Solid Database error 10033: Primary key unique constraint violation

The generic error handling method for a procedure can be declared with:

```

EXEC SQL WHENEVER SQLERROR [ROLLBACK [WORK],] ABORT;

```

When this statement is included in a stored procedure, all return values of executed SQL statements are checked for errors. If a statement execution returns an error, the procedure is automatically aborted and SQLERROR of the last cursor is returned. Optionally the transaction can also be rolled back.

The statement should be included before any EXEC SQL statements, directly following the DECLARE section of variables.

Below is an example of a complete procedure returning all table names from SYS_TABLES that start with 'SYS':

```
"CREATE PROCEDURE sys_tabs
RETURNS (tab VARCHAR)
BEGIN
-- abort on errors
EXEC SQL WHENEVER SQLERROR ROLLBACK, ABORT;
-- prepare the cursor
EXEC SQL PREPARE sel_tables
SELECT table_name
FROM sys_tables
WHERE table_name LIKE 'SYS%';
-- execute the cursor
EXEC SQL EXECUTE sel_tables INTO (tab);
-- loop through rows
EXEC SQL FETCH sel_tables;
WHILE sqlsuccess LOOP
RETURN ROW;
EXEC SQL FETCH sel_tables;
END LOOP
-- close and drop the used cursors
EXEC SQL CLOSE sel_tables;
EXEC SQL DROP sel_tables;
END";
```

Example: stored procedure with EXECDIRECT and cursors: The following stored procedure uses EXECDIRECT in one place and uses a cursor in another place.

```
"CREATE PROCEDURE p2
BEGIN

-- This variable holds an ID that we insert into the table.
DECLARE id INT;

-- Here are simple examples of EXECDIRECT.
EXEC SQL EXECDIRECT create table table1 (id_col INT);
EXEC SQL EXECDIRECT insert into table1 (id_col) values (1);

-- Here is an example of a cursor.
EXEC SQL PREPARE cursor1 INSERT INTO table1 (id_col) values (?);
id := 2;
WHILE id <= 10 LOOP
EXEC SQL EXECUTE cursor1 USING (id);
id := id + 1;
END LOOP;
EXEC SQL CLOSE cursor1;
EXEC SQL DROP cursor1;

END";
```

Calling other procedures

As calling a procedure forms a part of the supported SQL syntax, a stored procedure may be called from within another stored procedure. The default limit for levels of nested procedures is 16. When the maximum is exceeded, the transaction fails. The maximum nesting level is set in the MaxNestedProcedures parameter in the solid.ini configuration file. For details, see appendix "Configuration Parameters" in *solidDB Administration Guide*.

Like all SQL statements, a cursor should be prepared and executed like:

```
EXEC SQL PREPARE cp CALL myproc(?, ?);
EXEC SQL EXECUTE cp USING (var1, var2);
```

If procedure *myproc* returns one or more values, then subsequently a fetch should be done on the cursor *cp* to retrieve those values:

```
EXEC SQL PREPARE cp call myproc(?,?);
EXEC SQL EXECUTE cp USING (var1, var2) INTO
(ret_var1, ret_var2);
EXEC SQL FETCH cp;
```

Note that if the called procedure uses a *return row* statement, the calling procedure should utilize a WHILE LOOP construct to fetch all results.

Recursive calls are possible, but discouraged because cursor names are unique at connection level.

Positioned updates and deletes: In solidDB procedures it is possible to use positioned updates and deletes. This means that an update or delete will be done to a row, where a given cursor is currently positioned. The positioned updates and deletes can also be used within stored procedures using the cursor names used within the procedure.

The following syntax is used for positioned updates:

```
UPDATE table_name
SET column = value
WHERE CURRENT OF cursor_name
```

and for deletes

```
DELETE FROM table_name
WHERE CURRENT OF cursor_name
```

In both cases the *cursor_name* refers to a statement doing a SELECT on the table that is to be updated/deleted from.

Positioned cursor update is a semantically suspicious concept in SQL standard that may cause peculiarities also with a solidDB server.

Below is an example written with pseudo code that will cause an endless loop with a solidDB server (error handling, binding variables and other important tasks omitted for brevity and clarity):

```
"CREATE PROCEDURE ENDLESS_LOOP
BEGIN
EXEC SQL PREPARE MYCURSOR SELECT * FROM TABLE1;
EXEC SQL PREPARE MYCURSOR_UPDATE
  UPDATE TABLE1 SET COLUMN2 = 'new data';
  WHERE CURRENT OF MYCURSOR;"
EXEC SQL EXECUTE MYCURSOR;
EXEC SQL FETCH MYCURSOR;
WHILE SQLSUCCESS LOOP
  EXEC SQL EXECUTE MYCURSOR_UPDATE;
  EXEC SQL COMMIT WORK;
  EXEC SQL FETCH MYCURSOR;
END LOOP
END";
```

The endless loop is caused by the fact that when the update is committed, a new version of the row becomes visible in the cursor and it is accessed in the next FETCH statement. This happens because the incremented row version number is included in the key value and the cursor finds the changed row as the next greater key value after the current position. The row gets updated again, the key value is changed and again it will be the next row found.

In the above example, the updated COLUMN2 is not assumed to be part of the primary key for the table, and the row version number was the only part of the

index entry that changed. However, if a column value is changed that is part of the index through which the cursor has searched the data, the changed row may jump further forward or backward in the search set.

For these reasons, using positioned update is not recommended in general and searched update should be used instead whenever possible. However, sometimes the update logic may be too complex to be expressed in SQL WHERE clause and in such cases positioned update can be used as follows:

Positioned cursor update works deterministically in solidDB, when the WHERE clause is such that the updated row does not match the criteria and therefore does not reappear in the fetch loop. Constructing such a search criteria may require using additional column only for this purpose.

Note that in an open cursor, user changes do not become visible unless they are committed within the same database session.

Transactions: Stored procedures use transactions like any other interface to the database uses transactions. A transaction may be committed or rolled back either inside the procedure or outside the procedure. Inside the procedure a commit or roll back is done using the following syntax:

```
EXEC SQL COMMIT WORK;  
EXEC SQL ROLLBACK WORK;
```

These statements end the previous transaction and start a new one.

If a transaction is not committed inside the procedure, it may be ended externally using:

- solidDB SA
- Another stored procedure
- By autocommit, if the connection has AUTOCOMMIT switch set to ON

When a connection has autocommit activated, it does not force autocommit inside a procedure. The commit is done when the procedure exits.

Default cursor management: By default, when a procedure exits, all cursors opened in a procedure are closed. Closing cursors means that cursors are left in a prepared state and can be re-executed.

After exiting, the procedure is put in the procedure cache. When the procedure is dropped from the cache, all cursors are finally dropped.

The number of procedures kept in cache is determined by the `solid.ini` file setting:

```
[SQL]  
ProcedureCache = nbr_of_procedures
```

This means that, as long as the procedure is in the procedure cache, all cursors can be reused as long as they are not dropped. A solidDB server itself manages the procedure cache by keeping track of the cursors declared, and notices if the statement a cursor contains has been prepared.

As cursor management, especially in a heavy multi-user environment, can use a considerable amount of server resources, it is good practice to always close cursors

immediately and preferably also drop all cursors that are no longer used. Only the most frequently used cursors may be left non-dropped to reduce the cursor preparation effort.

Note that transactions are not related to procedures or other statements. Commit or rollback therefore does NOT release any resources in a procedure.

Notes on SQL:

- There is no restriction on the SQL statements used. Any valid SQL statement can be used inside a stored procedure, including DDL and DML statements.
- Cursors may be declared anywhere in a stored procedure. Cursors that are certainly going to be used are best prepared directly following the declare section.
- Cursors that are used inside control structures, and are therefore not always necessary, are best declared at the point where they are activated, to limit the amount of open cursors and hence the memory usage.
- The cursor name is an undeclared identifier, not a variable; it is used only to reference the query. You cannot assign values to a cursor name or use it in an expression.
- Cursors may be re-executed repeatedly without having to re-prepare them. Note that this can have a serious influence on performance; repetitively preparing cursors on similar statements may decrease the performance by around 40% in comparison to re-executing already prepared cursors!
- Any SQL statement will have to be preceded by the keywords EXEC SQL.

Functions for procedure stack viewing: The following functions may be included in stored procedures to analyze the current contents of the procedure stack:

- PROC_COUNT ()
This function returns the number of procedures in the procedure stack, including the current procedure.
- PROC_NAME (N)
This function returns the Nth procedure name in the stack. The first procedure is in position zero.
- PROC_SCHEMA (N)
This function returns the schema name of the Nth procedure in the procedure stack.

These functions allow for stored procedures that behave differently depending on whether they are called from an application or from a procedure.

Procedure privileges

Stored procedures are owned by the creator, and are part of the creator's schema. Users who need to run stored procedures in other schemas need to be granted EXECUTE privilege on the procedure:

```
GRANT EXECUTE ON procedure_name TO { USER | ROLE };
```

This function returns the schema name of the Nth procedure in the procedure stack.

All database objects accessed within the granted procedure, even subsequently called procedures, are accessed according to the rights of the owner of the procedure. No special grants are necessary.

Since the procedure is run with the privileges of the creator, the procedure not only has the creator's rights to access objects such as tables, but also uses the creator's schema and catalog. For example, suppose that user 'Sally' runs a procedure named 'Proc1' created by user 'Jasmine'. Suppose also that both Sally and Jasmine have a table named 'table1'. By default, the stored procedure Proc1 will use the table1 that is in Jasmine's schema, even if Proc1 was called by user Sally.

See also "ACCESS RIGHTS" on page 48 for more information about privileges and remote stored procedure calls.

Remote stored procedures in advanced replication configurations

In advanced replication configurations, stored procedures can be called locally or remotely. Remote procedures are stored procedures that one database server calls from another database server. Remote stored procedure calls use the following syntax:

```
CALL procedure_name AT node-ref;
```

where

node-ref indicates which database server the remote stored procedure is on

Remote stored procedures calls can only be made between two solidDB servers that have a master/replica relationship. The calls can be made in either direction; the master may call a stored procedure on the replica, or the replica may call a stored procedure on the master.

A remote stored procedure can be called from any context that allows a local procedure call. Thus, for example, you can call a remote stored procedure directly by using a CALL statement, or you can call the remote procedure from within a trigger, another stored procedure, or a START AFTER COMMIT statement.

A remotely-called stored procedure can contain any command that any other stored procedure may contain. All stored procedures are created using the same syntax rules. A single stored procedure may be called both locally and remotely at different times.

When called remotely, the stored procedure accepts parameters from the caller as if the call was local. However, a remote stored procedure cannot return a result set; it can only return an error code.

Both local and remote stored procedure calls are synchronous; regardless of whether the procedure is called locally or remotely, the caller waits until the value is returned; the caller does not continue on while the stored procedure executes in the background. If the stored procedure is called from inside a START AFTER COMMIT, the stored procedure call itself is synchronous, but the START AFTER COMMIT was not synchronous, so the stored procedure will execute as an asynchronous background process.

Transaction handling for remote stored procedures is different from transaction handling for local stored procedures. When a stored procedure is called remotely, the execution of the stored procedure is not a part of the transaction that contained the call. Therefore, you cannot roll back a stored procedure call by rolling back the transaction that called it.

CALL syntax for remote stored procedures

The full syntax of the command to call a remote stored procedure is:

```
CALL <proc-name>[(param [, param...])] AT node-def;  
node-def ::= DEFAULT | 'replica name' | 'master name'
```

DEFAULT is used only with the START AFTER COMMIT statement.

For example:

```
CALL MyProc('Smith', 750) AT replica1;  
CALL MyProcWithoutParameters AT replica2;
```

Note: You can only list one node definition per CALL. If you want to notify multiple replicas, for example, then you must call each of them separately. You can, however, create a stored procedure that contains multiple CALL statements, and then simply make a single call to that procedure.

Creating remote stored procedures

The remote stored procedure is always created on the server that executes the procedure, not on the server that calls the procedure. For example, if the master is going to call procedure foo() to execute on replica1, the procedure foo() must have been created on replica1. The master does not know the content of the stored procedure that it calls remotely. In fact, the master does not know anything about the stored procedure other than the information specified in the CALL statement itself.

For example, the following CALL statement includes the procedure's name, some parameter values, and the name of the replica on which the procedure is to be executed:

```
CALL foo(param1, param2) AT replica1
```

The stored procedure is not registered with the caller. This means that the caller calls the procedure blindly, without even knowing if it exists. If the caller tries to call a procedure that does not exist, the caller will get an error message.

Dynamic parameter binding is supported. For example, the following is legal:

```
CALL MYPROC(?, ?) AT MYREPLICA1;
```

Calls to the stored procedure are not buffered or queued. If you call the stored procedure and the procedure does not exist, the call does not persist, waiting until the stored procedure appears. Similarly, if the procedure exists, but the server that has that procedure is shut down or is disconnected from the network or inaccessible for any other reason, the call is not held open and retried when the server becomes accessible again. This is important to know when using the "Sync Pull Notify" (push synchronization) feature.

ACCESS RIGHTS: To call a stored procedure, the caller must have EXECUTE privilege on that procedure. (This is true for any stored procedure, whether it is called locally or remotely.)

When a procedure is called locally, it is executed with the privileges of the caller. When a procedure is called remotely, it may be executed either with the privileges of a specified user on the remote server, or with the privileges of the remote user who corresponds to the local caller. (The replica and master users must already be mapped to each other before the stored procedure is called. For more information about mapping replica users to master users, see *IBM solidDB Advanced Replication User Guide*.)

If a remote stored procedure was called from the replica (and is to be executed on the master), then you have the option of specifying which master user's privileges you would like the procedure to be executed with.

If the remote stored procedure was called from the master (and is to be executed on the replica), or if you do not specify which user's privileges to use, then the calling server will figure out which user's privileges should be used, based on which user called the stored procedure and the mapping between replica and master users.

These possibilities are explained in more detail below.

1. If the procedure was called from a replica (and will be executed on the master), then you may execute the SET SYNC USER statement to specify which master user's privileges to use. You must execute SET SYNC USER on the local server before calling the remote stored procedure. Once the sync user has been specified on the calling server, the calling server will send the user name and password to the remote server (the master server) each time a remote stored procedure is called. The remote server will try to execute the procedure using the user id and password that were sent with the procedure call. The user id and password must exist in the remote server, and the specified user must have appropriate access rights to the database and EXECUTE privilege on the called procedure.

The SET SYNC USER statement is valid only on a replica, so you can only specify the sync user when a replica calls a stored procedure on a master.

2. If the caller is a master, or if the call was made from a replica and you did not specify a sync user before the call, then the servers will attempt to determine which user on the remote server corresponds to the user on the local server.

If the calling server is a replica (R → M)

The calling server sends the following information to the remote server when calling a remote procedure:

Name of the master (SYS_SYNC_MASTERS.NAME).

Replica id (SYS_SYNC_MASTERS.REPLICA_ID).

Master user id (This master user id is the master user id that corresponds to the user id of the local user who called the procedure. Obviously, this local user must already be mapped to the corresponding master user.)

Note that this method of selecting the master user id is the same as the method used when a replica refreshes data — the replica looks up in the SYS_SYNC_USERS table to find the master user who is mapped to the current local replica user.

If the calling server is a master (M → R)

The calling server sends the following information to the remote server when calling a remote procedure:

Name of the master (SYS_SYNC_REPLICAS.MASTER_NAME).

Replica id (SYS_SYNC_REPLICAS.ID).

User name of the caller.

User id of the caller.

When the replica receives the master user id, the replica looks up the local user who is mapped to that master id. Since more than one replica user may be mapped to a single master user, the server will use the first local user it finds who is mapped to the specified master user and who has the privileges required to execute this stored procedure.

Before a master server can call a stored procedure on a replica server, the master must know the connect string of the replica. If a replica allows calls from a master, then the replica should define its own connect string information in the `solid.ini` file. This information is provided to the master (the replica includes a copy when it forwards any message to master). When the master receives the connect string from the replica, the master replaces the previous value (if the new value differs).

Example:

```
[Synchronizer]
ConnectStrForMaster=tcp replicahost 1316
```

It is also possible to inform the master of the replica's connect string by using the statement:

```
SET SYNC CONNECT <connect-info> TO REPLICA <replica-name>
```

This is useful if the master needs to call the replica but the replica has not yet provided its connect string to the master (i.e. has not yet forwarded any message to the master).

Deferred procedures in advanced replication configurations

Deferred procedures are stored procedures that are called after commit has been processed.

In advanced replication configurations, at the end of a committed transaction, you might want to perform a specific action. For example, if the transaction updated some data in a "master" publication, you might want to notify a replica that the master data was updated. solidDB allows the `START AFTER COMMIT` statement to specify an SQL statement that will be executed when the current transaction is committed. The specified SQL statement is called the "body" of the `START AFTER COMMIT`. The body is executed asynchronously in a separate connection.

For example, if you would like to call a stored procedure named `my_proc()` when the transaction commits, you would write:

```
START AFTER COMMIT NONUNIQUE CALL
    my_proc;
```

This statement may appear anywhere inside the transaction; it may be the first statement, the last statement, or any statement in between. Regardless of where the `START AFTER COMMIT` statement itself appears within the transaction, the body (the call to `my_proc`) will be executed only when the transaction is committed. In the example above, the body is on a separate line, but that is not required syntactically.

Because the body of the statement is not executed at the same time as the `START AFTER COMMIT` statement itself, there are two different phases to the `START AFTER COMMIT` command: the *definition* phase and the *execution* phase.

- In the definition phase of `START AFTER COMMIT`, you specify the body but do not execute it. The definition phase can occur anywhere inside a transaction; in other words, the statement "`START AFTER COMMIT ...`" can be placed in any order relative to other SQL statements in the same transaction.
- In the execution phase, the body of the `START AFTER COMMIT` statement is actually executed. The execution phase occurs when the `COMMIT WORK` statement for the transaction is executed. It is also possible to execute a `START AFTER COMMIT` in autocommit mode, but there is rarely a reason to do this.

The following example shows the use of a `START AFTER COMMIT` statement inside a transaction.

```
-- Any valid SQL statement(s)...  
...  
-- Creation phase. The function my_proc() is not actually called here.  
START AFTER COMMIT NONUNIQUE CALL my_proc(x, y);  
...  
-- Any valid SQL statement(s)...  
  
-- Execution phase: This ends the transaction and starts execution  
-- of the call to my_proc().  
COMMIT WORK;
```

A `START AFTER COMMIT` does not execute unless and until the transaction is successfully committed. If the transaction containing the `START AFTER COMMIT` is rolled back, the body of the `START AFTER COMMIT` is not executed. If you want to propagate the updated data from a replica to a master, this is an advantage because you only want the data propagated if it is committed. If you were to use triggers to start the propagation, the data would be propagated before it was committed.

The `START AFTER COMMIT` command applies only to the current transaction, that is, the one that the `START AFTER COMMIT` command was issued inside. It does not apply to subsequent transactions, or to any other transactions that are currently open in other connections.

The `START AFTER COMMIT` command allows you to specify only one SQL statement to be executed when the `COMMIT` occurs. However, that one SQL statement may be a call to a stored procedure, and that stored procedure can have many statements, including calls to other stored procedures. Furthermore, you can have more than one `START AFTER COMMIT` command per transaction. The body of each of these `START AFTER COMMIT` statements will be executed when the transaction is committed. However, these bodies will run independently and asynchronously; they will not necessarily execute in the same order as their corresponding `START AFTER COMMIT` statements, and they are likely to have overlapping execution (there is no guarantee that one will finish before the next one starts).

A common use of `START AFTER COMMIT` is to help implement "Sync Pull Notify" ("Push Synchronization") functionality in advanced replication setups.

If the body of your `START AFTER COMMIT` is a call to a stored procedure, that procedure can be local or it may be remote on one remote replica (or master).

If you are using Sync Pull Notify, you might want to call the same procedure on many replicas. To do this, you must use a slightly indirect method. The simplest method is to write one local procedure that calls many procedures on replicas. For example, if the body of the `START AFTER COMMIT` statement is "`CALL my_proc`", you could write `my_proc` to be similar to the following:

```
CREATE PROCEDURE my_proc  
BEGIN  
CALL update_inventory(x) AT replica1;  
CALL update_inventory(x) AT replica2;  
CALL update_inventory(x) AT replica3;  
END;
```

This approach works fine if your list of replicas is static. However, if you expect to add new replicas in the future, you might find it more convenient to update

"groups" of replicas based on their properties. This allows you to add new replicas with specific properties and then have existing stored procedures operate on those new replicas. This is done by making use of two features: the FOR EACH REPLICA clause in START AFTER COMMIT, and the DEFAULT clause in remote stored procedure calls.

If the FOR EACH REPLICA clause is used in START AFTER COMMIT, the statement will be executed once for each replica that meets the conditions in the WHERE clause. The statement is executed once FOR each replica, not once ON each replica. If there is no "AT node-ref" clause in the CALL statement, the stored procedure is called locally, that is, on the same server as the START AFTER COMMIT was executed on. To make sure that a stored procedure is called once on each replica, you must use the DEFAULT clause. The typical way to do this is to create a local stored procedure that contains a remote procedure calling that uses the DEFAULT clause.

For example, suppose that my_local_proc contains the following:
CALL update_sales_statistics AT DEFAULT;

The START AFTER COMMIT statement is the following:

```
START AFTER COMMIT FOR EACH REPLICA
WHERE region = 'north'
UNIQUE
CALL my_local_proc;
```

The WHERE clause is the following:

```
WHERE region = 'north'
```

Therefore, for each replica that has the properties region = 'north', you will call the stored procedure named my_local_proc.

That local procedure, in turn, executes the following:

```
CALL update_sales_statistics() AT DEFAULT
```

The keyword DEFAULT is resolved as the name of the replica. Each time that my_local_proc is called from inside the body of the START AFTER COMMIT, the DEFAULT keyword is the name of a different replica that has the property "region = 'north'".

Note: In this example:

- It is possible that not all replicas will have a procedure named update_sales_statistics(). If this is the case, the procedure will only be executed on those replicas that have the procedure. (The master will not send each replica a copy of the procedure; the master only calls existing procedures.)
- It is possible that not all replicas that have a procedure named update_sales_statistics() will have the SAME procedure. Each replica may have its own custom version of the procedure.

Before executing each statement on each replica, a connection to the replica is established.

When the START AFTER COMMIT command is used to call multiple replicas, this enables the use of the optional keyword "DEFAULT" in the syntax of the CALL command. For example, suppose that you use the following:

```

START AFTER COMMIT
  FOR EACH REPLICAS
  WHERE location = 'India'
  UNIQUE CALL push;

```

Then in the local procedure 'push', you can use the keyword "DEFAULT", which acts as a variable that contains the name of the replica in question.

```

CREATE PROCEDURE push
BEGIN
EXEC SQL EXECDIRECT CALL remoteproc AT DEFAULT;
END

```

Procedure 'push' will be called once for each replica that has a property named 'location' with value 'India'. Each time the procedure is called, "DEFAULT" will be set to the name of that replica. Thus, the following statement will call the procedure on that particular replica:

```
CALL remoteproc AT DEFAULT;
```

You can set the replica properties in the master with the statement:

```
SET SYNC PROPERTY propname = 'value' FOR REPLICAS replica_name;
```

For example

```
SET SYNC PROPERTY location = 'India' FOR REPLICAS asia_hq;
```

The statement specified in START AFTER COMMIT is executed as an independent transaction. It is not part of the transaction that contained the START AFTER COMMIT command. This independent transaction is run as though autocommit mode were on; in other words, you do not need an explicit COMMIT WORK to commit the work done in this statement.

In other respects, however, the execution of the statement is not much like a transaction. First, there is no guarantee that the statement will execute to completion. The statement is launched as an independent background task. If the server crashes, or if for some other reason the statement cannot be executed, then the statement disappears without being completely executed.

Second, because the statement is executed as a background task, there is no mechanism for returning an error. Third, there is no way to roll back the statement; if the statement execution is completed, the "transaction" statement is autocommitted regardless of whether any errors were detected. If the statement is a procedure call, then the procedure itself may contain COMMIT and ROLLBACK commands.

You can use the "RETRY" clause to try executing the statement more than once if it fails. The RETRY clause allows you to specify the number of times the server should attempt to retry the failed statement. You must specify the number of seconds to wait between each retry.

If you do not use the RETRY clause, the server attempts only once execute the statement, then the statement is discarded. If, for example, the statement tries to call a remote procedure, and if the remote server is down (or cannot be contacted due to a network problem), then the statement will not be executed and you will not get any error message.

Any statement, including the statement specified in a START AFTER COMMIT, executes in a certain "context". The context includes such factors as the default

catalog, the default schema, and so on. For a statement executed from within a START AFTER COMMIT, the statement's context is based on the context at the time that the START AFTER COMMIT is executed, not on the context at the time of the COMMIT WORK that actually causes the statement inside START AFTER COMMIT to run. In the example below, 'CALL FOO_PROC' is executed in the catalog foo_cat and schema foo_schema, not bar_cat and bar_schema.

```
SET CATALOG FOO_CAT;  
SET SCHEMA FOO_SCHEMA;  
START AFTER COMMIT UNIQUE CALL FOO_PROC;  
...  
SET CATALOG BAR_CAT;  
SET SCHEMA BAR_SCHEMA;  
COMMIT WORK;
```

The UNIQUE/NONUNIQUE keywords determine whether the server tries to avoid issuing the same command twice.

The UNIQUE keyword before <stmt> defines that the statement is executed only if there isn't identical statement under execution or "pending" for execution. Statements are compared with simple string compare. So for example 'call foo(1)' is different from 'call foo(2)'. Replicas are also taken into account in the comparison; in other words, UNIQUE does not prevent the server from executing the same trigger call on different replicas. Note that "unique" only blocks overlapping execution of statements; it does not prevent the same statement from being executed again later if it is called again after the current invocation has finished running.

NONUNIQUE means that duplicate statements can be executed simultaneously in the background.

Examples: The following statements are all considered different and are thus executed even though each contains the UNIQUE keyword. (Name is a unique property of replica.)

```
START AFTER COMMIT UNIQUE call myproc;  
START AFTER COMMIT FOR EACH REPLICa WHERE name='R1' UNIQUE call myproc;  
START AFTER COMMIT FOR EACH REPLICa WHERE name='R2' UNIQUE call myproc;  
START AFTER COMMIT FOR EACH REPLICa WHERE name='R3' UNIQUE call myproc;
```

But if the following statement is executed in the same transaction as the previous ones and if some of the replicas R1, R2, and R3 have the property "color='blue'", then the call is not executed for those replicas again.

```
START AFTER COMMIT FOR EACH REPLICa WHERE color='blue'  
UNIQUE call myproc;
```

Uniqueness does not prevent "automatic" execution from overlapping "manual" execution. For example, if you manually execute a command to refresh from a particular publication, and if the master also calls a remote stored procedure to refresh from that publication, the master will not "skip" the call because a manual refresh is already running. Uniqueness applies only to statements started by START AFTER COMMIT.

The START AFTER COMMIT statement can be used inside a stored procedure. For example, suppose that you want to post an event if and only if a transaction completed successfully. You could write a stored procedure that would execute a START AFTER COMMIT statement that would post the event if the transaction was committed (but not if it was rolled back). Your code might look similar to the following:

This sample also contains an example of "receiving" and then using an event parameter. See the stored procedure named "wait_on_event_e" in script #1.

```
-- To run this demo properly, you will need two users/connections.
-- This demo contains 5 separate "scripts", which must be executed
-- in the order shown below:
--     User1 executes the first script.
--     User2 executes the second script.
--     User1 executes the third script.
--     User2 executes the fourth script.
--     User1 executes the fifth script.
-- You may notice that there are some COMMIT WORK statements
-- in surprising places. These are to ensure that each user sees the
-- most recent changes of the other user. Without the COMMIT WORK
-- statements, in some cases one user would see an out-of-date
-- "snapshot" of the database.
--
-- Set autocommit off for both users/connections!
```

```
----- SCRIPT 1 (USER 1) -----
CREATE EVENT e (i int);
CREATE TABLE table1 (a int);

-- This inserts a row into table1. The value inserted into the is copied
-- from the parameter to the procedure.
"CREATE PROCEDURE inserter(i integer)
BEGIN
EXEC SQL PREPARE c_inserter INSERT INTO table1 (a) VALUES (?);
EXEC SQL EXECUTE c_inserter USING (i);
EXEC SQL CLOSE c_inserter;
EXEC SQL DROP c_inserter;
END";

-- This posts the event named "e".
"CREATE PROCEDURE post_event(i integer)
BEGIN
POST EVENT e(i);
END";

-- This demonstrates the use of START AFTER COMMIT inside a
-- stored procedure. After you call this procedure and
-- call COMMIT WORK, the server will post the event.
"CREATE PROCEDURE sac_demo
BEGIN
DECLARE MyVar INT;
MyVar := 97;
EXEC SQL PREPARE c_sacdemo START AFTER COMMIT NONUNIQUE CALL
  post_event(?);
EXEC SQL EXECUTE c_sacdemo USING (MyVar);
EXEC SQL CLOSE c_sacdemo;
EXEC SQL DROP c_sacdemo;
END";

-- When user2 calls this procedure, the procedure will wait until
-- the event named "e" is posted, and then it will call the
-- stored procedure that inserts a record into table1.
"CREATE PROCEDURE wait_on_event_e
BEGIN
-- Declare the variable that will be used to hold the event parameter.
-- Although the parameter was declared when the event was created, you
-- still need to declare it as a variable in the procedure that receives
-- that event.
DECLARE i INT;
WAIT EVENT
```



```

        WHEN e (i) BEGIN
            -- After we receive the event, insert a row into the table.
            EXEC SQL PREPARE c_call_inserter CALL inserter(?);
            EXEC SQL EXECUTE c_call_inserter USING (i);
            EXEC SQL CLOSE c_call_inserter;
            EXEC SQL DROP c_call_inserter;
        END EVENT
    END WAIT
END";

COMMIT WORK;

----- SCRIPT 2 (USER 2) -----
-- Make sure that user2 sees the changes that user1 made.
COMMIT WORK;

-- Wait until user1 posts the event.
CALL wait_on_event_e;
-- Do not commit work again (yet).

----- SCRIPT 3 (USER 1) -----
COMMIT WORK;

-- User2 should be waiting on event e, and should see the event after
-- we execute the stored procedure named sac_demo and then commit work.
-- Note that since START AFTER COMMIT statements are executed
-- asynchronously, there may be a slight delay between the COMMIT WORK
-- and the associated POST EVENT.
CALL sac_demo;
COMMIT WORK;

----- SCRIPT 4 (USER 2) -----
-- Commit the INSERT that we did earlier when we called inserter()
-- after receiving the event.
COMMIT WORK;

-----SCRIPT 5 (USER 1) -----
-- Ensure that we see the data that user2 inserted.
COMMIT WORK;

-- Show the record that user2 inserted.
SELECT * FROM table1;

COMMIT WORK;

```

Important considerations on START AFTER COMMIT

- When the body of the deferred procedure call (START AFTER COMMIT) is executed, it runs asynchronously in the background. This allows the server to immediately start executing the next SQL command in your program without waiting for the deferred procedure call statement to finish. It also means that you do not have to wait for completion before disconnecting from the server. In most situations, this is an advantage. However, in a few situations this may be a disadvantage. For example, if the body of the deferred procedure call locks records that are needed by subsequent SQL commands in your program, you may not appreciate having the body of the deferred procedure call run in the background while your next SQL command runs in the foreground and has to wait to access those same records.
- The statement to be executed will only be executed if the transaction is completed with a COMMIT, not a ROLLBACK. If the entire transaction is

explicitly rolled back, or if the transaction is aborted and thus implicitly rolled back (due to a failed connection, for example), then the body of the `START AFTER COMMIT` will not be executed.

- Although the transaction in which the deferred procedure call occurs can be rolled back (thus preventing the body of the deferred procedure call from running), the body of the deferred procedure call cannot itself be rolled back if it has executed. Because it runs asynchronously in the background, there is no mechanism for cancelling or rolling back the body once it starts executing.
- The statement in the deferred procedure call is not guaranteed to run until completion or to be run as an "atomic" transaction. For example, if your server crashes, then the statement will not resume executing the next time that the server starts, and any actions that were completed before the server crashed may be kept. To prevent inconsistent data in this type of situation, you must program carefully and make proper use of features like referential constraints to ensure data integrity.
- If you execute a `START AFTER COMMIT` statement in autocommit mode, then the body of the `START AFTER COMMIT` will be executed "immediately" (i.e. as soon as the `START AFTER COMMIT` is executed and automatically committed). At first, this might seem useless — why not just execute the body of the `START AFTER COMMIT` directly? There are a few subtle differences, however. First, a direct call to `my_proc` is synchronous; the server will not return control to you until the stored procedure has finished executing. If you call `my_proc` as the body of a `START AFTER COMMIT`, however, then the call is asynchronous; the server does not wait for the end of `my_proc` before allowing you to execute the next SQL statement. In addition, because `START AFTER COMMIT` statements are not truly executed "immediately" (i.e. at the time that the transaction is committed) but may instead be delayed briefly if the server is busy, you might or might not actually start running your next SQL statement before `my_proc` even starts executing. It is rare for this to be desirable behavior. However, if you truly want to launch an asynchronous stored procedure that will run in the background while you continue onward with your program, it is valid to do `START AFTER COMMIT` in autocommit mode.
- If more than one deferred procedure call was executed in the same transaction, then the bodies of all the `START AFTER COMMIT` statements will run asynchronously. This means that they will not necessarily run in the same order as you executed the `START AFTER COMMIT` statements within the transaction.
- The body of a `START AFTER COMMIT` must contain only one SQL statement. That one statement may be a procedure call, however, and the procedure may contain multiple SQL statements, including other procedure calls.
- The `START AFTER COMMIT` statement applies only to the transaction in which it is defined. If you execute `START AFTER COMMIT` in the current transaction, the body of the deferred procedure call will be executed only when the current transaction is committed; it will not be executed in subsequent transactions, nor will it be executed for transactions done by any other connections. `START AFTER COMMIT` statements do not create "persistent" behavior. If you would like the same body to be called at the end of multiple transactions, then you will have to execute a "`START AFTER COMMIT ... CALL my_proc`" statement in each of those transactions.
- The "result" of the execution of the body of the deferred procedure call (`START AFTER COMMIT`) statement is not returned in any way to the connection that ran the deferred procedure call. For example, if the body of the deferred procedure call returns a value that indicates whether an error occurred, that value will be discarded.

- Almost any SQL statement may be used as the body of a START AFTER COMMIT statement. Although calls to stored procedures are typical, you may also use UPDATE, CREATE TABLE, or almost anything else. (We do not advise putting another START AFTER COMMIT statements inside a START AFTER COMMIT, however.) Note that a statement like SELECT is generally useless inside an deferred procedure call because the result is not returned.
- Because the body is not executed at the time that the START AFTER COMMIT statement is executed inside the transaction, START AFTER COMMIT statements rarely fail unless the deferred procedure call itself or the body contains a syntax error or some other error that can be detected without actually executing the body.

What if you do not want the next SQL statement in your program to run until deferred procedure call statement has finished running? Here's a workaround:

1. At the end of the deferred procedure call statement (for example, at the end of the stored procedure called by the deferred procedure call statement), post an Event. (See *IBM solidDB Programmer Guide* for a description of events.)
2. Immediately after you commit the transaction that specified the deferred procedure call, call a stored procedure that waits on the event.
3. After the stored procedure call (to wait on the event), put the next SQL statement that your program wants to execute.

For example, your program might look like the following:

```
...
START AFTER COMMIT ... CALL myproc;
...
COMMIT WORK;
CALL wait_for_sac_completion;
UPDATE ...;
```

The stored procedure `wait_for_sac_completion` would wait for the event that `myproc` will post. Therefore, the UPDATE statement will not run until after the deferred procedure call statement finishes.

However, this workaround is slightly risky. Since deferred procedure call statements are not guaranteed to execute until completion, there is a chance that the stored procedure `wait_for_sac_completion` will never get the event that it is waiting for.

Why would anyone design a command that may or may not run to completion? The answer is that the primary purpose of the START AFTER COMMIT feature is to support "Sync Pull Notify". The Sync Pull Notify feature allows a master server to notify its replica(s) that data has been updated and that the replicas may request refreshes to get the new data. If this notification process fails for some reason, it would not result in data corruption; it would simply mean that there would be a longer delay before the replica refreshes the data. Since a replica is always given all the data since its last successful refresh operation, a delay in receipt of data does not cause the replica to permanently miss any data. For more details, see the section *Introduction to Sync Pull Notify* in the *IBM solidDB Advanced Replication User Guide*.

Note: The statement inside the body of the START AFTER COMMIT may be any statement, including SELECT. Remember, however, that the body of the START AFTER COMMIT does not return its results anywhere, so a SELECT statement is generally not useful inside a START AFTER COMMIT.

Note: If you are in auto-commit mode and execute `START AFTER COMMIT...`, then the given statement is started immediately in the background. "Immediately" here actually means "as soon as possible", because it's still executed asynchronously when the server has time to do it.

Example: Sync Pull Notify ("Push Synchronization"): This example showcases how to implement Sync Pull Notify (Master notifying all relevant Replicas that there is new data that they can request a refresh of) using the `START` and `CALL` statements. This example also uses triggers.

Consider a scenario where there is a Master M1 and Replicas R1 and R2.




Figure 1. Sync pull notify

To carry out Sync Pull Notify, follow the steps listed below:

1. Define a Procedure Pm1 in Master M1. In Procedure Pm1, include the statements:

```
EXECDIRECT CALL Pr1 AT R1;  
EXECDIRECT CALL Pr1 AT R2;
```

You will have one call for each interested Replica. Even though the replica name changes, typically the procedure name is the same on each replica.

2. Define a Procedure Pr1 in Replica R1. If a master is to invoke the Pr1 in more than one replica, then Pr1 should be defined for every replica that is of interest. See the replica procedure example in the example section below.
3. Define a Trigger for all relevant DML operations, such as
 - INSERT
 - UPDATE
 - DELETE

4. In each trigger body, embed the statement

```
EXECDIRECT START [UNIQUE] CALL Pm1;
```

5. Grant EXECUTE authority to the appropriate user on each replica. (A user Ur1 on the replica should already be mapped to a corresponding user Um1 on the master. The user Um1 must execute the

```
EXECDIRECT START [UNIQUE] CALL Pm1;
```

When Um1 calls the procedure remotely, the call will actually execute with the privileges of Ur1 when the call is executed on the replica.)

Example: Sliced replicas

In this example, a sales application has a table named `CUSTOMER`, which has a column named `SALESMAN`. The master database contains information for all salespersons. Each salesperson has her own replica database, and that replica has only a slice of the master's data; specifically, each salesperson's replica has the slice

of data for that salesperson. For example, salesperson Smith's replica has only the data for salesperson Smith. If the salesperson assigned to a particular customer changes, then the correct replicas should be notified. If XYZ corporation is reassigned from salesperson Smith to salesperson Jones, then salesperson Jones's replica database should add the data related to XYZ corporation, and salesperson Smith's replica should delete the data related to XYZ corporation. Here is the code to update both replica databases:

```
-- If a customer is reassigned to a different salesman, then we
-- must notify both the old and new salesmen.
-- NOTE: This sample shows only the "UPDATE" trigger, but in
-- reality, you'd also need to define INSERT and DELETE triggers.
CREATE TRIGGER T_CUST_AFTERUPDATE ON CUSTOMER
AFTER UPDATE
REFERENCING NEW SALESMAN AS NEW_SALESMAN,
REFERENCING OLD SALESMAN AS OLD_SALESMAN
BEGIN
IF NEW_SALESMAN <> OLD_SALESMAN THEN
EXEC SQL EXECDIRECT
START AFTER COMMIT
FOR EACH REPLICA WHERE NAME=OLD_SALESMAN
UNIQUE CALL CUST(OLD_SALESMAN);
EXEC SQL EXECDIRECT
START AFTER COMMIT
FOR EACH REPLICA WHERE NAME=NEW_SALESMAN
UNIQUE CALL CUST(NEW_SALESMAN);
ENDIF
END;
```

Suppose that in the application, the user assigns all customers in sales area 'CA' to salesperson Mike.

```
UPDATE CUSTOMER SET SALESMAN='Mike' WHERE SALES_AREA='CA';
COMMIT WORK;
```

The Master server has the following procedure:

```
CREATE PROCEDURE CUST(salesman VARCHAR)
BEGIN
EXEC SQL EXECDIRECT CALL CUST(salesman) AT salesman;
COMMIT WORK;
END
```

Each replica has the following procedure:

```
CREATE PROCEDURE CUST(salesman VARCHAR)
BEGIN
MESSAGE s BEGIN;
MESSAGE s APPEND REFRESH CUSTS(salesman);
MESSAGE s END;
COMMIT WORK;
MESSAGE s FORWARD TIMEOUT FOREVER;
COMMIT WORK;
END
```

In the procedure CUST(), we force the salesperson's replica to refresh from the data in the master. This procedure CUST() is defined on all the replicas. If we call the procedure on both the replica that the customer was reassigned to, and the replica that the customer was reassigned from, then the procedure updates both those replicas. Effectively, this will delete the out-of-date data from the replica that no longer has this customer, and will insert the data to the replica that is now responsible for this customer. If the publication and its parameters are properly defined, we do not need to write additional detailed logic to handle each possible operation, such as reassigning a customer from one salesperson to another; instead, we simply tell each replica to refresh from the most current data.

Note:

- It is possible to implement a Sync Pull Notify without triggers. The application may call appropriate procedures to implement SyncPull. Triggers are a way to achieve Sync Pull Notify in conjunction with the statement START AFTER COMMIT and remote procedure calls.
- Sometimes, in the Sync Pull Notify process, it is possible that a replica may have to exchange one extra round trip of messages unnecessarily. This could happen if the master invoked procedure tries to send a message to the replica that just sent the changes to the master, and that causes a change in the "hot data" in the master. But this can be avoided with careful usage of the START AFTER COMMIT statement. Be careful not to create an "infinite loop", where each update on the master leads to an immediate update on the replica, which leads to an immediate update on the master... The best way to avoid this is to be careful when creating triggers on the replica that might "immediately" send updated data to the master, which in turn "immediately" notifies the replica to refresh again.

Tracing the execution of background jobs: The START AFTER COMMIT statement returns a result-set with one INTEGER column. This integer is a unique "job" id that can be used to query the status of statements that failed to start for one reason or another (invalid SQL statement, no access rights, replica not available, and so on).

If a maximum number of uncommitted deferred procedure call statements is reached, then an error is returned when the deferred procedure call is issued. The maximum number is configurable in `solid.ini`. See *IBM solidDB Administrator Guide*.

If a statement cannot be started, the reason is logged to the system table `SYS_BACKGROUNDJOB_INFO`.

```
SYS_BACKGROUNDJOB_INFO
(
  ID INTEGER NOT NULL,
  STMT WVARCHAR NOT NULL,
  USER_ID INTEGER NOT NULL,
  ERROR_CODE INTEGER NOT NULL,
  ERROR_TEXT WVARCHAR NOT NULL,
  PRIMARY KEY(ID)
);
```

Only failed START AFTER COMMIT statements are logged into this table. If the statement (for example, a procedure call) starts successfully, no information is stored into the system tables.

User can retrieve the information from the table `SYS_BACKGROUNDJOB_INFO` using either SQL SELECT-query or calling a system procedure `SYS_GETBACKGROUNDJOB_INFO`. The input parameters is the jobID. The returned values are: ID INTEGER, STMT WVARCHAR, USER_ID INTEGER, ERROR_CODE INTEGER, ERROR_TEXT INTEGER.

Also an event `SYS_EVENT_SACFAILED` is posted when a statement fails to start.

```
CREATE EVENT SYS_EVENT_SACFAILED (ENAME WVARCHAR,
POSTSRVTIME TIMESTAMP,
UID INTEGER,
NUMDATAINFO INTEGER,
TEXTDATA WVARCHAR);
```

The NUMDATAINFO field contains the jobID. The application can wait for this event and use the jobID to retrieve the reason from the system table SYS_BACKGROUNDJOB_INFO.

The system table SYS_BACKGROUNDJOB_INFO can be emptied with the admin command **cleanbgjobinfo**. You need DBA privileges to execute this command, which means that only a DBA can delete the rows from the table.

Controlling background tasks: Background tasks can be controlled with the SSC API and admin commands (see *IBM solidDB Shared Memory Access and Linked Library Access User Guide* for details on the SSC API). The server uses the task type SSC_TASK_BACKGROUND for the tasks that execute statements started with START AFTER COMMIT. Note that there may be several of these tasks, but you cannot control them individually.

3.1.2 Stored procedures – external

solidDB supports external stored procedures written in C programming language. External stored procedures enable you to expand the server functionality by writing a C program.

The external stored procedures are registered with the CREATE PROCEDURE (external) statement.

The external procedures are accessed using the standard dynamic library interface as provided by the operating system (shared libraries in Linux and UNIX, DLLs in Windows). To enable use of external stored procedures, you must link your application dynamically to solidDB using shared memory access (SMA) or linked library access (LLA).

The external procedures are called with the same CALL statement syntax as SQL stored procedures; the control of the thread is passed to the external library routine until it returns. Furthermore, the database interface of the external procedures is similar to that of an ODBC application, except that the connection handle (hdbc) is passed to the external routine, and it can immediately start using it without establishing the connection.

The external procedures are removed with the DROP PROCEDURE statement.

Related reference:

A.16, “CREATE PROCEDURE (external),” on page 193

A.7, “CALL,” on page 165

A.35, “DROP PROCEDURE,” on page 222

3.2 Functions

A *function* is an operation that is denoted by a function name followed by a pair of parentheses enclosing the specification of arguments (there may be no arguments). In addition to built-in functions, solidDB supports both internal and external user-defined stored functions.

Built-in functions are provided within the database engine.

User-defined functions are registered to the database using the CREATE FUNCTION statement. There are two types of user-defined functions:

- User-defined *stored functions* can be written in the solidDB proprietary SQL procedure language.
- User-defined *external stored functions* can be written in C programming language. The external stored functions are loaded at runtime using standard dynamic library interfaces as provided by the operating system. External stored functions enable you to expand the server functionality by writing a C program.

Related reference:

Appendix B, “Functions,” on page 313

A *function* is an operation that is denoted by a function name followed by a pair of parentheses enclosing the specification of arguments (there may be no arguments). In addition to built-in functions, solidDB supports both internal and external user-defined stored functions.

A.12, “CREATE FUNCTION,” on page 173

A.13, “CREATE FUNCTION (external),” on page 175

3.3 Triggers

A trigger is a mechanism for executing a series of SQL statements when a particular action (an INSERT, UPDATE, or DELETE) occurs. The trigger contains a SQL statement that needs to be executed when the trigger is invoked. Triggers are created using solidDB proprietary stored procedure syntax.

You can create one or more triggers on a table, with each trigger defined to activate on a specific INSERT, UPDATE, or DELETE command. When a user modifies data within the table, the trigger that corresponds to the command is activated.

You can only use inline SQL or stored procedures with triggers. If you use a stored procedure in the trigger, the procedure must be created with the CREATE PROCEDURE command. A procedure invoked from a trigger body can invoke other triggers.

Triggers enable you to:

- Implement special integrity constraints, such as checking that certain conditions are maintained, for example, to prevent users from making incorrect or inconsistent data changes.
- Take action based on the value of a row before or after modification.
- Transfer much of the logic processing to the server, reducing the amount of work that your application needs to do as well as reducing network traffic.

3.3.1 Triggers – principles of operation

In solidDB's DML execution model, the solidDB server performs a number of validation checks before executing data manipulation statements (INSERT, UPDATE, or DELETE). Following is the execution order for data validation, trigger execution, and integrity constraint checking for a single DML statement.

1. Validate values if they are part of the statement (that is, not bound). This includes null value checking, data type checking (such as numeric), and so on.
2. Perform table level security checks.
3. Loop for each row affected by the SQL statement. For each row perform these actions in this order:
 - a. Perform column level security checks.
 - b. Fire BEFORE row trigger.

- c. Validate values if they are bound in. This includes null value checks, data type checking, and size checking (for example, checking if the character string is too long).

Note: Size checking is performed even for values that are not bound.

- d. Execute INSERT/UPDATE/DELETE
 - e. Fire AFTER ROW trigger
4. Commit statement
- a. Perform concurrency conflict checks.
 - b. Perform checks for duplicate values.
 - c. Perform referential integrity checks on invoking DML.

Note: A trigger itself can cause the DML to be executed, which applies to the steps shown in the above model.

3.3.2 Creating and modifying triggers

Use the CREATE TRIGGER statement to create a trigger.

You can disable an existing trigger or all triggers defined on a table by using the ALTER TRIGGER statement. The ALTER TRIGGER statement causes a solidDB server to ignore the trigger when an activating DML statement is issued. With this statement, you can also enable a trigger that is currently inactive.

To drop a trigger from the system catalog, use the DROP TRIGGER statement.

Related reference:

- A.5, "ALTER TRIGGER," on page 162
- A.24, "CREATE TRIGGER," on page 210
- A.46, "DROP TRIGGER," on page 228

3.3.3 Triggers and procedures

Triggers can call stored procedures and cause a solidDB server to execute other triggers. You can invoke procedures within a trigger body. You can also define a trigger body that contains only procedure calls. A procedure invoked from a trigger body can invoke other triggers.

When using stored procedures within the trigger body, you must first store the procedure with the CREATE PROCEDURE statement.

In a procedure definition, you can use COMMIT and ROLLBACK statements. However, in a trigger body, you cannot use COMMIT (including AUTOCOMMIT and COMMIT WORK) and ROLLBACK statements. You can use only the WHENEVER SQLERROR ABORT statement.

By default, you can nest triggers up to 16 levels deep. The limit can be changed using the **SQL.MaxNestedTriggers** parameter. If a trigger gets into an infinite loop, solidDB detects this recursive action when the maximum nesting level is reached, and returns an error to the user. For example, you could activate a trigger by attempting to insert into the table T1 and the trigger could call a stored procedure which also attempts to insert into T1, activating the trigger recursively.

If a set of nested triggers fails at any time, solidDB rolls back the statement which originally activated the triggers.

Setting default or derived columns

You can create triggers to set up default or derived column values in INSERT and UPDATE operations. When you create the trigger for this purpose using the CREATE TRIGGER command, the trigger must follow these rules:

- The trigger must be executed BEFORE the INSERT or UPDATE operation. Column values are modified with only a BEFORE trigger. Because the column value must be set before the INSERT or UPDATE operation, using the AFTER trigger to set column values is meaningless. Note also that the DELETE operation does not apply to modifying column values.
- For an INSERT and UPDATE operation, the REFERENCING clause must contain a NEW column value for modification. Note that modifying the OLD column value is meaningless.
- New column values can be set by simply changing the values of variables defined in the referencing section.

Using parameters and variables with triggers

When you update a record and that update invokes a trigger, the trigger itself can change the value of some columns within that record. In some situations, you might want to refer to both the "old" value and the "new" value within the trigger.

The REFERENCING clause allows you to create "aliases" for old and new values so that you can refer to either one within the same trigger. For example, assume there are two tables, one that holds customer information and one that holds invoice information. In addition to storing the amount of money billed for each invoice, the table contains a "total_bought" field for each customer; this "total_bought" field contains the cumulative total for all invoices ever sent to this customer. (This field might be used to identify high-volume customers.)

Any time the total_amount on an invoice is updated, the "total_bought" value for that customer's record in the customer table is also updated. To do this, the amount of the old value stored in the invoice is subtracted and the amount of the new value in the invoice is added. For example, if a customer's invoice used to be for \$100 and it is changed to \$150, then \$100 is subtracted and \$150 is added to the "total_bought" field. By properly using the REFERENCING clause, the trigger can "see" both the old value and the price column, thereby allowing the update of the total_bought column.

The column aliases created by the REFERENCING clause are valid only within the trigger.

Example: Trigger with referencing clause

```
-- This SQL sample demonstrates how to use the clause
-- "REFERENCING OLD AS old_col, REFERENCING NEW AS new_col"
-- to have simultaneous access to both the "OLD" and "NEW"
-- column values of the field while inside a trigger.
-- In this scenario, we have customers and invoices.
-- For each customer, we keep track of the cumulative total of
-- all purchases by that customer.
-- Each invoice stores the total amount of all purchases on
-- that invoice. If an total price on an invoice must be
-- adjusted, then the cumulative value of that customer's
-- purchases must also be adjusted.
-- Therefore, we update the cumulative total by subtracting
-- the "old" price on the invoice and adding the "new" price.
-- For example, if the amount on a customer's invoice was
-- changed from $100 to $150 (an increase of $50), then we
-- would update the customer's cumulative total by
-- subtracting $100 and adding $150 (a net increase of $50).
```

```

-- Drop the sample tables if they already exist.
DROP TABLE customers;
DROP TABLE invoices;
CREATE TABLE customers (
    customer_id INTEGER,    -- ID for each customer.
    total_bought FLOAT     -- The cumulative total price of
                          -- all this customer's purchases.
);
-- Each customer may have 0 or more invoices.
CREATE TABLE invoices (
    customer_id INTEGER,
    invoice_id INTEGER,    -- unique ID for each invoice
    invoice_total FLOAT    -- total price for this invoice
);
-- If the total_price on an invoice changes, then
-- update customers.total_bought to take into account
-- the change. Subtract the old invoice price and add the
-- new invoice price.
"CREATE TRIGGER old_and_new ON invoices
AFTER UPDATE
    REFERENCING OLD invoice_total AS old_invoice_total,
    REFERENCING NEW invoice_total AS new_invoice_total,
    -- If the customer_id doesn't change, we could use
    -- either the NEW or OLD customer_id.
    REFERENCING NEW customer_id AS new_customer_id
BEGIN
    EXEC SQL PREPARE upd_curs
    UPDATE customers
    SET total_bought = total_bought - ? + ?
    WHERE customers.customer_id = ?;
    EXEC SQL EXECUTE upd_curs
    USING (old_invoice_total, new_invoice_total,
          new_customer_id);
    EXEC SQL CLOSE upd_curs;
    EXEC SQL DROP upd_curs;
END";
-- When a new invoice is created, we update the total_bought
-- in the customers table.
"CREATE TRIGGER update_total_bought ON invoices
AFTER INSERT
    REFERENCING NEW invoice_total AS new_invoice_total,
    REFERENCING NEW customer_id AS new_customer_id
BEGIN
    EXEC SQL PREPARE ins_curs
    UPDATE customers
    SET total_bought = total_bought + ?
    WHERE customers.customer_id = ?;
    EXEC SQL EXECUTE ins_curs
    USING (new_invoice_total, new_customer_id);
    EXEC SQL CLOSE ins_curs;
    EXEC SQL DROP ins_curs;
END";
-- Insert a sample customer.
INSERT INTO customers (customer_id, total_bought)
VALUES (1000, 0.0);
-- Insert invoices for a customer; the INSERT trigger will
-- update the total_bought in the customers table.
INSERT INTO invoices (customer_id, invoice_id, invoice_total)
VALUES (1000, 5555, 234.00);
INSERT INTO invoices (customer_id, invoice_id, invoice_total)
VALUES (1000, 5789, 199.0);
-- Make sure that the INSERT trigger worked.
SELECT * FROM customers;
-- Now update an invoice; the total_bought in the customers
-- table will also be updated and the trigger that does
-- this will use the REFERENCING clauses
--     REFERENCING NEW invoice_total AS new_invoice_total,

```

```

--      REFERENCING OLD invoice_total AS old_invoice_total
UPDATE invoices SET invoice_total = 235.00
  WHERE invoice_id = 5555;
-- Make sure that the UPDATE trigger worked.
SELECT * FROM customers;
COMMIT WORK;

```

3.3.4 Triggers and transactions

Triggers require no commit from the invoking transaction in order to fire; DML statements alone cause triggers to fire. COMMIT WORK is also disallowed in a trigger body.

In a procedure definition, you can use COMMIT and ROLLBACK statements. But in a trigger body, you cannot use COMMIT and ROLLBACK statements. You can use only the WHENEVER SQLERROR ABORT statement. Note that if autocommit is on, then each statement inside the trigger is not treated as a separate statement and is not committed when it is executed; instead, the entire trigger body is executed as part of the INSERT, UPDATE, or DELETE statement that fired the trigger. Either the entire trigger (and the statement that fired it) is committed, or else the entire trigger (and the statement that fired it) is rolled back.

Recursion and concurrency conflict errors

If a DML statement updates/deletes a row that causes a trigger to be fired, you cannot update/delete the same row again within that trigger. In such cases an AFTER trigger event can cause a recursion error and a BEFORE trigger event can cause a concurrency conflict error.

The following sections explain these terms, provide some examples of triggers that create these problems, and provide a table (shown in “Summary of trigger cases” on page 69), that indicates the trigger situations that will and will not cause recursion errors or concurrency conflict errors.

Triggers and recursion: A piece of code is “recursive” if the code causes itself to execute again. For example, a stored procedure that calls itself is recursive. Recursion in stored procedures is occasionally useful. On the other hand, triggers can create a slightly more subtle type of recursion, which is invalid and prohibited by the solidDB server. A trigger that contains a statement that causes the same trigger to execute again on the same record is recursive. For example, a delete trigger would be recursive if it tries to delete the same record whose deletion fired the trigger.

If the database server were to allow recursion in triggers, then the server might go into an “infinite loop” and never finish executing the statement that fired the trigger. A concurrency conflict error occurs when a trigger executes an operation that “competes with” the statement that fired the trigger by trying to do the same type of action (for example, delete) within the same SQL statement. For example, if you create a trigger that is supposed to be fired when a record is deleted, and if that trigger tries to delete the same record whose deletion fired the trigger, then there are in essence two different “simultaneous” delete statements “competing” to delete the record; this results in a concurrency conflict. The following section provides an example of a defective delete trigger.

Examples of Defective Triggers Causing Recursion

The examples in this section explain just a few of the many restrictions and rules involving triggers.

In this scenario, an employee has resigned from a job and his or her medical coverage requires cancellation. The medical coverage also requires cancellation for the employee's dependents. A business rule for this situation is implemented by creating a trigger; the trigger is executed when an employee's record is deleted and the statements inside the trigger then delete the employee's dependents. (This example assumes that the employees and their dependents are stored in the same table; in the real world, dependents are normally kept in a separate table. This example also assumes that each family has a unique last name.)

```
CREATE TRIGGER do_not_try_this ON employees_and_dependents
AFTER DELETE
REFERENCING OLD last_name AS old_last_name
BEGIN
EXEC SQL PREPARE del_cursor
DELETE FROM employees_and_dependents
WHERE last_name = ?;
EXEC SQL EXECUTE del_cursor USING (old_last_name);
-- ... close and drop the cursor.
END;
```

Assume that an employee "John Smith" resigns and his medical coverage is deleted. When you delete "John Smith", the trigger is invoked immediately after John Smith is deleted and the trigger will try to delete ALL people named "John Smith", including not only the employee's dependents, but also the employee himself, since his name meets the criteria in the WHERE clause.

Every time an attempt is made to delete the employee's record, this action fires the trigger again. The code then recursively keeps trying to delete the employee by again firing the trigger, and again trying to delete. If the database server did not prohibit this or detect the situation, the server could go into an infinite loop. If the server detects this situation, it will give you an appropriate error, such as "Too many nested triggers."

A similar situation can happen with UPDATE. Assume that a trigger adds sales tax every time that a record is updated. Here's an example that causes a recursion error:

```
CREATE TRIGGER do_not_do_this_either ON invoice
AFTER UPDATE
REFERENCING NEW total_price AS new_total_price
BEGIN
-- Add 8% sales tax.
EXEC SQL PREPARE upd_curs1
UPDATE invoice SET total_price = 1.08 * total_price
WHERE ...;
-- ... execute, close, and drop the cursor...
END;
```

In this scenario, customer Ann Jones calls up to change her order; the new price (with sales tax) is calculated by multiplying the new subtotal by 1.08. The record is updated with the new total price; each time the record is updated, the trigger is fired, so updating the record once, causes the trigger to update it again and updates are repeated in an infinite loop.

If AFTER triggers can cause recursion or looping, what happens with BEFORE triggers? The answer is that, in some cases, BEFORE triggers can cause concurrency problems. Let's return to the first example of the trigger that deleted medical coverage for employees and their dependents. If the trigger were a BEFORE trigger (rather than an AFTER trigger), then just before the employee is deleted, we would execute the trigger, which in this case deletes everyone named John Smith. After the trigger is executed, the engine resumes its original task of

dropping employee John Smith himself, but the server finds either he isn't there or that his record cannot be deleted because it has already been marked for deletion — in other words, there is a concurrency conflict because there are two separate efforts to delete the same record.

Summary of trigger cases: In addition to the examples described in the previous section, the following table summarizes a number of additional cases, including those involving INSERTs, as well as UPDATEs and DELETEs.

The table is divided into the following five columns:

- *Trigger Mode* (that is, BEFORE or AFTER)
- *Operation* (INSERT, DELETE, or UPDATE)
- *Trigger Action* (what the trigger itself attempts to do, such as update the record that was just inserted)
- *Lock Type* ("optimistic" or "pessimistic")
- *Result* that you will see (for example, that the trigger action was successful, or that the trigger failed for a reason such as a recursion error like the one discussed in the previous section).

For details on interpreting a trigger entry in this table, see *Example Entry 1* later in this chapter.

Table 11. INSERT/UPDATE/DELETE operations for BEFORE/AFTER triggers

Trigger Mode	Operation	Trigger Action	Lock Type	Result
AFTER	INSERT	UPDATE the same row by adding a number to the value	Optimistic	Record is updated.
AFTER	INSERT	UPDATE the same row by adding a number to the value	Pessimistic	Record is updated.
BEFORE	INSERT	UPDATE the same row by adding a number to the value	Optimistic	Record is not updated since the WHERE condition of the UPDATE within the trigger body returns a NULL resultset (as the desired row is not yet inserted in the table).
BEFORE	INSERT	UPDATE the same row by adding a number to the value	Pessimistic	Record is not updated since the WHERE condition of the UPDATE within the trigger body returns a NULL resultset (as the desired row is not yet inserted in the table).
AFTER	INSERT	DELETE the same row that is being inserted	Optimistic	Record is deleted.
AFTER	INSERT	DELETE the same row that is being inserted	Pessimistic	Record is deleted.

Table 11. INSERT/UPDATE/DELETE operations for BEFORE/AFTER triggers (continued)

Trigger Mode	Operation	Trigger Action	Lock Type	Result
BEFORE	INSERT	DELETE the same row that is being inserted	Optimistic	Record is not deleted since the WHERE condition of the DELETE within the trigger body returns a NULL resultset (as the desired row is not yet inserted in the table).
BEFORE	INSERT	DELETE the same row that is being inserted	Pessimistic	Record is not updated since the WHERE condition of the UPDATE within the trigger body returns a NULL resultset (as the desired row is not yet inserted in the table).
AFTER	INSERT	INSERT a row	Optimistic	Too many nested triggers.
AFTER	INSERT	INSERT a row	Pessimistic	Too many nested triggers.
BEFORE	INSERT	INSERT a row	Optimistic	Too many nested triggers.
BEFORE	INSERT	INSERT a row	Pessimistic	Too many nested triggers.
AFTER	UPDATE	UPDATE the same row by adding a number to the value	Optimistic	Generates Solid Table Error: Too many nested triggers.
AFTER	UPDATE	UPDATE the same row by adding a number to the value	Pessimistic	Generates Solid Table Error: Too many nested triggers.
BEFORE	UPDATE	UPDATE the same row by adding a number to the value.	Optimistic	Record is updated, but does not get into a nested loop because the WHERE condition in the trigger body returns a NULL resultset and no rows are updated to fire the trigger recursively.
BEFORE	UPDATE	UPDATE the same row by adding a number to the value.	Pessimistic	Record is updated, but does not get into a nested loop because the WHERE condition in the trigger body returns a NULL resultset and no rows are updated to fire the trigger recursively.
AFTER	UPDATE	DELETE the same row that is being updated.	Optimistic	Record is deleted.

Table 11. INSERT/UPDATE/DELETE operations for BEFORE/AFTER triggers (continued)

Trigger Mode	Operation	Trigger Action	Lock Type	Result
AFTER	UPDATE	DELETE the same row that is being updated.	Pessimistic	Record is deleted.
BEFORE	UPDATE	DELETE the same row that is being updated.	Optimistic	Concurrency conflict error.
BEFORE	UPDATE	DELETE the same row that is being updated.	Pessimistic	Concurrency conflict error.
AFTER	DELETE	INSERT a row with the same value.	Optimistic	Same record is inserted after deleting.
AFTER	DELETE	INSERT a row with the same value.	Pessimistic	Hangs at the time of firing the trigger.
BEFORE	DELETE	INSERT a row with the same value.	Optimistic	Same record is inserted after deleting
BEFORE	DELETE	INSERT a row with the same value.	Pessimistic	Hangs at the time of firing the trigger.
AFTER	DELETE	INSERT a row with the same value.	Optimistic	Record is deleted.
AFTER	DELETE	UPDATE the same row by adding a number to the value.	Pessimistic	Record is deleted.
BEFORE	DELETE	UPDATE the same row by adding a number to the value.	Optimistic	Record is deleted.
BEFORE	DELETE	UPDATE the same row by adding a number to the value	Pessimistic	Record is deleted.
AFTER	DELETE	DELETE same row	Optimistic	Too many nested triggers.
AFTER	DELETE	DELETE same record	Pessimistic	Too many nested triggers
BEFORE	DELETE	DELETE same record	Optimistic	Concurrency conflict error.
BEFORE	DELETE	DELETE same record	Pessimistic	Concurrency conflict error.

Here's an example entry from the table and an explanation of that entry:

Table 12. Example Entry 1

Trigger	Operation	Trigger Action	Lock Type	Result
AFTER	INSERT	UPDATE the same row by adding a number to the value	Optimistic	Record is updated.

In this situation, we have a trigger that fires AFTER an INSERT operation is done. The body of the trigger contains statements that update the same row as was inserted (that is, the same row as the one that fired the trigger). If the lock type is "optimistic", then the result will be that the record gets updated. (Because there is no conflict, the locking [optimistic versus pessimistic] does not make a difference).

Note that in this case there is no recursion issue, even though we update the same row that we just inserted. The action that "fires" the trigger is not the same as the action taken inside the trigger, and so we do not create a recursive/looping situation.

Here's another example from the table:

Table 13. Example entry 2

Trigger	Operation	Trigger Action	Lock Type	Result
BEFORE	INSERT	UPDATE the same row by adding a number to the value	Optimistic	Record is not updated since the WHERE condition of the UPDATE within the trigger body returns a NULL resultset (as the desired row is not yet inserted in the table).

In this case, we try to insert a record, but before the insertion takes place the trigger is run. In this case, the trigger tries to update the record (for example, to add sales tax to it). Since the record is not yet inserted, however, the UPDATE command inside the trigger does not find the record, and never adds the sales tax. Thus the result is the same as if the trigger had never fired. There is no error message, so you may not realize immediately that your trigger does not do what you intended.

Flawed trigger: Flawed trigger logic occurs in the following example in which the same row is deleted in a BEFORE UPDATE trigger; this causes solidDB to generate a concurrency conflict error.

Flawed Trigger

```

DROP EMP;
COMMIT WORK;

CREATE TABLE EMP(C1 INTEGER);
INSERT INTO EMP VALUES (1);
COMMIT WORK;

"CREATE TRIGGER TRIG1 ON EMP
  BEFORE UPDATE
  REFERENCING OLD C1 AS OLD_C1
  BEGIN

```



```

EXEC SQL WHENEVER SQLERROR ABORT;
EXEC SQL PREPARE CUR1 DELETE FROM EMP WHERE C1 = ?;
EXEC SQL EXECUTE CUR1 USING (OLD_C1);
END";

UPDATE EMP SET C1=200 WHERE C1 = 1;
SELECT * FROM EMP;

ROLLBACK WORK;

```

Note:

If the row that is updated/deleted were based on a unique key, instead of an ordinary column (as in the example above), solidDB generates the following error message: *1001: key value not found.*

To avoid recursion and concurrency conflict errors, be sure to check the application logic and take precautions to ensure the application does not cause two transactions to update or delete the same row.

Error handling: If a procedure returns an error to a trigger, the trigger causes its invoking DML command to fail with an error. To automatically return errors during the execution of a DML statement, you must use `WHENEVER SQLERROR ABORT` statement in the trigger body. Otherwise, errors must be checked explicitly within the trigger body after each procedure call or SQL statement.

For any errors in the user written business logic as part of the trigger body, users must use the `RETURN SQLERROR` statement. For details, see 3.3.6, "Raising errors from inside triggers," on page 74.

If `RETURN SQLERROR` is not specified, then the system returns a default error message when the SQL statement execution fails. Any changes to the database due to the current DML statement are undone and the transaction is still active. In effect, transactions are not rolled back if a trigger execution fails, but the current executing statement is rolled back.

Note:

Triggered SQL statements are a part of the invoking transaction. If the invoking DML statement fails due to either the trigger or another error that is generated outside the trigger, all SQL statements within the trigger are rolled back along with the failed invoking DML command.

It is the responsibility of the invoking transaction to commit or rollback any DML statements executed within the trigger's procedure. However, this rule does not apply if the DML command invoking the trigger fails as a result of the associated trigger. In this case, any DML statements executed within that trigger's procedure are automatically rolled back.

The `COMMIT` and `ROLLBACK` statements must be executed outside the trigger body and cannot be executed within the trigger body. If one executes `COMMIT` or `ROLLBACK` within the trigger body or within a procedure called from the trigger body or another trigger, the user will get a run-time error.

Nested and recursive triggers: If a trigger gets into an infinite loop, a solidDB server detects this recursive action when the 16-level nesting (or `MaxNestedTriggers` system parameter maximum is reached). For example, an insert attempt on table T1 activates a trigger and the trigger could call a stored procedure

which also attempts to insert into Table T1, recursively activating the trigger. A solidDB server returns an error on a user's insert attempt.

If a set of nested triggers fails at any time, a solidDB server rolls back the command which originally activated the triggers.

3.3.5 Trigger privileges and security

Because triggers can be activated by a user's attempt to INSERT, UPDATE, or DELETE data, no privileges are required to execute them.

When a user invokes a trigger, the user assumes the privileges of the owner of the table on which the trigger is defined. The action statements are executed on behalf of the table owner, not the user who activates the trigger. However, to create a trigger which uses a stored procedure requires that the creator of the trigger meet one of the following conditions:

- You have DBA privileges.
- You are the owner of the table on which the trigger is being defined.
- You were granted all privileges on the table.

If the creator has DBA authority and creates a table for another user, a solidDB server assumes that unqualified names specified in the TRIGGER command belong to the user. For example, the following command is executed under DBA authority:

```
CREATE TRIGGER A.TRIG ON EMP BEFORE UPDATE
```

Since the EMP table is unqualified, the solidDB server assumes that the qualified table name is A.EMP, not DBA.EMP.

3.3.6 Raising errors from inside triggers

At times, it is possible to receive an error in executing a trigger. The error may be due to execution of SQL statements or business logic. If a trigger returns an error, it causes its invoking DML command to fail. To automatically return errors during the execution of a DML statement, you must use the WHENEVER SQLERROR ABORT statement in the trigger body. Otherwise, errors must be checked explicitly within the trigger body after each procedure call or SQL statement.

For any errors in the user-written business logic as part of the trigger body, users can receive errors in a procedure variable using the RETURN SQLERROR *error_string* or RETURN SQLERROR *char_variable* statements.

The error is returned in the following format:

```
User error: error_string
```

If a user does not specify the RETURN SQLERROR statement in the trigger body, all trapped SQL errors are raised with a default *error_string* determined by the system.

Note: Triggered SQL statements are a part of the invoking transaction. If the invoking DML statement fails due to either the trigger or another error that is generated outside the trigger, all SQL statements within the trigger are rolled back along with the failed invoking DML command.

Below is an example of using WHENEVER SQLERROR ABORT to make sure that the trigger catches an error in a stored procedure that it calls.

```

-- If you return an SQLERROR from a stored procedure, the error is
-- displayed. However, if the stored procedure is called from inside
-- a trigger, then the error is not displayed unless you use the
-- SQL statement WHENEVER SQLERROR ABORT.

CREATE TABLE table1 (x INT);
CREATE TABLE table2 (x INT);

"CREATE PROCEDURE stproc1
BEGIN
    RETURN SQLERROR 'Here is an error!';
END";
COMMIT WORK;

"CREATE TRIGGER displays_error ON table1 BEFORE INSERT
BEGIN
    EXEC SQL WHENEVER SQLERROR ABORT;
    EXEC SQL EXECEDIRECT CALL stproc1();
END";
COMMIT WORK;

"CREATE TRIGGER does_not_display_error ON table2 BEFORE INSERT
BEGIN
    EXEC SQL EXECEDIRECT CALL stproc1();
END";
COMMIT WORK;

-- This shows that the error is returned if you execute the stored procedure.
CALL stproc1();

-- Displays an error because the trigger had WHENEVER SQL ERROR ABORT.
INSERT INTO table1 (x) values (1);
-- Does not display an error.
INSERT INTO table2 (x) values (1);

```

3.3.7 Obtaining trigger information

You obtain trigger information in the following ways:

- Use the B.7, “Trigger functions,” on page 322.
- Perform a query on the trigger system table E.1.36, “SYS_TRIGGERS,” on page 362.

3.3.8 Trigger parameter settings

Setting nested trigger maximum

Triggers can invoke other triggers or a trigger can invoke itself (recursive trigger). The maximum number of nested or recursive triggers can be set with the **SQL.MaxNestedTriggers** parameter.

```

[SQL]
MaxNestedTriggers = n;

```

where *n* is the maximum number of nested triggers.

The default number for nested triggers is 16.

Setting the trigger cache

Triggers are cached in a separate cache. Each user has a separate cache for triggers. As the triggers are executed, the trigger procedure logic is cached in the trigger cache and reused when the trigger is executed again.

You can set the size of the trigger cache using the **SQL.TriggerCache** parameter.

```
[SQL]
TriggerCache = n;
```

where *n* is the number of triggers being reserved for the cache.

3.3.9 Trigger example Trigger Example

This example shows how simple triggers work. It contains some triggers that work correctly and some triggers that contain errors. For the successful triggers in the example, a table (named *trigger_test*) is created and six triggers are created on that table. Each trigger, when fired, inserts a record into another table (named *trigger_output*). After performing the DML statements (INSERT, UPDATE, and DELETE) that fire the triggers, the results of the triggers are displayed by selecting all records from the *trigger_output* table.

```
DROP TABLE TRIGGER_TEST;
DROP TABLE TRIGGER_ERR_TEST;
DROP TABLE TRIGGER_ERR_B_TEST;
DROP TABLE TRIGGER_ERR_A_TEST;
DROP TABLE TRIGGER_OUTPUT;
COMMIT WORK;
-- Create a table that has a column for each of the possible trigger
-- types (for example, BI = a trigger that is on Insert
-- operations and that executes as a "Before" trigger).
CREATE TABLE TRIGGER_TEST(
    XX VARCHAR,
    BI VARCHAR, -- BI = Before Insert
    AI VARCHAR, -- AI = After Insert
    BU VARCHAR, -- BU = Before Update
    AU VARCHAR, -- AU = After Update
    BD VARCHAR, -- BD = Before Delete
    AD VARCHAR -- AD = After Delete
);
COMMIT WORK;

-- Table for 'before' trigger errors
CREATE TABLE TRIGGER_ERR_B_TEST(
    XX VARCHAR,
    BI VARCHAR,
    AI VARCHAR,
    BU VARCHAR,
    AU VARCHAR,
    BD VARCHAR,
    AD VARCHAR
);

INSERT INTO TRIGGER_ERR_B_TEST VALUES('x','x','x','x','x',
    'x','x');
COMMIT WORK;

-- Table for 'after X' trigger errors
CREATE TABLE TRIGGER_ERR_A_TEST(
    XX VARCHAR,
    BI VARCHAR, -- Before Insert
    AI VARCHAR, -- After Insert
    BU VARCHAR, -- Before Update
    AU VARCHAR, -- After Update
    BD VARCHAR, -- Before Delete
    AD VARCHAR -- After Delete
);

INSERT INTO TRIGGER_ERR_A_TEST VALUES('x','x','x','x','x',
```

```

        'x','x');
COMMIT WORK;

CREATE TABLE TRIGGER_OUTPUT(
        TEXT VARCHAR,
        NAME VARCHAR,
        SCHEMA VARCHAR
);
COMMIT WORK;

-----
-- Successful triggers
-----
-- Create a "Before" trigger on insert operations. When a record is
-- inserted into the table named trigger_test, then this trigger is
-- fired. When this trigger is fired, it inserts a record into the
-- "trigger_output" table to show that the trigger actually executed.

"CREATE TRIGGER TRIGGER_BI ON TRIGGER_TEST
    BEFORE INSERT
    REFERENCING NEW BI AS NEW_BI
BEGIN
    EXEC SQL PREPARE BI INSERT INTO TRIGGER_OUTPUT VALUES(
        'BI', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE BI;
    SET NEW_BI = 'TRIGGER_BI';
END";
COMMIT WORK;

"CREATE TRIGGER TRIGGER_AI ON TRIGGER_TEST
    AFTER INSERT
    REFERENCING NEW AI AS NEW_AI
BEGIN
    EXEC SQL PREPARE AI INSERT INTO TRIGGER_OUTPUT VALUES(
        'AI', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE AI;
    SET NEW_AI = 'TRIGGER_AI';
END";
COMMIT WORK;

"CREATE TRIGGER TRIGGER_BU ON TRIGGER_TEST
    BEFORE UPDATE
    REFERENCING NEW BU AS NEW_BU
BEGIN
    EXEC SQL PREPARE BU INSERT INTO TRIGGER_OUTPUT VALUES(
        'BU', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE BU;
    SET NEW_BU = 'TRIGGER_BU';
END";
COMMIT WORK;

"CREATE TRIGGER TRIGGER_AU ON TRIGGER_TEST
    AFTER UPDATE
    REFERENCING NEW AU AS NEW_AU
BEGIN
    EXEC SQL PREPARE AU INSERT INTO TRIGGER_OUTPUT VALUES(
        'AU', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE AU;
    SET NEW_AU = 'TRIGGER_AU';
END";
COMMIT WORK;

"CREATE TRIGGER TRIGGER_BD ON TRIGGER_TEST
    BEFORE DELETE
    REFERENCING OLD BD AS OLD_BD
BEGIN
    EXEC SQL PREPARE BD INSERT INTO TRIGGER_OUTPUT VALUES(

```

```

    'BD', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE BD;
    SET OLD_BD = 'TRIGGER_BD';
END";
COMMIT WORK;

"CREATE TRIGGER TRIGGER_AD ON TRIGGER_TEST
    AFTER DELETE
    REFERENCING OLD AD AS OLD_AD
BEGIN
    EXEC SQL PREPARE AD INSERT INTO TRIGGER_OUTPUT VALUES(
    'AD', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE AD;
    SET OLD_AD = 'TRIGGER_AD';
END";
COMMIT WORK;

-----
-- This attempt to create a trigger will fail. The statement
-- specifies the wrong data type for the error variable named
-- ERRSTR.
-----

"CREATE TRIGGER TRIGGER_ERR_AU ON TRIGGER_ERR_A_TEST
    AFTER UPDATE
    REFERENCING NEW AU AS NEW_AU
BEGIN
    -- The following line is incorrect; ERRSTR must be declared
    -- as VARCHAR, not INTEGER;
    DECLARE ERRSTR INTEGER;
    -- ...
    RETURN SQLERROR ERRSTR;
END";
COMMIT WORK;

-----
-- Trigger that returns an error message.
-----

"CREATE TRIGGER TRIGGER_ERR_BI ON TRIGGER_ERR_B_TEST
    BEFORE INSERT
    REFERENCING NEW BI AS NEW_BI
BEGIN
    -- ...
    RETURN SQLERROR 'Error in TRIGGER_ERR_BI';
END";
COMMIT WORK;

-----
-- Success trigger tests. These Insert, Update, and Delete
-- statements will force the triggers to fire. The SELECT
-- statements will show you the records in the trigger_test and
-- trigger_output tables.
-----

INSERT INTO TRIGGER_TEST(XX) VALUES ('XX');
COMMIT WORK;

-- Show the records that were inserted into the trigger_test
-- table. (The records for trigger_output are shown later.)

SELECT * FROM TRIGGER_TEST;
COMMIT WORK;

UPDATE TRIGGER_TEST SET XX = 'XX updated';
COMMIT WORK;

```

```

-- Show the records that were inserted into the trigger_test
-- table. (The records for trigger_output are shown later.)

SELECT * FROM TRIGGER_TEST;
COMMIT WORK;

DELETE FROM TRIGGER_TEST;
COMMIT WORK;

SELECT * FROM TRIGGER_TEST;

-- Show that the triggers did run and did add values to the
-- trigger_output table. You should see 6 records one for
-- each of the triggers that executed. The 6 triggers are:
--     BI, AI, BU, AU, BD, AD.

SELECT * FROM TRIGGER_OUTPUT;
COMMIT WORK;

-----
-- Error trigger test
-----

INSERT INTO TRIGGER_ERR_B_TEST(XX) VALUES ('XX');
COMMIT WORK;

```

3.4 Sequences

Sequencer objects are objects that are used to get sequence numbers. The sequence objects are used to get sequence numbers in an efficient manner.

The syntax is:

```
CREATE [DENSE] SEQUENCE sequence_name
```

Depending on how the sequence is created, there may or may not be holes in the sequence (the sequence can be sparse or dense). Dense sequences guarantee that there are no holes in the sequence numbers. The sequence number allocation is bound to the current transaction. If the transaction rolls back, the sequence number allocations are also rolled back. The drawback of dense sequences is that the sequence is locked out from other transactions until the current transaction ends.

If there is no need for dense sequences, a sparse sequence can be used. A sparse sequence guarantees uniqueness of the returned values, but it is not bound to the current transaction. If a transaction allocates a sparse sequence number and later rolls back, the sequence number is simply lost.

A sequence object can be used, for example, to generate primary key numbers. The advantage of using a sequence object instead of a separate table is that the sequence object is specifically fine-tuned for fast execution and requires less overhead than normal update statements.

Both dense and sparse sequence numbers start from 1.

After creating the sequence with the CREATE SEQUENCE statement, you can access the sequence object values by using the following constructs in SQL statements:

- *sequence_name*.CURRVAL which returns the current value of the sequence
- *sequence_name*.NEXTVAL which increments the sequence by one and returns the next value.

An example of creating unique identifiers automatically for a table is given below:
`INSERT INTO ORDERS (id, ...) VALUES (order_seq.NEXTVAL, ...);`

The sequence values can be set dynamically with the SET SEQUENCE statement.
`SET SEQUENCE <sequence_name> VALUE <BIGINT_value>`

Sequences in stored procedures

Sequences can also be used inside stored procedures.

Table 14. Sequences in stored procedures

Stored procedure statement	Usage
<code>EXEC SEQUENCE <i>sequence_name</i>.CURRENT INTO <i>variable</i></code>	Retrieves a new sequence value
<code>EXEC SEQUENCE <i>sequence_name</i>.NEXT INTO <i>variable</i></code>	Retrieves the current sequence value
<code>EXEC SEQUENCE <i>sequence_name</i> SET VALUE USING <i>variable</i></code>	Sets sequence values Tip: Alternatively, you can use the SET SEQUENCE <name> VALUE <bigint value> statement.

The following example is a stored procedure that retrieves a new sequence number:

```
"CREATE PROCEDURE get_my_seq
RETURNS (val INTEGER)
BEGIN
EXEC SEQUENCE my_sequence.NEXT INTO (val);
END";
```

3.5 Events

Event alerts are database objects that are used to signal events in solidDB databases. Together with stored procedures, the events can be used for automating administrative. You can make your application use event alerts instead of polling, which uses more resources.

The events mechanism is based on one connection *waiting* on an event until another connection *posts* that event. More than one connection may wait on the same event. If multiple connections wait on the same event, all waiting connections are notified when the event is posted. A connection may also wait on multiple events, in which case it will be notified when any of those events are posted.

There are two types of events:

- User-defined events

User-defined events can only be used within stored procedures.

User-defined event objects are created with the CREATE EVENT statement and removed with the DROP EVENT statement.

The CREATE EVENT defines the event name and a set of parameters.

```
CREATE EVENT event_name [(parameter_name datatype)]
```

The name can be any user-specified alphanumeric string.

Event parameters must be local variables or parameters in the stored procedure where the event is triggered.

- System events

System events can be used within stored procedures; additionally, you can also wait on a system event outside a stored procedure using the ADMIN EVENT command.

The system event names and parameters are described in Appendix H, "System events," on page 401.

When an application calls a stored procedure that waits for a specific event to happen, the application is blocked until the event is posted and received. In multi-threaded environments, separate threads and connections can be used to access the database during the event wait. All clients that are waiting for the posted event will receive the event.

Events are always posted (sent) inside stored procedures using the POST EVENT stored procedure statement. User events are also received inside stored procedures.

```
post_statement ::= POST EVENT event_name [( parameters)]
```

Each connection has its own event queue. The events to be collected in the event queue are specified (registered) with the REGISTER EVENT stored procedure statement and removed from the queue with the UNREGISTER EVENT statement.

```
wait_register-statement ::= REGISTER EVENT | UNREGISTER EVENT event_name
```

Procedures are made to wait for an event to happen with the WAIT EVENT stored procedure statement:

```
wait_event_statement ::=
  WAIT EVENT
  [event_specification...]
  END WAIT
event_specification ::=
  WHEN event_name [(parameters)] BEGIN
    statements
  END EVENT
```

You may wait on either system-defined events or user-defined events.

Tip: With system events, you may also wait on an event without using a stored procedure by using the ADMIN EVENT statement. However, you cannot post events using ADMIN EVENT.

```
ADMIN EVENT 'register sys_event_hsbstateswitch';
ADMIN EVENT 'wait';
```

If you want to stop the stored procedure waiting for an event, you can use ODBC function SQLCancel() called from a separate thread in the client application. This function cancels executing statements. Alternatively, you can create a specific user event and send it. The waiting stored procedure must be modified to wait for this additional event. The client application recognizes this event and exits the waiting loop.

Access rights and events

The creator of an event or the database administrator can grant and revoke access rights on that event. Access rights can be granted to users and roles. If a user has "SELECT" access right on an event, then the user has the right to wait on that event. If a user has the INSERT access right on an event, then the user may post that event.

Related reference:

A.2, "ADMIN EVENT," on page 156

A.11, "CREATE EVENT," on page 172

A.31, "DROP EVENT," on page 220

A.15, "CREATE PROCEDURE," on page 180

3.5.1 Using events - Example 1

This example uses a pair of SQL scripts to show how to use events.

This example contains two scripts. Script 1 waits on an event and Script 2 posts the event. Once the event has been posted, the event that is waiting will finish waiting and move on to the next command.

To execute this example code, you will need two consoles so that you can start the WaitOnEvent.sql script and then run the PostEvent.sql script while WaitOnEvent.sql is waiting.

In this example, the stored procedure that waits does not actually do anything after the event has posted; the script merely finishes the wait and returns to the caller. The caller can then proceed to do whatever it wants, which in this case is to SELECT the record that was inserted while we were waiting.

This example waits for only a single event, which is called "record_was_inserted". Later in this chapter we will have another script that waits for multiple events using a single "WAIT".

```
===== SCRIPT 1=====
-- SCRIPT NAME: WaitOnEvent.sql
-- PURPOSE:
-- This is one of a set of scripts that demonstrates posting events
-- and waiting on events. The sequence of steps is shown below:
--
-- THIS SCRIPT (WaitOnEvent.sql) PostEvent.sql script
-----
-- CREATE EVENT.
-- CREATE TABLE.
-- WAIT ON EVENT.
--   Insert a record into table.
--   Post event.
-- SELECT * FROM TABLE.
--
-- To perform these steps in the proper order, start running this
-- script FIRST, but remember that this script does not finish running
-- until after the post_event script runs and posts the event.
-- Therefore, you will need two open consoles so that you can leave
-- this running/waiting in one window while you run the other script
-- (post_event) in the other window.
-- Create a simple event that has no parameters.
-- Note that this event (like any event) does not have any
-- commands or data; the event is just a label that allows both the
-- posting process and the waiting process to identify which event has
-- been posted (more than one event may be registered at a time).
-- As part of our demonstration of events, this particular event
-- will be posted by the other user after he or she inserted a record.
CREATE EVENT record_was_inserted;
-- Create a table that the other script will insert into.
CREATE TABLE table1 (int_col INTEGER);
-- Create a procedure that will wait on an event
-- named "record_was_inserted".
-- The other script (PostEvent.sql) will post this event.
"CREATE PROCEDURE wait_for_event
BEGIN
-- If possible, avoid holding open a transaction. Note that in most
-- cases it's better to do the COMMIT WORK before the procedure,
```

```

-- not inside it. See "Waiting on Events" at the end of this example.
EXEC SQL COMMIT WORK;
-- Now wait for the event to be posted.
WAIT EVENT
  WHEN record_was_inserted BEGIN
  -- In this demo, we simply fall through and return from the
  -- procedure call, and then we continue on to the next
  -- statement after the procedure call.
  END EVENT
END WAIT;
END";
-- Call the procedure to wait. Note that this script will not
-- continue on to the next step (the SELECT) until after the
-- event is posted.
CALL wait_for_event();
COMMIT WORK;
-- Display the record inserted by the other script.
SELECT * FROM table1;

```

Guidelines for Committing Transaction in Script 1 (WaitOnEvent.sql)

Whenever possible, complete any current transaction before waiting on an event. If you execute a WAIT inside a transaction, then the transaction will be held open until the event occurs and the next COMMIT or ROLLBACK is executed. This means that during the wait, the server will hold locks, which may lead to excessive bonsai tree growth. For details on the Bonsai Tree and preventing its growth, read the section "Reducing Bonsai Tree Size by Committing Transactions," in *solidDB Administration Guide*.

In this example, we have put COMMIT WORK inside the procedure immediately before the WAIT. However, this is not usually a good solution; putting the COMMIT or ROLLBACK inside the "wait" procedure means that if the procedure is called as part of another transaction, then the COMMIT or ROLLBACK will terminate that enclosing transaction and start a new transaction, which is probably not what you want. If, for example, you were entering data into a "child" table with a referential constraint and you are waiting for the referenced data to be entered into the "parent" table, then breaking the transaction into two transactions would simply cause the insert of the "child" record to fail because the parent would not have been inserted yet.

The best strategy is to design your program so that you do not need to WAIT inside a transaction; instead, your "wait" procedure should be called between transactions if that is possible. By using events/waits, you have some control over the order in which things are done and you can use this to help ensure that dependencies are met without actually putting everything into a single transaction. For example, in an "asynchronous" situation you might be waiting for both a child and a parent record to be inserted, and if your database server did not have the "events" feature, then you might require that both records be inserted in the same transaction so that you could ensure referential integrity.

By using events/waits, you can ensure that the insertion of the parent is done first; you can then put the insertion of the child record in a second transaction because you can guarantee that the parent will always be present when the child is inserted. (To be more precise, you can ALMOST guarantee that the parent will be present when the child is inserted. If you break up the insertions into two different transactions, then even if you ensure that the parent is inserted before the child, there is a slight chance that the parent would be deleted before the program tried to insert the child record.)

```

===== SCRIPT 2=====
-- SCRIPT NAME: PostEvent.sql
-- PURPOSE:
-- This script is one of a set of scripts that demonstrates posting
-- events and waiting on events. The sequence of steps is shown below:
--
-- WaitOnEvent.sql THIS SCRIPT (PostEvent.sql)
-----
-- Create event.
-- Create table.
-- Wait on event.
--   INSERT A RECORD INTO TABLE.
--   POST THE EVENT.
-- Select * from table.
-- Insert a record into the table.
INSERT INTO table1 (int_col) VALUES (99);
COMMIT WORK;
-- Create a stored procedure to post the event.
"CREATE PROCEDURE post_event
BEGIN
  -- Post the event.
  POST EVENT record_was_inserted;
END";
-- Call the procedure that posts the event.
CALL post_event();
DROP PROCEDURE post_event;
COMMIT WORK;

```

3.5.2 Using events - Example 2

This example shows how to write a stored procedure that will wait on multiple events and that will finish the wait when any one of those events is posted.

This example contains two scripts. Script 1 waits on more than one events. Script 2 is a stored procedure.

```

===== SCRIPT 1=====
-- SCRIPT NAME: MultiWaitExamplePart1.sql
-- PURPOSE:
-- This code shows how to wait on more than one event.
-- If you run this demonstration, you will see that a "wait" lasts only
-- until one of the events is received. Thus a wait on multiple events
-- is like an "OR" (rather than an "AND"); you wait until event1 OR
-- event2 OR ... occurs.
--
-- This demo uses 2 scripts, one of which waits for an event(s) and one
-- of which posts an event.
-- To run this example, you will need 2 consoles.
-- 1) Run this script (MultiWaitExamplePart1.sql) in one window. After
-- this script reaches the point where it is waiting for the event, then
-- start Step 2.
-- 2) Run the script MultiWaitExamplePart2.sql in the other window.
-- This will post one of the events.
-- After the event is posted, the first script will finish.
-- Create the 3 different events on which we will wait.
CREATE EVENT event1;
CREATE EVENT event2(i INTEGER);
CREATE EVENT event3(i INTEGER, c CHAR(4));
-- When an event is received, the process that is waiting on the event
-- will insert a record into this table. That lets us see which events
-- were received.
CREATE TABLE event_records(event_name CHAR(10));
-- This procedure inserts a record into the event_records table.
-- This procedure is called when an event is received.
"CREATE PROCEDURE insert_a_record(event_name_param CHAR(10))
BEGIN
  EXEC SQL PREPARE insert_cursor

```

```

INSERT INTO event_records (event_name) VALUES (?);
EXEC SQL EXECUTE insert_cursor USING (event_name_param);
EXEC SQL CLOSE insert_cursor;
EXEC SQL DROP insert_cursor;
END";
-- This procedure has a single "WAIT" command that has 3 subsections;
-- each subsection waits on a different event.
-- The "WAIT" is finished when ANY of the events occur, and so the
-- event_records table will hold only one of the following:
-- "event1",
-- "event2", or
-- "event3".
"CREATE PROCEDURE event_wait(i1 INTEGER)
RETURNS (eventresult CHAR(10))
BEGIN
  DECLARE i INTEGER;
  DECLARE c CHAR(4);
  -- The specific values of i and c are irrelevant in this example.
  i := i1;
  c := 'mark';
  -- Set eventresult to an empty string.
  eventresult := '';
  -- Will we exit after any of these 3 events are posted, or must
  -- we wait until all of them are posted? The answer is that
  -- we will exit after any one event is posted and received.
  WAIT EVENT
  -- When the event named "event1" is received...
  WHEN event1 BEGIN
    eventresult := 'event1';
    -- Insert a record into the event_records table showing that
    -- this event was posted and received.
    EXEC SQL PREPARE call_cursor
    CALL insert_a_record(?);
    EXEC SQL EXECUTE call_cursor USING (eventresult);
    EXEC SQL CLOSE call_cursor;
    EXEC SQL DROP call_cursor;
    RETURN;
  END EVENT
  WHEN event2(i) BEGIN
    eventresult := 'event2';
    EXEC SQL PREPARE call_cursor2
    CALL insert_a_record(?);
    EXEC SQL EXECUTE call_cursor2 USING (eventresult);
    EXEC SQL CLOSE call_cursor2;
    EXEC SQL DROP call_cursor2;
    RETURN;
  END EVENT
  WHEN event3(i, c) BEGIN
    eventresult := 'event3';
    EXEC SQL PREPARE call_cursor3
    CALL insert_a_record(?);
    EXEC SQL EXECUTE call_cursor3 USING (eventresult);
    EXEC SQL CLOSE call_cursor3;
    EXEC SQL DROP call_cursor3;
    RETURN;
  END EVENT
  END WAIT
END";
COMMIT WORK;
-- Call the procedure that waits until one of the events is posted.
CALL event_wait(1);
-- See which event was posted.
SELECT * FROM event_records;

===== SCRIPT 2 =====
-- SCRIPT NAME: MultiWaitExamplePart2.sql
-- PURPOSE:

```

```

-- This is script 2 of 2 scripts that show how to wait for multiple
-- events. See the instructions at the top of MultiWaitExamplePart1.sql.
-- Create a stored procedure to post an event.
"CREATE PROCEDURE post_event1
BEGIN
  -- Post the event.
  POST EVENT event1;
END";
--Create a stored procedure to post the event.
"CREATE PROCEDURE post_event2(param INTEGER)
BEGIN
  -- Post the event.
  POST EVENT event2(param);
END";
--Create a stored procedure to post the event.
"CREATE PROCEDURE post_event3(param INTEGER, s CHAR(4))
BEGIN
  -- Post the event.
  POST EVENT event3(param, s);
END";
COMMIT WORK;
-- Notice that to finish the "wait", only one event needs to be posted.
-- You may execute any one of the following 3 CALL commands to post an
-- event.
-- We've commented out 2 of them; you may change which one is de
-- commented.
CALL post_event1();
--CALL post_event2(2);
--CALL post_event3(3, 'mark');

```

3.5.3 Using events - Example 3

This example shows a simple usage of the REGISTER EVENT and UNREGISTER EVENT commands.

Example 1 and Example 2 do not use REGISTER EVENT, yet their WAIT commands succeeded anyway. The reason for this is that when you wait on an event, you will be registered implicitly for that event if you did not already explicitly register for it. Thus you only need to explicitly register events if you want them to start being queued now but you do not want to start WAITING for them until later.

```

CREATE EVENT e0;
CREATE EVENT e1 (param1 int);
COMMIT WORK;

-- Create a procedure to register the events to that when they occur
-- they are put in this connection's event queue.
"CREATE PROCEDURE eeregister
BEGIN
  REGISTER event e0;
  REGISTER EVENT e1;
END";

CALL eeregister;
COMMIT WORK;

-- Create a procedure to post the events.
"CREATE PROCEDURE eepost
BEGIN
  DECLARE x int;
  x := 1;
  POST EVENT e0;
  POST EVENT e1(x);
END";

```

```

COMMIT WORK;

-- Post the events. Even though we haven't yet waited on the events,
-- they will be stored in our queue because we registered for them.
CALL eepost;
COMMIT WORK;

-- Now create a procedure to wait for the events.
"CREATE PROCEDURE eewait
  RETURNS (whichEvent VARCHAR(100))
  BEGIN
  DECLARE i INT;

  WAIT EVENT
    WHEN e0 BEGIN
      whichEvent := 'event0';
    END EVENT

    WHEN e1(i) BEGIN
      whichEvent := 'event1';
    END EVENT

  END WAIT

END";

COMMIT WORK;

-- Since we already registered for the 2 events and we already
-- posted the 2 events, when we call the eewait procedure twice
-- it should return immediately, rather than waiting.
CALL eewait;
CALL eewait;
COMMIT WORK;

-- Unregister for the events.
"CREATE PROCEDURE eeunregister
  BEGIN
  UNREGISTER event e0;
  UNREGISTER EVENT e1;
END";

CALL eeunregister;
COMMIT WORK;
CREATE EVENT e0;
CREATE EVENT e1 (param1 int);
COMMIT WORK;

-- Create a procedure to register the events to that when they occur
-- they are put in this connection's event queue.
"CREATE PROCEDURE eeregister
  BEGIN
  REGISTER event e0;
  REGISTER EVENT e1;
END";

CALL eeregister;
COMMIT WORK;

-- Create a procedure to post the events.
"CREATE PROCEDURE eepost
  BEGIN
  DECLARE x int;
  x := 1;
  POST EVENT e0;
  POST EVENT e1(x);

```

```

END";

COMMIT WORK;

-- Post the events. Even though we haven't yet waited on the events,
-- they will be stored in our queue because we registered for them.
CALL eepost;
COMMIT WORK;

-- Now create a procedure to wait for the events.
"CREATE PROCEDURE eewait
  RETURNS (whichEvent VARCHAR(100))
  BEGIN
  DECLARE i INT;

  WAIT EVENT
  WHEN e0 BEGIN
    whichEvent := 'event0';
  END EVENT

  WHEN e1(i) BEGIN
    whichEvent := 'event1';
  END EVENT

  END WAIT

END";

COMMIT WORK;

-- Since we already registered for the 2 events and we already
-- posted the 2 events, when we call the eewait procedure twice
-- it should return immediately, rather than waiting.
CALL eewait;
CALL eewait;
COMMIT WORK;

-- Unregister for the events.
"CREATE PROCEDURE eeunregister
  BEGIN
  UNREGISTER event e0;
  UNREGISTER EVENT e1;
END";

CALL eeunregister;
COMMIT WORK;

```

4 Using SQL for database administration

You can manage the database and its users and schema using SQL statements. The SQL-based management tasks include managing roles and privileges, tables, indexes, transactions, catalogs, and schemas.

4.1 Using solidDB SQL syntax

The solidDB SQL syntax is based on the ANSI X3H2-1989 (SQL-89) level 2 standard including important SQL-92 and SQL-99 extensions.

To execute SQL statements, you can use the solidDB SQL Editor (**solsql**) or solidDB SQL Remote Control (**solcon**), or various ODBC or JDBC compliant tools. Additionally, you can execute SQL statements from a file, for example, to automate tasks. You can also use such files for rerunning your SQL statements later or as a document of your users, tables, and indexes.

Remember: When using **solsql**, SQL statements must be terminated with a semicolon (;). Otherwise, terminating SQL statements with a semicolon leads to a syntax error.

Related reference:

Appendix A, “IBM solidDB SQL statements,” on page 141

This section provides descriptions of the solidDB SQL statements, including examples.

4.1.1 solidDB SQL data types

solidDB SQL supports data types specified in the SQL-92 Standard Entry Level specifications, as well as important Intermediate Level enhancements. Refer to Appendix C, “Data types,” on page 323, for a complete description of the supported data types.

You can also define some data types with the optional length, scale, and precision parameters. In that case, the default properties of the corresponding data type are not used.

4.1.2 solidDB ADMIN COMMANDS

The solidDB SQL includes the extension **ADMIN COMMAND** *'command [command_args]'* that is used for performing basic administrative tasks, such as backups, performance monitoring, and shutdown.

To access a short description of available ADMIN COMMANDs, execute **ADMIN COMMAND 'help'**.

Related information:

A.1, “ADMIN COMMAND,” on page 141

4.1.3 Scalar functions

The solidDB server supports the use of direct function names.

For example:

```
SELECT substring(line, 1,4) FROM test;
```

However, functions whose names match reserved words must be enclosed with escape characters (double quotation marks). For example:

```
SELECT "left"(line,4) FROM test;
```

or:

```
SELECT {fn left(line,4)} FROM test;
```

The latter example corresponds to the ODBC implementation-independent syntax. It can be used in all interfaces.

4.2 Managing user privileges and roles

You can use solidDB teletype tools, and many ODBC compliant SQL tools to modify user privileges. Users and roles are created and deleted using SQL statements or commands. A file consisting of several SQL statements is called an SQL script.

In the `samples/sql` directory, you will find the SQL script `sample.sql`, which gives an example of creating users and roles. You can run it by using `solsql`. To create your own users and roles, you can make your own script describing your user environment.

4.2.1 User privileges

When using solidDB databases in a multi-user environment, you may want to apply user privileges to hide certain tables from some users. For example, you may not want an employee to see the table in which employee salaries are listed, or you may not want other users to change your test tables.

You can apply five different kinds of user privileges. A user may be able to view, delete, insert, update or reference information in a table or view. Any combination of these privileges may also be applied. A user who has none of these privileges to a table is not able to use the table at all.

Note: Once user privileges are granted, they take effect when the user who is granted the privileges logs on to the database. If the user is already logged on to the database when the privileges are granted, they take effect only if the user:

- accesses for the first time the table or object on which the privileges are set, or
- disconnects and then reconnects to the database.

4.2.2 User roles

Privileges can also be granted to an entity called a role.

A *role* is a group of privileges that can be granted to users as one unit. You can create roles and assign users to certain roles. A single user may have more than one role assigned, and a single role may have more than one user assigned.

Note:

- The same string cannot be used both as a user name and a role name.
- Once a user role is granted, it takes effect when the user who is granted the role logs on to the database. If the user is already logged on to the database when the role is granted, the role takes effect when the user disconnects and then reconnects to the database.

System roles

solidDB offers the following system roles. The system role names are reserved user names.

Table 15. System roles

Reserved Names	Description
PUBLIC	This role grants privileges to all users. When user privileges to a certain table are granted to the role <i>PUBLIC</i> , all current and future users have the specified user privileges to this table. This role is granted automatically to all users.
SYS_ADMIN_ROLE	This is the default role for the database administrator. This role has administration privileges to all tables, indexes and users, as well as the right to use solidDB Remote Control. This is also the database creator role.
_SYSTEM	This is the schema name of all system tables and views.
SYS_CONSOLE_ROLE	This role has the right to use solidDB Remote Control, but does not have other administration privileges.
SYS_SYNC_ADMIN_ROLE	This is the administrator role for data synchronization functions.
SYS_SYNC_REGISTER_ROLE	This role is only for registering and unregistering a replica database to the master.

4.2.3 Creating users

```
CREATE USER username IDENTIFIED BY password;
```

Only an administrator has the privilege to execute this statement. The following example creates a new user named CALVIN with the password HOBBS.

```
CREATE USER CALVIN IDENTIFIED BY HOBBS;
```

4.2.4 Deleting users

```
DROP USER username;
```

Only an administrator has the privilege to execute this statement. The following example deletes the user named CALVIN.

```
DROP USER CALVIN;
```

4.2.5 Changing a password

You can change your password with the ALTER USER command. Administrators can also change passwords of other users.

Procedure

Use the following syntax to change the password of a given user:

```
ALTER USER username IDENTIFIED BY new_password
```

Example

The following statement changes the password of the user CALVIN to GUBBES
ALTER USER CALVIN IDENTIFIED BY GUBBES

Related tasks:

4.2.13, “Changing administrator's user name and password,” on page 93
The database system administrator's password can be changed with the ALTER USER command. The administrator's user name cannot be changed with the ALTER USER command; the user name can be changed by granting the SYS_ADMIN_ROLE to another user and then dropping the original administrator user account.

4.2.6 Creating roles

```
CREATE ROLE rolename;
```

The following example creates a new user role named GUEST_USERS.

```
CREATE ROLE GUEST_USERS;
```

4.2.7 Deleting roles

```
DROP ROLE role_name;
```

The following example deletes the user role named GUEST_USERS.

```
DROP ROLE GUEST_USERS;
```

4.2.8 Granting privileges to a user or a role

```
GRANT user_privilege ON table_name TO username or role_name ;
```

The possible user privileges on tables are SELECT, INSERT, DELETE, UPDATE, REFERENCES and ALL. ALL provides a user or a role all five privileges mentioned above. A new user has no privileges until they are granted.

The following example grants INSERT and DELETE privileges on a table named TEST_TABLE to the GUEST_USERS role.

```
GRANT INSERT, DELETE ON TEST_TABLE TO GUEST_USERS;
```

The EXECUTE privilege provides a user the right to execute a stored procedure:

```
GRANT EXECUTE ON procedure_name TO username or role_name ;
```

The following example grants EXECUTE privilege on a stored procedure named SP_TEST to user CALVIN.

```
GRANT EXECUTE ON SP_TEST TO CALVIN;
```

4.2.9 Granting privileges to a user by giving the user a role

```
GRANT role_name TO username ;
```

The following example gives the user CALVIN the privileges that are defined for the GUEST_USERS role.

```
GRANT GUEST_USERS TO CALVIN;
```

4.2.10 Revoking privileges from a user or a role

```
REVOKE user_privilege ON table_name FROM username or role_name ;
```

The following example revokes the INSERT privilege on the table named TEST_TABLE from the GUEST_USERS role.

```
REVOKE INSERT ON TEST_TABLE FROM GUEST_USERS;
```

4.2.11 Revoking privileges by revoking the role of a user

```
REVOKE role_name FROM username ;
```

The following example revokes the privileges that are defined for the GUEST_USERS role from CALVIN.

```
REVOKE GUEST_USERS FROM CALVIN;
```

4.2.12 Granting administrator privileges to a user

```
GRANT SYS_ADMIN_ROLE TO username ;
```

The following example grants administrator privileges to CALVIN, who now has all privileges to all tables.

```
GRANT SYS_ADMIN_ROLE TO CALVIN;
```

You may also want to grant a user the right to perform data synchronization operations. To do this, execute the command:

```
GRANT SYS_SYNC_ADMIN_ROLE TO HOBBS
```

Note:

If the autocommit mode is set OFF, you need to commit your work. To commit your work use the following SQL statement: COMMIT WORK; If the autocommit mode is set ON, the transactions are committed automatically.

4.2.13 Changing administrator's user name and password

The database system administrator's password can be changed with the ALTER USER command. The administrator's user name cannot be changed with the ALTER USER command; the user name can be changed by granting the SYS_ADMIN_ROLE to another user and then dropping the original administrator user account.

About this task

The database system administrator account is created when the solidDB database is created. The creator of the database has the SYS_ADMIN_ROLE user role.

Important: You must remember your user name and password to be able to connect to solidDB. There are no default user names. The administrator user name you enter when creating the database is the only user name available for connecting to the new database for the first time. If you cannot connect to solidDB because you have forgotten your system administrator user name or password, contact IBM Software Support.

Procedure

- **Changing administrator's password**

To change the administrator's password, issue the following command:

```
ALTER USER username identified by new_password
```

- **Changing administrator's user name**

To change the administrator's user name:

1. Create a user with the following command:

```
CREATE USER username IDENTIFIED BY password
```
2. Grant the new user the SYS_ADMIN_ROLE role with the following command:

```
GRANT SYS_ADMIN_ROLE TO username
```
3. Drop the original administrator account with the following command:

```
DROP USER original_dba_username
```

Related tasks:

4.2.5, "Changing a password," on page 91

You can change your password with the ALTER USER command. Administrators can also change passwords of other users.

Related reference:

A.6, "ALTER USER," on page 162

A.25, "CREATE USER," on page 216

A.52, "GRANT," on page 234

A.47, "DROP USER," on page 229

4.3 Managing tables

The solidDB database tables are managed using SQL commands.

In the solidDB server installation directory, you can find an SQL script named `sample.sql`, which gives an example of managing tables. You can run the script using, for example, solidDB SQL Editor (**solsql**).

The SQL statements for creating, deleting, and altering tables are described in details in Appendix A, "IBM solidDB SQL statements," on page 141.

4.3.1 Accessing system tables

The solidDB system tables store solidDB server information, including user information. Your ability to access specific system tables depends on your user's role and access rights. For example, DBAs can view all information about all stored procedures, including the procedure definition text (i.e. the CREATE PROCEDURE statement). Normal users can see the stored procedures, including the procedure definition text, for procedures that they have created. Normal users who have execute access on a stored procedure, but who did not create that stored procedure, may look at some information about that stored procedure but may not see the procedure definition text. For a list of system tables, refer to Appendix E, "Database system tables and system views," on page 345.

The table below provides the viewing access and/or object granting privileges for specific system tables and their data by user role and user access rights.

Note that a "User with access rights" in this table refers to a normal user who has any one of the following rights: INSERT, UPDATE, DELETE, or SELECT access.*

Table 16. Viewing tables and granting access

Tasks	DBA	Owner	User with access rights*	User with no access rights
Viewing SYS_TABLES	All (no restrictions)	All (no restrictions)	All (no restrictions)	All (no restrictions)

Table 16. Viewing tables and granting access (continued)

Tasks	DBA	Owner	User with access rights*	User with no access rights
Viewing User tables in SYS_TABLES	All (no restrictions)	Restricted to the owners' tables only	All tables to which the user has INSERT, UPDATE, DELETE, SELECT, or REFERENCES access rights.	No tables can be viewed.
Viewing SYS_COLUMNS	All (no restrictions)	Columns in the owner's tables	Columns in tables to which the user has INSERT, UPDATE, DELETE, SELECT, or REFERENCES access rights.	No columns can be viewed.
Viewing SYS_PROCEDURES (excluding the procedure definition text — i.e. the text of the CREATE PROCEDURE statement)	All (no restrictions)	Those procedures created by the user (owner).	Those procedures in which the user has execute access.	No procedures can be viewed.
Viewing Procedure definition text in SYS_PROCEDURES	All (no restrictions)	Those procedures created by the user (owner)	Note that execute access does not allow the user to see the procedure definition text.	No procedures or procedure definition text can be viewed.
Ability to Grant Access rights on procedures	Yes	Yes	No	No
Viewing SYS_TRIGGERS	All (no restrictions)	Those triggers created by the user (owner)	None	No triggers can be viewed.
Viewing Trigger definition text in SYS_TRIGGERS	All (no restrictions)	Those triggers created by the user (owner)	None	No triggers can be viewed.

4.3.2 Creating tables

```
CREATE TABLE table_name (column_name column_type
    [, column_name column_type]...);
```

All users have privileges to create tables.

The following example creates a new table named TEST with the column I of the column type INTEGER and the column TEXT of the column type VARCHAR.

```
CREATE TABLE TEST (I INTEGER, TEXT VARCHAR);
```

4.3.3 Removing tables

```
DROP TABLE table_name;
```

Only the creator of the particular table or users having SYS_ADMIN_ROLE have privileges to remove tables.

The following example removes the table named TEST.

```
DROP TABLE TEST;
```

Note:

For catalogs and schemas: The ANSI standard for SQL defines the keywords RESTRICT and CASCADE. When dropping a catalog or a schema, if you use the keyword RESTRICT, then you cannot drop a catalog or schema if it contains other database objects (for example, tables). Using the keyword CASCADE allows you to drop a catalog or schema that still contains database objects — the database objects that it contains will automatically be dropped. The default behavior (if you do not specify either RESTRICT or CASCADE) is RESTRICT.

For database objects other than Catalogs and Schemas: The keywords RESTRICT and CASCADE are not accepted as part of most DROP statements in solidDB SQL. Furthermore, for these database objects, the rules are more complex than simply "pure CASCADE" or "pure RESTRICT" behavior, but generally objects are dropped with drop behavior RESTRICT. For example, if you try to drop table1 but table2 has a foreign key dependency on table1, or if there are publications that reference table1, then you will not be able to drop table1 without first dropping the dependent table or publication. However, the server does not use RESTRICT behavior for all possible types of dependency. For example, if a view or a stored procedure references a table, the referenced table can still be dropped, and the view or stored procedure will fail the next time that it tries to reference that table. Also, if a table has a corresponding synchronization history table, that synchronization history table will be dropped automatically. For more information about synchronization history tables, see *solidDB Advanced Replication Guide*.

4.3.4 Adding columns to a table

```
ALTER TABLE table_name ADD COLUMN column_name column_type;
```

Only the creator of the particular table or users having SYS_ADMIN_ROLE have privileges to add or delete columns in a table.

The following example adds the column C of the column type CHAR(1) to the table TEST.

```
ALTER TABLE TEST ADD COLUMN C CHAR(1);
```

4.3.5 Deleting columns from a table

```
ALTER TABLE table_name DROP COLUMN column_name;
```

A column cannot be dropped if it is part of a unique constraint or primary key. For details on primary keys, read 4.4, "Managing indexes," on page 97.

The following example statement deletes the column C from the table TEST.

```
ALTER TABLE TEST DROP COLUMN C;
```

Note:

If the autocommit mode is set OFF, you need to commit your work before you can modify the data in the table you altered. To commit your work after altering a table, use the following SQL statement:

```
COMMIT WORK;
```

If the autocommit mode is set ON, then all statements, including DDL (Data Definition Language) statements, are committed automatically.

4.3.6 Viewing table names and definitions

To view the table names and definitions in your database, use the LIST and DESCRIBE statements.

Procedure

- To list all user-created tables, issue the following command:

```
LIST TABLES
```

- To list all system and user-created tables, issue the following command:

```
SELECT * FROM TABLES
```

TABLES is a system view that contains all the tables. In the output, the table names can be found in the column TABLE_NAME.

- To view the basic table information, issue the following command:

```
DESCRIBE TABLE <table_name> NICE
```

DESCRIBE TABLE *table_name* with the keyword NICE prints the table name, the storage type used, list of parent and child tables, columns, column types, column precisions, column nullability, whether column is primary key and whether column is sec key. In RAW format, the create statement is printed along with create statements of associated indexes and constraints.

- To view the table definition, issue the following command:

```
DESCRIBE TABLE <table_name> RAW
```

DESCRIBE TABLE *table_name* with the keyword RAW prints the CREATE TABLE create statement for the table, including the CREATE statements of associated indexes and constraints.

4.4 Managing indexes

Indexes are used to speed up access to tables. The database engine uses indexes to access the rows in a table directly. Without indexes, the engine would have to search the whole contents of a table to find the desired row. You can create as many indexes as you like on a single table; however, adding indexes does slow down write operations, such as inserts, deletes, and updates on that table. For details on creating indexes to improve performance, read 7.3, “Using indexes to improve query performance,” on page 135.

There are two kinds of indexes: non-unique indexes and unique indexes. A unique index is an index where all key values are unique. A unique index is always created, when the UNIQUE restraint is used when creating an index.

You can create and delete indexes using the SQL statements.

4.4.1 Primary key indexes

To retrieve a single specific record from a table, the record must be uniquely identified. The solidDB server uses *primary keys* to uniquely identify each record in each table. A primary key is a column or combination of columns that contains a unique value or combination of values. Each table has a primary key — either explicit or implicit.

The solidDB server creates a *primary key index* automatically based on the field or fields of the primary key. A primary key index, like any index, speeds up access to data in the table. Unlike other indexes, however, a primary key index also controls

the order in which records are stored in the database. This is called *clustering*. Records are stored in ascending order based on the primary key values.

If the creator of the table does not specify a primary key, the server creates a primary key for the table automatically. To ensure uniqueness in that primary key, the server uses a hidden internal row identifier. The value of that row identifier may be retrieved and used in queries by way of a symbolic pseudo column name ROWID.

Note: Always create a primary key at table creation. In solidDB, you cannot add an explicit primary key after the table has been created. If a primary key is not specified by a user, the most efficient query method is not available (unless ROWID is used) for that table. Also, such a table cannot be used in referential integrity constraints as a referenced table.

Once a primary key is defined (whether by the table creator or by the server), the server prevents rows with duplicate primary key values from being inserted into the table.

Primary key indexes cannot be dropped.

4.4.2 Secondary key indexes

Since indexes speed up searches, it is often helpful for a table to have one index for each attribute (or combination of attributes) that is used frequently in searches. All indexes other than the primary index are called "secondary indexes".

A table may have as many indexes as you like, as long as each index has a unique combination of columns, order of columns, and order of values (ASCending, DESCending). For example, in the code shown below, the third index duplicates the first and will either generate an error message or will waste disk space with duplicate information.

```
CREATE INDEX i1 ON TABLE t1 (col1, col2);
-- The following is ok because although the columns are the same as in
-- index i1, the order of the columns is different.
CREATE INDEX i2 ON TABLE t1 (col2, col1);
-- The following is not ok because index i3 would be exactly the
-- same as index i1.
CREATE INDEX i3 ON TABLE t1 (col1, col2); -- ERROR.
-- The following is ok because although the columns and
-- column order are the same, the order of the index values
-- (ASCending vs. DESCending) is different.
CREATE INDEX i3b ON TABLE t1 (col1, col2) DESC;
```

Note that if one index is a "leading subset" of another (meaning that the columns, column order, and value order of all N columns in index2 are exactly the same as the first N column(s) of index1), then you only need to create the index that is the superset. For example, suppose that you have an index on the combination of DEPARTMENT + OFFICE + EMP_NAME. This index can be used not only for searches by department, office and emp_name together, but also for searches of just the department, or just the department and office together. So there is no need to create a separate index on the department name alone, or on the department and office alone. The same is true for ORDER BY operations; if the ORDER BY criterion matches a subset of an existing index, then the server can use that index.

Keep in mind that if you defined a primary key or unique constraint, that key or constraint is implemented as an index. Thus you never need to create an index that is a "leading subset" of the primary key or of an existing unique constraint; such an index would be redundant.

Note that when searching using a secondary index, if the server finds all the requested data in the index key, the server doesn't need to look up the complete row in the table. (This applies only to "read" operations, i.e. SELECT statements. If the user updates values in the table, then the data rows in the table as well as the values in the index(es) must be updated.)

4.4.3 Protection against duplicate indexes

solidDB contains a protection against duplicate indexes. Occasionally, the recreation of an index (DROP/CREATE) can fail if other indexes were created whereby the original index became a duplicate index.

For example, you could create a table containing five columns, named A, B, C, D, E, and create the following indexes on the table:

- B
- AB
- BCE
- ABC

Index B is used for searching or filtering column B. Index BCE starts with column B. Therefore, queries that use an index for locating column B can use index BCE. The same is the case with indexes AB and ABC. Thus, indexes B and AB are duplicate indexes.

Duplicate indexes have, for example, the following adverse effects:

- The required storage space increases.
- The update performance decreases.
- Backup time increases.

If you attempt to create duplicate indexes, index creation fails and solidDB issues the following error:

```
SOLID Table Error 13199: Duplicate index definition
```

For more information, see Appendix, *Error Codes*, in *IBM solidDB Administrator Guide*.

4.4.4 Creating an index on a table

```
CREATE [UNIQUE] INDEX index_name ON base_table_name  
  column_identifier [ASC | DESC]  
  [, column_identifier [ASC | DESC]] ...
```

Only the creator of the particular table or users having SYS_ADMIN_ROLE have privileges to create or drop indexes.

The following example creates an index named X_TEST on column I of the table TEST.

```
CREATE INDEX X_TEST ON TEST (I);
```

4.4.5 Creating a unique index on a table

```
CREATE UNIQUE INDEX index_name ON table_name (column_name);
```

The following example creates a unique index named UX_TEST on column I of the table TEST.

```
CREATE UNIQUE INDEX UX_TEST ON TEST (I);
```

4.4.6 Deleting an index

```
DROP INDEX index_name;
```

The following example deletes the index named X_TEST.

```
DROP INDEX X_TEST;
```

Note:

After creating or dropping an index, you must commit (or roll back) your work before you can modify the data in the table on which you created or dropped the index.

4.5 Referential integrity

Referential integrity is a concept for ensuring that relationships between database tables remain consistent. The tables between which referential integrity must exist are called *referenced tables* and *referencing tables*.

Referential integrity is enforced by using *foreign keys*. A foreign key is a field in the referencing table that matches a *primary key* column (or other similar unique column) of the referenced table. The foreign key can be used to represent a conceptual relationship of the type 1-to-n, such as, 'an employee belongs to a department'. For example, referential integrity ensures that if the referencing table has a foreign key to the referenced table, you cannot add a record to the referencing table unless there is a corresponding record in the referenced table.

Foreign keys are maintained with referential constraint definitions. The constraints specify what *referential actions* solidDB must take when the constraint is violated. This can happen, for example, when a row with a referenced primary key is deleted from the referenced table.

4.5.1 Primary keys and candidate keys

To enforce referential integrity, a referenced table must contain a primary key (preferable) or candidate keys.

A *primary key* is a column or combination of columns that has the same properties as a unique constraint. Because primary keys are used to identify a row in a table, they must be unique, and must have the NOT NULL attribute. A table cannot have more than one primary key, but it can have multiple unique keys. Primary keys are optional, and can be defined when a table is created or altered.

A primary key is defined with the primary key constraint syntax in the CREATE TABLE statement. For example:

```
CREATE TABLE customers (  
  cust_id INTEGER PRIMARY KEY,  
  name CHAR(24),  
  city CHAR(40));
```

Alternatively, you can define a unique index on a column or a group of columns and enact the NOT NULL constraint for them. Effectively, this will produce a *candidate key*. Using an explicit primary key is however preferable; it can provide performance gains while deriving joins.

4.5.2 Foreign keys

A foreign key is a column (or group of columns) within a table that refers to (or "relates to") a unique value in a referenced table. Each value in the foreign key column must have a matching value in the referenced table.

To ensure that each record in the referencing table references exactly one record in the referenced table, the referenced column(s) in the referenced table must have a primary key constraint or have both unique and not-null constraints. Having a unique index is not sufficient.

Example 1:

In a banking environment, one table might hold customer information (Customers), and another table might hold account information (Accounts). Each account is related to a particular customer, and each customer is identified with a unique ID (CUST_ID). Some customers can have more than one account. The CUST_ID can then serve as the primary key of the Customers table. The Accounts table also contains the CUST_ID information to identify which customer owns a particular account; this makes it possible to look up customer information based on account information. The copy of the CUST_ID in the Accounts table is a foreign key; it references the matching value in the primary key of the Customers table.

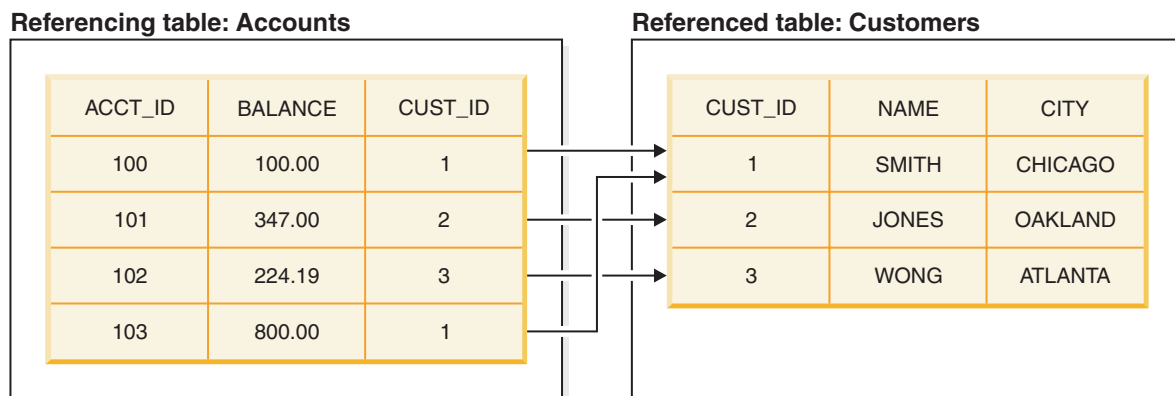


Figure 2. Example: Tables with referential constraints

In the above example, the referencing table Accounts can be created with the following statement:

```
CREATE TABLE accounts (
    acct_id INTEGER PRIMARY KEY,
    balance FLOAT,
    cust_id INTEGER REFERENCES customers);
```

In the REFERENCES clause, only the referenced table is specified, with no referenced column. By default, the primary key is assumed. This is a preferred way that helps avoiding errors while specifying the referenced columns.

In the above example, the primary key and foreign key use a single column. However, primary and foreign keys can be composed of more than one column. Since each foreign key value must exactly match the corresponding primary key value, the foreign key must contain the same number and data type of columns as the primary key, and these key columns must be in the same order.

A foreign key can also have different column names than the primary key. The foreign key and primary key can also have different default values. However, since values in the referenced table must be unique, default values are not much used and are rarely used for columns that are part of a primary key. Default values are also not used very often for foreign key columns.

Although primary key values must be unique, foreign key values are not required to be unique. For example, a single customer at a bank might have multiple accounts. The account ID (ACCT_ID) that appears in the primary key column in the Customers table must be unique; however, the same CUST_ID might occur multiple times in the foreign key column in the ACCOUNTS table. As you can see in the illustration above, customer SMITH has more than one account, and therefore her CUST_ID appears more than once in the foreign key column of the ACCOUNTS table.

Example 2:

In some cases, a foreign key in a table can refer to a primary key in the same table – in such a case, the same table is the referenced table and the referencing table. For example, in a table of employees, each employee record might have a field that contains the ID of the manager (MGR_ID) of that employee. The managers themselves might be stored in the same table. Thus the MGR_ID of that table can be a foreign key that refers to the employee ID (EMP_ID) of that same table. The following diagram illustrates this.

A self-referential table

EMP_ID	MGR_ID	EMP_NAME
1	NULL	ANNAN
10	1	WONG
20	1	SMITH
147	10	JONES
162	20	RAMA

Figure 3. Self-referential constraint

In this example, Rama's manager is Smith (Rama's MGR_ID is 20, and Smith's EMP_ID is 20). Smith reports to Annan (Smith's MGR_ID is 1, and Annan's EMP_ID is 1.) Jones' manager is Wong, and Wong's manager is Annan. If Annan is the president of the company, she does not have a manager, and the value of the foreign key (MGR_ID) is NULL.

If a primary key is composed of more than one column, it should be defined after the columns. For example:

```
CREATE TABLE DEPT (  
    DIVNO INTEGER,  
    DEPTNO INTEGER,  
    DNAME VARCHAR,  
    PRIMARY KEY (DIVNO, DEPTNO));
```

A similar syntax may be used for foreign keys. However, foreign keys should always be defined with the CONSTRAINT syntax that also includes a constraint name. If you have defined a constraint name, you can remove the constraint dynamically with ALTER TABLE statements after tables have been created.

Example of creating a table with a CONSTRAINT name (emp_fk1):

```
CREATE TABLE EMP (  
    EMPNO INTEGER PRIMARY KEY,  
    DIVNO INTEGER,  
    DEPTNO INTEGER,  
    ENAME VARCHAR,  
    CONSTRAINT emp_fk1 FOREIGN KEY (DIVNO, DEPTNO) REFERENCES DEPT);
```

Note: Similarly to other integrity constraints, you can name referential integrity constraints (foreign keys) and manipulate them (drop or add) dynamically, with the ALTER TABLE statement. For more information, see 4.5.4, “Dynamic constraint management,” on page 104.

Defining a foreign key always creates an index on the foreign key column(s). Each time when a referenced record is updated or deleted, the server checks that there are no referencing records that are left without a reference. Foreign key indexes improve the performance of foreign key checking.

Related reference:

A.23, “CREATE TABLE,” on page 207

A.4, “ALTER TABLE,” on page 157

4.5.3 Referential actions

solidDB maintains referential integrity by taking action when referential constraints are violated. Referential constraints are violated, for example, in the following cases:

- A row that contains an invalid foreign key value is inserted in the referencing table.
- A foreign key in the referencing table is updated to an invalid value..
- A row with a referenced primary key is deleted from the referenced table..
- A referenced primary key is updated in the referenced table.

The referential action is defined as part of the table_constraint_definition in a CREATE TABLE or ALTER TABLE statement. The following actions (options) are available when the constraint is violated:

NO ACTION

This option restricts the operation, or rolls back the operation that violates the referential integrity constraint. . The NO ACTION option allows changes to a table that temporarily violate a referential constraint; if you have a requirement that the table state must never violate any constraint even temporarily, use the RESTRICT option.

CASCADE

In the case of operations performed on the referenced table, this option cascades the operations on the referenced table down to the referencing tables. This includes deleting all the referencing rows (a cascading delete) and updating all the referencing foreign-key values (a cascading update).

No cycles are allowed in CASCADE referential actions. An effort to create a cycle composed of foreign keys that have cascading actions results in an error.

For any two tables, at most one CASCADE UPDATE path between them can be defined. The restriction does not apply to CASCADE DELETE.

SET DEFAULT

In the case of operations performed on the referenced table, this option sets the referencing columns to a predefined default value.

SET NULL

In the case of operations performed on the referenced table, this option sets the referencing columns to null.

RESTRICT

This option restricts all operations violating the constraint.

Referential integrity actions sometimes allow changes to a table that temporarily violate a referential constraint. The NO ACTION option allows such violations.

If no option is specified, the default NO ACTION is used.

Related reference:

A.23, "CREATE TABLE," on page 207

A.4, "ALTER TABLE," on page 157

4.5.4 Dynamic constraint management

Constraints can be managed dynamically with the ALTER TABLE clause. The sub-clauses that can be used are:

- ADD CONSTRAINT. This clause adds a named constraint to a table.
- DROP CONSTRAINT. This clause removes a named constraint from a table.

Note:

In solidDB, when the keyword CONSTRAINT is used, the constraint name is mandatory.

- CHECK. This constraint allows you to specify rules to your tables or table columns. Each rule is a condition that must not be false for any row in the table on which it is defined. Otherwise the table cannot be updated.

The rules are Boolean expressions. The rule can check, for example, a range of values, equity, or the rule can be a simple comparison. You can use several checks in one statement. The following expressions and operators are available:

Table 17. Expressions and operators

Expression	Explanation
<	less than
>	greater than

Table 17. Expressions and operators (continued)

Expression	Explanation
=	equal to
<=	less than or equal to
>=	greater than or equal to
<>	not equal to
AND	conjunction
ANY	in the list that follows or in the table specified
BETWEEN	between
IN	in the list that follows or in the table specified
MAX	maximum value
MIN	minimum value
NOT	negation
OR	disjunction
XOR	exclusive or

- **UNIQUE.** The UNIQUE constraint requires that no two rows in a table contain the same value in a given column or list of columns. You can create a unique constraint at either the table level or the column level. Note: primary keys contain the unique constraint.
- **FOREIGN KEY.** The FOREIGN KEY constraint requires that each value in the foreign key column must have a matching value in the referenced table.

Note:

solidDB automatically generates names for unnamed constraints. If you want to view the names, use the command **solidd -x hidddenames**.

For constraint syntax information and examples, see the CREATE TABLE and ALTER TABLE sections in Appendix A, “IBM solidDB SQL statements,” on page 141.

4.6 Managing database objects

In the solidDB server, you use catalogs and schemas to organize your data. The use of schemas in the solidDB server conforms to the SQL standard, but the use of catalogs is an extension to the SQL standard. In general, a catalog can be thought of as a logical database. A schema typically corresponds to a user.

Catalogs and schemas help you to group database objects (tables, sequences, and so on) in a hierarchical way. This allows you to put related items into the same group. For example, all the tables related to the accounting system might be in one

group (for example, a catalog), while all the tables related to the human resources system might be in another group. You can also group database objects by user. For example, all of the tables used by Jane Smith might be in a single schema.

Catalogs are the highest (broadest) level of the hierarchy. Schema names are the mid-level. Specific database objects, such as tables, are the lowest (narrowest) level of the hierarchy. Thus, a single catalog may contain multiple schemas, and each of those schemas may contain multiple tables.

Object names must be unique within a group, but do not have to be unique across groups. Thus, for example, Jane Smith's schema and Robin Trower's schema might each have a table named "bills". These two tables have nothing to do with each other. They may have different structures and different data, even though they have the same name. Similarly, the catalog "accounting_catalog" and the catalog "human_resources_catalog" might each have a schema named "david_jones". Those schemas are unrelated to each other, even though they have the same name.

If you want to specify a particular table and that table name is not unique in the database, you can identify it by specifying the catalog, schema, and table name, for example:

```
accounting_catalog.david_jones.bills
```

If you do not specify the complete name (that is, if you omit the schema, or the schema and the catalog), the server uses the current/default catalog and schema name to determine which table to use.

4.6.1 Catalogs

A physical database file may contain more than one logical database. Each logical database is a complete, independent group of database objects, such as tables, indexes, procedures, triggers, and so on. Each logical database is a catalog. Note that a solidDB catalog is not just limited to indexes (as in the traditional sense of a library card catalog, which serves to locate an item without containing the full contents of the item).

Catalogs allow you to logically partition databases so you can:

- Organize your data to meet the needs of your business, users, and, applications.
- Specify multiple master or replica databases (by using logical databases) for synchronization within one physical database server. For more details on implementing synchronization in multi-master environments, read "Multi-master synchronization model" in *IBM solidDB Advanced Replication User Guide*.

4.6.2 Schemas

A catalog may contain one or more schemas. A schema is a persistent database object that provides a definition for part or all of the database. It represents a collection of database objects associated with a specific schema name. These objects include tables, views, indexes, stored procedures, triggers, and sequences. Schemas allow you to provide each user with his or her own database objects (such as tables) within the same logical database (that is, a single catalog). If no schema is specified with a database object, the default schema is the user id of the user creating the object.

4.6.3 Uniquely identifying objects within catalogs and schemas

Schemas make it possible for two different users to create tables with the same names in the same physical database or even in the same logical database. For example, assume in a single physical database, there are two separate catalogs, `employee_catalog` and `inventory_catalog`. Assume also that each catalog contains two separate schemas, named `smith` and `jones`, and that the same Smith owns both "smith" schemas and the same Jones owns both "jones" schemas. If Smith and Jones create a table named `books` in each of their schemas, then we have a total of 4 tables named "books", and these tables are accessible as:

```
employee_catalog.smith.books  
employee_catalog.jones.books  
inventory_catalog.smith.books  
inventory_catalog.jones.books
```

As you can see, the catalog name and schema name can be used to "qualify" (uniquely identify) the name of a database object such as a table. Object names can be qualified in all DML statements by using the syntax:

```
catalog_name.schema_name.database_object
```

or

```
catalog_name.user_id.database_object
```

For example:

```
SELECT cust_name FROM accounting_dept.smith.overdue_bills;
```

You can qualify one or more database objects with a schema name, whether or not you specify a catalog name. The syntax is:

```
schema_name.database_object_name
```

or

```
user_id.database_object_name
```

For example,

```
SELECT SUM(sales_tax) FROM jones.invoices;
```

To use a schema name with a database object, you must have already created the schema.

By default, database objects that are created without schema names are qualified using the user ID of the database object's creator. For example:

```
user_id.table_name
```

Catalog and schema contexts are set using the `SET CATALOG` or `SET SCHEMA` statement.

If a catalog context is not set using `SET CATALOG`, then all database object names are resolved by using the default catalog name.

Note: When creating a new database or converting an old database to a new format, the user is prompted to specify a default catalog name for the database system catalog. Users can access the default catalog name without knowing this specified default catalog name. For example, users can specify the following syntax to access the system catalog:

```
""._SYSTEM.table
```

solidDB translates the empty string ("") specified as a catalog name to the default catalog name. solidDB also provides for automatic resolution of _SYSTEM schema to the system catalog, even when users provide no catalog name.

The following SQL statements provide examples of creating catalogs and schemas. Refer to Appendix A, “IBM solidDB SQL statements,” on page 141, for a formal definition of the solidDB SQL statements.

4.6.4 Creating a catalog

```
CREATE CATALOG catalog_name
```

Only the creator of the database or users having SYS_ADMIN_ROLE have privileges to create or drop catalogs.

The following example creates a catalog named C and assumes the userid is SMITH

```
CREATE CATALOG C;
SET CATALOG C;
CREATE TABLE T (i INTEGER);
SELECT * FROM T;
--The name T is resolved to C.SMITH.T
```

4.6.5 Setting a catalog and schema context

The following example sets a catalog context to C and the schema context to S.

```
SET CATALOG C;
SET SCHEMA S;
CREATE TABLE T (i INTEGER);
SELECT * FROM T;
-- The name T is resolved to C.S.T
```

4.6.6 Deleting a catalog

```
DROP CATALOG catalog_name
```

The following example deletes the catalog named C.

```
DROP CATALOG C;
```

4.6.7 Creating a schema

```
CREATE SCHEMA schema_name
```

Any database user can create a schema; however, the user must have permission to create the objects that pertain to the schema (for example, CREATE PROCEDURE, CREATE TABLE, and so on).

Note that creating a schema does not implicitly make that new schema the current/default schema. You must explicitly set that schema with the SET SCHEMA statement if you want the new schema to become the current schema.

The following example creates a schema named FINANCE and assumes the user id is SMITH:

```
CREATE SCHEMA FINANCE;
CREATE TABLE EMPLOYEE (EMP_ID INTEGER);
-- NOTE: The employee table is qualified to SMITH.EMPLOYEE, not
```

```
-- FINANCE.EMPLOYEE. Creating a schema does not implicitly make that
-- new schema the current/default schema.
SET SCHEMA FINANCE;
CREATE TABLE EMPLOYEE (ID INTEGER);
SELECT ID FROM EMPLOYEE;
-- In this case, the table is qualified to FINANCE.EMPLOYEE
```

4.6.8 Deleting a schema

```
DROP SCHEMA schema_name
```

The following example deletes the schema named FINANCE.

```
DROP SCHEMA FINANCE;
```

5 Managing transactions

This section explains how to manage transactions, how to deal with concurrency control and locking, and how to choose the durability level.

5.1 Defining read-only and read-write transactions

To define a transaction to be read-only or read-write, use the following SQL commands:

```
SET TRANSACTION { READ ONLY | READ WRITE }
```

The following options are available with this command.

- **READ ONLY**
Use this option for a read only transaction.
- **READ WRITE**
Use this option for a read and write transaction. This option is the default.

Note: To detect conflicts between transactions, use the standard ANSI SQL command `SET TRANSACTION ISOLATION LEVEL` to define the transaction with a Repeatable Read or Serializable isolation level. For details, read section *Choosing transaction isolation levels* in *IBM solidDB Administrator Guide*.

Transactions must be ended with the `COMMIT WORK` or `ROLLBACK WORK` commands unless autocommit is used.

5.2 Concurrency control and locking

The purpose of *concurrency control* is to prevent two different users (or two different connections by the same user) from trying to update the same data at the same time. Concurrency control can also prevent one user from seeing out-of-date data while another user is updating the same data.

The following examples explain why concurrency control is needed. For both examples, suppose that your checking account contains \$1,000. During the day you deposit \$300 and spend \$200 from that account. At the end of the day your account should have \$1,100.

- **Example 1: No concurrency control**

1. At 11:00 AM, bank teller #1 looks up your account and sees that you have \$1,000. The teller subtracts the \$200 check, but is not able to save the updated account balance (\$800) immediately.
2. At 11:01 AM, another teller #2 looks up your account and still sees the \$1,000 balance. Teller #2 then adds your \$300 deposit and saves your new account balance as \$1,300.
3. At 11:09 AM, bank teller #1 returns to the terminal, finishes entering and saving the updated value that is calculated to be \$800. That \$800 value writes over the \$1300.
4. At the end of the day, your account has \$800 when it should have had \$1,100 (\$1000 + 300 - 200).

- **Example 2: Concurrency control**

1. When teller #1 starts working on your account, a *lock* is placed on the account.
2. When teller #2 tries to read or update your account while teller #1 is updating your account, teller #2 will not be given access and gets an error message.
3. After teller #1 has finished the update, teller #2 can proceed.
4. At the end of the day, your account has \$1,100 (\$1000 - 200 + 300).

In Example 1, the account updates are done simultaneously rather than in sequence and one update write overwrites another update. In Example 2, to prevent two users from updating the data simultaneously (and potentially writing over each other's updates), the system uses a *concurrency control* mechanism.

solidDB offers two different concurrency control mechanisms, *pessimistic concurrency control* and *optimistic concurrency control*.

The pessimistic concurrency control mechanism is based on *locking*. A *lock* is a mechanism for limiting other users' access to a piece of data. When one user has a lock on a record, the lock prevents other users from changing (and in some cases reading) that record. Optimistic concurrency control mechanism does not place locks but prevents the overwriting of data by using timestamps.

5.2.1 PESSIMISTIC vs. OPTIMISTIC concurrency control

solidDB offers two different concurrency control mechanisms, *pessimistic* and *optimistic*.

- **Pessimistic concurrency control** (or *pessimistic locking*) is called "pessimistic" because the system assumes the worst — it assumes that two or more users will want to update the same record at the same time, and then prevents that possibility by locking the record, no matter how unlikely conflicts actually are.

The locks are placed as soon as any piece of the row is accessed, making it impossible for two or more users to update the row at the same time. Depending on the lock mode (*shared*, *exclusive*, or *update*), other users might be able to read the data even though a lock has been placed. For more details on the lock modes, see "Lock modes: shared, exclusive, and update" on page 115.

- **Optimistic concurrency control** (or *optimistic locking*) assumes that although conflicts are possible, they will be very rare. Instead of locking every record every time that it is used, the system merely looks for indications that two users actually did try to update the same record at the same time. If that evidence is found, then one user's updates are discarded and the user is informed.

For example, if User1 updates a record and User2 only wants to read it, then User2 simply reads whatever data is on the disk and then proceeds, without checking whether the data is locked. User2 might see slightly out-of-date information if User1 has read the data and updated it, but has not yet committed the transaction.

Optimistic locking is available on disk-based tables (D-tables) only.

The solidDB implementation of optimistic concurrency control uses multiversioning.

1. Each time that the server reads a record to try to update it, the server makes a copy of the version number of the record and stores that copy for later reference.
2. When it is time to commit the transaction, the server compares the original version number that it read against the version number of the currently committed data.

- If the version numbers are the same, then no one else changed the record and the system can write the updated value.
- If the originally read value and the current value on the disk are not the same, then someone has changed the data since it was read, and the current operation is probably out-of-date. Thus the system discards the version of the data, aborts the transaction, and returns an error message.
- The step of checking the version numbers is called *validation*. The validation can be performed at the commit time (normal validation) or at the time of writing each statement (early validation). In solidDB, early validation is the default method (**General.TransactionEarlyValidate=yes**).

Each time a record is updated, the version number is updated as well.

solidDB can store multiple versions of each data row temporarily, rather than giving each user the version of data is on the disk at the moment it is read. Each user's transaction sees the database as it was at the time that the transaction started. This way the data that each user sees is consistent throughout the transaction, and users are able to concurrently access the database. For more details about multiversioning, see *Multithread processing* in the *IBM solidDB Getting Started Guide*.

Note: Even though the optimistic concurrency control mechanism is sometimes called optimistic locking, it is not a true locking scheme—the system does not place any locks when optimistic concurrency control is used. The term locking is used because optimistic concurrency control serves the same purpose as pessimistic locking by preventing overlapping updates.

When you use optimistic locking, you do not find out that there is a conflict until just before you write the updated data. In pessimistic locking, you find out there is a conflict as soon as you try to read the data.

To use an analogy with banks, pessimistic locking is like having a guard at the bank door who checks your account number when you try to enter; if someone else (a spouse, or a merchant to whom you wrote a check) is already in the bank accessing your account, then you cannot enter until that other person finishes her transaction and leaves. Optimistic locking, on the other hand, allows you to walk into the bank at any time and try to do your business, but at the risk that as you are walking out the door the bank guard will tell you that your transaction conflicted with someone else's and you will have to go back and do the transaction again.

With pessimistic locking, the first user to request a lock, gets it. Once you have the lock, no other user or connection can override your lock. In solidDB, the lock lasts until the end of the transaction or in the case of long table locks, the lock lasts until you explicitly release it.

Default concurrency control mechanisms

The default concurrency control mechanism depends on the table type:

- Disk-based tables (D-tables) are by default optimistic.
- Main-memory tables (M-tables) are always pessimistic.

You can override optimistic locking and specify pessimistic locking instead. You can do this at the level of individual tables. One table might follow the rules of optimistic locking while another table follows the rules of pessimistic locking. Both

tables can be used within the same transaction and even the same statement; solidDB handles this internally.

Locking and performance

Optimistic locking allows fast performance and high concurrency (access by multiple users), at the cost of occasionally refusing to write data that was initially accepted but was found at the last second to conflict with another user's changes.

Pessimistic locking requires overhead for every operation, whether or not two or more users are actually trying to access the same record. The overhead is small but adds up because every row that is updated requires a lock. Furthermore, every time that a user tries to access a row, the system must also check whether the requested row(s) are already locked by another user or connection.

For example, if two bank tellers are accessing the same record around the same time and bank teller #1 gets a lock, teller #2 must check for that lock, no matter how unlikely it is that teller #2 will want to work on the same record exactly at the same time as teller #1. Checking every record that is used will take time. Furthermore, it is important that during the checking, no other teller tries to run the same check as teller #2 (otherwise they might both see at 10:59:59 that record X is not in use, and then they might both try to lock it at 11:00:00). Thus even checking a lock can itself require another lock to prevent two users from changing the locks at the time.

Choosing concurrency control mechanism

In most scenarios, optimistic concurrency control is more efficient and offers higher performance. When choosing between pessimistic and optimistic locking, consider the following:

- Pessimistic locking is useful if there are a lot of updates and relatively high chances of users trying to update data at the same time.

For example, if each operation can update a large number of records at a time (the bank might add interest earnings to every account at the end of each month), and two applications are running such operations at the same time, they will have conflicts.

Pessimistic concurrency control is also more appropriate in applications that contain small tables that are frequently updated. In the case of these so-called *hotspots*, conflicts are so probable that optimistic concurrency control wastes effort in rolling back conflicting transactions.

If you are migrating an application from another DBMS that uses pessimistic locking, you should use the pessimistic mode with solidDB too. Using pessimistic mode in solidDB means that you do not have make changes to the application.

- Optimistic locking is useful if the possibility for conflicts is very low – there are many records but relatively few users, or very few updates and mostly read-type operations.

5.2.2 Locks and lock modes

A *lock* is a mechanism for preventing two or more users from doing conflicting operations at the same time. Operations conflict if at least one of the operations involves updating the data (via UPDATE, DELETE, INSERT, ALTER TABLE, and so on). If all the operations are read-only operations (such as SELECT), then there is no conflict.

solidDB does not allow users to specify row-level locks explicitly. There is no LOCK RECORD command; the server does all row-level locking for you. The server also does table-level locking for you. If you need to set table-level locks explicitly, you may do so using the LOCK TABLE command.

Table-level vs. row-level locks

solidDB allows both table-level locks and row-level locks.

Row-level locks

Row-level locks are placed on single records (rows) that the statements in a transaction define. The locks are placed as soon as any piece of the row is accessed.

Row-level locks are always implicit, solidDB sets the locks when necessary. You cannot lock or unlock row-level locks manually.

Table-level locks

Table-level locks can be thought of as metadata locks; they prevent concurrent users from making schema changes (DDL operations) simultaneously or while records within the table are being changed.

For example, if you are updating a customer's home phone number, you do not want another user to drop the telephone number column at the same time. If the other user was allowed to drop the telephone number column before you were finished, your transaction would try to write an updated telephone number to a column that no longer exists, thus resulting in data corruption.

Most table-level locks are implicit; the server itself sets those locks when necessary. For example, when the server recognizes that a particular operation (such as an UPDATE statement without a where clause) will affect every record in the table, the server itself can lock the entire table if it thinks that would be most efficient, and if no conflicting locks on the table already exist. Also, when you acquire a lock on a record in a table, you also implicitly acquire a lock (usually a shared lock) on the entire table. This prevents prevent one user from dropping the table or modifying the structure of the table while another user is updating data in the table.

You can also lock and unlock table-level locks manually using the LOCK TABLE and UNLOCK TABLE commands.

Table-level locks are always pessimistic; the server puts a real lock on the table rather than just looking at versioning information. This is true even if the table is set to optimistic locking.

In setups using advanced replication, table-level locks are typically used with *Maintenance Mode* operations. For more details, see *Introduction to Maintenance Mode* in the *IBM solidDB Advanced Replication User Guide*.

Lock modes: shared, exclusive, and update

Depending on the lock mode, when one user has a lock on a record, the lock prevents other users from changing or even reading that record.

There are three lock modes:

- SHARED

Row-level shared locks allow multiple users to read data, but do not allow any users to change that data.

Table-level shared locks allow multiple users to perform read and write operations on the table, but do not allow any users to perform DDL operations. Multiple users can hold shared locks simultaneously.

- **EXCLUSIVE**

An exclusive lock allows only one user/connection to update a particular piece of data (insert, update, and delete). When one user has an exclusive lock on a row or table, no other lock of any type may be placed on it.

- **UPDATE**

Update locks are always row-level locks. When a user accesses a row with the `SELECT... FOR UPDATE` statement, the row is locked with an update mode lock. This means that no other user can read or update the row and ensures the current user can later update the row.

Update locks are similar to exclusive locks. The main difference between the two is that you can acquire an update lock when another user already has a shared lock on the same record. This lets the holder of the update lock read data without excluding other users. However, once the holder of the update lock changes the data, the update lock is converted into an exclusive lock.

Also, update locks are asymmetric with respect to shared locks. You can acquire an update lock on a record that already has a shared lock, but you cannot acquire a shared lock on a record that already has an update lock. Because an update lock prevents subsequent read locks, it is easier to convert the update lock to an exclusive lock.

Shared and exclusive locks cannot be mixed. If User1 has an exclusive lock on a record, User2 cannot get a shared lock or an exclusive lock on that same record.

All locks within a particular category (such as shared locks) are equal.

- All users regardless the user privileges are equal: locks placed by a DBA are no more and no less strong than locks placed by any other user.
- All ways of executing statements that place locks are equal: the lock can be executed as part of ,
- It does not matter whether the lock was executed as part of an interactively typed statement, called from a compiled remote application, or called from within the local application when using solidDB with shared memory access or linked library access, or if the lock was placed as a result of a statement inside a stored procedure or trigger.

Some locks can be escalated. For example, if you are using a scroll cursor and you acquire a shared lock on a record, and then later within that same transaction you update that record, your shared lock may be upgraded to an exclusive lock.

Getting an exclusive lock is only possible if there are no other locks (shared or exclusive) on the table; if you and another user both have shared locks on the same record, then the server cannot upgrade your shared lock to an exclusive lock until the other user drops her shared lock.

Lock modes for table-level locks

The `EXCLUSIVE` and `SHARED` lock modes are used for both pessimistic and optimistic tables. By default, optimistic and pessimistic tables are locked in shared mode; unless you are altering the table, the locks on tables are usually shared locks.

When you execute an ALTER TABLE operation, you get a shared lock on that table. That allows other users to continue to read data from the table, but prevents them from making changes to the table. If other users want to do DDL operations (such as ALTER TABLE) on the same table at the same time, they will either have to wait or will get an error message.

Also, in advanced replication setups, some solidDB statements (such as REFRESH or MESSAGE EXECUTE) that can be run with the optional PESSIMISTIC keyword, use EXCLUSIVE table-level locks even when the tables are optimistic.

Lock duration and timeout

By default, a lock is held from the time it is acquired until the end of the transaction (completed with commit or rollback). If you try to get an exclusive lock on a record that another user has already locked (shared or exclusive), you cannot get a lock; instead, your transaction will fail with an error. You can define whether solidDB should fail your transaction immediately or, before failing, wait and try again for a specified number of seconds. This is controlled with a *lock timeout* setting.

The *lock timeout* setting is the time in seconds that the engine waits for a lock to be released. By default, solidDB lock timeout is set to 30 seconds. If transactions tend to be very short, a brief wait allows you to continue activities that otherwise would have been blocked by locks.

When the lock timeout interval is reached, solidDB terminates the timed-out statement. For example, if User1 is querying a specific row in a table and User2 tries to update data in the same row, the update will not go through until the User1's query is completed (or times out). If the query of User1 is completed and the User2 query has not timed out yet, a lock is issued for the update transaction of User2. If User1 does not finish before the query of User2 times out, the server terminates User2's statement.

The default lock timeout is controlled with the **General.LockWaitTimeOut** parameter. In advanced replication setups, you might also want to set the default lock time for table-level locks with the **General.TableLockWaitTimeout** parameter.

The default timeout can be overridden with the following transaction or section specific commands:

- LOCK TABLE WAIT – sets the timeout for table-level locks for specific tables (D-tables only)
- SET LOCK TIMEOUT – sets the timeout for both table-level and row-level locks
SET LOCK TIMEOUT does not change the timeout for those tables for which the table-level timeout has been set with LOCK TABLE WAIT.

Note: The LOCK TABLE WAIT mechanism does not apply to M-tables. For example, if in Session1 you lock table DEPARTMENT (LOCK TABLE DEPARTMENT IN EXCLUSIVE MODE, an attempt to insert values into the table in Session2 (INSERT INTO DEPARTMENT VALUES ...) will return error 10014 Resource is locked. immediately.

The wait mechanism in lock timeout applies only to pessimistic locking. There is no such thing as "waiting for an optimistic lock". If someone else changed the data since the time that you read it, no amount of waiting will prevent a conflict that has already occurred. In fact, since optimistic concurrency methods do not place locks, there is no "optimistic lock" to wait on.

LONG exclusive locks

solidDB allows you to prevent exclusive locks from being released when the locking transaction commits. These type of *long* exclusive locks are set with the LONG option in the LOCK TABLE command.

For example:

```
LOCK TABLE emp IN LONG EXCLUSIVE MODE
```

If the locking transaction aborts or is rolled back, all locks are released, including LONG locks. You must unlock long locks explicitly using the UNLOCK command. LONG duration locks are allowed only in EXCLUSIVE mode. LONG shared locks are not supported.

Transaction isolation levels and lock duration

Update locks and exclusive locks are always held until the time that the transaction completes. Shared locks ("read locks") are also held until the end of the transaction but the transaction isolation level can affect how shared locks behave. For example, SERIALIZABLE isolation level does additional checks. It checks also that no new rows are added to the result set that the transaction should have seen. In other words, it prevents other users from inserting rows that would have qualified for the result set that is in the transaction.

Example:

If a SERIALIZABLE transaction has an update command like UPDATE customers SET x = y WHERE area_code = 415;, solidDB does not allow other users to enter records with area_code=415 until the serializable transaction is committed.

Note: solidDB's implementation of holding shared locks until the end of transaction differs from some other servers. Some servers will release shared locks before the end of a transaction, if the transaction isolation level is low enough. Other database servers might also allow you to extend the duration of read/shared locks to ensure that within a single transaction, data looks the same every time you view it.

Also, in other servers, transaction isolation level might affect not only how long you lock a record, but also what you see. For example, on systems that allow both READ UNCOMMITTED (sometimes called "dirty read") and READ COMMITTED, your isolation level affects what you see, not just what other users can or cannot see because you have locked certain records.

5.2.3 Setting concurrency control

The concurrency control method and lock modes can be controlled with solidDB SQL statements and configuration parameters.

Setting the concurrency (locking) mode to optimistic or pessimistic

The concurrency mode of disk-based tables can be set to optimistic or pessimistic for all tables or for specific tables. In-memory tables are always pessimistic.

By default D-tables use optimistic locking.

- To set the concurrency mode for a specific table, use the ALTER TABLE <table_name> SET OPTIMISTIC|PESSIMISTIC command.

For example:

```
ALTER TABLE MyTable1 SET PESSIMISTIC;  
ALTER TABLE MyTable2 SET OPTIMISTIC;
```

- To control the default concurrency mode for all tables, set the **General.Pessimistic** parameter to 'yes' or 'no' (default is 'no').

For example:

```
[General]  
Pessimistic=yes
```

The **General.Pessimistic** parameter takes effect only at the time that the server starts. If you edit the `solid.ini` file manually, the change will not be visible until the server restarts.

Since the value of the **General.Pessimistic** can change, the concurrency control for a table may change. It is possible for a table to use optimistic concurrency control during one instance of the server and pessimistic during another.

When you set the **General.Pessimistic** parameter to 'yes', the server defaults to pessimistic locking for

- any new tables that are created, and
- for any existing tables whose concurrency control method has never been set explicitly with the ALTER TABLE command.

If you set a table's locking mode by using the ALTER TABLE command, the ALTER TABLE command takes precedence.

Related topics

- “Setting mixed concurrency control”

Setting mixed concurrency control

With D-tables, you can use mixed concurrency control methods. Mixed concurrency control is available by setting individual tables to optimistic or pessimistic.

By default, solidDB uses optimistic concurrency control for D-tables. M-tables are always pessimistic.

Mixed concurrency control is a combination of row-level pessimistic locking and optimistic concurrency control. By turning on row-level locking table-by-table, you can specify that a single transaction use both concurrency control methods simultaneously. This can be set for both read-only and read-write transactions.

To set individual tables for optimistic or pessimistic concurrency, use the following command:

```
ALTER TABLE base_table_name SET {OPTIMISTIC | PESSIMISTIC}
```

Note:

When using solidDB with advanced replication, pessimistic table-level locks in shared mode are possible with tables that are synchronized. This functionality provides users with the option to run some operations for synchronization in pessimistic mode even with optimistic tables. For example, when a REFRESH is executed in pessimistic mode in a replica, solidDB locks all tables in shared mode; later, if necessary, the server can "promote" these locks to exclusive table locks. This is done in a few synchronization statements when optional keyword PESSIMISTIC is specified. Read operations do not use any locks.

Reading the concurrency mode

The method for reading the concurrency mode of a table depends on how the concurrency mode has been set.

- If the concurrency mode of the table has been set explicitly with the ALTER TABLE command, the concurrency mode is recorded in a SYS_TABLEMODES system table.

Check the values in the SYS_TABLEMODES with the following command:

```
SELECT SYS_TABLEMODES.ID, SYS_TABLEMODES.MODE, SYS_TABLES.TABLE_NAME
FROM SYS_TABLEMODES, SYS_TABLES
WHERE SYS_TABLEMODES.ID = SYS_TABLES.ID
AND SYS_TABLES.TABLE_NAME = '<table_name>';
```

For example:

```
SELECT SYS_TABLEMODES.ID, SYS_TABLEMODES.MODE, SYS_TABLES.TABLE_NAME
FROM SYS_TABLEMODES, SYS_TABLES
WHERE SYS_TABLEMODES.ID = SYS_TABLES.ID
AND SYS_TABLES.TABLE_NAME = 'TESTTABLE2';
```

ID	MODE	TABLE_NAME
--	----	-----
10002	PESSIMISTIC	TESTTABLE2

1 rows fetched.

If the concurrency mode has not been set using the ALTER TABLE command, the SYS_TABLEMODES system table does not contain information about the concurrency mode of the table.

- If the concurrency mode of the table has not been set with the ALTER TABLE command, check the setting of the **General.Pessimistic** parameter with the following command:

```
ADMIN COMMAND 'describe parameter General.Pessimistic';
```

If the value in the solid.ini file has not been changed since the server started, and if the value has not been overridden by an ADMIN COMMAND, you can also check the parameter setting in the solid.ini file.

Setting lock timeout

The lock timeout setting can be modified with the SET LOCK TIMEOUT and LOCK TABLE WAIT commands. By default, lock timeout is set to 30 seconds.

- Use LOCK TABLE WAIT <timeout_in_seconds> to set timeout for table-level locks.

Note: The LOCK TABLE WAIT command is effective on disk-based tables only.

- Use SET LOCK TIMEOUT <timeout_in_seconds> to set the lock timeout for both table-level and row-level locks in a session.

Note: SET LOCK TIMEOUT does not change the timeout for those tables for which the table-level timeout has been set with LOCK TABLE WAIT.

By default, the granularity for the timeout is in seconds. The lock timeout can be set at millisecond granularity by adding "MS" after the value, for example:

```
LOCK TABLE emp,dept IN SHARED MODE WAIT 10MS;
```

or

```
SET LOCK TIMEOUT 10MS;
```

Without the "MS", the lock timeout is in seconds.

Note: The maximum timeout is 1000 seconds (a little over 15 minutes). The server will not accept a longer value.

Setting lock timeout for optimistic tables

When you use `SELECT FOR UPDATE`, the selected rows become locked even if the locking mode of the table is optimistic. If the statement is successful, you can execute consecutive updates and deletes on those rows without a fear that they would fail on a conflict.

Statements that are competing in concurrent transactions for the write access to the rows (such as `SELECT FOR UPDATE`, `UPDATE` or `DELETE`) will fail immediately if the lock cannot be granted. That corresponds to zero timeout. You can introduce lock waits by setting the optimistic lock timeout to a non-zero value:

```
SET OPTIMISTIC LOCK TIMEOUT {seconds| milliseconds MS}
```

When the timeout is set within a session, the `SELECT FOR UPDATE` statement will block (wait) until the timeout expires, or until the lock is released.

The timeout does not affect the `DELETE` or `UPDATE` statements that will return immediately – either with a success or a conflict error.

The `SET OPTIMISTIC LOCK TIMEOUT` statement takes effect immediately. If needed, the default behavior can be reset by setting the timeout value to zero.

5.3 Choosing transaction durability level

solidDB provides two durability levels, strict and relaxed. *Strict durability* means that as soon as a transaction is committed, the server writes the information to the transaction log file. *Relaxed durability* means that the server may not write the information as soon as the transaction is committed; instead, the server might wait, for example, until it is less busy or until it can write multiple transactions in a single write operation.

If you can afford to lose a small amount of recent data, and if performance is crucial to you, use relaxed durability. Relaxed durability is appropriate when each individual transaction is not crucial. For example, if you are monitoring system performance and you want to store data on response times, you may only be interested in average response times, which will not be significantly affected if you are missing a few pieces of data. Since measuring performance will itself affect performance (by using up resources such as processor capacity and I/O bandwidth), you probably want your performance tracking operations themselves to have high performance (low cost) rather than high precision. Relaxed durability is appropriate in this situation.

On the other hand, if you are tracking financial data, such as bill payments, you probably want to ensure that 100% of your committed data is stored and recoverable. In this situation, use strict durability.

You should use relaxed durability only when you can afford to lose a few of the most recent transactions. Otherwise, use strict durability. If you are not sure whether strict or relaxed durability is appropriate, use strict durability.

5.3.1 Setting transaction durability level

There are four ways to set the transaction durability level, described below in a descending order of precedence.

1. SET TRANSACTION DURABILITY – transaction-level setting

```
SET TRANSACTION DURABILITY { RELAXED | STRICT }
```

For example:

```
SET TRANSACTION DURABILITY STRICT;
```

The SET TRANSACTION DURABILITY command specifies the transaction durability for a transaction. You must issue the SET TRANSACTION DURABILITY command in the beginning of a statement. The durability setting applies only to the current transaction, until the transaction is committed or aborted.

2. SET DURABILITY – session-level setting

```
SET DURABILITY { RELAXED | STRICT }
```

For example:

```
SET DURABILITY STRICT;
```

The SET DURABILITY command specifies the transaction durability for a session. A *session* is the time between connecting and disconnecting to the server. Each user has a separate session, even if the sessions overlap in time. (A single user may establish more than one session, for example, by running multiple copies of **solsql**, or by writing a program that makes multiple connections to the same server). When you use the SET DURABILITY statement, you specify the transaction durability only for the session in which the command is issued. Your choice will not affect any other user, any other open session that you yourself currently have, or any future session that you may have. Each user session may set its own transaction durability level, based on how important it is for the session not to lose any data.

The effect of this statement lasts until the end of the session, or until another SET DURABILITY command is issued.

3. **Logging.DurabilityLevel** - server default setting

```
[Logging]  
DurabilityLevel=3
```

The parameter setting affects all users.

The **Logging.DurabilityLevel** parameter can be changed dynamically. If you want to change the default setting while the server is running, issue the following command:

```
ADMIN COMMAND 'parameter Logging.DurabilityLevel={1 | 2 | 3}'
```

The command is effective immediately.

4. **Logging.DurabilityLevel** - factory value for server

If you do not set the transaction durability level using any of the methods above, the server will use relaxed durability (**Logging.DurabilityLevel=1**).

6 Diagnostics and troubleshooting for SQL

solidDB provides the following diagnostic tools for troubleshooting SQL:

- SQL info facility and the EXPLAIN PLAN FOR statement used to tune your application and identify inefficient SQL statements in your application
- Tracing facilities for stored procedures and triggers

You can use these facilities to observe performance, troubleshoot problems, and produce high quality problem reports. These reports let you pinpoint the source of your problems by isolating them under product categories (such as solidDB ODBC API, solidDB ODBC Driver, or solidDB JDBC Driver).

6.1 Observing performance

You can use the SQL Info facility to provide information about an SQL statement and the SQL statement EXPLAIN PLAN FOR to show the execution graph that the SQL optimizer selected for a given SQL statement. Typically, if you need to contact IBM Software Support, you will be asked to provide the SQL statement, EXPLAIN PLAN output, and SQL Info output from the EXPLAIN PLAN run with info level 8 for more extensive trace output.

6.1.1 SQL Info facility

The SQL Info facility generates information for each SQL statement processed by solidDB.

To generate the SQL Info, you run your application with the SQL Info facility enabled. The SQL Info facility can be enabled in the following ways:

- `Sql.Info=<info_level>` parameter
- `ADMIN COMMAND 'trace on info <info_level>'` command
- `SET SQL INFO ON LEVEL info_level FILE file_name` statement

The tracing level (*info_level*) is defined as an integer between 0 (no tracing) and 8 (solidDB info from every fetched row).

Table 18. SQL Info levels

Info level	Description
0	no output
1	table, index, and view info in SQL format
2	SQL execution graphs (technical support use only)
3	some SQL estimate info, solidDB selected key name
4	all SQL estimate info, solidDB selected key info
5	solidDB info also from discarded keys
6	solidDB table level info

Table 18. SQL Info levels (continued)

Info level	Description
7	SQL info from every fetched row
8	solidDB info from every fetched row

The trace information is output by default to the `soltrace.out` file in the solidDB working directory. You can also specify the output file using the **SQL.InfoFileName** parameter. This is recommended since the `soltrace.out` file may contain information from several sources.

Examples

```
[SQL]
Info = 1
InfoFileName = solidsql_trace.txt
```

The following command turns on the SQL Info facility on level 3, outputting the trace information to a `my_query.txt` file in the working directory. This SQL Info facility is turned on only for the client that executes the statement.

```
SET SQL INFO ON LEVEL 1 FILE 'my_query.txt'
```

The following SQL statement turns off the SQL Info facility:

```
SET SQL INFO OFF
```

6.1.2 EXPLAIN PLAN FOR statement

The EXPLAIN PLAN FOR statement is used to show the execution plan that the SQL optimizer has selected for a given SQL statement.

The syntax of the EXPLAIN PLAN FOR statement is:

```
EXPLAIN PLAN FOR sql_statement
```

An execution plan is a series of primitive operations, and an ordering of these operations, that solidDB performs to execute the statement. Each operation in the execution plan is called a unit.

Table 19. EXPLAIN PLAN FOR units

Unit	Description
JOIN UNIT*	Join unit joins two or more tables. The join can be done by using loop join or merge join.
TABLE UNIT	The table unit is used to fetch the data rows from a table or index.
ORDER UNIT	Order unit is used to order rows for grouping or to satisfy ORDER BY. The ordering can be done in memory or using an external disk sorter.
GROUP UNIT	Group unit is used to do grouping and aggregate calculation (SUM, MIN, and so on).
UNION UNIT*	Union unit performs the UNION operation. The unit can be done by using loop join or merge join.

Table 19. EXPLAIN PLAN FOR units (continued)

Unit	Description
INTERSECT UNIT*	Intersect unit performs the INTERSECT operation. The unit can be done by using loop join or merge join.
EXCEPT UNIT*	Except unit performs the EXCEPT operation. The unit can be done by using loop join or merge join.

*This unit is generated also for queries that reference only a single table. In that case no join is executed in the unit; it simply passes the rows without manipulating them.

The table returned by the EXPLAIN PLAN FOR statement contains the following columns.

Table 20. EXPLAIN PLAN FOR table columns

Column Name	Description
ID	The output row number, used only to guarantee that the rows are unique.
UNIT_ID	This is the internal unit id in the SQL interpreter. Each unit has a different id. The unit id is a sparse sequence of numbers, because the SQL interpreter generates unit ids also for those units that are removed during the optimization phase. If more than one row has the same unit id it means that those rows belong to the same unit. For formatting reasons the info from one unit may be divided into several different rows.
PAR_ID	Parent unit id for the unit. The parent id number refers to the id in the UNIT_ID column.
JOIN_PATH	For join, union, intersect, and except units there is a join path which specifies which tables are joined in the unit and the join order for tables. The join path number refers to the unit id in the UNIT_ID column. It means that the input to the unit comes from that unit. The order in which the tables are joined is the order in which the join path is listed. The first listed table is the outermost table in a loop join.
UNIT_TYPE	Unit type is the execution graph unit type.
INFO	Info column is reserved for additional information. It may contain, for example, index usage, the database table name and constraints used in solidDB to select rows. Note that the constraints listed here may not match those constraints given in the SQL statement.

The following texts may exist in the INFO column for different types of units.

Table 21. Texts in the unit INFO column

Unit type	Text in Info column	Description
TABLE UNIT	<i>tablename</i>	The table unit refers to table <i>tablename</i> .

Table 21. Texts in the unit INFO column (continued)

Unit type	Text in Info column	Description
TABLE UNIT	<i>constraints</i>	The constraints that are passed to the database engine are listed. If for example in joins the constraint value is not known in advance, the constraint value is displayed as NULL.
TABLE UNIT	SCAN TABLE	Full table scan is used to search for rows.
TABLE UNIT	SCAN <i>indexname</i>	Index <i>indexname</i> is used to search for rows. If all selected columns are found from an index, sometimes it is faster to scan the index instead of the entire table because the index has fewer disk blocks.
TABLE UNIT	PRIMARY KEY	The primary key is used to search rows. This differs from SCAN in that the whole table is not scanned because there is a limiting constraint to the primary key attributes.
TABLE UNIT	INDEX <i>indexname</i>	Index <i>indexname</i> is used to search for rows. For every matching index row, the actual data row is fetched separately.
TABLE UNIT	INDEX ONLY <i>indexname</i>	Index <i>indexname</i> is used to search for rows. All selected columns are in the index, so the actual data rows are not fetched separately by reading from the table.
JOIN UNIT	MERGE JOIN	Merge join is used to join the tables.
JOIN UNIT	3-MERGE JOIN	A 3-merge join is used to merge the tables.
JOIN UNIT	LOOP JOIN	Loop join is used to join the tables.
ORDER UNIT	NO ORDERING REQUIRED	No ordering is required, the rows are retrieved in correct order from solidDB.
ORDER UNIT	EXTERNAL SORT	External sorter is used to sort the rows.
ORDER UNIT	FIELD <i>n</i> USED AS PARTIAL ORDER	For distinct result sets, an internal sorter (in-memory sorter) is used for sorting and the rows retrieved from solidDB are partially sorted with column number <i>n</i> . The partial ordering helps the internal sorter avoid multiple passes over the data.
ORDER UNIT	<i>n</i> FIELDS USED FOR PARTIAL SORT	An internal sorter (in-memory sorter) is used for sorting and the rows retrieved from solidDB are partially sorted with <i>n</i> fields. The partial ordering helps the internal sorter to avoid multiple passes over the data.

Table 21. Texts in the unit INFO column (continued)

Unit type	Text in Info column	Description
ORDER UNIT	NO PARTIAL SORT	Internal sorter is used for sorting. The rows are retrieved in random order from solidDB to the sorter.
UNION UNIT	MERGE JOIN	Merge join is used to join the tables.
UNION UNIT	3-MERGE JOIN	A 3-merge join is used to merge the tables.
UNION UNIT	LOOP JOIN	Loop join is used to join the tables.
INTERSECT UNIT	MERGE JOIN	Merge join is used to join the tables.
INTERSECT UNIT	3-MERGE JOIN	A 3-merge join is used to merge the tables.
INTERSECT UNIT	LOOP JOIN	Loop join is used to join the tables.
EXCEPT UNIT	MERGE JOIN	Merge join is used to join the tables.
EXCEPT UNIT	3-MERGE JOIN	A 3-merge join is used to merge the tables.
EXCEPT UNIT	LOOP JOIN	Loop join is used to join the tables.

Example 1

```
EXPLAIN PLAN FOR SELECT * FROM TENKTUP1 WHERE
UNIQUE2_NI BETWEEN 0 AND 99;
```

Table 22. EXPLAIN PLAN FOR, Example 1

ID	UNIT_ID	PAR_ID	JOIN_PATH	UNIT_TYPE	INFO
1	2	1	3	JOIN UNIT	
2	3	2	0	TABLE UNIT	TENKTUP1
3	3	2	0		FULL SCAN
4	3	2	0		UNIQUE2_NI <= 99
5	3	2	0		UNIQUE2_NI >= 0
6	3	2	0		

Execution graph

JOIN UNIT 2 gets input from TABLE UNIT 3

TABLE UNIT 3 for table TENKTUP1 does a full table scan with constraints UNIQUE2_NI <= 99 and UNIQUE2_NI >= 0




Figure 4. Execution graph 1

Example 2

```
EXPLAIN PLAN FOR SELECT * FROM TENKTUP1, TENKTUP2
WHERE TENKTUP1.UNIQUE2 > 4000 AND TENKTUP1.UNIQUE2 < 4500
AND TENKTUP1.UNIQUE2 = TENKTUP2.UNIQUE2;
```

Table 23. EXPLAIN PLAN FOR, Example 2

ID	UNIT_ID	PAR_ID	JOIN_PATH	UNIT_TYPE	INFO
1	6	1	9	JOIN UNIT	MERGE JOIN
2	6	1	10		
3	9	6	0	ORDER UNIT	NO ORDERING REQUIRED
4	8	9	0	TABLE UNIT	TENKTUP2
5	8	9	0		PRIMARY KEY
6	8	9	0		UNIQUE2 < 4500
7	8	9	0		UNIQUE2 > 4000
8	8	9	0		
9	10	6	0	ORDER UNIT	NO ORDERING REQUIRED
10	7	10	0	TABLE UNIT	TENKTUP1
11	7	10	0		PRIMARY KEY
12	7	10	0		UNIQUE2 < 4500
13	7	10	0		UNIQUE2 > 4000
14	7	10	0		

Execution graph

JOIN UNIT 6 the input from order units 9 and 10 are joined using merge join algorithm

ORDER UNIT 9 orders the input from TABLE UNIT 8. Since the data is retrieved in correct order, no real ordering is needed

ORDER UNIT 10 orders the input from TABLE UNIT 7. Since the data is retrieved in correct order, no real ordering is needed

TABLE UNIT 8: rows are fetched from table TENKTUP2 using primary key. Constraints UNIQUE2 < 4500 and UNIQUE2 > 4000 are used to select the rows

TABLE UNIT 7: rows are fetched from table TENKTUP1 using primary key. Constraints UNIQUE2 < 4500 and UNIQUE2 > 4000 are used to select the rows




Figure 5. Execution graph 2

6.2 Tracing facilities for stored procedures and triggers

When debugging a stored procedure or a trigger, you may want to add "trace" commands to see which parts of the code are executing. Or you may want to trace every statement within the procedure or trigger. The following two sections explain how to do these things.

6.2.1 User-definable trace output from procedure code

From inside your stored procedure or trigger, you can send "trace" output to the `sqltrace.out` file by using the following command:

```
WRITETRACE (entry VARCHAR)
```

You can turn the output on or off by using the command:

```
ADMIN COMMAND 'usertrace { on | off }  
user username { procedure | trigger | table } entity_name'
```

The "entity_name" is the name of the procedure, trigger, or table for which you want to turn tracing on or off. If the keyword "table" is specified, then all triggers on that table are traced.

You may turn on (or off) tracing for a specified procedure, a specified trigger, or for all triggers on a specified table.

Trace is activated only when the specified user calls the procedure / trigger. This is useful, for example, when tracing propagated procedure calls in a advanced replication master.

Turning on tracing turns it on in all procedure/trigger calls by this user, not just calls from the connection that switched the trace on. If you have multiple connections that use the same username, then all of the calls in all of those connections will be traced. Furthermore, the tracing will be done on calls propagated to (executed on) the master, as well as the calls executed on the replica.

6.2.2 Procedure execution trace

If you must trace EVERY statement in your stored procedure or trigger, then you do not want to spend time to write a WRITETRACE statement for every SQL statement. Instead, you can simply turn on "PROCTRACE", which traces every statement inside the specified stored procedure or trigger. As with USERTRACE, you can turn proctrace on for a specified procedure, a specified trigger, or for all triggers associated with a particular table. The syntax is:

```
ADMIN COMMAND 'proctrace { on | off }  
user username { procedure | trigger | table } entity_name'
```

The "entity_name" is the name of the procedure, trigger, or table for which you want to turn tracing on or off.

Trace is activated only when the specified user calls the procedure / trigger. This is useful, for example, when tracing propagated procedure calls in a advanced replication master.

Turning on tracing turns it on in all procedure/trigger calls by this user, not just calls from the connection that switched the trace on. If you have multiple connections that use the same username, then all of the calls in all of those connections will be traced. Furthermore, the tracing will be done on calls propagated to (executed on) the master, as well as the calls executed on the replica.

If the keyword "table" is specified, then all triggers on that table are traced.

Example:

```
"create procedure trace_sample(i integer)  
returns(j integer)  
begin  
    j := 2*i;  
    return row;  
end";  
commit work;
```

```
admin command 'proctrace on user DBA procedure TRACE_SAMPLE';  
call trace_sample(2);
```

OUTPUT FROM EXAMPLE:

```
23.01 17:25:17 ---- PROCEDURE 'DBA.DBA.TRACE_SAMPLE' TRACE BEGIN ----  
0001:CREATE PROCEDURE TRACE_SAMPLE(I INTEGER)  
0002:RETURNS(J INTEGER)  
0003:BEGIN  
    --> I:=2  
    --> J:=NULL  
    --> SQLSUCCESS:=1  
    --> SQLERRNUM:=NULL  
    --> SQLERRSTR:=NULL  
    --> SQLROWCOUNT:=NULL  
  
0004:          J := 2*I;  
    --> J:=4
```

```
0005:          RETURN ROW;
0006:END
23.01 17:25:17 ---- PROCEDURE 'DBA.DBA.TRACE_SAMPLE' TRACE END ----
```

6.3 Measuring and improving performance of START AFTER COMMIT statements

6.3.1 Tuning performance of START AFTER COMMIT statements

Background tasks can be controlled with SSC-API and admin commands (see *IBM solidDB Shared Memory Access and Linked Library Access User Guide* for details). The task type `SSC_TASK_BACKGROUND` is used for the tasks that execute statements started with `START AFTER COMMIT`. You can give this task type higher priority or lower priority, or you may suspend this task type.

Note that there may be more than one of these tasks, but you cannot control them individually. In other words, if you call `SSCSuspendTaskClass` for `SSC_TASK_BACKGROUND`, it will suspend all the background tasks.

6.3.2 Analyzing failures in START AFTER COMMIT statements

There is a limit on the number of uncommitted `START AFTER COMMIT` statements that may exist simultaneously. (By "uncommitted", we mean that the transaction in which the `START AFTER COMMIT` statement was executed has not yet been committed. At this point, the body of the `START AFTER COMMIT` statement — for example, the procedure call — has not yet even started to execute.) If the maximum is reached, then an error is returned when the next `START AFTER COMMIT` is issued. The maximum number is configurable in `solid.ini` using the parameter named `MaxStartStatements` (for details, see the description of this parameter in *IBM solidDB Administrator Guide*).

If a statement cannot be started, the reason for it is logged into the system table `SYS_BACKGROUNDJOB_INFO`. Only failed `START AFTER COMMIT` statements are logged into this table. For more details about this table, see E.1.4, "`SYS_BACKGROUNDJOB_INFO`," on page 347.

The user can retrieve the information from the table `SYS_BACKGROUNDJOB_INFO` using either an SQL `SELECT` statement or by calling the system procedure `SYS_GETBACKGROUNDJOB_INFO`. The stored procedure `SYS_GETBACKGROUNDJOB_INFO` returns the row that matches the given jobid of the `START AFTER COMMIT` statement. For more details about `SYS_GETBACKGROUNDJOB_INFO`, see G.2.1, "`SYS_GETBACKGROUNDJOB_INFO`," on page 400.

If you want to be notified when a statement fails to start, you can wait on the system event `SYS_EVENT_SACFAILED`. See its description in H.1, "Miscellaneous events," on page 401 for details about this event. The application can wait for this event and use the jobid to retrieve the error message from the system table `SYS_BACKGROUNDJOB_INFO`.

7 Tuning performance with SQL

This section discusses SQL-related techniques that you can use to improve the performance of solidDB.

For tips on optimizing advanced replication data synchronization, see *Performance monitoring and tuning* in the *IBM solidDB Advanced Replication User Guide*.

7.1 Tuning SQL statements and applications

Tuning the SQL statements, especially in applications where complex queries are involved, is generally the most efficient means of improving the database performance.

Be sure to tune your application *before* tuning the RDBMS because:

- during application design you have control over the SQL statements and data to be processed
- you can improve performance even if you are unfamiliar with the internal working of the RDBMS you are going to use
- if your application is not tuned well, it will not run well even on a well-tuned RDBMS

You should know what data your application processes, what are the SQL statements used, and what operations the application performs on the data. For example, you can improve query performance when you keep SELECT statements simple, avoiding unnecessary clauses and predicates.

7.1.1 Evaluating application performance

To isolate areas where performance is lacking in your application, the solidDB provides the following diagnostic tools for observing database performance:

- SQL info facility
- EXPLAIN PLAN FOR statement

These tools are helpful in tuning your application and identifying any inefficient SQL statements in it. Read 6, “Diagnostics and troubleshooting for SQL,” on page 123 for additional information about how to use these tools.

In addition, the following commands provide useful information for evaluating performance.

- **ADMIN COMMAND 'status'**

This command returns statistics information from the server. For details, read about this command in *IBM solidDB Administrator Guide*.

- **ADMIN COMMAND 'perfmon'**

The command returns detailed performance statistics from the server. For details on producing performance reports with this command, see section *Performance counters (perfmon)* in the *IBM solidDB Administrator Guide*.

- **ADMIN COMMAND 'trace'**

This command switches tracing on for SQL statements and network communication. For complete syntax, see the trace option syntax under A.1, "ADMIN COMMAND," on page 141.

7.1.2 Using stored procedure language

Using stored procedures can speed up some operations in two ways:

- Statements in stored procedures are parsed and compiled once and then stored in compiled form. Statements outside stored procedures are re-parsed and compiled every time that they are executed. Thus, putting statements in stored procedures reduces overhead (parsing and compiling) if the statements are executed more than once.
- If you have multiple statements inside a single stored procedure, calling that stored procedure once may use fewer network "trips" than passing each statement individually from the client to the server.

7.1.3 Executing prepared statements (prepare once, execute many)

You can improve performance by preparing statements that are expected to be executed multiple times.

About this task

In your application design, identify statements that are executed multiple times. Ensure that the application prepares each SQL statement only once and uses the same prepared SQL statement for executes. Additionally, use dynamic parameter binding to ensure that the same SQL statement can be executed with different parameter values.

Using prepared SQL statements improves performance for the following reasons:

- An extra call to the server to prepare the SQL statement is omitted.
- The following type of overhead associated with the SQL prepare operation can be avoided: allocating memory for the statement, checking access rights, and validating column and table names.
- The SQL interpreter and optimizer can create an optimized access plan for data retrieval and reuse it when the statement is executed again.

Procedure

To prepare and execute an SQL statement in separate steps:

1. Prepare the SQL statement expression by calling the prepare function (for example, `SQLPrepare()`) and passing the statement or expression as the statement argument.
2. Call the bind parameter function (`SQLBindParameter()`) to bind any parameter markers you have in the SQL statement. The solidDB server supports dynamic parameter markers represented by a question mark (?).
3. Execute the prepared statement by calling the execute function (`SQLExecute()`).

What to do next

To re-execute the prepared statement, call the execute function again.

7.2 Optimizing single-table SQL queries

solidDB provides a simple SQL optimization feature that increases performance with specific types of single-table SQL queries. Performance improvements apply to SELECT, DELETE, and UPDATE statements. The feature does not apply to INSERT statements.

Simple SQL optimization is enabled/disabled with the **SQL.SimpleSQLOpt** parameter. By default, this feature is turned on and the **SQL.SimpleSQLOpt** parameter does not appear in the `solid.ini` file. To disable the feature, you must add the following lines to the `solid.ini` file:

```
[SQL]
SimpleSQLOpt=No
```

Once you have added these lines to the file, you can always enable the feature by specifying **SimpleSQLOpt=Yes** or removing the parameter from the [SQL] section.

When simple SQL optimization is turned on, solidDB automatically optimizes single-table SQL queries that meet the following conditions:

- The statement accesses only a single table.
- The statement does not contain a view, subquery, UNION, INTERSECT, or so on.
- The statement does not use ROWNUM.
- The statement does not use a solidDB sequence object that is used to retrieve sequence numbers.

As with other optimization techniques, the simple SQL optimization feature speeds up most queries, but reduces performance for a few types of queries. If particular queries run more slowly when you are using simple SQL optimization, you can turn off the feature.

7.3 Using indexes to improve query performance

You can use indexes to improve the performance of queries. A query that references an indexed column in its WHERE clause can use the index. If the query selects only the indexed column, the query can read the indexed column value directly from the index, rather than from the table.

If all the fields in the SELECT list of a query are in an index, then the solidDB optimizer can simply use that index, rather than doing an extra lookup to read the complete record. Similarly, if all the fields of a WHERE clause are in an index, then the optimizer can use that index — if the information in the index is enough to prove that the record will not qualify for the WHERE clause, then the optimizer can avoid looking up the complete record.

For example, suppose that we have a WHERE clause that refers to two or more columns, for example:

```
WHERE col1 = x AND col2 >= a AND col2 <=b
```

Suppose further that we have an index that contains both col1 and col2, and that has either col1 or col2 as the leading column of the key. For example, if we have an index on col2 + col3 + col1 then this index contains both columns, and one of those columns (col2) is the leading column in the key. If the user's query is

```
SELECT col1, col4
FROM table1
WHERE col1 = x AND col2 >= a AND col2 <=b;
```

then we do not need to look up the complete record unless the search criteria are met. After all, if the search criteria are not met, then we do not care what value col4 has and so we do not need to look up the full record.

If a table has a primary key, solidDB orders the rows on disk in the order of the values of the primary key. Since the rows are physically in order by the primary key, the primary key itself serves as an index, and optimization tips that apply to indexes also apply to the primary key.

If the table does not have a user-specified primary key, then the rows are ordered using the ROWID. The ROWID is assigned to each row when it is inserted, and each record gets a larger ROWID than the record inserted before it. Thus, in tables without user-specified primary keys, the records are stored in the order in which those rows were inserted. For more information about primary keys, read 4.4.1, "Primary key indexes," on page 97.

Searches with row value constructor constraints are optimized to use an index if an index is available. For efficiency, solidDB uses an index to resolve row value constructor constraints of the form (A, B, C) >= (1, 2, 3), where the operator may be any of the following: <, <=, >= and >. (The server does not use an index to resolve row value constructor constraints that contain the operators =, !=, or <>. The server may use an index to resolve other types of constraints that use =, !=, or <>.) For more information about row value constructors, see 2.5.5, "Row value constructors," on page 19.

Indexes improve the performance of queries that select a small percentage of rows from a table. You should consider using indexes for queries that select less than 15% of table rows.

7.3.1 Full table scan

If a query cannot use an index, solidDB must perform a full table scan to execute the query. This involves reading all rows of a table sequentially. Each row is examined to determine whether it meets the criteria of the query's WHERE clause. Finding a single row with an indexed query can be substantially faster than finding the row with a full table scan. On the other hand, a query that selects more than 15% of a table's rows may be performed faster by a full table scan than by an indexed query.

You should check every query using the EXPLAIN PLAN statement. Use your real data when doing this because the best plan will depend upon the actual amount of data and the characteristics of that data. The output from the EXPLAIN PLAN statement allows you to detect whether an index is really used and, if necessary, you can redo the query or the index. Full table scans often cause slow response time for SELECT queries as well as excessive disk activity.

To diagnose performance degradation problems, you can request statistics on file operations using ADMIN COMMAND 'perfmon' as described in section *Performance counters (perfmon)* in the *IBM solidDB Administrator Guide*.

To perform a full table scan, every block in the table is read. For each block, every row stored in the block is read. To perform an indexed query, the rows are read in the order in which they appear in the index, regardless of which blocks contain them. If a block contains more than one selected row it may be read more than once. So, there are cases when a full table scan requires less I/O than an indexed query, if the result set is relatively large.

7.3.2 Concatenated indexes

An index can be made up of more than one column. Such an index is called a concatenated index. We recommend using concatenated indexes when possible.

Whether or not an SQL statement uses a concatenated index is determined by the columns contained in the WHERE clause of the SQL statement. A query can use a concatenated index if it references a leading portion of the index in the WHERE clause. A leading portion of an index refers to the first column or columns specified in the CREATE INDEX statement.

Example:

```
CREATE INDEX job_sal_deptno ON emp(job, sal, deptno);
```

This index can be used by these queries:

```
SELECT * FROM emp WHERE job = 'clerk' and sal =  
800 and deptno = 20;  
SELECT * FROM emp WHERE sal = 1250 and job = salesman;  
SELECT job, sal FROM emp WHERE job = 'manager';
```

The following query does not contain the first column of the index in its WHERE clause and therefore cannot use the index:

```
SELECT * FROM emp WHERE sal = 6000;
```

Choosing columns to index

The following list gives guidelines in choosing columns to index:

- You should create indexes on columns that are used frequently in WHERE clauses.
- You should create indexes on columns that are used frequently to join tables.
- You should create indexes on columns that are used frequently in ORDER BY clauses.
- You should create indexes on columns that have few of the same values or unique values in the table.
- You should not create indexes on small tables (tables that use only a few blocks) because a full table scan may be faster than an indexed query.
- If possible, choose a primary key that orders the rows in the most appropriate order.
- If only one column of the concatenated index is used frequently in WHERE clauses, place that column first in the CREATE INDEX statement.
- If more than one column in a concatenated index is used frequently in WHERE clauses, place the most selective column first in the CREATE INDEX statement.

7.4 Waiting on events

In many programs, you may have to wait for a particular condition to occur before you can perform a certain task. In some cases, you may use a "while" loop to check whether the condition has occurred. solidDB provides Events, which in some cases allow you to avoid wasting CPU time spinning in a loop waiting for a condition.

One (or more) clients or threads can wait on an event, and another client or thread can post that event. For example, several threads might wait for a sensor to get a new piece of data. Another thread (working with that sensor) can post an event indicating that the data is available. For more information about events, see 3.5,

“Events,” on page 80 and various sections of Appendix A, “IBM solidDB SQL statements,” on page 141, including A.11, “CREATE EVENT,” on page 172.

7.5 Optimizing batch inserts and updates

You should design a database schema that supports running a batch insert in primary key order. Data in the database file is stored physically in the order defined by the primary key of the table. If no primary key is defined, data is stored in the database file in the order it is written to the database. Database operations (that is, reads and writes) always access data at the page level. The page size (block size) of the database is defined with the **IndexFile.BlockSize** parameter.

If the batch write operations are performed in the order that supports the primary key, the caching algorithms of the server are able to group the database file write operations. In this way, a larger number of rows are written to the disk in one physical disk I/O operation. In the worst case, if the insert order is different from the primary key order, each insert or delete operation requires rewriting a database page where only one row has changed.

For these reasons, it makes sense to ensure that tables of a batch write operation have primary keys that match the access order of the batch write operation. This type of database schema can make a significant difference in the performance of the operation.

For example, create a table with the following command:

```
CREATE TABLE USAGE_EVENT (  
  EVENT_ID INTEGER NOT NULL PRIMARY KEY,  
  DEVICE_ID INTEGER NOT NULL,  
  EVENT_DATA VARCHAR NOT NULL);
```

In this table, `EVENT_ID` is a sequence number. The insert and delete operations are done in the order specified by the `EVENT_ID` column, allowing for maximum efficiency.

The performance of batch write operations on this same table can be significantly worse if the first column of the primary key were `DEVICE_ID`, but data was written to the database in the `EVENT_ID` order. In this scenario, the number of file I/O operations needed to complete the batch write operation increases when the size of the table grows.

7.5.1 Increasing speed of batch inserts and updates

You can optimize the speed for large batch inserts and updates to solidDB. Following are guidelines for increasing speed:

1. Check that you are running the application with the `AUTOCOMMIT` mode set off.
solidDB ODBC Driver's default setting is `AUTOCOMMIT`. This is the standard setting according to the ODBC specification. To set your application with `AUTOCOMMIT` off, call the `SQLSetConnectOption` function as in the following example:

```
rc = SQLSetConnectOption  
(hdbc, SQL_AUTOCOMMIT, SQL_AUTOCOMMIT_OFF);
```
2. Do not use large transactions. Five hundred (500) rows is recommended as the initial transaction size. The optimal value for the transaction size is dependent on the particular application; you may need to experiment.

3. To make batch inserts faster, you can turn logging off. This, however, increases the risk of data loss during system failure. In some environments, this trade-off is tolerable.

Number 1 and 2 of these guidelines are the most important actions you can take to increase the speed of batch inserts. The actual rate of insertions also depends on your hardware, on the amount of data per row, and on the existing indices for the table.

7.6 Using optimizer hints

Hints are an extension of SQL that provide directives to the SQL optimizer for determining the query execution plan that is used. Hints are specified through embedded pseudo comments within query statements. The optimizer detects these directives or hints and bases its query execution plan accordingly. Optimizer hints allow applications to be optimized under various conditions to the data, query type, and the database. They not only provide solutions to performance problems occasionally encountered with queries, but shift control of response times from the system to the user.

Hints are needed because due to various conditions with the data, user query, and database, the SQL optimizer is not always able to choose the best possible execution plan. For example, you might want to force a merge join because you know, unlike the optimizer, that your data is already sorted. Also, sometimes specific predicates in queries can cause performance problems that the optimizer cannot eliminate. The optimizer may be using an index that you know is not optimal. In this case, you may want to force the optimizer to use one that produces faster results.

Hints are available for:

- Selecting merge or nested loop join
- Using a fixed join order as given in the from list
- Selecting internal or external sort
- Selecting a particular index
- Selecting a table scan over an index scan
- Selecting sorting before or after grouping

You can place hints in SQL statements as static strings, just after a SELECT, UPDATE, or DELETE keyword. Hints are not allowed after the INSERT keyword.

When there is an error in a hint specification, the whole SQL statement fails with an error message.

Enabling and disabling hints

Hints are enabled and disabled using the `SQL.EnableHints` parameter. By default, hints are enabled (`SQL.EnableHints=yes`).

Related reference:

A.55, "HINT," on page 236

A.55.1, "Hints supported by solidDB," on page 240

7.7 Diagnosing poor performance

There are different areas in solidDB that can result in performance degradation. In order to remedy performance problems, you need to determine the underlying cause. Following is a table that lists common symptoms of poor performance, possible causes, and directs you to the section in this chapter for the remedy.

Table 24. Diagnosing poor performance

Symptoms	Diagnosis	Solution
Slow response time for a single query. Other concurrent access to the database is affected. Disk may be busy.	<ul style="list-style-type: none"> Inefficient usage of indexes in the query. Non-optimal decision from the Optimizer. External sorting is not defined and a large internal sorting is causing excessive swapping to disk. 	<p>If index definitions are missing, create new indices or modify existing ones to match the indexing requirements of the slow query. For more details, read <i>Using indexes to improve query performance</i>.</p> <p>Run the EXPLAIN PLAN FOR statement for the slow query and verify whether the query optimizer is using the indices. For more details, see <i>EXPLAIN PLAN FOR statement</i>.</p> <p>If the Optimizer is not choosing the optimal query execution plan, override the Optimizer decision by using optimizer hints. For more details, see <i>Using optimizer hints</i>.</p>
Slow response time is experienced for all queries. An increase in the number of concurrent users deteriorates the performance more than linearly. When all users are thrown out and then reconnected, performance still does not improve.	Insufficient cache size.	Increase the cache size. Allocate for cache at least 0.5 MB per concurrent user or 2-5% of the database size. For more details, read the section <i>Defining database cache size</i> in <i>IBM solidDB Administrator Guide</i> .
Slow response time is experienced for all queries and write operations. When all users are thrown out and are connected, performance only improves temporarily. The disk is very busy.	The Bonsai Tree is too large to fit into the cache.	Make sure that there are no unintentionally long-running transactions. Verify that all transactions (also read-only transactions) are committed in a timely manner. For more details, read <i>Reducing Bonsai Tree size by committing transactions</i> in <i>IBM solidDB Administrator Guide</i> .
Slow performance during batch write operation as the database size increases. There is an excessive amount of disk I/O.	<ul style="list-style-type: none"> The data is committed to the database in batches that are too small. Data is written to disk in an order that is not supported by the primary key of the table. 	<p>Make sure that the autocommit is switched off and the write operations are committed in batches of at least 100 rows per transaction.</p> <p>Modify the primary keys or batch write processes so that write operations occur in the primary key order. For more details, read <i>Optimizing batch inserts and update</i>.</p>
The server process footprint grows excessively and causes the operating system to swap. The disk is very busy. The ADMIN COMMAND 'report' output shows a long list of currently active statements.	SQL statements have not been closed and dropped after use.	Make sure that the statements that are no longer in use by the client application are closed and dropped in a timely manner.

Appendix A. IBM solidDB SQL statements

This section provides descriptions of the solidDB SQL statements, including examples.

solidDB uses SQL syntax based on the ANSI X3H2 and IEC/ISO 9075 SQL standards. The SQL-89 Level 2 and SQL-92 Entry level standards are fully supported, including important ANSI X3H2-1992 (SQL-92) extensions. Many features of full SQL-92, SQL-99, and SQL-2003 standards are also supported.

A.1 ADMIN COMMAND

```
ADMIN COMMAND 'command_name'
```

```
command_name ::= ABORT | ASSETEXIT | BACKUP |  
BACKGROUNDJOB | BACKUPLIST | CHECKPOINTING | CLEANBGJOBINFO |  
CLOSE | DESCRIBE | ERRORCODE | ERROREXIT | ERRORMESSAGE | FILESPEC |  
GETREADONLYFLAG | HELP | HOTSTANDBY | INDEXUSAGE | INFO | LOGMESSAGE |  
LOGREADER | MAKECP | MEMORY | MESSAGES | MONITOR | NETBACKUP |  
NETBACKUPLIST | NETSTAT | NOTIFY | OPEN | PARAMETER | PASSTHROUGH STATUS |  
PERFMON | PERFMON DIFF | PERFMON TIMERS | PID | PROCTRACE | PROTOCOLS |  
REPORT | RUNMERGE | SAVE | SHUTDOWN | SQLLIST | STARTMERGE | STATUS |  
THROWOUT | TID | TRACE | TRACEMESSAGE | USERID | USERLIST |  
USERTRACE | VERSION
```

Usage

The ADMIN COMMAND is a SQL extension specific to solidDB server. You use ADMIN COMMANDs to execute administrative operations.

Using ADMIN COMMAND with solidDB SQL Editor (solsql)

When used with the solidDB SQL Editor (**solsql**), the *command_name* must be given with single quotation marks. For example:

```
ADMIN COMMAND 'backup'
```

If you use double quotation marks, the *command_name* is not recognized and the command fails.

Using ADMIN COMMAND with solidDB Remote Control (solcon)

When used with the solidDB Remote Control (**solcon**), the ADMIN COMMAND syntax includes the *command_name* only, without the quotation marks. For example:

```
backup
```

Abbreviations

Abbreviations for ADMIN COMMANDs are also available. For example:

```
ADMIN COMMAND 'bak'
```

To access a list of abbreviated commands, execute the following command:

```
ADMIN COMMAND 'help'
```

Return values

The result set contains two columns: RC and TEXT:

- The RC (return code) column is a command return code. If the execution of the command was successful, value 0 is returned.
- The TEXT column is the command reply.

Help

To access a list of abbreviated commands, execute the following command:

```
ADMIN COMMAND 'help'
```

To access the options and syntax description for a specific command, execute the following command:

```
ADMIN COMMAND 'command_name help'
```

Important:

- **All options of the ADMIN COMMAND are not transactional and cannot be rolled back.**

- **ADMIN COMMANDs and starting transactions**

Although ADMIN COMMANDs are not transactional, they will start a new transaction if one is not already open. (They do not commit or roll back any open transaction.) This effect is usually insignificant. However, it may affect the 'start time' of a transaction, and that may occasionally have unexpected effects. The concurrency control in solidDB is based on a versioning system; you see a database as it was at the time that your transaction started.

For example, if you issue an ADMIN COMMAND without another commit and then leave for an hour; when you return, your next SQL command may see the database as it was an hour ago, that is, when you first started the transaction with the ADMIN COMMAND.

- **Error codes**

Error codes in ADMIN COMMANDs return an error only if the command syntax or parameter values are incorrect. If only the requested operation may be started, the command returns SQLSUCCESS (0). The outcome of the operation itself is written into a result set. The result set has two columns: RC and TEXT. The RC (return code) column contains the return code of the operation: it is "0" for success, and different numeric values for errors. It is thus necessary to check both the codes of the ADMIN COMMAND statement and of the operation.

Table 25. ADMIN COMMAND syntax and options

Option syntax	Description
ADMIN COMMAND 'abort [backup netbackup]'	Aborts the active local or network backup process. The backup operation is not guaranteed to be atomic, therefore the cancelled operation might produce an incomplete backup file to the backup directory until the next backup takes place. If the option is not entered, the command defaults to ADMIN COMMAND 'abort backup'.
ADMIN COMMAND 'assertexit' Abbreviation: asex	Terminates the server immediately without a proper shut down.

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
<pre>ADMIN COMMAND 'backgroundjob' [LIST [-1] [user]] [ABORT {jobid user ALL}] [DELETE ERRORINFO {jobid user ALL}]' user ::= USER {username userid}</pre> <p>Abbreviation: bgjob</p>	<p>Lists and optionally aborts running background jobs, that is, SQL statements that have been started by using the START AFTER COMMIT statement.</p> <ul style="list-style-type: none"> • LIST option lists running jobs for all users or a specified user. • -1 option refers to a long list (similar to ADMIN COMMAND 'userlist -1'). • ABORT option aborts either jobs by job identification number or all jobs by user identification number. If you give the ABORT without arguments, it aborts all jobs from all users. • DELETE ERRORINFO option deletes error information from the SYS_BACKGROUNDJOB_INFO system table, where the errors encountered by background jobs are stored. This option performs the same operation as the deprecated ADMIN COMMAND 'CLEANBGJOBINFO' command.
<pre>ADMIN COMMAND 'backup [-s] [backup_directory]'</pre> <p>Abbreviation: bak</p>	<p>Makes a backup of the database. The operation can be performed in a synchronized or an asynchronous (default) manner. The synchronized operation is specified by using the -s option.</p> <p>The default backup directory is defined with the General.BackupDirectory. The backup directory can also be given as an argument. For example, backup abc creates a backup in directory abc. All directory definitions are relative to the solidDB working directory.</p>
<pre>ADMIN COMMAND 'backuplist'</pre> <p>Abbreviation: bls</p>	<p>Displays a status list of last local backups.</p>
<pre>ADMIN COMMAND 'checkpointing {ON OFF}'</pre> <p>Abbreviation: cp</p>	<p>Turns checkpointing on or off.</p>
<pre>ADMIN COMMAND 'cleanbgjobinfo'</pre> <p>Abbreviation: cleanbgi</p>	<p>Note: This command has been deprecated. Use ADMIN COMMAND 'backgroundjob' instead.</p> <p>Cleans the table SYS_BACKGROUNDJOB_INFO containing status data of background procedures.</p>
<pre>ADMIN COMMAND 'close'</pre> <p>Abbreviation: clo</p>	<p>Closes the server to new connections; no new connections are allowed.</p>
<pre>ADMIN COMMAND 'describe parameter param'</pre> <p>Abbreviation: des</p>	<p>Returns a description of all parameters or a parameter specified with <i>param</i>.</p> <p><i>param</i> must be given in the format section_name.param_name. The section and parameter names are case-insensitive.</p> <p>The following example describes parameter Com.Trace = y/n:</p> <pre>ADMIN COMMAND 'des parameter com.trace' RC TEXT -- ---- 0 Trace 0 If set to 'yes', trace information of the network messages is written to a file 0 BOOL 0 RW/STARTUP 0 0 0 No 7 rows fetched.</pre>
<pre>ADMIN COMMAND 'errorcode {all SOLID_error_code}'</pre> <p>Abbreviation: ec</p>	<p>Returns a description of all error codes or a specific error code.</p> <p><i>SOLID_error_code</i> is the code number, for example 10034.</p> <pre>ADMIN COMMAND 'errorcode 10034'; RC TEXT -- ---- 0 Code: DBE_ERR_SEQEXIST (10034) 0 Class: Database 0 Type: Error 0 Text: Sequence already exists 4 rows fetched.</pre>
<pre>ADMIN COMMAND 'errorexit <number>'</pre> <p>Abbreviation: erex</p>	<p>Forces the server into an immediate process exit with the given process exit code.</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'errormsg <string>' Abbreviation: errmsg	Outputs the user-defined <string> to the error message log (solerror.out).
ADMIN COMMAND 'filespec [-d -a <file_name> <max_file_size_in_bytes> [<device_number>"]' Abbreviation: fs	<p>Displays or modifies database (index) file specifications defined with the IndexFile.FileSpec parameter as well file sizes and current fill ratios (percentage).</p> <ul style="list-style-type: none"> • -d deletes the database file specified with <file_name max_file_size_in_bytes> [device_number] • -a adds a new database file specification as specified with <file_name> >max_file_size_in_bytes> [<device_number>] <p>For example: ADMIN COMMAND 'fs -a "solid3.db 3000M"; RC TEXT -- ---- 0 Added: FileSpec_3 = solid3.db 3145728000</p> <p>The database file specification changes are stored in the solid.ini configuration file at shutdown.</p>
ADMIN COMMAND 'getreadonlyflag' Abbreviation: grof	<p>Returns the read-only status of the database.</p> <p>You can set the database set to read-only mode with the General.Readonly parameter. Alternatively, if the server runs out of disk space, it switches to read-only mode automatically.</p>
ADMIN COMMAND 'help' Abbreviation: ?	Displays available commands.
ADMIN COMMAND 'hotstandby [option]' Abbreviation: hsb	<p>A HotStandby command.</p> <p>For a list of options, see the <i>IBM solidDB High Availability User Guide</i>.</p>
ADMIN COMMAND 'indexusage' Abbreviation: idxu	Displays the indexes, showing the number of times each index has been used.

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'info [options]' Abbreviation: info	<p>Returns server information.</p> <p>The output consists of 25 rows of data.</p> <p><i>options</i> are as follows:</p> <ul style="list-style-type: none"> • numusers - Number of current users. • maxusers - Maximum number of users. • sernum - Server serial number. • dbsize - Database size (KB). • logsize - Size of log files (KB). • uptime - Server startup timestamp. • bcktime - Timestamp of last successfully completed local backup. • cptime - Timestamp of last successfully completed checkpoint. • tracestate - Current trace state - see ADMIN COMMAND 'trace' for information on tracing. • monitorstate - Current monitor state, shown as the number of users who have SQL monitoring currently enabled (see ADMIN COMMAND 'monitor' for information on SQL monitoring). <p>If all users have SQL monitoring enabled, the value is -1.</p> <ul style="list-style-type: none"> • openstate - Current state for accepting connections. Open means that the database server accepts new connections. • nummerges - Number of merges. • numlocks - Number of locks. • numcursors - Number of open cursors. • numtransactions - Number of open transactions. • memtotal - Total amount of allocated memory (bytes). • dbfreesize - Amount of free space remaining in database (KB). • dbpagesize - Database page size (KB). • imdbsize - Amount of space used by in-memory tables (including temporary tables and transient tables) and the indexes on those tables. The return value is in kilobytes (KB) and is in the form of a VARCHAR. • name - Server name. You can set the server name with the solidDB startup option <code>-n name</code>. • primarystarttime - The time the Primary role has started. • secondarystarttime - The time the Secondary role has started. • dbconfigsize - The configured database size (MB), as set with the IndexFile.FileSpec parameter. • dbcreatetime dbcreationtime - Database creation timestamp. • processsize psize - System-level virtual process size (KB). <p>More than one option can be used per command. Values are returned in the same order as requested, one row for each value.</p> <p>Example:</p> <pre>ADMIN COMMAND 'info dbsize logsize'; RC TEXT -- ---- 0 851968 0 573440 2 rows fetched.</pre>
ADMIN COMMAND 'logmessage <string>' Abbreviation: logmsg	<p>Outputs the user-defined <string> to the message log (solmsg.out).</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'logreader stop [all <partition_id>]' Abbreviation: lr	<p>This command stops the transmission of log records on active log reader connections.</p> <p>When this command is issued, the active log reader applications reach the end of the result set (SQLSTATE 0200, No data found) when fetching the next row of the SYS_LOG table.</p> <p>If the form LOGREADER STOP or LOGREADER STOP ALL is used, all log record transmissions are stopped. If a <PARTITION_ID> is given, the command affect only the log reader operation on that partition.</p> <p>To access the log again, the application needs to reconnect. The log reading may be resumed without any loss of information if the last read position is known. If the SYS_LOG table is accessed without specifying the log position, the reading starts from the live data.</p> <p>Important: The stopping of the log transmission is effective immediately, regardless of the fact that there might be records in the log awaiting transmission.</p> <p>If the server is running in the relaxed durability mode (default), do not execute LOGREADER STOP before all the records are written to the log, if those records are meant to be seen in the log reader. With the default logging settings, it is safe to wait for 5 seconds after the last write operation.</p>
ADMIN COMMAND 'makecp [-s]' Abbreviation: mcp	<p>Makes a checkpoint.</p> <p>Only users with SYS_ADMIN_ROLE privilege can execute this command.</p> <p>By default, the checkpoint is asynchronous. With the option -s, the command returns only after the checkpoint has completed.</p>
ADMIN COMMAND 'memory' Abbreviation: mem	<p>Returns the server process memory size, that is, the amount of memory allocated by the server based on internal solidDB memory counters, including the memory used by data in the in-memory tables.</p> <p>Note: The reported process memory size can differ from the process size reported by your operating system.</p>
ADMIN COMMAND 'messages [{ warnings errors }] [count]' Abbreviation: mes	<p>Displays server messages. Optionally, you can also specify the severity and message numbers of the output.</p> <p>For example, ADMIN COMMAND 'messages warnings 100' displays last 100 warnings.</p>
ADMIN COMMAND 'monitor { on off } [user { username userid }]' Abbreviation: mon	<p>Sets server monitoring on and off.</p> <p>When set to on, user activity and SQL calls are logged into the sol trace.out file.</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'netbackup [options] [DELETE_LOGS KEEP_LOGS] [connect connect str] [dir backup dir]' Abbreviation: nbak	<p>Makes a network backup of the database. The operation can be performed as a synchronized or an asynchronous (default) manner.</p> <p>options can be</p> <ul style="list-style-type: none"> • -s Synchronized execution • -I Executes a full database integrity check • -i Executes a database index integrity check <p>DELETE_LOGS KEEP_LOGS defines whether backup logs are deleted or kept in the source server. Default is DELETE_LOGS.</p> <p>Note:</p> <ul style="list-style-type: none"> • DELETE_LOGS is referred to as <i>Full backup</i>. • KEEP_LOGS is referred to as <i>Copy backup</i>. Using KEEP_LOGS corresponds to setting the General.NetBackupDeleteLog parameter to no. <p>connect <i>connect str</i> specifies the connection to the NetBackup Server. If <i>connect str</i> is omitted, it must be specified in the <code>solid.ini</code> configuration file. For the full connect string syntax, see Format of the connect string.</p> <p><i>dir backup dir</i> defines the backup directory in the NetBackup Server. The path can be either absolute or relative to the netbackup root directory.</p> <p>The default connect string and the default netbackup directory are defined with the General.NetBackupConnect and the General.NetBackupDirectory parameters.</p> <p>The options that are entered with this command override the values specified in the <code>solid.ini</code> file.</p> <p>Directory definitions are relative to the solidDB working directory.</p>
ADMIN COMMAND 'netbackuplist' Abbreviation: nbls	Displays a status list of the most recently made network backups of the database server.
ADMIN COMMAND 'netstat' Abbreviation: net	Displays server settings and the network status.
ADMIN COMMAND 'notify user {username user id ALL } message' Abbreviation: not	<p>This command sends an event to a given user with event identifier NOTIFY. This identifier is used to cancel an event-waiting thread when the statement timeout is not long enough for a disconnect or to change the event registration.</p> <p>The following example sends a notify message to a user with user id 5; the event then gets the value of the message parameter.</p> <p>ADMIN COMMAND 'notify user 5 Canceled by admin'</p>
ADMIN COMMAND 'open' Abbreviation: ope	Opens server for new connections; new connections are allowed.

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'parameter [-r] [-t] [name[= [* value] [temporary]]' Abbreviation: par	<p>Displays and sets server parameter values.</p> <p>If you run the command without any options, all parameters are displayed.</p> <p>The output can contain three columns. For example:</p> <pre>0 PassThrough SqlPassthroughRead Force Conditional None</pre> <ul style="list-style-type: none"> • First column shows the current value (Force) that might have been changed dynamically. • Second column shows the value set in the .ini file at startup. (Conditional) • Third column shows the factory value. (None) • -r means that only the current parameter values are returned. • -t means that the changed value is not stored in the solid.ini file (same as temporary). • name may be a section name or a parameter name prefaced by a section name (section_name.parameter_name). There must be a period between the section name and the parameter name. • = [* value] [temporary] <ul style="list-style-type: none"> - If you assign a parameter value with an asterisk (*), the parameter will be set to its factory value. - If value is not specified, the parameter will be set to its startup value. - temporary means that the changed value is not stored in the solid.ini file. <p>For example:</p> <ul style="list-style-type: none"> • 'parameter general' displays all parameters from section [General]. • 'parameter general.readonly' displays the parameter ReadOnly in the [General] section. • 'parameter com.trace=yes' sets communication trace on. • 'parameter com.trace=' sets communication trace to its startup value. • 'parameter com.trace=*' sets communication trace to its factory value.
ADMIN COMMAND 'passthrough status' Abbreviation: pt	<p>Provides the following status information about the SQL passthrough connections:</p> <ul style="list-style-type: none"> • NO REMOTE SERVER - no remote server object defined • NOT CONNECTED - not connected, no errors • CONNECTED - connected • LOGIN FAILED - failed at login • CONNECTION BROKEN - connection broken

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'perfmon [- c - r] [print_options] [name_prefix_list] [†] Abbreviation: pmon	<p>Returns server performance counters for the past few minutes at approximately one minute intervals. Most values are shown as the average number of events per second. Counters that cannot be expressed as events per second (for example, database size) are expressed in absolute values.</p> <ul style="list-style-type: none"> • -c - prints actual counter values for each snapshot. • -r - prints counter values in raw mode, which includes only the latest counter values without any formatting. The counter names are not printed. This option is useful if actual monitoring is performed using some other external program that retrieves the counter values from the server. You can retrieve the counter names with the --xnames option. • print_options <ul style="list-style-type: none"> - -xtime - prints the time in seconds - -xtimediff - prints the difference to the last pmon call in milliseconds - -xnames - prints out the column names for the output - -xdiff - indicates the difference to the last ADMIN COMMAND 'perfmon' execution instead of the absolute value • name_prefix_list - limits the output to specific counter types, as indicated by the first word in the counter name. For example, to print all File related counters, the name_prefix_list should be file. You can also specify multiple prefixes. <p>The following example returns all information:</p> <p>ADMIN COMMAND 'perfmon'</p> <p>The following example returns all values for counters whose name starts with prefix File and Cache.</p> <p>ADMIN COMMAND 'perfmon -c file cache'</p>
ADMIN COMMAND 'perfmon diff [start stop] [filename][interval]' Abbreviation: pmon diff	<p>Starts a server task that prints out all perfmon counters with specified intervals to a file.</p> <ul style="list-style-type: none"> • filename is the name of the output file. The performance data is output in comma-separated value format; the first row contains the counter names, and each subsequent row contains the performance data per each sampling time. The default file name is pmondiff.out. • interval is the interval in milliseconds at which performance data is collected. The default interval is 1000 milliseconds. <p>The following command starts a task that outputs performance data to myd.csv file on 500 milliseconds interval:</p> <p>ADMIN COMMAND 'pmon diff start myd.csv 500'</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'perfmon timers [start stop list clear]' Abbreviation: pmon timers	<p>Produces information about execution times of database operations such as SQL execute and file operations for each user.</p> <p>start starts the timers and clears the existing counter values.</p> <p>stop stops the timers and keeps the current counter values.</p> <p>list lists the current counter values.</p> <p>clear clears the current counter values.</p> <p>The timer information is given in seconds. The values are cumulative since last perfmon timers start or perfmon timers clear.</p> <p>The output can be viewed in the console window (perfmon timers list) or printed into a report file with ADMIN COMMAND 'report report_name'. In the report file, the timer information is listed under the section PERFORMANCE TIMERS. The output lists the execution times for each user, identified with the userid.</p> <p>For more information and examples of the output, see ADMIN COMMAND 'perform timers'.</p>
ADMIN COMMAND 'pid' Abbreviation: pid	Returns server process id.
ADMIN COMMAND 'proctrace { on off } user <i>username</i> { procedure trigger table } <i>entity_name</i> ' Abbreviation: ptrc	<p>This turns on tracing in stored procedures and triggers.</p> <p><i>username</i> is the name of the user whose procedure calls (or triggers) you want to trace. If multiple connections are using the same username, calls from all of those connections will be traced. Furthermore, if you are using advanced replication, the tracing will be done not only for calls on the replica, but also calls that are propagated to the master and then executed on the master.</p> <p><i>entity_name</i> is the name of the procedure, trigger, or table for which you want to turn tracing on or off. If you specify a procedure or trigger name, then it will generate output for every statement in the specified procedure or trigger. If you specify a table name, then it will generate output for all triggers on that table. Trace is activated only when the specified username calls the procedure / trigger.</p> <p>For more details about proctrace, see section Tracing facilities for stored procedures and triggers in <i>IBM solidDB SQL Guide</i>.</p> <p>See also ADMIN COMMAND 'usertrace'.</p>
ADMIN COMMAND 'protocols' Abbreviation: prot	<p>Returns a list of available communication protocols.</p> <p>Example (Windows environments):</p> <pre>ADMIN COMMAND 'protocols'; RC TEXT -- ---- 0 NmPipe np 0 TCP/IP tc 2 rows fetched.</pre>
ADMIN COMMAND 'report <i>filename</i> ' Abbreviation: rep	Generates a report of server information and statistics to a file defined with <i>filename</i> .
ADMIN COMMAND 'runmerge' Abbreviation: rm	Runs an index merge.
ADMIN COMMAND 'save parameters [<i>filename</i>]' Abbreviation: save	Saves the set of current configuration parameter values to a file. If no file name is given, the default <i>solid.ini</i> file is rewritten. This operation is performed implicitly at each checkpoint.
ADMIN COMMAND 'shutdown [<i>force</i>]' Abbreviation: sd	<p>Stops the server process.</p> <p>If the force option is used, the active transactions are aborted and the users are disconnected forcefully.</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'sqllist [top number_of_statements]'	Prints out a list of the longest running SQL statements among the currently running statements. You must specify the number of statements you want to list.
ADMIN COMMAND 'startmerge' Abbreviation: sm	Starts and waits for completion of merge.
ADMIN COMMAND 'status' Abbreviation: sta	<p>Displays statistics for the server since the startup.</p> <p>The output provides the following information:</p> <ul style="list-style-type: none"> • Server startup timestamp • Working directory • Configuration file location and name • Memory statistics – Amount of memory allocated by the server based on internal solidDB memory counters, including the memory used by data in the in-memory tables (same as ADMIN COMMAND 'memory' output value) • Process size statistics in KB <ul style="list-style-type: none"> – Resident set size - Actual process size in memory as reported by the operating system – Virtual size - System-level virtual process size (same as ADMIN COMMAND 'info processsize' output value) • Transaction count statistics: <ul style="list-style-type: none"> – Commit – Number of committed transactions – Abort – Number of system-aborted transactions – Rollback – Number of transactions rolled back by user – Total – Total number of committed, aborted, and rolled back transactions – Read-only – Number of read-only transactions – Trxbuf – Number of transactions in transaction buffer – Active – Number of active transactions (same as performance counter <i>Trans active</i>) – Validate – Number of active transactions being validated at commit phase (same as performance counter <i>Trans validate</i>) • Cache count statistics: <ul style="list-style-type: none"> – Hit rate – Percentage of successful bufferpool cache hits (disk access avoided) – Find – Number of searches in cache – Read – Number of read operations on disk – Write – Number of write operations from cache to disk • Database statistics: <ul style="list-style-type: none"> – Index writes – Number of write operations – (Index writes) After last merge – Number of write operations since last merge – Log writes – Number of log write operations – (Log writes) After last cp – Number of log write operations since last checkpoint – Active searches – Number of active searches on database engine level – (Active searches) Average – Average number of active searches on database engine level – Database size – Log size • User count statistics <ul style="list-style-type: none"> – Current – Number of current connected users – Maximum – Number of concurrently connected users since startup – Total – Number of connected users since startup <p>For more information, see Checking database status.</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'status backup netbackup' Abbreviation: sta backup netbackup	Displays status of the last started local or network backup. The status can be one of the following: <ul style="list-style-type: none"> • If the last backup was successful or no backups have been requested, the output is 0 SUCCESS. • If the backup is in process (for example, started but not ready yet), then the output is 14003 ACTIVE. • If the backup is being finalized, the output is 14003 STOPPING. • If the last backup failed, the output is: <i>errorcode</i> ERROR where the <i>errorcode</i> shows the reason for the failure.
ADMIN COMMAND 'throwout {username userid all}' Abbreviation: to	Exits all or specific users from solidDB. To exit a specified user, give the username or user id as an argument. To throw out all users, use the keyword ALL as an argument.
ADMIN COMMAND 'tid' Abbreviation: tid	This command returns the ID (4-digit code) of the current user thread (in the server).
ADMIN COMMAND 'trace { on off } sql est estplans rpc sync flowplans rexec batch logreader passthrough xa hac info <level> func proc all active' Abbreviation: tra	Sets server trace on or off. The name of the default trace file is soltrace.out. The tracing options are: <ul style="list-style-type: none"> • sql - SQL messages • est - SQL estimator information • estplans - SQL execution plan • rpc - Network communications • sync - synchronization messages • flowplans - plans of SQL statements related to advanced replication • rexec - remote procedure call information • batch - background job and deferred procedure call information • logreader - logs the following information into the trace file soltrace.out. <ul style="list-style-type: none"> – Logreader read started. – Errors in logreader cursor start. Total of 14 different error conditions are printed. – Logreader read stopped. – Abnormal read stop after certain system changes. – High level information of number of returned log records and read progress. Each information is tagged with user id so operations from different users can be separated. <ul style="list-style-type: none"> • passthrough - provides tracing information about the SQL passthrough connections and the loading of the ODBC driver as follows: <ul style="list-style-type: none"> – Loading of the ODBC driver: driver name and status of the load – Status of connections to the back-end: connect/reconnect/disconnect/broken • xa - distributed transaction information • hac - High Availability Controller (HAC); trace information is output to hactrace.out in the HAC working directory Note: To start tracing on HAC, you must issue the command on a HAC connection. For example, connect to HAC with solsql using the port defined with the HACController.Listen parameter in the solidhac.ini configuration file. • info <level> - SQL execution trace (level can be 0...8) • func - function execution information • proc - stored procedure execution information • all - both SQL messages and network communications messages are written to the trace file. • active - lists all active traces

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'tracemessage <string>' Abbreviation: trcmsg	Outputs the user-defined <string> to the trace message log (soltrace.out).
ADMIN COMMAND 'userid' Abbreviation: uid	<p>Returns the user identification number of the current connection.</p> <p>The lifetime of an Id is that of the user session. After a user logs out, the number may be reused.</p> <pre>ADMIN COMMAND 'userid' RC TEXT -- ---- 0 8 1 rows fetched.</pre> <p>For example, the userid can be used in the ADMIN COMMAND 'throwout' command to disconnect a specific user.</p>

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
ADMIN COMMAND 'userlist [-1] [name id]' Abbreviation: ul	<p>This command displays a list of users that are currently logged into the database, as well as information about various database operations and settings for each user. The option -1 (long) displays a more detailed output.</p> <p>Without the -1 option, the following information is displayed: <i>User name</i>, <i>User Id</i>, <i>Tid</i>, <i>Type</i>, <i>Machine Id</i>, <i>Login time</i>, <i>Client version</i>, and <i>Appinfo</i> (if available).</p> <p>With the -1 option, the following information is displayed:</p> <ul style="list-style-type: none"> • <i>User name</i> - The user name of the connected user. • <i>User Id</i> - The user session identification number (userid) within the database. The lifetime of the userid is that of the user session. After the user logs out, the number can be reused. • <i>Tid</i> - The identification number as a 4-digit code of the current user thread in the server. • <i>Type</i> - Client type. Possible values are: <ul style="list-style-type: none"> - <i>Java</i>, which refers to a client using JDBC - <i>ODBC</i>, which refers to a client using ODBC, including solidDB SQL Editor (solsql) - <i>Solcon</i>, which refers to solidDB Remote Control (solcon) • <i>Machine id</i> - The client computer name (host name) and its IP address, if available • <i>Login time</i> - The client computer login timestamp • <i>Client version</i> - The version of the JDBC or ODBC client, as of V7.0.0.2 Interim Fix 2. <p>Note:</p> <ul style="list-style-type: none"> - The client version information is not available for clients prior to V7.0.0.2 Interim Fix 2 or for solidDB Remote Control (solcon) connections. - For solidDB SQL Editor (solsql) connections, the ODBC client version is shown. • <i>Appinfo</i> - The value of the client computer's environmental variable SOLAPPINFO (ODBC), or the value of JDBC connection property solid_appinfo. • <i>Last activity</i> - The time when the client last time sent a request to the server. • <i>Autocommit</i> - Value 0 means that the autocommit mode is switched off; the current transaction is open until a COMMIT or ROLLBACK statement is issued. Value 1 means that the autocommit mode is switched on; each statement is automatically committed. • <i>RPC compression</i> - Indicates whether the data transmission compression is on or off. • <i>Transparent failover</i> - This field indicates if Transparent Failover (TF) is in use (HotStandby configurations). Because solidDB tools do not support TF, you will only see a "no" value in this field when using solsql or solcon. • <i>Transparent cluster</i> - Transparent cluster indicates whether the load balancing feature (in HSB) is enabled for this connection or not. • <i>Transaction active</i> - This field indicates whether there is an open, uncommitted transaction on the connections (value 1) or not (value 0). When the connection is set for Autocommit, the value is, most of the time, 0. • <i>Transaction duration</i> - This field indicates the duration of the currently open transaction. After COMMIT or ROLLBACK, the value becomes 0. • <i>Transaction isolation</i> - This field indicates the transaction isolation level for the transactions. The isolation level decides how data which is a part of an ongoing transaction is made visible to other transactions.

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
<p>..continued..</p> <p>ADMIN COMMAND 'userlist [-1] [name id]'</p> <p>Abbreviation: ul</p>	<ul style="list-style-type: none"> • <i>Transaction durability</i> - This field indicates the durability of the currently open transaction. • <i>Transaction safeness</i> - This field indicates the safeness of the currently open transaction (set with HotStandby.SafenessLevel). • <i>Transaction autocommit</i> - This field indicates whether the currently open transaction is automatically committed. If the transaction autocommit for the current transaction is switched off (value 0), the current transaction is open until a COMMIT or ROLLBACK statement is issued. After that, a new statement starts a new transaction. If the autocommit mode is switched on for the current transaction (value 1), each statement is automatically committed. • <i>Current catalog</i> - Indicates the current catalog name. • <i>Current schema</i> - Indicates the current schema name. • <i>Sortgroubby</i> - Indicates how the GROUP BY statement is performed if explicit information about the number of result groups is not available. There are two possible values: <ul style="list-style-type: none"> - ADAPTIVE - GROUP BY input is pre-sorted if the real number of result groups exceeds the number of rows that fit into the central memory array for GROUP BY. - <ul style="list-style-type: none"> STATIC - GROUP BY input is pre-sorted whenever there are at least two items in the GROUP BY list. Otherwise, the GROUP BY input is not pre-sorted. • <i>Simple optimizer rules</i> - Indicates whether simple optimizer rules are in use (SQL.SimpleOptimizerRules) Possible values are Yes/No/Default. • <i>Statement max time</i> - Indicates the connection-specific statement maximum execution time in seconds. This setting is effective until a new maximum time is given. Zero time indicates that there is no maximum time. This is the default value. • <i>Lock timeout</i> - Indicates the timeout set by using the SET LOCK TIMEOUT statement. • <i>Optimistic lock timeout</i> - Indicates the timeout set by using the SET OPTIMISTIC LOCK TIMEOUT statement. • <i>Idle timeout</i> - Indicates the timeout set by using the SET IDLE TIMEOUT statement. • <i>Join Path Span</i> - Indicates the join path span value set by using the SET SQL JOINPATHSPAN statement. • <i>RPC seqno</i> - Internal protocol message sequence number. • <i>SQL sortarray</i> - The size of user-specific internal sort array. • <i>SQL unionsfromors</i> - The value tells how many (at most) OR operators may be converted to UNIONS. Unions are faster but require more memory to execute. • <i>EVENT QUEUE LENGTH</i> - Indicates the number of posted events in the event queue. • <i>Connection idle timeout</i> - Indicates the connection idle timeout setting • <i>Stmt id</i> - The current statement identification number. The numbers are session specific and they are assigned for each different statement. • <i>Stmt state</i> - An internal statement execution state. • <i>Stmt rowcount</i> - The number of rows retrieved or inserted in the current statement. • <i>Stmt start time</i> - The current statement start date and time. • <i>Stmt last activity time</i> - The timestamp of the most recent statement. • <i>Stmt duration</i> - Internal statement duration in seconds. Note: this value has no relevance to the externally visible statement latency. Typically, the statement duration is much longer than latency. • <i>Stmt SQL str</i> - The current SQL statement string.

Table 25. ADMIN COMMAND syntax and options (continued)

Option syntax	Description
<pre>ADMIN COMMAND 'usertrace { on off } user <i>username</i> { procedure trigger table } <i>entity_name</i>' Abbreviation: utrc</pre>	<p>This turns on user tracing in stored procedures and triggers. This command will generate output for every WRITETRACE statement in the specified procedure or trigger.</p> <ul style="list-style-type: none"> <i>username</i> is the name of the user whose procedure calls (or triggers) you want to trace. If multiple connections are using the same username, then calls from all of those connections will be traced. Furthermore, if you are using advanced replication, the tracing will be done not only for calls on the replica, but also calls that are propagated to the master and then executed on the master. <i>entity_name</i> is the name of the procedure, trigger, or table for which you want to turn tracing on or off. If you specify a table name, it will generate output for all triggers on that table. Trace is activated only when the specified user calls the procedure / trigger. <p>For more details about usertrace, see section Tracing facilities for stored procedures and triggers in <i>IBM solidDB SQL Guide</i>.</p> <p>See also ADMIN COMMAND 'proctrace'.</p>
<pre>ADMIN COMMAND 'version' Abbreviation: ver</pre>	<p>Displays server version information and information related to the solidDB software licence in use.</p>

A.2 ADMIN EVENT

```
ADMIN EVENT 'command'
command_name ::=
    REGISTER { event_name [ , event_name ... ] | ALL } |
    UNREGISTER { event_name [ , event_name ... ] | ALL } |
    WAIT
event_name ::= the name of a system event
```

Usage

The ADMIN EVENT statement is a solidDB-specific extension to SQL that allows you to register for and wait for system-generated events without writing and calling a stored procedure. You cannot wait for user events using ADMIN EVENT. If you want to wait for user events, you must write and call a stored procedure.

When using events outside stored procedures, you must register for and wait for the event explicitly. You must register for the event before you wait for it.

Tip: The ADMIN EVENT works differently from the way that WAIT works in stored procedures. In stored procedures, explicit registration is optional.

- You cannot register to synchronization events (starting with "SYNC_") with this command, because the ADMIN EVENT 'wait' command is not able to return variable result sets. Instead, you must use stored procedures to handle synchronization events.
- Once the connection starts to wait for an event, the connection will not be able to do anything else until the event is posted.
- You may register for multiple events. When you wait, you cannot specify which type of event to wait for. The wait will continue until you have received any of the events for which you have registered.
- The ADMIN EVENT command does not provide an option to post an event.
- To use ADMIN EVENT, you must have administrator privileges or be granted the role SYS_ADMIN_ROLE.

Examples

```
ADMIN EVENT 'register sys_event_hsbstateswitch';
ADMIN EVENT 'wait';
ADMIN EVENT 'unregister sys_event_hsbstateswitch';
```

After the event is posted by the system, you will see something similar to the following:

ENAME	POSTSRVTIME	UID	NUMDATAINFO	TEXTDATA
SYS_EVENT_HSBSTATESWITCH	2003-10-28 18:10:14	-1	NULL	PRIMARY ACTIVE

1 rows fetched.

Related reference:

A.15, “CREATE PROCEDURE,” on page 180

Related information:

3.5, “Events,” on page 80

A.3 ALTER REMOTE SERVER

```
ALTER REMOTE SERVER SET USERNAME <username> | PASSWORD <password>
```

Usage

The ALTER REMOTE SERVER statement modifies the back-end login information in the SYS_SERVER system table. The login data is used for SQL passthrough purposes in Universal Cache.

By default, username and password are stored in uppercase letters. To retain case sensitivity, enter the username and password in single quotation marks.

Examples

```
ALTER REMOTE SERVER SET USERNAME dba PASSWORD dba
CREATE REMOTE SERVER PASSWORD 'PwD123'
```

See also

A.18, “CREATE [OR REPLACE] REMOTE SERVER,” on page 201

A.38, “DROP REMOTE SERVER,” on page 223

A.4 ALTER TABLE

```
ALTER TABLE base_table_name
{
  ADD [COLUMN] column_identifier data_type
  [DEFAULT literal | NULL] [NOT NULL] |
  ADD CONSTRAINT constraint_name dynamic_table_constraint |
  DROP CONSTRAINT constraint_name |
  ALTER [ COLUMN ] column_name
  {DROP DEFAULT | {SET DEFAULT literal | NULL} } |
  {{ADD | DROP} NOT NULL }
  DROP [COLUMN] column_identifier |
  RENAME [COLUMN]
  column_identifier column_identifier |
  MODIFY [COLUMN] column_identifier data-type |
  MODIFY SCHEMA schema_name |
  SET HISTORY COLUMNS (c1, c2, c3) |
  SET {OPTIMISTIC | PESSIMISTIC} |
```

```

SET STORE {DISK | MEMORY} |
SET [NO]SYNCHISTORY |
SET TABLE NAME new_base_table_name
SET POOLSIZE percentage
ADD CACHE SEGMENT segment_name [WHERE range_specification]
DROP CACHE SEGMENT segment_name

}

where
dynamic_table_constraint::=
{FOREIGN KEY (column_identifier [, column_identifier] ...)
REFERENCES table_name [(column_identifier [, column_identifier] ) ...]}
[referential_triggered_action] |
CHECK (check_condition) | UNIQUE (column_identifier)
referential_triggered_action::=
ON {UPDATE | DELETE} {CASCADE | SET NULL | SET DEFAULT |
RESTRICT | NO ACTION}
range_specification ::= <column> < <value>
                        <column> > <value>
                        <column> = <value>
                        <column> <> <value>
                        <column> BETWEEN <value> AND <value>
                        <column> OLDER THAN <value> DAYS
                        <column> OLDER THAN <value> SECONDS
                        <column> NEWER THAN <value> DAYS
                        <column> NEWER THAN <value> SECONDS

```

Usage

You can use the ALTER TABLE statement for the following:

- Modify the structure of a table; add, remove, modify, and rename columns.

Note: A column cannot be dropped if it is part of a unique or primary key.

- Modify the column size.

You can change the width of a column using the ALTER TABLE command. You can increase the column width at any time (that is, whether a table is empty [no rows] or non-empty). You cannot decrease the column width when the table is non-empty; a table must be empty to decrease the column width.

For supported data types and column size restrictions, see Appendix C, “Data types,” on page 323.

- Set tables to use optimistic or pessimistic concurrency control.

Disk-based tables (D-tables) can be set to optimistic or pessimistic with the statement ALTER TABLE *base_table_name* SET {OPTIMISTIC | PESSIMISTIC}. By default, all D-tables are optimistic. M-tables are always pessimistic. A database-wide default for D-tables can be set with the parameter

General.Pessimistic.

- Change the schema and owner for the table.

The owner of a table can be changed using the ALTER TABLE *base_table_name* MODIFY SCHEMA *schema_name* statement. This statement gives all rights, including creator rights, to the new owner. The old owner's access rights to the table, excluding the creator rights, are preserved.

- Change whether the table is stored in memory or on disk.

A table may be changed from disk-based to in-memory based or vice-versa. This may be done only if the table is empty. If you try to change a table to the same

storage mode that it already uses (for example, if you try to change an in-memory table to use in-memory storage), then the command has no effect, and no error message is issued.

- Change the referential constraint settings.
For more information, see 4.5, “Referential integrity,” on page 100.
- Define synchronization setting for tables.
For more information, see A.4.1, “ALTER TABLE ... SET HISTORY COLUMNS” and A.4.2, “ALTER TABLE ... SET SYNCHISTORY,” on page 160.
- Work with cache segments:
 - Change cache segment size (SET POOLSIZE *percentage*).

Note: You must restart the server to make the cache segment size change effective immediately. If you do not restart the server, the cache segment size change is gradual.

- Assign a table (ADD CACHE SEGMENT) or a part of a table (ADD CACHE SEGMENT ... WHERE) to a cache segment.

When defining a data range, the OLDERTHAN and NEWERTHAN keywords specify date and time ranges; the data type used must be of type DATE or TIMESTAMP.

If the value cannot be converted to the column type, the operation fails.

- Drop a table from a cache segment (DROP CACHE SEGMENT).

Examples

```
ALTER TABLE table1 ADD x INTEGER;
ALTER TABLE table1 RENAME COLUMN old_name new_name;
ALTER TABLE table1 MODIFY COLUMN xyz SMALLINT;
ALTER TABLE table1 DROP COLUMN xyz;
ALTER TABLE table1 SET STORE MEMORY;
ALTER TABLE table1 SET PESSIMISTIC;
ALTER TABLE table2 ADD COLUMN col_new CHAR(8) DEFAULT 'VACANT' NOT NULL;
ALTER TABLE table2 ALTER COLUMN col_new SET DEFAULT 'EMPTY';
ALTER TABLE table2 ALTER COLUMN col_new DROP DEFAULT;
ALTER TABLE dept_tab1 ADD CONSTRAINT div_check CHECK(division_id < 12);
ALTER TABLE dept_tab1 DROP CONSTRAINT div_check;

ALTER TABLE HISTORY_DATA ADD CACHE SEGMENT HISTORY_CACHE
ALTER TABLE HISTORY_DATA ADD CACHE SEGMENT HISTORY_CACHE WHERE DT OLDERTHAN 7 DAYS
```

Related reference:

A.23, “CREATE TABLE,” on page 207

A.9, “CREATE CACHE SEGMENT,” on page 169

A.4.1, “ALTER TABLE ... SET HISTORY COLUMNS”

A.4.2, “ALTER TABLE ... SET SYNCHISTORY,” on page 160

A.4.1 ALTER TABLE ... SET HISTORY COLUMNS

```
ALTER TABLE table_name SET HISTORY COLUMNS ( col1, col2, colN ...)
```

Usage

To further optimize the synchronization history process, after you set tables for synchronization history, you can use the SET HISTORY COLUMNS statement to specify which column updates in the master and its corresponding synchronized table cause entries to the history table. If you do not use this statement to specify particular columns, then all update operations (on all columns) in the master database cause a new entry to the history table when the corresponding synchronized table is updated. Generally, we recommend using ALTER TABLE ...

SET HISTORY COLUMNS for columns that are used for search criteria or for joining.

Usage in master

Use SET SYNCHISTORY and SET HISTORY COLUMNS in the master to enable incremental publications on a table.

Usage in replica

Use SET SYNCHISTORY and SET HISTORY COLUMNS in the replica to enable incremental REFRESH on a table.

Note: In order for ALTER TABLE ... SET HISTORY COLUMNS to succeed, the statement ALTER TABLE ... SET SYNCHISTORY has to be executed first. Executing ALTER TABLE ... SET NOSYNCHISTORY removes also the effect of ALTER TABLE ... SET HISTORY COLUMNS.

Examples

```
ALTER TABLE myLargeTable SET HISTORY COLUMNS (accountid);
```

Return values

Table 26. ALTER TABLE SET HISTORY COLUMNS return values

Error code	Description
13047	No privilege for operation
13100	Illegal table mode combination
13134	Table is not a base table
25038	Table is referenced in publication <i>publication_name</i> ; drop or alter operations are not allowed
25039	Table is referenced in subscription to publication <i>publication_name</i> ; drop or alter operations are not allowed.

Related reference:

A.4.2, "ALTER TABLE ... SET SYNCHISTORY"

A.4.2 ALTER TABLE ... SET SYNCHISTORY

```
ALTER TABLE table_name SET {SYNCHISTORY | NOSYNCHISTORY}
```

Usage

The SET SYNCHISTORY / NOSYNCHISTORY clause tells the server to use the incremental publications mechanism of solidDB advanced replication architecture for the specified table. By default, SYNCHISTORY is not on. When this statement is set to SYNCHISTORY for a specified table, a shadow table is automatically created to store old versions of updated or deleted rows of the main table. The shadow table is called a *synchronization history table* or simply a *history table*.

The data in a history table is referred to when a replica gets an incremental REFRESH from a publication in the master. For example, if you delete the record

with Ms. Smith's telephone bill from the main table, a copy of the record is stored in the synchronization history table. When the replica refreshes, the master checks the history table and tells the replica that Ms. Smith's record was deleted. The replica can then delete that record also. If the percentage of records that were deleted or changed is fairly small, an incremental update is faster than downloading the entire table from the master. If you perform a full REFRESH, rather than an incremental REFRESH, the history table is not used. The data in the table on the master is simply copied to the replica.

Versioned data is automatically deleted from the database when there are no longer any replicas that need the data to fulfill REFRESH requests.

You must use this command to turn on synchronization history before a table can participate in master/replica synchronization. You can use this command on a table even if data currently exists in that table; however ALTER TABLE SET SYNCHISTORY can only be used if the specified table is not referenced by an existing publication.

SET SYNCHISTORY must be specified in the tables of both master and replica databases.

To check if SYNCHISTORY is on for a table, check the SYS_TABLEMODES system table. The MODE column contains the SYNCHISTORY information.

For example, use the following query:

```
SELECT mode
FROM SYS_TABLES, SYS_TABLEMODES
WHERE table_name = 'MY_TABLE' AND SYS_TABLEMODES.ID = SYS_TABLES.ID;
MODE
----
SYNCHISTORY
1 rows fetched.
```

SYS_TABLEMODES only shows the mode of tables for which the mode was explicitly set. SYS_TABLEMODES does not show the mode of tables that are in the the default mode. If SYNCHISTORY (or NOSYNCHISTORY) is not set for the table, the query returns an empty result set.

Usage in master

Use SET SYNCHISTORY in the master to enable incremental publications on a table.

Usage in replica

Use SET SYNCHISTORY in the replica to enable incremental REFRESHES on a table.

Note: If the Replica is read only (no changes are done to the replicated parts of the publication), the statement ALTER TABLE ... SET SYNCHISTORY is not needed. However, at the same time, the SYS_SYNC_KEEPLOCALCHANGES parameter must be set to yes with the following command:

```
set sync parameter SYS_SYNC_KEEPLOCALCHANGES 'Yes';
```

Examples

```
ALTER TABLE myLargeTable SET SYNCHISTORY;
ALTER TABLE myVerySmallTable SET NOSYNCHISTORY;
```

Return values

Table 27. ALTER TABLE SET SYNCHISTORY return values

Error code	Description
13047	No privilege for operation
13100	Illegal table mode combination
13134	Table is not a base table
25038	Table is referenced in publication <i>publication_name</i> ; drop or alter operations are not allowed
25039	Table is referenced in subscription to publication <i>publication_name</i> ; drop or alter operations are not allowed.

Related reference:

A.4.1, “ALTER TABLE ... SET HISTORY COLUMNS,” on page 159

A.5 ALTER TRIGGER

```
ALTER TRIGGER trigger_name_attr SET {ENABLED | DISABLED}
trigger_name_attr ::= [ catalog_name. [ schema_name. ] ] trigger_name
```

Usage

The ALTER TRIGGER statement modifies trigger attributes. The valid attributes are ENABLED and DISABLED trigger.

The ALTER TRIGGER DISABLED statement causes solidDB to ignore the trigger when an activating DML statement is issued. With this command, you can also enable a trigger that is currently inactive or disable a trigger that is currently active.

To alter a trigger on a table, you must be the owner of the table, or a user with DBA authority.

Examples

```
ALTER TRIGGER trig_on_employee SET ENABLED;
```

A.6 ALTER USER

```
ALTER USER username IDENTIFIED {BY new_password | EXTERNALLY}
[CONNECTIONS integer]
```

Usage

The ALTER USER statement modifies the password of a user or changes the authentication type of a user.

BY *new_password*

The ALTER USER *username* IDENTIFIED BY *new_password* statement modifies the password of a specified user.

IDENTIFIED EXTERNALLY

The ALTER USER *username* IDENTIFIED EXTERNALLY changes the authentication of the user to use external authentication. When using external authentication, the *username* must match the user name of an operating system user account.

Additionally, to use external authentication, you must enable the use of IBM Global Security Kit (GSKit). For more details, see Enabling encryption with IBM Global Security Kit (GSKit).

Note: When using external authentication on Windows systems, your environment might require you to provide the user name in the two-part format of *domain_name\username* or *username@domain_name*. You can specify the *domain_name* using the **General.DefaultDomainName** parameter. The solidDB server uses the value of the **General.DefaultDomainName** parameter to resolve the two-part user ID as expected by the Windows operating system.

CONNECTIONS *integer*

The CONNECTIONS keyword defines the maximum number of connections for a user.

integer is a positive integer that defines the number of connections allowed for the specified user. The default value is 0, which means that there is no limit on concurrent connections for the specified user.

Once the maximum number of connections have been made, new connections fail with the following error message:

Server Error 14559: Exceeded maximum number of allowed connections for user. Close existing connections or contact administrator to increase the limit.

Only users with administrator rights can change the password or authentication type of other users.

Examples

```
ALTER USER MANAGER IDENTIFIED BY 02CPTG;  
ALTER USER solid1 IDENTIFIED EXTERNALLY;  
ALTER USER solid1 IDENTIFIED EXTERNALLY CONNECTIONS 7;
```

Related tasks:

4.2.5, “Changing a password,” on page 91

You can change your password with the ALTER USER command. Administrators can also change passwords of other users.

4.2.13, “Changing administrator's user name and password,” on page 93

The database system administrator's password can be changed with the ALTER USER command. The administrator's user name cannot be changed with the ALTER USER command; the user name can be changed by granting the SYS_ADMIN_ROLE to another user and then dropping the original administrator user account.

A.6.1 ALTER USER (replica)

```
ALTER USER replica_user SET MASTER master_name USER user_specification
```

where:

```
user_specification ::= { master_user IDENTIFIED BY master_password | NONE }
```

```
ALTER USER username SET {PUBLIC | PRIVATE}
```

Usage

The ALTER USER *replica_user* ... statement is used to map replica user ids to specified master user ids. It is always a replica user id that is mapped to a master user id.

The mapping of user ids is used for implementing security in a multi-master or multi-tier synchronization environment. In such environments, it is difficult to maintain the same username and passwords in separate, geographically dispersed databases.

Only a user with DBA authority or SYS_SYNC_ADMIN_ROLE can map users. To implement mapping, an administrator must know the master user name and password. If NONE is specified, the mapping is removed.

All replica databases are responsible for subscribing to the SYNC_CONFIG system publication to update user information. Public master user names and passwords are downloaded during this process to a replica database using the MESSAGE APPEND SYNC_CONFIG command. Through mapping of the replica user id with the master user id, the system determines the currently active master user, based on the local user id that is logged to the replica database. If during SYNC_CONFIG loading the system does not detect mapping, it determines the currently active master user through the matching user id and password in the master and the replica.

For more details on using mapping for security, see *Implementing security through access rights and roles* in the *IBM solidDB Advanced Replication User Guide*.

You can also limit which master user accounts are downloaded to the replica during SYNC_CONFIG loading. This is done by altering the status of users as private or public with the following command:

```
ALTER USER username SET PRIVATE | PUBLIC
```

The default is PUBLIC. If the PRIVATE option is set for a user, that user's information is not included in a SYNC_CONFIG subscription, even if they are specified in a SYNC_CONFIG request. Only a user with DBA authority or SYS_SYNC_ADMIN_ROLE can alter a user's status.

This allows administrators to ensure that no user ids with administration rights are sent to a replica. For example, for security reasons, administrators might want to ensure that DBA passwords are never public.

Usage in master

In a master database, you can set user ids to PUBLIC or PRIVATE.

Usage in replica

In a replica database, you can map a replica user id to a master user id.

Examples

The following statement maps a replica user id *smith_1* to a master user id *dba* with the password of the user DBA.

```
ALTER USER SMITH_1 SET MASTER MASTER_1 USER DBA IDENTIFIED BY DBA
```

The following statement sets the status of the user *admin1* of the master database to PRIVATE, defining that the user account of *admin1* should not be downloaded to any replica.

```
ALTER USER admin1 SET PRIVATE;
```

The following statement sets the status of the user *salesman* of the master database to PUBLIC, defining that the user account of *salesman* should be downloaded to all replicas.

```
ALTER USER salesman SET PUBLIC;
```

Return codes

Table 28. ALTER USER return values

Error code	Description
13047	No privilege for operation
13060	User name <i>xxx</i> not found
25020	Database is not a master database
25062	User <i>user_id</i> is not mapped to master <i>user_id</i>
25063	User <i>user_id</i> is already mapped to master <i>user_id</i>

A.7 CALL

```
CALL procedure_name [( argument [, argument ...])] [AT node-def]
```

where

```
argument ::= parameter_name=expression | expression
```

```
node-def ::= DEFAULT | <replica_name> | <master_name>
```

Usage

The CALL statement is used for calling stored procedures that have been created with the CREATE PROCEDURE statement. To call a stored procedure, the caller must have EXECUTE privilege on the procedure.

argument

parameter-name

Name of the parameter to which the argument is assigned.

When an argument is assigned to a parameter by name, all the arguments that follow it must also be assigned by name. A named argument can be specified only once.

expression

Each specification of expression is an argument of the CALL. The expression must be a literal or a parameter marker (?)

The *n*th unnamed argument of the CALL statement corresponds to the *n*th parameter defined in the CREATE PROCEDURE statement for the procedure. If the procedure definition specified default values for the parameters, you do not have to give values for all the parameters.

You can enter arguments in either of two ways:

Named parameters

Use a parameter name (in the form *parameter_name = expression*), even if the arguments are not in the same order as the parameters.

Important: If you specify a *parameter_name* for one argument, you must use parameter names for all the arguments that follow it.

Positional parameters

Use the parameter position only, with no *parameter_name*, specifying each expression is in the same order as the parameter in the procedure definition to which the argument corresponds.

You can omit arguments, that is, you call a procedure with fewer arguments than the procedure definition specified. The solidDB server initializes missing arguments to their corresponding default values.

The following rules apply to omitting parameters:

- You can omit only parameters for which a default value is specified. If the missing arguments have no default values, solidDB issues an error.
- Positional parameters can be omitted only in the order specified in the procedure definition, starting from the last parameter.

AT node_ref

The AT node_ref clause is supported only in advanced replication configurations that use remote stored procedures; it can be executed from a master node to one of its replica nodes or vice-versa.

DEFAULT

DEFAULT means that the *current replica context* is used. The current replica context is only defined when the procedure call is started in the background using the START AFTER COMMIT statement with the FOR EACH REPLICA option. If the default is not set, the error Default node not defined is returned. DEFAULT can be used inside stored procedures and in a statement started with START AFTER COMMIT.

replica_name

Defines the name of the replica node where the remote procedure is executed.

For example, to call the procedure remote_proc on the replica node called replica2, issue the following statement

```
CALL remote_proc AT replica2;
```

master_name

Defines the name of the master node where the remote procedure is executed.

A remote stored procedure cannot return a result set; it can only return an error code.

A single call statement can only call a single procedure on a single node. If you want to call more than one procedure on a single node, you must execute multiple CALL statements. If you want to execute the same procedure (with the same procedure name) on more than one node, you have to either use the START AFTER COMMIT FOR EACH REPLICA statement or execute multiple calls.

Transactions in remote procedure calls

A remote procedure call (whether or not it was started by a START AFTER COMMIT) is executed in a separate transaction from the transaction that it was called from. The caller cannot roll back or commit the remote procedure call. The procedure that is executing in the called node is responsible for issuing its own commit or rollback statement.

Also, remote procedure calls are not durable. If the server goes down right after issuing the remote procedure call, the call is lost. It will not be executed in recovery phase.

Return values from the remote procedure

When you call a remote stored procedure, you cannot get a complete result set returned. Instead, the server returns the return value of the stored procedure (a single value) or an error code.

Note: If the remote procedure is executed in the background (using START AFTER COMMIT), no return value is returned to the user. Even error codes are not returned.

A procedure call is executed synchronously; it returns after the call is executed. However, if the procedure call is executed using START AFTER COMMIT (for example, START AFTER COMMIT UNIQUE CALL FOO AT REPLICA1), the procedure call is executed asynchronously in the background. The asynchronous execution is a characteristic of the START AFTER COMMIT statement, not of the procedure call.

Example: Calling a stored procedure without default parameter values

The procedure participants has been created in the database as follows:

```
"CREATE PROCEDURE participants( adults integer ,
children integer ,
pets integer )
BEGIN
END"
```

The following procedure calls return identical results:

```
CALL participants(adults = 5, children = 2, pets = 3)
CALL participants(5, 2, 3)
```

Note: The following syntax fails because the parameter name has been defined for only one argument (adults). If a name is given to a parameter, all parameters following it must also have a name.

```
CALL participants(adults = 5, 2, 3)
```

The following syntax succeeds because the positional parameters

Example: Calling stored procedures with default parameter values

The procedure paying_guests has been created in the database as follows:

```
"CREATE PROCEDURE paying_guests( adults integer = 1,
children integer = '0',
pets integer = '0')
BEGIN
END"
```

The following procedure calls use value 2 for parameter adults and default values for parameters children and pets:

```
CALL paying_guests(adults = 2)
CALL paying_guests(2)
```

The following procedure call uses value 3 for parameter pets and default values for parameters adults and children:

```
CALL paying_guests(pets = 3)
```

Example: Procedure call with dynamic parameters

The following procedure call uses parameter markers (?):

```
CALL RemoteProc(?,?) AT MyReplica1;
```

Example: START AFTER COMMIT FOR EACH REPLICA statement

```
START AFTER COMMIT FOR EACH REPLICA WHERE NAME LIKE 'REPLICA%'
UNIQUE CALL MYPROC AT DEFAULT.
```

Related reference:

A.15, "CREATE PROCEDURE," on page 180

A.52, "GRANT," on page 234

Related information:

"Using parameters" on page 24

"Procedure privileges" on page 46

A.7.1, "Access rights for remote stored procedure calls"

A.7.1 Access rights for remote stored procedure calls

When a stored procedure is called remotely, you must take into account the access rights — i.e. does the caller have the right to execute this procedure on the remote server?

CASE 1: If the Sync user is set with the command SET SYNC USER.

The caller sends the user name and password of the "sync user" to the remote server, and the remote server tries to execute the procedure using that user name and password. In this case, the username and password must exist in the remote server (i.e. the server that the stored procedure will be executed on) and the user must have appropriate access rights to the database and the called procedure.

CASE 2: If the Sync user is not set:

The caller sends the following information to the remote server when calling a remote procedure:

If the caller is the master and the remote server is the replica (M → R):

- Name of the master (SYS_SYNC_REPLICAS.MASTER_NAME).
- Replica id (SYS_SYNC_REPLICAS.ID).
- User name of the caller.
- User id of the caller.

If the caller is the replica and the remote procedure is the master (R → M):

- Name of the master (SYS_SYNC_MASTERS.NAME).
- Replica id (SYS_SYNC_MASTERS.REPLICA_ID).

- Master user id (The same user id is used as when a replica refreshes data. There has to be a mapping from the local replica user to a master user in SYS_SYNC_USERS table.)

The following actions are performed in the called node:

If the remote node is a replica (M → R):

- Get the master id from table SYS_SYNC_MASTERS according to the master name received from the caller (master itself doesn't know it's id in the replica). From the table SYS_SYNC_USERMAPS get the replica user ids according to master user name and master id. Select the first user that has access rights to the procedure.
- If there are no matching rows in SYS_SYNC_USERMAPS, then get NAME and PASSWD from the table SYS_SYNC_USERS according to master id and master user name received from the caller and try to execute the procedure using them.

If the remote node is a master (R → M)

- Try to execute the procedure using the user id received from the replica.

If the replica allows calls from any master it should define its own connect string information in the `solid.ini` file, for example:

```
[Synchronizer]
ConnectStrForMaster=tcp replicahost 1316
```

The replica sends that connect string automatically to the master when it forwards any message to the master. When the master receives the connect string from the replica, it replaces any previous value (if it differs).

The master can set the connect string to the replica (if the replica has not done any messaging and the master needs to call it and knows that the connect string has changed) using the following statement:

```
SET SYNC CONNECT <connect-info> TO REPLICA <replica-name>
```

A.8 COMMIT WORK

```
COMMIT [WORK]
```

Usage

The COMMIT or COMMIT WORK statement makes changes in the database permanent, terminating the transaction. To discard the changes, use the ROLLBACK command.

If you do not commit a transaction explicitly and the program does not make the commit for you automatically, the transaction will be rolled back. For example, the solidDB SQL Editor (**solsql**) does not commit transactions by default.

Related reference:

A.67, "ROLLBACK WORK," on page 275

A.9 CREATE CACHE SEGMENT

```
CREATE CACHE SEGMENT name POOLSIZE percentage
```

Usage

The CREATE CACHE SEGMENT statement divides the database buffer pool (cache) into segments. The cache segment size is defined as a percentage of the total cache size. The total size for the assigned segments cannot be more than 80 % of the total cache size. The total cache size is defined with the **IndexFile.CacheSize** parameter.

After you have created cache segments with the CREATE CACHE SEGMENT statement, assign tables or data ranges to the segment with the ALTER TABLE ... ADD CACHE SEGMENT statement.

Parameters

name Identifies the cache segment name. The name must be a valid SQL identifier.

percentage Defines the size of the cache segment as a percentage of the total cache size. The segment size must be 1 - 80. The total size of all cache segments cannot exceed 80.

To check the names and sizes of cache segments, query the system table SYS_CACHESEGMENTS.

Example

To create a cache segment that is called HISTORY_CACHE that uses 20 % percentage of the database cache, use the following statement:

```
CREATE CACHE SEGMENT HISTORY_CACHE POOLSIZE 20
```

A.10 CREATE CATALOG

```
CREATE CATALOG catalog_name
```

Usage

The CREATE CATALOG creates a new catalog in the solidDB database. The use of catalogs in solidDB is an extension to the SQL standard.

Catalogs allow you to logically partition databases so you can organize your data to meet the needs of your business or application.

A solidDB database file may contain more than one *logical database*. Each logical database is a complete, independent group of database objects, such as tables, indexes, triggers, stored procedures, and so on. Each logical database is implemented as a database catalog. Thus, solidDB can have one or more catalogs.

When creating a new database or converting an old database to a new format, you are prompted for a default catalog name. This default catalog name allows for backward compatibility of solidDB databases prior to version 3.x.

A catalog can have zero or more schema names. The default schema name is the user ID of the user who creates the catalog.

A schema can have zero or more database object names. A database object can be qualified by a schema or user ID.

The catalog name is used to qualify a database object name.

Important: The catalog name must not contain spaces.

Database object names can be qualified in all DML statements as:

```
catalog_name.schema_name.database_object
```

or

```
catalog_name.user_id.database_object
```

If you use the catalog name, you must also use the schema name. The converse is not true; you may use the schema name without using the catalog name (if you have already done an appropriate SET CATALOG statement to specify the default catalog).

```
catalog_name.database_object -- Illegal  
schema_name.database_object -- Legal
```

Only a user with DBA authority (SYS_ADMIN_ROLE) can create a catalog for a database.

Creating a catalog does not automatically make that catalog the current default catalog. If you have created a new catalog and want your subsequent commands to execute within that catalog, you must also execute the A.71.3, "SET CATALOG," on page 282 statement. For example:

```
CREATE CATALOG MyCatalog;  
CREATE SCHEMA smith; -- not in MyCatalog  
SET CATALOG MyCatalog;  
CREATE SCHEMA jones; -- in MyCatalog
```

To use schemas, a schema name must be created with the A.20, "CREATE SCHEMA," on page 203 statement before creating the database object name. However, a database object name can be created without a schema name. In such cases, database objects are qualified using user_id only.

A catalog context can be set in a program using the A.71.3, "SET CATALOG," on page 282 statement.

A catalog can be dropped from a database using the A.30, "DROP CATALOG," on page 219. When dropping a catalog name, all objects associated with the catalog name must be dropped prior to dropping the catalog.

Rules for resolving catalog names

- A fully qualified name (*catalog_name.schema_name.database_object_name*) does not need any name resolution, but will be validated.
- If a catalog context is not set using SET CATALOG, all database object names are resolved always using the default catalog name as the catalog name. The database object name is resolved using schema name resolution rules. For more information about these rules, see A.20, "CREATE SCHEMA," on page 203.
- If a catalog context is set and the catalog name cannot be resolved using the *catalog_name* in the context, *database_object_name* resolution fails.
- To access a database system catalog, users do not need to know the system catalog name. Users can specify ""._SYSTEM.table". solidDB translates the empty string " that is used as a catalog name to the default catalog name. solidDB also provides automatic resolution of _SYSTEM schema to the system catalog, even when the catalog name is not provided.

Examples

```
CREATE CATALOG C;
SET CATALOG C;
CREATE SCHEMA S;
SET SCHEMA S;
CREATE TABLE T (i INTEGER);
SELECT * FROM T;
-- the name T is resolved to C.S.T

-- Assume the userid is SMITH
CREATE CATALOG C;
SET CATALOG C;
CREATE TABLE T (i INTEGER);
SELECT * FROM T;
--The name T is resolved to C.SMITH.T

-- Assume there is no Catalog context set.
-- Meaning the default catalog name is BASE or the setting
-- of the base catalog.
CREATE SCHEMA S;
SET SCHEMA S;
CREATE TABLE T (i INTEGER);
SELECT * FROM T;
--The name T is resolved to <BASE>.S.T

CREATE CATALOG C1;
SET CATALOG C1;
CREATE SCHEMA S1;
SET SCHEMA S1;
CREATE TABLE T1 (c1 INTEGER);

CREATE CATALOG C2;
SET CATALOG C2;
CREATE SCHEMA S2;
SET SCHEMA S2;
CREATE TABLE T1 (c2 INTEGER)

SET CATALOG BASE;
SET SCHEMA USER;
SELECT * FROM T1;
-- This select will give an error as it
-- cannot resolve the T1.
```

A.11 CREATE EVENT

```
CREATE EVENT event_name
  [(parameter_name datatype
    [parameter_name data_type...])]
```

where

event_name is any user-defined alphanumeric string that identifies the event

parameter_name is a user-defined name of the parameter

data_type is an SQL data type

Usage

The CREATE EVENT statement creates user-defined event objects in the solidDB databases.

After you have created a user-defined event, you must register, post, and wait on the event using stored procedures.

Example

```
CREATE EVENT ALERT1(I INTEGER, C CHAR(4));
```

Related reference:

A.2, “ADMIN EVENT,” on page 156

A.31, “DROP EVENT,” on page 220

A.61, “POST EVENT,” on page 270

A.63, “REGISTER EVENT,” on page 273

A.79.2, “data_type,” on page 305

Related information:

3.5, “Events,” on page 80

A.12 CREATE FUNCTION

```
CREATE FUNCTION function_name [(parameter_definition
    [, parameter_definition ...])]
    RETURN[S] data_type
    [SQL_data_access_indication]
    [LANGUAGE SQL]
    BEGIN
        function_body
    END
```

```
parameter_definition ::= parameter_name
                        data_type [=literal]
```

```
SQL_data_access_indication ::= CONTAINS SQL | READS SQL DATA | NO SQL
```

Usage

The CREATE FUNCTION statement creates user-defined *stored functions*. Stored functions are written in the solidDB proprietary SQL procedure language.

You can use the stored SQL functions to expand the server functionality by writing functions that behave in the same way as built-in functions.

Stored functions accept literals, variables, variable markers (“?”) or column names as input parameters. Output parameters are not supported.

Parameters

function_name

Identifies the function name. The function name can also include the catalog name and schema:

```
[catalog_name [.schema].]function_name
```

The name must be a valid SQL identifier. The function name must be unique within the schema. Although there is no prohibition against it, the user-defined function should not be given the same name as a built-in function.

parameter_definition ::= *parameter_name data_type* [=literal]

Identifies the number of input parameters of the function, and specifies the name and data type, and optional default value of each parameter. One entry in the list must be specified for each parameter that the function expects to receive. Up to 90 parameters can be specified.

Example:

```
CREATE FUNCTION week_number (s TIMESTAMP)
...
```

You can create a function that has no parameters. For example:

```
CREATE FUNCTION howdy() RETURNS VARCHAR20
...
```

data_type must be a valid solidDB A.79.2, “data_type,” on page 305.
literal defines the default value of the parameter. The value must be a literal.

Stored functions accept literals, variables, variable markers (“?”) or column names as input parameters. If a column name is used, the column value is retrieved from the current row variable instance (for example in the WHERE clause). Output parameters are not supported.

RETURN[S] data_type

Identifies the data type of for the output of the function.

data_type must be a valid solidDB A.79.2, “data_type,” on page 305.

SQL_data_access_indication

Defines whether the function issues any SQL statements and, if so, what type.

CONTAINS SQL

Indicates that SQL statements that do not read or modify SQL data can be executed by the function.

READS SQL DATA

Indicates that only read-only SQL statements can be executed by the function.

NO SQL

Indicates that the function cannot execute any SQL statements.

Note: The validity of the contents of the function are not checked against the declared data access indication.

LANGUAGE SQL

This optional clause indicates that the CREATE FUNCTION statement is being used to register a new function written in solidDB SQL procedure language. See A.13, “CREATE FUNCTION (external),” on page 175 for details about how to create external functions written in C programming language.

function_body

The function body must be a valid solidDB SQL procedure statement. See A.15, “CREATE PROCEDURE,” on page 180 for more details.

For valid expressions, see A.79.3, “expression,” on page 306.

Return values

- SQL_ERROR
- SQL_SUCCESS

Restrictions

- The returned value cannot be a complex expression; it can only be a variable or a literal.
- COMMIT is not allowed within a function.
- RETURNS NULL ON NULL INPUT is not supported. If any actual parameter is NULL, a NULL function value is forced automatically.

Example 1: Creating and using a simple function hello

Creating the function:

```
CREATE FUNCTION hello (s VARCHAR(20))
RETURNS VARCHAR(50)
BEGIN
    RETURN 'Hello';
END
```

Using the function:

```
SELECT hello('world');

HELLO('world')
-----
Hello
```

Example 2: Creating and using a function to return week numbers

Creating the function:

```
CREATE FUNCTION week_number (s TIMESTAMP)
RETURNS INTEGER
BEGIN
    DECLARE week INTEGER;
    week := floor(( dayofyear(s) - dayofweek(s) + 7)/7);
    RETURN week;
END
```

Using the function:

```
SELECT quota, week_number(sales_date)AS week, week_number(now())
AS current_week
FROM sales ORDER BY quota DESC LIMIT 1;

QUOTA WEEK    CURRENT_WEEK
-----
467    23          43

SELECT quota FROM sales
WHERE current_week(sales_date) = 23 ORDER BY quota DESC LIMIT 1;

QUOTA
-----
467

INSERT INTO report VALUES (?, week_number(now()), week_number(?));
```

Related reference:

A.13, “CREATE FUNCTION (external)”

A.32, “DROP FUNCTION,” on page 220

A.52, “GRANT,” on page 234

A.13 CREATE FUNCTION (external)

```
CREATE FUNCTION function_name [(parameter_definition
    [, parameter_definition ...])]
RETURN[S] data_type
[SQL_data_access_indication]
LANGUAGE C
EXTERNAL NAME 'external_function_library_identifier'
```

parameter_definition::= *parameter_name data_type*[=*literal*]

SQL_data_access_indication::= CONTAINS SQL | READS SQL DATA | NO SQL

Usage

The CREATE FUNCTION (external) statement registers user-defined *external stored functions*. The solidDB external stored functions can be written in C programming language.

You can use the external functions to expand the server functionality by writing a C program.

The external functions are used in the same way as built-in or user-defined functions. solidDB accesses the external functions by using the standard dynamic library interface (DLLs in Windows, shared libraries in Linux and UNIX). The control of the thread is passed to the external library routine until it returns. Furthermore, the database interface of the external functions is similar to that of an ODBC application, except that the connection handle (hdbc) is passed to the external routine, and it can immediately start using it without establishing the connection.

To enable use of external stored functions, you must link your application dynamically to solidDB using shared memory access (SMA) or linked library access (LLA).

Stored functions accept literals, variables, variable markers ("?",) or column names as input parameters. Output parameters are not supported.

The external functions are scalar functions, which return a single value each time it is called.

Parameters

function_name

Identifies the function name. The function name can also include the catalog name and schema:

[catalog_name[.schema].]function_name

The name must be a valid SQL identifier. The function name must be unique within the schema. Although there is no prohibition against it, the user-defined function should not be given the same name as a built-in function.

parameter_definition:= *parameter_name data_type* [= *literal*]

Identifies the number of input parameters of the function, and specifies the name and data type, and optional default value of each parameter. One entry in the list must be specified for each parameter that the function expects to receive. Up to 90 parameters can be specified.

Example:

```
CREATE FUNCTION week_number (s TIMESTAMP)
...
```

You can create a function that has no parameters. For example:

```
CREATE FUNCTION howdy() RETURNS VARCHAR20
...
```

data_type must be a valid solidDB A.79.2, "data_type," on page 305.

literal defines the default value of the parameter. The value must be a literal.

Stored functions accept literals, variables, variable markers ("?",) or column names as input parameters. If a column name is used, the column value is

retrieved from the current row variable instance (for example in the WHERE clause). Output parameters are not supported.

RETURN[S] *data_type*

Identifies the data type of for the output of the function.

data_type must be a valid solidDB A.79.2, "data_type," on page 305.

SQL_data_access_indication

Defines whether the function issues any SQL statements and, if so, what type.

CONTAINS SQL

Indicates that SQL statements that do not read or modify SQL data can be executed by the function.

READS SQL DATA

Indicates that only read-only SQL statements can be executed by the function.

NO SQL

Indicates that the function cannot execute any SQL statements.

Note: The validity of the contents of the function are not checked against the declared data access indication.

LANGUAGE C

Indicates that the CREATE FUNCTION statement is being used to register a new function based on code written in C programming language.

EXTERNAL NAME '*external_function_library_identifier*'

Identifies the name of the user-written code which implements the function being defined. The name must be enclosed by single quotation marks. Extraneous blanks are not permitted.

external_function_library_identifier::=[*path*]*library_name*!*C_routine_name*

[*path*]*library_name*

Identifies the absolute path (optional) and name of the function. If you specify an absolute path, you need to append the file extension of the library file. If you do not specify the file extension, you might need to add the location of the library to the LD_LIBRARY_PATH or LIBPATH (Linux and UNIX) or PATH (Windows) environment variable.

For example:

- On UNIX systems, '/u/jchui/mylib/myfunc.so' would cause solidDB to look in /u/jchui/mylib for the myfunc.so library.
- On Windows systems, 'd:\mylib\myfunc.dll' would cause solidDB to load the file myfunc.dll from the d:\mylib directory.

Tip: If you do not specify the file extension, you can port and use the library on several platforms.

!*C_routine_name*

Identifies the entry point name of the C routine (function) to be invoked. The exclamation point (!) serves as a delimiter between the library name and the routine name.

'!proc8' would direct solidDB to look for the library in the location specified by [*path*]*library_name*, and to use entry point *proc8* within that library.

Note: On Windows 32-bit systems, use the calling convention `_cdecl` for calling the C routines. This is to allow checking and handling of varying argument counts. Also, ensure that your DLL exports the routines you want to call: for example, include the routines in the corresponding `.def` file in your DLL project.

Note: The library (and the function within the library) does not need to exist when the `CREATE FUNCTION` statement is issued. However, when the function is used, the library and function within the library must exist and be accessible.

Return values

- `SQL_ERROR`
- `SQL_SUCCESS`

Restrictions

- The call interface of the external functions is assumed to be non-fenced and thread-safe.
- All functions are assumed non-deterministic. The keyword `[NOT] DETERMINISTIC` is allowed in the syntax but the setting `DETERMINISTIC` is not effective.
- `COMMIT [WORK]` is not allowed within a function.
- Only one parameter passing convention is available; the return value is passed as an implicit output parameter that comes after the input parameters to the called C function.
- You should not start new threads that execute in parallel with the called function; the provided connection handle is not multithread safe in the same manner as normal ODBC connection handle is. If new threads are launched, they must create new connections. If the new thread tries to share the same connection as the function itself, the `SQLException: Invalid Handle` error is returned.

Example

Tip: For more examples, see the `samples/procedures` directory in your `solidDB` installation directory.

The following example registers the external function `EXTFUNCCUBEVOLUME`, which calculates the volume of a cube using the length of the edge of the cube as an input parameter:

```
CREATE FUNCTION EXTFUNCCUBEVOLUME (edge float)
RETURN FLOAT
LANGUAGE C
EXTERNAL NAME 'example!lib!extfunccubevolume'
```

C program for the external function EXTFUNCCUBEVOLUME

```
#include <soliddbc3.h>

int extfunccubevolume(
    SQLHDBC hdbc, /* not used in this function */
    char sqlst[6],
    char qualName[],
    char diagMsg[71],
    SQLULEN nArgs,
    SQLLEN nullind[2],
    double* edge, /* input */
    double* retval, /* return value */
```

```

{
  if (nullind[0] == SQL_NULL_DATA) {
    /* NULL input implies NULL return value */
    nullind[1] = SQL_NULL_DATA;
    return (SQL_SUCCESS);
  }
  if (*edge < 0.0) {
    /* error: edge cannot be < 0 */
    strcpy(diagMsg, "extfunccubevolume: parameter 'edge' cannot be < 0");
    strcpy(sqlst, "22003"); /* Numeric values out of range */
    return (SQL_ERROR);
  }
  *retval = (*edge) * (*edge) * (*edge);
  nullind[1] = sizeof(double);
  return (SQL_SUCCESS);
}

```

Using the external function EXTFUNCCUBEVOLUME

```
SELECT extfunccubevolume(edge) FROM cubic_containers WHERE material = 'STEEL';
```

Related reference:

A.12, “CREATE FUNCTION,” on page 173

A.32, “DROP FUNCTION,” on page 220

A.52, “GRANT,” on page 234

A.14 CREATE INDEX

```

CREATE [UNIQUE] INDEX index_name
  ON base_table_name
  (column_identifier [ASC | DESC]
  [, column_identifier [ASC | DESC]] ...)

```

Usage

The CREATE INDEX statement creates an index for a table based on the given columns.

The keyword UNIQUE specifies that the column(s) being indexed must contain unique values. If more than one column is specified, the combination of columns must have a unique value, but the individual columns do not need to have unique values.

For example, if you create an index on the combination of LAST_NAME and FIRST_NAME, the following data values are acceptable because although there are duplicate first names and duplicate last names, no 2 rows have the same value for both first name and last name.

```

SMITH, PATTI
SMITH, DAVID
JONES, DAVID

```

Keywords ASC and DESC specify whether the given columns should be indexed in ascending or descending order. If neither ASC nor DESC is specified, ascending order is used.

If the *index_name* starts with a digit, the *index_name* must be given in double quotation marks.

Examples

```

CREATE UNIQUE INDEX UX_TEST ON TEST (I);
CREATE INDEX X_TEST ON TEST (I DESC, J DESC);
CREATE INDEX "053_TEST" ON TEST (I DESC, J DESC);

```

Related reference:

A.17, “CREATE [OR REPLACE] PUBLICATION,” on page 197

A.15 CREATE PROCEDURE

```
CREATE PROCEDURE procedure_name [(parameter_definition [, parameter_definition ...])]
  [RETURNS (output_column_definition [, output_column_definition ...])]
  [SQL_data_access_indication]
  [LANGUAGE SQL] BEGIN procedure_body END;

parameter_definition ::= [parameter_mode] parameter_name data_type
output_column_definition::= column_name column_type

SQL_data_access_indication::= CONTAINS SQL | READS SQL DATA | MODIFIES SQL DATA | NO SQL

procedure_body ::= [declare_statement; ...][procedure_statement; ...]

parameter_mode ::= IN | OUT | INOUT

declare_statement ::= DECLARE variable_name data_type

procedure_statement ::= prepare_statement | execute_statement |
  fetch_statement | control_statement | post_statement |
  wait_event_statement | wait_register_statement | exec_direct_statement |
  writetrace_statement | sql_dml_or_ddl_statement

prepare_statement ::= EXEC SQL PREPARE
  { cursor_name | CURSORNAME( { string_literal | variable } ) }
  sql_statement

execute_statement ::=
  EXEC SQL EXECUTE cursor_name
    [USING (variable [, variable ...])]
    [INTO (variable [, variable ...])] |
  EXEC SQL CLOSE cursor_name |
  EXEC SQL DROP cursor_name |
  EXEC SQL {COMMIT | ROLLBACK} WORK |
  EXEC SQL SET TRANSACTION {READ ONLY | READ WRITE} |
  EXEC SQL WHENEVER SQLERROR {ABORT | ROLLBACK [WORK], ABORT}
  EXEC SEQUENCE sequence_name.CURRENT INTO variable |
  EXEC SEQUENCE sequence_name.NEXT INTO variable |
  EXEC SEQUENCE sequence_name SET VALUE USING variable

fetch_statement ::= EXEC SQL FETCH cursor_name

cursor_name ::= literal

post_statement ::= POST EVENT event_name [(parameters)]

wait_event_statement ::=
  WAIT EVENT
  [event_specification ...]
  END WAIT

event_specification ::=
  WHEN event_name [(parameters)]
  BEGIN statements
  END EVENT

wait_register_statement ::=
  REGISTER EVENT event_name |
  UNREGISTER EVENT event_name

writetrace_statement ::= WRITETRACE(string)

control_statement ::=
  SET variable_name = value | variable_name ::= value |
  WHILE expression
  LOOP procedure_statement... END LOOP |
```

```

    LEAVE |
    IF expression THEN procedure_statement ...
      [ ELSEIF procedure_statement ... THEN] ...
      ELSE procedure_statement ... END IF |
    RETURN | RETURN SQLERROR OF cursor_name | RETURN ROW |
RETURN NOROW

exec_direct_statement ::=
EXEC SQL [USING (variable [, variable ...])]
[CURSORNAME(variable)]
EXECDIRECT sql_dml_or_ddl_statement |
EXEC SQL cursor_name
[USING (variable [, variable ...])]
[INTO (variable [, variable ...])]
[CURSORNAME(variable)]
EXECDIRECT sql_dml_or_ddl_statement

```

A.15.1 Usage

Stored procedures are simple programs, or procedures, that are executed in the server. The user can create a procedure that contains several SQL statements or a whole transaction and execute it with a single call statement. Usage of stored procedures reduces network traffic and allows more strict control to access rights and database operations.

Procedures are created with the statement `CREATE PROCEDURE name body` and dropped with `DROP PROCEDURE name`.

Procedures are called with the statement `CALL name [parameter ...]`.

The stored procedure syntax is a proprietary syntax modeled from SQL-99 specifications and dynamic SQL. Procedures contain control statements and SQL statements.

It is possible to create a stored procedure with an empty body.

All SQL DML and DDL statements can be used in procedures. Thus the procedure can, for example, create tables or commit a transaction. Each SQL statement in the procedure is atomic.

HotStandby configurations

In HotStandby configurations, all SQL stored procedures are executed in the Primary unless they are specified as read-only procedures by way of the SQL standard clause *SQL Data Access Indication* in the procedure declaration.

```

<SQL-data-access-indication> ::=
    NO SQL |
    READS SQL DATA |
    CONTAINS SQL |
    MODIFIES SQL DATA

```

To avoid unnecessary handovers of read-only procedures and functions, one of the following values can be declared:

- NO SQL
- READS SQL DATA
- CONTAINS SQL

Only MODIFIES SQL DATA (which is the default) inflicts transaction handover.

The <SQL-data-access-indication> clause comes between the (optional) RETURNS clause and the procedure body.

For example:

```
"CREATE PROCEDURE PHONEBOOK_SEARCH
(IN FIRST_NAME VARCHAR, LAST_NAME VARCHAR)
RETURNS (PHONE_NR NUMERIC, CITY VARCHAR)
READS SQL DATA
BEGIN
-- procedure_body
END";

---
```

Access rights

Procedures are owned by the creator of the procedure. Specified access rights can be granted to other users. When the procedure is run, it has the creator's access rights to database objects.

Commit in stored procedures

The "autocommit" functionality works differently for statements inside a stored procedure than for statements outside a stored procedure. For SQL statements outside a stored procedure, each individual statement is implicitly followed by a COMMIT WORK operation when autocommit is on. For a stored procedure, however, the implicit COMMIT WORK is executed after the stored procedure has returned to the caller. Note that this does not imply that a stored procedure is "atomic". As indicated above, a stored procedure may contain its own COMMIT and ROLLBACK commands. The implicit COMMIT WORK executed after the procedure returns will commit only that portion of the stored procedure statements that were executed since:

- the last COMMIT WORK inside the procedure
- the last ROLLBACK WORK inside the procedure
- the start of the procedure (if no COMMIT or ROLLBACK commands were executed during the procedure)

If one stored procedure is called from inside another, the implicit COMMIT WORK is done only after the end of the OUTERMOST procedure call. There is no implicit COMMIT WORK done after nested procedure calls.

For example, in the following script, the implicit COMMIT WORK is executed only after the **CALL outer_proc();** statement:

```
"CREATE PROCEDURE inner_proc
BEGIN
...
END";
CREATE PROCEDURE outer_proc
BEGIN
...
EXEC SQL PREPARE cursor1 CALL inner_proc();
EXEC SQL EXECUTE cursor1;
...
END";
CALL outer_proc();
```

A.15.2 parameter_modes

parameter_mode ::= IN | OUT | INOUT

Stored procedures provide for three different parameter modes: input parameters, output parameters, and input/output parameters.

- Input parameters are passed to the stored procedure from the calling program. The *parameter_mode* value is IN. This is the default behaviour.
- Output parameters are returned to the calling program from the stored procedure. The *parameter_mode* value is OUT.
- Input/output parameters pass values into the procedure and return a value back to the calling procedure. The *parameter_mode* is INOUT.

See the table below for a comparison of the parameter modes:

Table 29. Comparison of the parameter modes

Feature	IN	OUT	INOUT
Default/specified	Default.	Must be specified.	Must be specified.
Operation	Passes values to a subprogram.	Returns values to the caller.	Passes initial values to a subprogram; returns updated values to the caller.
Action	Formal parameter, acts like a constant.	Formal parameter, acts like an uninitialised variable.	Formal parameter, acts like an initialised variable.
Value assignation	Formal parameter, cannot be assigned a value.	Formal parameter, cannot be used in an expression; must be assigned a value.	Formal parameter, should be assigned a value.
Parameter type	Actual parameter, can be a constant, initialised variable, literal, or expression.	Actual parameter, must be a variable.	Actual parameter, must be a variable.

At programming interfaces, the output parameters are bound to variables as follows:

- In JDBC, with the method `CallableStatement.registerOutParameter()`.
- In ODBC, with the function `SQLBindParameter()`, where the third argument, `InputOutputType`, may be of type:
 - `SQL_PARAM_INPUT`
 - `SQL_PARAM_OUTPUT`
 - `SQL_PARAM_INPUT_OUTPUT`

A.15.3 prepare_statement

```
prepare_statement ::= EXEC SQL PREPARE
  { cursor_name | CURSORNAME( { string_literal | variable } ) }
  sql_statement
```

The SQL statements are first prepared with the EXEC SQL PREPARE statement

The *cursor_name* specification is a cursor name that must be given. It can be any unique cursor name inside the transaction. If the procedure is not a complete transaction, other open cursors outside the procedure may have conflicting cursor names.

A.15.4 execute_statement

```
execute_statement ::=  
EXEC SQL EXECUTE cursor_name  
  [USING (variable [, variable ...])] |  
  [INTO (variable [, variable ...])] |  
EXEC SQL CLOSE cursor_name |  
EXEC SQL DROP cursor_name |  
EXEC SQL {COMMIT | ROLLBACK} WORK |  
EXEC SQL SET TRANSACTION {READ ONLY | READ WRITE} |  
EXEC SQL WHENEVER SQLERROR {ABORT | ROLLBACK [WORK], ABORT}  
EXEC SEQUENCE sequence_name.CURRENT INTO variable |  
EXEC SEQUENCE sequence_name.NEXT INTO variable |  
EXEC SEQUENCE sequence_name SET VALUE USING variable
```

```
cursor_name ::= literal
```

Executing prepared SQL statements

The SQL *statement* is executed with the statement

```
EXEC SQL EXECUTE cursor [opt_using] [opt_into]
```

where the optional *opt-using* specification has the syntax:

```
USING (variable_list)
```

where *variable_list* contains a list of procedure variables or parameters separated by a comma. These variables are input parameters for the SQL statement. The SQL input parameters are marked with the standard question mark syntax in the prepare statement. If the SQL statement has no input parameters, the USING specification is ignored.

The optional *opt_into* specification has the syntax

```
INTO (variable_list)
```

where *variable_list* contains the variables that the column values of the SQL SELECT statement are stored into. The INTO specification is effective only for SQL SELECT statements.

After the execution of UPDATE, INSERT and DELETE statements an additional variable is available to check the result of the statement. Variable SQLROWCOUNT contains the number of rows affected by the last statement.

Using transactions

The following statement is used to terminate transactions:

```
EXEC SQL {COMMIT | ROLLBACK} WORK
```

The following statement is used to control the type of transactions:

```
EXEC SQL SET TRANSACTION {READ ONLY | READ WRITE}
```

Closing and dropping cursors

When you are finished using a cursor, you can either CLOSE the cursor or CLOSE and DROP the cursor.

If you are likely to reuse the cursor and want to improve performance, you should only CLOSE the cursor. When you close a cursor, all the memory allocated during the execute phase is released but the cursor is kept in prepared state.

Dropping a cursor frees all allocated resources. The next time you use the dropped cursor, it needs to be prepared.

Checking for errors

The result of each EXEC SQL statement executed inside a procedure body is stored into the variable SQLSUCCESS. This variable is automatically generated for every procedure. If the previous SQL statement was successful, a value one is stored into SQLSUCCESS. After a failed SQL statement, a value zero is stored into SQLSUCCESS.

```
EXEC SQL WHENEVER SQLERROR {ABORT | [ROLLBACK [WORK], ABORT]}
```

is used to decrease the need for IF NOT SQLSUCCESS THEN tests after every executed SQL statement in a procedure. When this statement is included in a stored procedure all return values of executed statements are checked for errors. If statement execution returns an error, the procedure is automatically aborted. Optionally the transaction can be rolled back.

The error string of latest failed SQL statement is stored into variable SQLERRSTR.

A.15.5 fetch_statement

```
fetch_statement ::= EXEC SQL FETCH cursor_name
```

The EXEC SQL FETCH statement fetches rows. If the fetch completed successfully, the column values are stored into the variables defined in the *opt_into* specification of the EXECUTE or EXECDIRECT statement.

A.15.6 post_statement

```
post_statement ::= POST EVENT event_name [( parameters) ]  
                [UNIQUE | DATA UNIQUE]
```

The POST EVENT statement is used for posting system events and user-defined events.

The keyword UNIQUE means that only last post is kept in event queue for each user and for each event. For example after POST EVENT EV(1) and POST EVENT EV(2) only EV(2) is in event queue if EV(1) is not processed before EV(2) is posted. Event EV(1) is discarded.

The keyword DATA UNIQUE means that also event parameters must be unique. So after calls POST EVENT EV(1), POST EVENT EV(2) and POST EVENT EV(2) events EV(1) and EV(2) are kept in event queue. First EV(2) is discarded.

The POST EVENT statement is allowed only inside stored procedures.

Related information:

3.5, "Events," on page 80

A.15.7 wait_register_statement

```
wait_register_statement ::=  
    REGISTER EVENT | UNREGISTER EVENT event_name
```

The REGISTER EVENT statement tells the server that you would like to be notified of all future occurrences of the specified event, even if you are not yet waiting for

it. By separating the REGISTER EVENT and WAIT EVENT commands, you can start queuing events immediately, while waiting until later to actually start processing them.

You do not need to register for every event before waiting for it. When you wait on an event, you will be registered implicitly for that event if you did not already explicitly register for it. Thus you only need to explicitly register events if you want them to start being queued now but you do not want to start waiting for them until later.

The UNREGISTER EVENT statement unregisters the specified event.

The REGISTER EVENT and UNREGISTER EVENT statement are allowed only inside stored procedures.

Related information:

3.5, “Events,” on page 80

A.15.8 wait_event_statement

```
wait_event_statement ::=
    WAIT EVENT
    [event_specification ...]
    END WAIT

event_specification ::=
    WHEN event_name [(parameters)]
    BEGIN statements
    END EVENT
```

The WAIT EVENT statement makes procedures to wait for an event to happen. You can use WAIT EVENT with both system-defined events and user-defined events.

If you want to stop the stored procedure waiting for an event, you can use ODBC function SQLCancel() called from a separate thread in the client application. This function cancels executing statements. Alternatively, you can create a specific user event and send it. The waiting stored procedure must be modified to wait for this additional event. The client application recognizes this event and exits the waiting loop.

Tip: With system events, you may also wait on an event without using a stored procedure by using the ADMIN EVENT statement. However, you cannot post events using ADMIN EVENT.

```
ADMIN EVENT 'register sys_event_hsbstateswitch';
ADMIN EVENT 'wait';
```

Related information:

3.5, “Events,” on page 80

A.15.9 control_statement

```
control_statement ::=
    SET variable_name = value | variable_name ::= value |
    WHILE expression
        LOOP procedure_statement... END LOOP |
    LEAVE |
    IF expression THEN procedure_statement ...
```

```

[ ELSEIF procedure_statement ... THEN] ...
ELSE procedure_statement ... END IF |
RETURN | RETURN SQLERROR OF cursor_name | RETURN ROW |
RETURN NOROW

```

The following control statements are available in the procedures:

Table 30. Control statements

Control Statement	Description
<code>set <i>variable</i> = <i>expression</i></code>	Assigns a value to a variable. The value can be either a literal value (for example, 10 or 'text') or another variable. Parameters are considered as normal variables.
<code><i>variable</i> ::= <i>expression</i></code>	Alternate syntax for assigning values to variables.
<pre> while <i>expr</i> loop <i>statement-list</i> end loop </pre>	Loops while expression is true.
<code>leave</code>	Leaves the innermost while loop and continues executing the procedure from the next statement after the keyword end loop.
<pre> if <i>expr</i> then <i>statement-list1</i> else <i>statement-list2</i> end if </pre>	Executes <i>statements-list1</i> if expression <i>expr</i> is true; otherwise, executes <i>statement-list2</i> .
<pre> if <i>expr1</i> then <i>statement-list1</i> elseif <i>expr2</i> then <i>statement-list2</i> end if </pre>	If <i>expr1</i> is true, executes <i>statement-list1</i> . If <i>expr2</i> is true, executes <i>statement-list2</i> . The statement can optionally contain multiple elseif statements and also an else statement.
<code>return</code>	Returns the current values of output parameters and exits the procedure. If a procedure has a <i>return row</i> statement, return behaves like <i>return norow</i> .
<code>return sqlerror of <i>cursor-name</i></code>	Returns the sqlerror associated with the cursor and exits the procedure.
<code>return row</code>	Returns the current values of output parameters and continues execution of the procedure. Return row does not exit the procedure and return control to the caller.
<code>return norow</code>	Returns the end of the set and exits the procedure.

Note: A procedure defined with the RETURNS clause returns always at least one row. This happens even if the row is NULL-filled.

To not return empty rows, use RETURN NOROW in the procedure body. It returns immediately without a result set (SQL_NO_DATA). The statement RETURN NOROW can be used after a sequence of RETURN ROW statements whereby the procedure returns the rows generated with RETURN ROW.

A.15.10 writetrace_statement

```
writetrace_statement ::= WRITETRACE(string)
```

The WRITETRACE statement function sends trace output (string) to the soltrace.out trace file. This can be useful when debugging problems in stored procedures.

The output will only be written if you turn tracing on.

For more information about tracing in stored procedures and how to turn on tracing, see 6.2, “Tracing facilities for stored procedures and triggers,” on page 129.

A.15.11 exec_direct_statement

```
exec_direct_statement ::=  
EXEC SQL [USING (variable [, variable ...])]  
[CURSORNAME(variable)]  
EXECDIRECT sql_dml_or_ddl_statement |  
EXEC SQL cursor_name  
[USING (variable [, variable ...])]  
[INTO (variable [, variable ...])]  
[CURSORNAME(variable)]  
EXECDIRECT sql_dml_or_ddl_statement
```

The EXECDIRECT statement allows you to execute statements inside stored procedures without preparing those statements first. This reduces the amount of code required. If the statement is a cursor, you still need to close and drop it; only the PREPARE statement can be skipped.

Considerations for EXECDIRECT usage:

- If the statement specifies a cursor name, the cursor must be dropped with the EXEC SQL DROP statement.
- If a cursor name is not specified, you do not need to drop the statement.
- If the statement is a fetch cursor, the INTO... clause must be specified.
- If the INTO clause is specified, the cursor name must be specified; otherwise the FETCH statement is able to specify which cursor name the row should be fetched from. You may have more than one open cursor at a time.

A.15.12 Procedure stack functions

The following functions may be used to analyze the current contents of the procedure stack: PROC_COUNT(), PROC_NAME(N), PROC_SCHEMA(N).

PROC_COUNT() returns the number of procedures in the procedure stack. This includes the current procedure.

PROC_NAME(N) returns the Nth procedure name in the stack. First procedure position is zero.

PROC_SCHEMA(N) returns the schema name of the Nth procedure in procedure stack.

A.15.13 Dynamic cursor names

```
CURSORNAME(  
    prefix -- VARCHAR  
)
```

The CURSORNAME() function allows you to dynamically generate a cursor name rather than hard-coding the cursor name.

Note:

Strictly speaking, CURSORNAME() is not a function, despite the syntactic similarity. CURSORNAME(arg) does not actually return anything; instead it sets the name of the current statement's cursor based on the given argument. However, it is convenient to refer to it as a function, and therefore we will do so.

Cursor names must be unique within a connection. This causes problems in recursive stored procedures because each invocation uses the same cursor name(s). When the recursive procedure calls itself, the second invocation will find that the first invocation has already created a cursor with the same name as the second invocation wants to use.

To get around this problem, we must generate unique cursor names dynamically, and we must be able to use those names when we declare and use cursors. To enable us to generate unique names and use them as cursors, we use 2 functions:

- GET_UNIQUE_STRING
- CURSORNAME

The GET_UNIQUE_STRING function does just what its name suggests — it generates a unique string. The CURSORNAME function (actually a pseudo-function) allows you to use a dynamically generated string as part of a cursor name.

Note that GET_UNIQUE_STRING returns a different output each time it is called, even if the input is the same. CURSORNAME, on the other hand, returns the same output each time if the input is the same each time.

Below is an example of using GET_UNIQUE_STRING and CURSORNAME to dynamically generate a cursor name. The dynamically generated cursorname is assigned to the placeholder "cname", which is then used in each statement after the PREPARE.

```
DECLARE autoname VARCHAR;  
Autoname := GET_UNIQUE_STRING('CUR ');  
EXEC SQL PREPARE cname CURSORNAME(autoname) SELECT * FROM TABLES;  
EXEC SQL EXECUTE cname USING(...) INTO(...);  
EXEC SQL FETCH cname;  
EXEC SQL CLOSE cname;  
EXEC SQL DROP cname;
```

CURSORNAME() can only be used in PREPARE statements and EXECDIRECT statements. It cannot be used in EXECUTE, FETCH, CLOSE, DROP, and so on.

By using the CURSORNAME() feature and the GET_UNIQUE_STRING() function, you can generate unique cursor names in recursive stored procedures. If the procedure calls itself, then each time that this function is called within the stored procedure, this function will return a unique string that can be used as the cursor name in a PREPARE statement. See below for some examples of code that you could use inside a stored procedure.

Note that each call to CURSORNAME(autoname) returns the same value — i.e. the same cursor name, as long as the input (autoname) does not change.

A.15.14 Examples of CREATE PROCEDURE statement

The examples in this section illustrate typical CREATE PROCEDURE statements and clauses.

CREATE PROCEDURE

```
"create procedure test2(tableid integer)
  returns (cnt integer)
begin
  exec sql prepare c1 select count(*) from sys_tables where id > ?;
  exec sql execute c1 using (tableid) into (cnt);
  exec sql fetch c1;
  exec sql close c1;
  exec sql drop c1;
end";
```

Using the explicit RETURN statement

This example uses the explicit RETURN statement to return multiple rows, one at a time.

```
"create procedure return_tables
  returns (name varchar)
begin
  exec sql execdirect create table table_name (lname char (20));
  exec sql whenever sqlerror rollback, abort;
  exec sql prepare c1 select table_name from sys_tables;
  exec sql execute c1 into (name);
  while sqlsuccess loop
    exec sql fetch c1;
    if not sqlsuccess
      then leave;
    end if
    return row;
  end loop;
  exec sql close c1;
  exec sql drop c1;
end";
```

Using EXECDIRECT

```
-- This example shows how to use "execdirect".
"CREATE PROCEDURE p
BEGIN
  DECLARE host_x INT;
  DECLARE host_y INT;

  -- Examples of execdirect without a cursor. Here we create a table
  -- and insert a row into that table.
  EXEC SQL EXECDIRECT create table foo (x int, y int);
  EXEC SQL EXECDIRECT insert into foo(x, y) values (1, 2);

  SET host_x = 1;

  -- Example of execdirect with cursor name.
  -- In this example, "c1" is the cursor name; "host_x" is the
  -- variable whose value will be substituted for the "?";
  -- "host_y" is the variable into which we will store the value of the
  -- column y (when we fetch it).
  -- Note: although you do not need a "prepare" statement, you still
  -- need close/drop.
  EXEC SQL c1 USING(host_x) INTO(host_y) EXECDIRECT
    SELECT y from foo where x=?;
  EXEC SQL FETCH c1;
  EXEC SQL CLOSE c1;
  EXEC SQL DROP c1;
END";
```

Using CURSORNAME

This example shows the usage of the CURSORNAME() pseudo-function. This shows only part of the body of a stored procedure, not a complete stored procedure.

```
-- Declare a variable that will hold a unique string that we can use
-- as a cursor name.
DECLARE autoname VARCHAR ;
Autoname := GET_UNIQUE_STRING('CUR_') ;
EXEC SQL PREPARE curs_name CURSORNAME(autoname) SELECT * FROM TABLES;
EXEC SQL EXECUTE curs_name USING(...) INTO(...);
EXEC SQL FETCH curs_name;
EXEC SQL CLOSE curs_name;
EXEC SQL DROP curs_name;
```

Using GET_UNIQUE_STRING and CURSORNAME

This example uses the GET_UNIQUE_STRING and CURSORNAME functions in a recursive stored procedure.

The stored procedure below demonstrates the use of these two functions in a recursive procedure. The cursor name "curs1" appears to be hardcoded, but in fact it has been mapped to the dynamically generated name.

```
-- Demonstrate GET_UNIQUE_STRING and CURSORNAME functions in a
-- recursive stored procedure.
-- Given a number N greater than or equal to 1, this procedure
-- returns the sum of the numbers 1 - N. (This can also be done in a loop,
-- but the purpose of the example is to show the use of the
-- CURSORNAME function in a recursive procedure.)
"CREATE PROCEDURE Sum1ToN(n INT)
RETURNS (SumSoFar INT)
BEGIN
    DECLARE SumOfRemainingItems INT;
    DECLARE nMinusOne INT;
    DECLARE autoname VARCHAR;

    SumSoFar := 0;
    SumOfRemainingItems := 0;
    nMinusOne := n - 1;

    IF (nMinusOne > 0) THEN
        Autoname := GET_UNIQUE_STRING('CURSOR_NAME_PREFIX_') ;
        EXEC SQL PREPARE curs1 CURSORNAME(autoname) CALL Sum1ToN(?);
        EXEC SQL EXECUTE curs1 USING(nMinusOne) INTO(SumOfRemainingItems);
        EXEC SQL FETCH curs1;
        EXEC SQL CLOSE curs1;
        EXEC SQL DROP curs1;
    END IF;

    SumSoFar := n + SumOfRemainingItems;
END";
```

Using EXECDIRECT in CREATE PROCEDURE

```
CREATE TABLE table1 (x INT, y INT);
INSERT INTO table1 (x, y) VALUES (1, 2);

"CREATE PROCEDURE FOO
RETURNS (r INT)
BEGIN
    DECLARE autoname VARCHAR;
    Autoname := GET_UNIQUE_STRING('CUR_');
    EXEC SQL curs_name INTO(r) CURSORNAME(autoname) EXECDIRECT
        SELECT y FROM TABLE1 WHERE x = 1;
    EXEC SQL FETCH curs_name;
    EXEC SQL CLOSE curs_name;
```

```
EXEC SQL DROP curs_name;
END";
```

```
CALL foo();
SELECT * FROM table1;
```

Creating a unique name for a synchronization message

Creating a unique name for a synchronization message:

```
DECLARE Autoname VARCHAR;
DECLARE Sqlstr VARCHAR;
Autoname := get_unique_string('MSG_') ;
Sqlstr := 'MESSAGE' + autoname + 'BEGIN';
EXEC SQL EXECDIRECT Sqlstr;
...
Sqlstr := 'MESSAGE' + autoname + 'FORWARD';
EXEC SQL EXECDIRECT Sqlstr;
```

Using GET_UNIQUE_STRING

This example demonstrates how to use the GET_UNIQUE_STRING() function to generate unique message names from within a recursive stored procedure.

```
CREATE TABLE table1 (i int, beginMsg VARCHAR, endMsg VARCHAR);
```

```
-- This is a simplified example of recursion.
-- Note that the messages that are composed are not actually used! This is
-- not a true example of synchronization; it is only an example of
-- generating unique message names. The "count" parameter is the
-- number of times that you want this function to call
-- itself (not including the initial call).
"CREATE PROCEDURE repeater(count INT)
```

```
BEGIN
```

```
DECLARE Autoname VARCHAR;
DECLARE MsgBeginStr VARCHAR;
DECLARE MsgEndStr VARCHAR;
```

```
Autoname := GET_UNIQUE_STRING('MSG_');
MsgBeginStr := 'MESSAGE ' + Autoname + ' BEGIN';
MsgEndStr := 'MESSAGE ' + Autoname + ' END';
```

```
EXEC SQL c1 USING (count, MsgBeginStr, MsgEndStr) EXECDIRECT
    INSERT INTO table1 (i, beginMsg, endMsg) VALUES (?, ?, ?);
EXEC SQL CLOSE c1;
EXEC SQL DROP c1;
```

```
-- Once you have composed the SQL statement as a string,
-- you can execute it one of two ways:
-- 1) by using the EXECDIRECT feature or
-- 2) by preparing and executing the SQL statement.
-- In this example, we use EXECDIRECT.
EXEC SQL EXECDIRECT MsgBeginStr;
EXEC SQL EXECDIRECT MsgEndStr;
-- Do something useful here.
```

```
-- The recursive portion of the function.
IF (count > 1) THEN
    SET count = count - 1;
    -- Note that we can also use our unique name as a cursor name,
    -- as shown below.
    EXEC SQL Autoname USING (count) EXECDIRECT CALL repeater(?);
    EXEC SQL CLOSE Autoname;
    EXEC SQL DROP Autoname;
END IF
```

```
RETURN;
```



```

END";

CALL repeater(3);

-- Show the message names that we composed.
SELECT * FROM table1;

```

The output from this SELECT statement would look similar to the following:

```

I  BEGINMSG                                ENDMSG
-----
1  MESSAGE MSG_019 BEGIN                   MESSAGE MSG_019 END
2  MESSAGE MSG_020 BEGIN                   MESSAGE MSG_020 END
3  MESSAGE MSG_021 BEGIN                   MESSAGE MSG_021 END

```

A.16 CREATE PROCEDURE (external)

```

CREATE PROCEDURE procedure_name [(parameter_definition [, parameter_definition ...])]
  [RETURNS (output_column_definition [, output_column_definition ...])]
  [SQL_data_access_indication]
  LANGUAGE C
  EXTERNAL NAME 'external_procedure_library_identifier'

```

parameter_definition ::= [*parameter_mode*] *parameter_name* *data_type*
parameter_mode ::= IN | OUT | INOUT

output_column_definition::= *column_name* *column_type*

SQL_data_access_indication::= CONTAINS SQL | READS SQL DATA | MODIFIES SQL DATA | NO SQL

Usage

The CREATE PROCEDURE (external) statement defines an external procedure written in C programming language.

The external procedures are accessed using the standard dynamic library interface as provided by the operating system (shared libraries in Linux and UNIX, DLLs in Windows). To enable use of external stored procedures, you must link your application dynamically to solidDB using shared memory access (SMA) or linked library access (LLA).

The external procedures are called with the same CALL statement syntax as SQL stored procedures; the control of the thread is passed to the external library routine until it returns. Furthermore, the database interface of the external procedures is similar to that of an ODBC application, except that the connection handle (hdbc) is passed to the external routine, and it can immediately start using it without establishing the connection.

The output of the external procedure can be returned as a result set of a single row or as output parameters.

Parameters

procedure_name

Identifies the procedure name. The procedure name can also include the catalog name and schema:

[catalog_name[.schema].]procedure_name

The name must be a valid SQL identifier. The procedure name must be unique within the schema.

parameter_definition ::= [parameter_mode] parameter_name data_type

Identifies the parameters of the procedure, and specifies the mode, optional parameter name, and data type of the parameter. One entry in the list must be specified for each parameter that the procedure will expect.

parameter_mode ::= IN | OUT | INOUT

Identifies the parameter mode.

IN Identifies the parameter as an input parameter to the procedure. Any changes made to the parameter within the procedure are not available to the calling SQL application when control is returned. The default is IN.

OUT

Identifies the parameter as an output parameter for the procedure.

INOUT

Identifies the parameter as both an input and output parameter for the procedure.

Note: If the RETURNS keyword is used, the parameter mode can only be IN.

parameter_name

Specifies the name of the parameter. The parameter name must be unique for the procedure.

data_type

Specifies the data type of the parameter. The data type must be a valid solidDB A.79.2, "data_type," on page 305.

RETURNS (*output_column_definition*)

Identifies the output of the procedure.

output_column_definition ::= column_name column_type

column_name

Specifies the name of the output column.

column_type

Specifies the data type of the output column. The data type must be a valid solidDB A.79.2, "data_type," on page 305.

Note: If the RETURNS keyword is used, the parameter mode can only be IN.

SQL_data_access_indication

Defines whether the procedure issues any SQL statements and, if so, what type.

CONTAINS SQL

Indicates that the procedure contains SQL statements but these statements do not read or modify SQL data.

READS SQL DATA

Indicates that the procedure contains SQL statements that read SQL data but do not modify any data.

MODIFIES SQL DATA

Indicates that the procedure contains SQL statements modify SQL data .

NO SQL

Indicates that the procedure does not execute any SQL statements.

Note: The validity of the contents of the procedure are not checked against the declared data access indication.

LANGUAGE C

Indicates that the CREATE PROCEDURE statement is being used to register a new procedure based on code written in C programming language.

EXTERNAL NAME '*external_procedure_library_identifier*'

Identifies the name of the user-written code which implements the procedure being defined. The name must be enclosed by single quotation marks. Extraneous blanks are not permitted.

`external_procedure_library_identifier ::= [path]
library_name!C_routine_name`

[*path*] *library_name*

Identifies the full path (optional) and name of the procedure. If you specify an absolute path, you need to append the file extension of the library file. If you do not specify the file extension, you might need to add the location of the library to the LD_LIBRARY_PATH or LIBPATH (Linux and UNIX) or PATH (Windows) environment variable.

For example:

- On UNIX systems, '/u/jchui/mylib/myproc.so' would cause solidDB to look in /u/jchui/mylib for the myproc.so library.
- On Windows systems, 'd:\mylib\myproc.dll' would cause solidDB to load the file myproc.dll from the d:\mylib directory.

Tip: If you do not specify the file extension, you can port and use the library on several platforms.

!*C_routine_name*

Identifies the entry point name of the C routine (procedure) to be invoked. The exclamation point (!) serves as a delimiter between the library name and the routine name.

'!proc8' would direct solidDB to look for the library in the location specified by [*path*] *library_name*, and to use entry point *proc8* within that library.

Note: On Windows 32-bit systems, use the calling convention `_cdecl` for calling the C routines. This is to allow checking and handling of varying argument counts. Also, ensure that your DLL exports the routines you want to call: for example, include the routines in the corresponding `.def` file in your DLL project.

Note: The library (and the procedure within the library) does not need to exist when the CREATE PROCEDURE statement is issued. However, when the procedure is called, the library and procedure within the library must exist and be accessible.

Return values

- SQL_ERROR
- SQL_SUCCESS

Restrictions

- The call interface of the external procedures is assumed to be non-fenced and thread-safe.

- All procedures are assumed non-deterministic; the procedure depends on some state values that affect the results.
- You should not start new threads that execute in parallel with the called procedure; the provided connection handle is not multithread safe in the same manner as normal ODBC connection handle is. If new threads are launched, they cannot access the connection concurrently or after the original external procedure has returned.

Example

Tip: For more examples, see the `samples/procedures` directory in your solidDB installation directory.

The following example registers an external stored procedure `EXTPROCEXAMPLE1`.

```
CREATE PROCEDURE EXTPROCEXAMPLE1(IN P1 INTEGER, INOUT P2 INTEGER, OUT P3 CHAR(10))
READS SQL DATA
LANGUAGE C
EXTERNAL NAME 'example1lib1!extprocexample1'
```

C program for the procedure

```
#include <solidodbc3.h>
/* UTF-8 sequence for a char can be up to 4 bytes */
#define CHARMAXBYTELEN(charlen) (4 * (charlen) + 1)
int extprocexample1(SQLHDBC hdbc, /* implicit input argument: 'hot' connection handle */
                   char sqlst[6], /* implicit output argument: sqlstate */
                   char qualName[], /* implicit input argument: procedure name */
                   char diagMsg[71], /* implicit output argument: diagnostics message */
                   SQLULEN nArgs, /* implicit input: actual number of args */
                   SQLLEN nullind[3], /* implicit input/output: null/size indicator variables
                                       (when nArgs > 0) */
                   int* p1, int* p2, char p3[CHARMAXBYTELEN(10)]) /* explicit arguments */
{
    SQLHSTMT hstmt = SQL_NULL_STMT;
    SQLRETURN rc;
    int retcode;

    if (nArgs != 3) {
        goto argcount_failure;
    }
    nullind[2] = SQL_NULL_DATA;
    if (nullind[0] == SQL_NULL_DATA || nullind[1] == SQL_NULL_DATA) {
        goto null_input_detected;
    }
    rc = SQLAllocHandle(SQL_HANDLE_STMT, hdbc, &hstmt);
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    rc = SQLPrepare(hstmt, "SELECT P2,P3 FROM EXTPROCEXAMPLE1_TAB "
                    "WHERE P1 = ? AND P2 > ?", SQL_NTS);
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    rc = SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT,
                        SQL_C_SLONG, SQL_INTEGER,
                        0, 0, p1, sizeof(*p1), NULL);
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    rc = SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT,
                        SQL_C_SLONG, SQL_INTEGER,
                        0, 0, p2, sizeof(*p2), NULL);
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    rc = SQLBindColumn(hstmt, 1, SQL_C_SLONG, p2, 0, NULL);
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    rc = SQLBindColumn(hstmt, 2, SQL_C_CHAR, p3, CHARMAXBYTELEN(10), NULL);
    if (rc != SQL_SUCCESS) {
```

```

        goto failure;
    }
    rc = SQLExecute(hstmt);
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    rc = SQLFetch(hstmt);
    if (rc == SQL_NO_DATA) {
        goto no_data_found;
    }
    if (rc != SQL_SUCCESS) {
        goto failure;
    }
    nullind[1] = sizeof(int);
    nullind[2] = sizeof(int);
    retcode = SQL_SUCCESS;
cleanup:;
    if (hstmt != SQL_NULL_HANDLE) {
        SQLFreeHandle(SQL_HANDLE_STMT, hstmt);
    }
    return retcode;
null_input_detected:;
no_data_found:;
    retcode = SQL_SUCCESS;
    nullind[1] = SQL_NULL_DATA;
    nullind[2] = SQL_NULL_DATA;
    goto cleanup;
failure:;
    retcode = SQL_ERROR;
    if (hstmt != SQL_NULL_HANDLE) {
        rc = SQLError(SQL_NULL_HANDLE, SQL_NULL_HANDLE, hstmt,
            sqlst, &nativeerror, diagMsg, 71, NULL);
        if (rc == SQL_NO_DATA) {
            goto generic_failure;
        }
    }
    else {
        rc = SQLError(SQL_NULL_HANDLE, hdbc, SQL_NULL_HANDLE,
            sqlst, &nativeerror, diagMsg, 71, NULL);
        if (rc == SQL_NO_DATA) {
            goto generic_failure;
        }
    }
    goto exit_cleanup;
generic_failure:;
    memcpy(sqlst, "1C000", 6);
    memcpy(diagMsg, "Generic error", 14);
    goto exit_cleanup;
argcount_failure:;
    memcpy(sqlst, "1C000", 6);
    strncpy(diagMsg, "wrong number of arguments, 3 expected", 71);
    goto exit_cleanup;
}

```

Related reference:

A.15, "CREATE PROCEDURE," on page 180

A.35, "DROP PROCEDURE," on page 222

A.17 CREATE [OR REPLACE] PUBLICATION

```

"CREATE [OR REPLACE] PUBLICATION publication_name
  [(parameter_definition [,parameter_definition...])]
BEGIN
  main_result_set_definition...
END";

```

```

main_result_set_definition ::=
RESULT SET FOR main_replica_table_name

```

```

BEGIN
  SELECT select_list
  FROM master_table_name
  [ WHERE search_condition ] ;
  [ [DISTINCT] result_set_definition... ]

```

```

END

result_set_definition ::=
RESULT SET FOR replica_table_name

BEGIN
    SELECT select_list
    FROM master_table_name
    [ WHERE search_condition ] ;
    [[ DISTINCT] result_set_definition...]
END

```

where

search_condition can reference *parameter_definitions* and/or columns of replica tables defined on previous (higher) levels.

Usage

The CREATE [OR REPLACE] PUBLICATION statement creates or replaces *publications*. Publications define the sets of data that can be REFRESHed from the master to the replica database. A publication is always transactionally consistent, that is, its data has been read from the master database in one transaction and the data is written to the replica database in one transaction.

Important: The data read from the publication is internally consistent unless the master is using the READ COMMITTED isolation level.

Search conditions of a SELECT clause can contain input arguments of the publication as parameters. The parameter name must have a colon as a prefix.

Publications can contain data from multiple tables. The tables of the publication can be independent or there can be relations between the tables. If there is a relation between tables, the result sets must be nested. The WHERE clause of the SELECT statement of the inner result set of the publication must refer to a column of the table of the outer result set.

If the relation between outer and inner result set of the publication is a N-1 relationship, the keyword DISTINCT must be used in the result set definition.

The *replica_table_name* can be different from the *master_table_name*. The publication definition provides the mapping between the master and replica tables. If you have multiple replicas, all the replicas should use the same name, even if that name is different from the name used in the master. Column names in the master and replica tables must be the same.

The initial download is always a *full publication*, which means that all data contained in the publication is sent to the replica database. Subsequent downloads (refreshes) for the same publication may be *incremental publications*, which means that they contain only the data that has been changed since the prior REFRESH. To enable usage of incremental publications, SYNCHISTORY has to be set ON for tables included in the publication in both the master and replica databases. For details, see A.4.2, "ALTER TABLE ... SET SYNCHISTORY," on page 160 and A.37, "DROP PUBLICATION REGISTRATION," on page 222.

If the optional keywords "OR REPLACE" are used, and the publication already exists, it will be replaced with the new definition. Since the publication was not dropped and re-created, replicas do not need to reregister, and subsequent

REFRESHes from that publication can be incremental rather than full, depending upon exactly what changes were made to the publication.

To avoid having a replica refresh from a publication while you are updating that publication, you may temporarily set the catalog's sync mode to Maintenance mode. However, using maintenance mode is not absolutely required when replacing a publication.

If you replace an existing publication, the new definition of the publication will be sent to each replica the next time that replica requests a refresh. The replica does not need to explicitly reregister itself to the publication.

When you replace an existing publication with a new definition, you can change the resultset definitions. You cannot change the parameters of the publication. The only way to change the parameters is to drop the publication and create a new one, which also means that the replicas must reregister and the replicas will get a full refresh rather than an incremental refresh the next time that they request a refresh.

When you replace an existing publication, the privileges related to that publication are left unchanged. You do not need to re-create them.

The CREATE OR REPLACE PUBLICATION command can be executed in any situation where it is valid to execute the CREATE PUBLICATION command.

Important: If you use CREATE OR REPLACE PUBLICATION to alter the contents of an existing advanced replication publication, you have to take care of removing invalid rows from Replica.

Usage in master

You define the publication in the master database to enable the replicas to get refreshes from it.

Usage in replica

There is no need to define the publications in the replicas. Publication subscription functionality depends on the definitions only at the master database. If the CREATE OR REPLACE PUBLICATION command is executed in a replica, it will store the publication definition to the replica, but the publication definition is not used for anything.

Note: If a database is both a replica (for a master above it) and a master (to a replica below it), you might want to create a publication definition in the database.

Examples

Example 1:

The following example publication retrieves data from the customer table using the area code of the customer as search criterion. For each customer, the orders and invoices of the customer (1-N relation) as well as the dedicated salesman of the customer (1-1 relation) are also retrieved.

```
"CREATE PUBLICATION PUB_CUSTOMERS_BY_AREA
  (IN_AREA_CODE VARCHAR)
BEGIN
```

```

RESULT SET FOR CUSTOMER
BEGIN
    SELECT * FROM CUSTOMER
    WHERE AREA_CODE = :IN_AREA_CODE;
    RESULT SET FOR CUST_ORDER
    BEGIN
    SELECT * FROM CUST_ORDER
    WHERE CUSTOMER_ID = CUSTOMER.ID;
    END
    RESULT SET FOR INVOICE
    BEGIN
    SELECT * FROM INVOICE
    WHERE CUSTOMER_ID = CUSTOMER.ID;
    END
    DISTINCT RESULT SET FOR SALESMAN
    BEGIN
    SELECT * FROM SALESMAN
    WHERE ID = CUSTOMER.SALESMAN_ID;
    END
END
END";

```

Note:

The colon (:) in :IN_AREA_CODE is used to designate a reference to a publication parameter with the same name.

Example 2:

Developers decided to add a new column C in table T, which is referenced in publication P. The modification must be made to the master database and all replica databases.

The tasks to execute in the master database are:

```

-- Prevent other users from doing concurrent synchronization operations
-- to this catalog.
SET SYNC MAINTENANCE MODE ON;
ALTER TABLE T ADD COLUMN C INTEGER;
COMMIT WORK;
CREATE OR REPLACE PUBLICATION P ... (column C added also to publication)
COMMIT WORK;
SET SYNC MAINTENANCE MODE OFF;

```

The tasks to execute in all replica databases are:

```

-- Prevent other users from doing concurrent synchronization operations
-- to this catalog.
SET SYNC MAINTENANCE MODE ON;
ALTER TABLE T ADD COLUMN C INTEGER;
COMMIT WORK;
SET SYNC MAINTENANCE MODE OFF;

```

Return values

Table 31. CREATE PUBLICATION return values

Error Code	Description
13047	No privilege for operation. You do not have the privileges required to drop this publication or create a publication.
13120	The name is too long for the publication

Table 31. CREATE PUBLICATION return values (continued)

Error Code	Description
25015	Syntax error: <i>error_message</i> , line <i>line_number</i>
25021	Database is not master or replica database. Publications can be created only in a master or replica database. (As a practical matter, they should only be created in a master database.)
25033	Publication <i>publication_name</i> already exists
25049	Referenced table <i>table_name</i> not found in subscription hierarchy
25061	Where condition for table <i>table_name</i> must refer to an outer table of the publication

A.18 CREATE [OR REPLACE] REMOTE SERVER

```
CREATE [OR REPLACE] REMOTE SERVER [USERNAME <username> PASSWORD <password>]
```

Usage

The CREATE REMOTE SERVER statement creates back-end login information in the SYS_SERVER system table for SQL passthrough purposes in Universal Cache.

The CREATE OR REPLACE REMOTE SERVER statement modifies login information in the SYS_SERVER system table, overwriting an existing entry.

In most cases, the InfoSphere® CDC technology is used for transferring backend login data to the frontend. When mirroring or a refresh is started on the first subscription from solidDB server to the backend data server, the InfoSphere CDC for solidDB instance retrieves the login data from the backend InfoSphere CDC instance and stores it in the solidDB system table SYS_SERVER with the statement CREATE REMOTE SERVER. The password stored in the SYS_SERVER table is hidden.

The InfoSphere CDC technology does not transfer the backend login data in the cases described below.

- If the InfoSphere CDC instance for the backend is running under a user ID that automatically has access to the database, the login data does not need to be stored.
- If the backend data server is DB2® for z/OS® or DB2 for iSeries®, the login data cannot be fetched. To avoid errors, you need to set the InfoSphere CDC for solidDB system parameter **retrieve_credentials** to FALSE.
- If you have upgraded your InfoSphere CDC for solidDB installation and subscription from V6.3, the 6.3 version of InfoSphere CDC for solidDB has not stored the backend login data in the SYS_SERVER table.

In the latter two cases, use the CREATE REMOTE SERVER (or ALTER REMOTE SERVER) statement to define the login data manually.

By default, username and password are stored in uppercase letters. To retain case sensitivity, enter the username and password in single quotation marks.

Examples

```
CREATE REMOTE SERVER USERNAME dba PASSWORD dba
CREATE REMOTE SERVER USERNAME 'AdMin' PASSWORD 'PwD123'
CREATE OR REPLACE REMOTE SERVER USERNAME dba PASSWORD 'PwD123'
```

See also

A.3, “ALTER REMOTE SERVER,” on page 157

A.38, “DROP REMOTE SERVER,” on page 223

A.19 CREATE ROLE

```
CREATE ROLE role_name
```

Usage

The CREATE ROLE statement creates a new user role.

A *role* is a group of privileges that can be granted to users as one unit. You can create roles and assign users to certain roles. A single user may have more than one role assigned, and a single role may have more than one user assigned.

Note:

- The same string cannot be used both as a user name and a role name.
- Once a user role is granted, it takes effect when the user who is granted the role logs on to the database. If the user is already logged on to the database when the role is granted, the role takes effect when the user disconnects and then reconnects to the database.

solidDB offers the following system roles. The system role names are reserved user names.

Table 32. System roles

Reserved Names	Description
PUBLIC	This role grants privileges to all users. When user privileges to a certain table are granted to the role <i>PUBLIC</i> , all current and future users have the specified user privileges to this table. This role is granted automatically to all users.
SYS_ADMIN_ROLE	This is the default role for the database administrator. This role has administration privileges to all tables, indexes and users, as well as the right to use solidDB Remote Control. This is also the database creator role.
_SYSTEM	This is the schema name of all system tables and views.
SYS_CONSOLE_ROLE	This role has the right to use solidDB Remote Control, but does not have other administration privileges.
SYS_SYNC_ADMIN_ROLE	This is the administrator role for data synchronization functions.
SYS_SYNC_REGISTER_ROLE	This role is only for registering and unregistering a replica database to the master.

Example

```
CREATE ROLE GUEST_USERS;
```

A.20 CREATE SCHEMA

```
CREATE SCHEMA schema_name
```

Usage

The CREATE SCHEMA statement creates a new schema.

Schemas are a collection of database objects, such as tables, views, indexes, events, triggers, sequences, and stored procedures for a database user. The default schema name is the user id; there is one default schema for each user. The use of schemas in solidDB conforms to the SQL standard.

The schema name is used to qualify a database object name. Database object names are qualified in all DML statements as:

```
catalog_name.schema_name.database_object_name
```

or

```
user_id.database_object_name
```

To logically partition a database, users can create a catalog before they create a schema. For details on creating a catalog, read A.10, “CREATE CATALOG,” on page 170. When creating a new database or converting an old database to a new format, you are prompted for a default catalog name.

To use schemas, a schema name must be created before creating the database object name (such as a table name or procedure name). However, a database object name can be created without a schema name. In such cases, database objects are qualified using *user_id* only.

You can specify the database object names in a DML statement explicitly by fully qualifying them or implicitly by setting the schema name context using:

```
SET SCHEMA schema_name
```

Creating a schema does not automatically make that schema the current default schema. If you have created a new schema and want your subsequent commands to execute within that schema, you must also execute the A.71.9, “SET SCHEMA,” on page 284 statement. For example:

```
CREATE SCHEMA MySchema;  
CREATE TABLE t1; -- not in MySchema  
SET SCHEMA MySchema;  
CREATE TABLE t2; -- in MySchema
```

A schema can be dropped from a database using the A.41, “DROP SCHEMA,” on page 224 statement. When dropping a schema name, all objects associated with the schema name must be dropped prior to dropping the schema.

A schema context can be removed using the SET SCHEMA USER statement.

Rules for resolving schema names:

- A fully qualified name (*schema_name.database_object_name*) does not need any name resolution, but will be validated.

- If a schema context is not set using SET SCHEMA, all database object names are resolved always using the user id as the schema name.
- If the database object name cannot be resolved from the schema name, then the database object name is resolved from all existing schema names.
- If name resolution finds either zero matching or more than one matching database object name, then a solidDB server issues a name resolution conflict error.

Examples

```
-- Assume the userID is SMITH.
CREATE SCHEMA FINANCE;
CREATE TABLE EMPLOYEE (EMP_ID INTEGER);
SET SCHEMA FINANCE;
CREATE TABLE EMPLOYEE (ID INTEGER);
SELECT ID FROM EMPLOYEE;
-- In this case, the table is qualified to FINANCE.EMPLOYEE
SELECT EMP_ID FROM EMPLOYEE;
-- This will give an error as the context is with FINANCE and
-- table is resolved to FINANCE.EMPLOYEE

--The following are valid schema statements: one with a schema context,
--the other without.
SELECT ID FROM FINANCE.EMPLOYEE;
SELECT EMP_ID FROM SMITH.EMPLOYEE
--The following statement will resolve to schema SMITH without a schema
--context
SELECT EMP_ID FROM EMPLOYEE;
```

A.21 CREATE SEQUENCE

```
CREATE [DENSE] SEQUENCE sequence_name
[START WITH constant] [INCREMENT BY constant]
```

Usage

Sequencer objects are objects that are used to get sequence numbers.

DENSE sequences guarantee that there are no holes in the sequence numbers. The sequence number allocation is bound to the current transaction. If the transaction rolls back, then the sequence number allocations are also rolled back. The drawback of dense sequences is that the sequence is locked out from other transactions until the current transaction ends.

Sparse sequences guarantee uniqueness of the returned values, but they are not bound to the current transaction. If a transaction allocates a sparse sequence number and later rolls back, the sequence number is simply lost.

Both dense and sparse sequence numbers start from 1, unless specified otherwise with START WITH *constant*.

Sequence numbers are 8-byte values. Sequence values can be stored in BIGINT, BINARY, or INT data types.

- BIGINT is recommended.
- 8-byte BINARY values can store a complete sequence number, but BINARY values are not always as convenient to work with as integer data types.
- Sequence values stored in INT variables lose information because an 8-byte sequence number will not fit in a 4-byte integer.

When an 8-byte number sequence number is stored as a 4-byte integer (in a stored procedure or in an application program), the highest four bytes are omitted. This will cause loss of data when the sequence number is above $2^{31} - 1$ (=2147483647).

The following example demonstrates this behavior:

```
CREATE SEQUENCE seq1;

-- Set the sequence number to 2^31 - 1,
-- then return that value and the "next" value (2^31).
"CREATE PROCEDURE set_seq1_to_2G
RETURNS (x INT, y INT)
BEGIN
DECLARE int1 INTEGER;
int1 := 2147483647;
EXEC SEQUENCE seq1 SET VALUE USING int1;
EXEC SEQUENCE seq1 CURRENT INTO x;
EXEC SEQUENCE seq1 NEXT INTO y;
END";

COMMIT WORK;

CALL set_seq1_to_2G();
```

The return values from the call are:

x	y
2147483647	-2147483648

The value for x is correct, but the value for y is a negative number instead of the correct positive number.

The advantage of using a sequencer object instead of a separate table is that the sequencer object is specifically fine-tuned for fast execution and requires less overhead than normal update statements.

Sequence values can be incremented and used within SQL statements. The increment is defined with `INCREMENT BY constant`.

The following constructs can be also used in SQL:

```
sequence_name.CURRVAL
sequence_name.NEXTVAL
```

Select access rights are required to retrieve the current sequence value. Update access rights are required to allocate new sequence values. These access rights are granted and revoked in the same way as table access rights.

Sequences in stored procedures

Sequences can also be used inside stored procedures.

Table 33. Sequences in stored procedures

Stored procedure statement	Usage
<code>EXEC SEQUENCE <i>sequence_name</i>.CURRENT INTO <i>variable</i></code>	Retrieves a new sequence value
<code>EXEC SEQUENCE <i>sequence_name</i>.NEXT INTO <i>variable</i></code>	Retrieves the current sequence value
<code>EXEC SEQUENCE <i>sequence_name</i> SET VALUE USING <i>variable</i></code>	Sets sequence values Tip: Alternatively, you can use the <code>SET SEQUENCE <name> VALUE <bigint value> statement</code> .

Examples

```
CREATE DENSE SEQUENCE SEQ1;  
INSERT INTO ORDER (id) VALUES (SEQ1.NEXTVAL);  
INSERT INTO ORDER (id) VALUES (SEQ1.NEXTVAL);  
"CREATE PROCEDURE get_my_seq  
  RETURNS (val INTEGER)  
  BEGIN  
    EXEC SEQUENCE my_sequence.NEXT INTO (val);  
  END";
```

A.22 CREATE SYNC BOOKMARK

```
CREATE SYNC BOOKMARK bookmark_name
```

Usage

The CREATE SYNC BOOKMARK statement creates a bookmark in a master database. The statement is valid only in solidDB advanced replication setups.

A *bookmark* represents a user-defined version of the database. It is a persistent snapshot of a solidDB database, which provides a reference for performing specific synchronization tasks. Bookmarks are used typically to export data from a master for import into a replica using the EXPORT SUBSCRIPTION command. Exporting and importing data allows you to create a replica from a master more efficiently if you have databases larger than 2 GB.

To create a bookmark, you must have administrative DBA privileges or SYS_SYNC_ADMIN_ROLE. There is no limit to the number of bookmarks you can create in a database. A bookmark is created only in a master database. The system issues an error if you attempt to create a bookmark in a replica database.

If a table is set up for synchronization history with the A.4.2, "ALTER TABLE ... SET SYNCHISTORY," on page 160 command, a bookmark retains history information for the table. For this reason, use the A.44, "DROP SYNC BOOKMARK," on page 227 statement to drop bookmarks when they are no longer needed. Otherwise, extra history data will increase disk space usage.

When you create a new bookmark, the system associates other attributes, such as creator of the bookmark, creation data and time, and a unique bookmark ID. This metadata is maintained in the system table E.2.10, "SYS_SYNC_BOOKMARKS," on page 368.

Usage in master

Use the CREATE SYNC BOOKMARK statement to create a bookmark in a master database.

Usage in replica

The CREATE SYNC BOOKMARK statement cannot be used in a replica database.

Example

```
CREATE SYNC BOOKMARK BOOKMARK_AFTER_DATALOAD;
```

Return values

Table 34. CREATE SYNC BOOKMARK return values

Error Code	Description
25066	Bookmark already exists
13047	No privilege for operation

A.23 CREATE TABLE

```
CREATE [ { [GLOBAL] TEMPORARY | TRANSIENT } ] TABLE base_table_name
(column_element [, column_element] ...) [STORE {MEMORY | DISK}]
[LIKE table_name]
```

```
base_table_name ::= base_table_identifier | schema_name.base_table_identifier |
catalog_name.schema_name.base_table_identifier
```

```
column_element ::= column_definition | table_constraint_definition
```

```
column_definition ::= column_identifier
data_type [DEFAULT literal | NULL] [NOT NULL]
[column_constraint_definition [column_constraint_definition] ...]
```

```
column_constraint_definition ::= [CONSTRAINT constraint_name]
UNIQUE | PRIMARY KEY |
REFERENCES ref_table_name [(referenced_columns)] |
CHECK (check_condition)
```

```
table_constraint_definition ::= [CONSTRAINT constraint_name]
UNIQUE (column_identifier [, column_identifier] ...) |
PRIMARY KEY (column_identifier [, column_identifier] ...) |
CHECK (check_condition) |
{FOREIGN KEY (column_identifier [, column_identifier] ...)
REFERENCES table_name [(referenced_columns)]
[referential_triggered_action] }
```

```
referential_triggered_action:: =
ON {UPDATE | DELETE} {CASCADE | SET NULL | SET DEFAULT |
RESTRICT | NO ACTION}
```

Usage

Tables are created through the CREATE TABLE statement. The CREATE TABLE statement requires a list of the columns created, the data types, and, if applicable, sizes of values within each column, in addition to other options, such as creating primary keys.

If the table name or column names starts with a digit, it must be given in double quotation marks.

STORE (table type)

The STORE clause indicates whether the table should be stored in memory or on disk. The CREATE TABLE statement takes precedence over the storage location specified by the **General.DefaultStoreIsMemory** parameter.

In-memory tables can be persistent (normal) tables, temporary tables, or transient tables.

All temporary tables and transient tables must be in-memory tables. You do not need to specify the `STORE MEMORY` clause; temporary tables and transient tables will automatically be created as in-memory tables if you omit the `STORE` clause. For temporary tables and transient tables, the **General.DefaultStoreIsMemory** parameter is ignored. You will get an error if you try to explicitly create temporary tables or transient tables as disk-based tables.

The keyword "GLOBAL" is included to comply with the SQL:1999 standard for temporary tables. In solidDB, all temporary tables are global, whether or not the GLOBAL keyword is used.

Row size of disk-based and in-memory tables

In a disk-based table, the maximum size of a row (excluding BLOBs) is approximately 1/3 of the page size. In an in-memory table, the maximum size of a row (including BLOBs) is approximately the page size. (There is a small amount of overhead used in both disk-based and in-memory pages, so all of the page is not available for user data.) The page sized (block size) is defined with the **IndexFile.BlockSize** parameter (default is 16 KB).

The server does not use simple rules to determine BLOB storage, but as a general rule each BLOB occupies 256 bytes from the page where the row resides, and the rest of the BLOB goes to separate BLOB pages. If the BLOB is shorter than 256 bytes, then it is stored entirely in the main disk page, not BLOB pages.

Each row is limited to 1000 columns.

Referential integrity

You should always define a primary key when you create a table. If you do not define a primary key, solidDB will create one automatically. This will lead to unexpected data order on the disk and may cause performance degradation. An appropriate primary key speeds up queries using the primary key.

Constraint definitions are available for both the column and table level. For the column level, constraints defined with `NOT NULL` specify that a non-null value is required for a column insertion. `UNIQUE` specifies that no two rows are allowed to have the same value. `PRIMARY KEY` ensures that the column(s), which is (are) a primary key, does not permit two rows to have the same value and does not permit any `NULL` values; `PRIMARY KEY` is thus equivalent to the combination of `UNIQUE` and `NOT NULL`. The `REFERENCES` clause with `FOREIGN KEY` specifies a table name and a list of columns for a referential integrity constraint. This means that when data is inserted or updated in this table, the data must match the values in the referenced tables and columns.

The `CHECK` keyword restricts the values that can be inserted into a column (for example, restricting the values with a specific integer range). When defined, the check constraint performs a validation check for any data that is inserted or updated in that column. If the data violates the constraint, then the modification is prohibited. For example:

```
CREATE TABLE table1 (salary DECIMAL CHECK (salary >= 0.0));
```

The `check_condition` is a boolean expression that specifies the check constraints for the column. Check constraints are defined with the predicates `>`, `<`, `=`, `<>`, `<=`, `>=` and the keywords `BETWEEN`, `IN`, `LIKE` (which may contain wildcard characters),

and IS [NOT] NULL. The expression (similar to the syntax of a WHERE clause) can be qualified with keywords AND and OR. For example:

```
...CHECK (col1 = 'Y' OR col1 = 'N')...  
...CHECK (last_name IS NOT NULL)...
```

UNIQUE and PRIMARY KEY constraints can be defined at the column level or the table level. They also automatically create a unique index on the specified columns.

FOREIGN KEY is used to specify that the column(s) listed are foreign keys in this table.

REFERENCES specifies the table and those column(s) that are references of the foreign key(s). Although column-level constraints can use a REFERENCES clause, only table-level constraints can use the FOREIGN KEY ... REFERENCES clause.

To use the REFERENCES constraint with FOREIGN keys, a foreign key must always include enough columns in its definition to uniquely identify a row in the referenced table. A foreign key must contain the same number and type (data type) of columns as the primary key in the referenced table as well as be in the same order; however, a foreign key can have different column names and default values than the primary key.

The following rules apply on the constraints:

- The *check_condition* cannot contain subqueries, aggregate functions, host variables, or parameters.
- Column check constraints can reference only the columns on which they are defined.
- Table check constraints can reference any columns in the table, that is if all columns in the table have been defined earlier in the statement.
- A table may have only one primary key constraint, but may have multiple unique constraints.
- The UNIQUE and PRIMARY KEY constraints in the CREATE TABLE statement can be used to create indexes. In the ALTER TABLE statement, a column cannot be dropped if it is part of a unique or primary key. You may want to use the CREATE INDEX statement to create an index instead – this way the index will have a name and you can drop it. The CREATE INDEX statement also offers some additional features, such as the ability to create non-unique indexes and to specify if the indexes are sorted in ascending or descending order.
- The referential integrity rules for persistent, transient, and temporary table types are different.
 - A temporary table may reference another temporary table, but may not reference any other type of table (transient or persistent). No other type of table may reference a temporary table.
 - Transient tables can reference other transient tables and persistent tables. They cannot reference temporary tables. Neither temporary tables nor persistent tables can reference a transient table.

Creating tables like existing tables

Instead of providing the table definition in the CREATE TABLE statement, you can create a table with the same definition as in an existing table by using the CREATE *new_table_name* LIKE *existing_table_name* construction.

For example:

```
CREATE TABLE MY_TABLE1(A INTEGER NOT NULL PRIMARY KEY, B INT) STORE MEMORY;
CREATE TABLE MY_TABLE2 LIKE MY_TABLE1;
```

The following information is copied to the new table from the existing table:

- Columns
- Column data types
- Default values
- NOT NULL constraint
- Primary key
- Table type (in-memory or disk-based table)

The following information is not copied:

- Data
- Unique and secondary keys
- Foreign keys
- Check constraints

To create tables like existing tables, you must have access rights to the existing table and rights to create tables to the current catalog and schema.

Examples

```
CREATE TABLE DEPT (DEPTNO INTEGER NOT NULL, DNAME VARCHAR, PRIMARY KEY(DEPTNO));
CREATE TABLE DEPT2 (DEPTNO INTEGER NOT NULL PRIMARY KEY, DNAME VARCHAR);
CREATE TABLE DEPT3 (DEPTNO INTEGER NOT NULL UNIQUE, DNAME VARCHAR);
CREATE TABLE DEPT4 (DEPTNO INTEGER NOT NULL, DNAME VARCHAR, UNIQUE(DEPTNO));
CREATE TABLE EMP (DEPTNO INTEGER, ENAME VARCHAR, FOREIGN KEY (DEPTNO)
REFERENCES DEPT (DEPTNO)) STORE DISK;
CREATE TABLE EMP2 (DEPTNO INTEGER, ENAME VARCHAR, CHECK (ENAME IS NOT NULL),
FOREIGN KEY (DEPTNO) REFERENCES DEPT (DEPTNO)) STORE MEMORY;
CREATE GLOBAL TEMPORARY TABLE T1 (C1 INT);
CREATE TRANSIENT TABLE T2 (C1 INT);
CREATE TABLE "053TEST" ("5DEPT" INTEGER NOT NULL UNIQUE, "6DEPT" VARCHAR);
```

Related information:

Appendix C, "Data types," on page 323

A.24 CREATE TRIGGER

```
CREATE TRIGGER trigger_name ON table_name time_of_operation
   triggering_event [REFERENCING column_reference]
   BEGIN trigger_body END
```

where:

```
trigger_name ::= literal
table_name ::= literal
time_of_operation ::= BEFORE | AFTER
triggering_event ::= INSERT | UPDATE | DELETE
column_reference ::= {OLD | NEW} column_name [AS] col_identifier
   [, REFERENCING column_reference ]

trigger_body ::=
   [ declare_statement; ... ]
   [ trigger_statement; ... ]

old_column_name ::= literal
new_column_name ::= literal
col_identifier ::= literal
```

Usage

A trigger provides a mechanism for executing a series of SQL statements when a particular action (an INSERT, UPDATE, or DELETE) occurs. The body of the trigger contains the SQL statement(s) that the user wants to execute. The body of the trigger is written using the stored procedure language (see A.15, “CREATE PROCEDURE,” on page 180).

You can create one or more triggers on a table, with each trigger defined to activate on a specific INSERT, UPDATE, or DELETE command. When a user modifies data within the table, the trigger that corresponds to the command is activated.

You can only use inline SQL or stored procedures with triggers. If you use a stored procedure in the trigger, the procedure must be created with the CREATE PROCEDURE command. A procedure invoked from a trigger body can invoke other triggers.

To create a trigger, you must be a DBA or owner of the table on which the trigger is being defined.

Triggers are created with the CREATE TRIGGER statement and dropped from the system catalog with the DROP TRIGGER statement. Triggers can also be disabled with the ALTER TRIGGER statement.

Example

```
"CREATE TRIGGER TRIGGER_BI ON TRIGGER_TEST
  BEFORE INSERT
  REFERENCING NEW BI AS NEW_BI
BEGIN
  EXEC SQL PREPARE BI INSERT INTO TRIGGER_OUTPUT VALUES (
    'BI', TRIG_NAME(0), TRIG_SCHEMA(0));
  EXEC SQL EXECUTE BI;
  SET NEW_BI = 'TRIGGER_BI';
END";
```

Related reference:

A.5, “ALTER TRIGGER,” on page 162

A.46, “DROP TRIGGER,” on page 228

Related information:

3.3.3, “Triggers and procedures,” on page 64

3.3.9, “Trigger example,” on page 76

A.24.1 trigger_name

trigger_name identifies the trigger. It can contain up to 254 characters.

A.24.2 BEFORE | AFTER clause

The BEFORE | AFTER clause specifies whether to execute the trigger before or after the invoking DML statement. In some circumstances, the BEFORE and AFTER clauses are interchangeable. However, there are some situations where one clause is preferred over the other.

- It is more efficient to use the BEFORE clause when performing data validation, such as domain constraint and referential integrity checking.

- When you use the AFTER clause, table rows which become available due to the invoking DML statement are processed. Conversely, the AFTER clause also confirms data deletion after the invoking DELETE statement.

You can define up to six triggers per table, one for each combination of action (INSERT, UPDATE, DELETE) and time (BEFORE and AFTER):

- BEFORE INSERT
- BEFORE UPDATE
- BEFORE DELETE
- AFTER INSERT
- AFTER UPDATE
- AFTER DELETE

The following example shows trigger trig01 defined BEFORE INSERT ON table1.

```
"CREATE TRIGGER TRIG01 ON table1
  BEFORE INSERT
  REFERENCING NEW COL1 AS NEW_COL1
BEGIN
  EXEC SQL PREPARE CUR1
    INSERT INTO T2 VALUES (?);
  EXEC SQL EXECUTE CUR1 USING (NEW_COL1);
  EXEC SQL CLOSE CUR1;
  EXEC SQL DROP CUR1;
END"
```

Following are examples (including implications and advantages) of using the BEFORE and AFTER clause of the CREATE TRIGGER command for each DML operation:

- UPDATE operation

The BEFORE clause can verify that modified data follows integrity constraint rules before processing the UPDATE. If the REFERENCING NEW AS *new_column_identifier clause* is used with the BEFORE UPDATE clause, then the updated values are available to the triggered SQL statements. In the trigger, you can set the default column values or derived column values before performing an UPDATE.

The AFTER clause can perform operations on newly modified data. For example, after a branch address update, the sales for the branch can be computed.

If the REFERENCING OLD AS *old_column_identifier clause* is used with the AFTER UPDATE clause, then the values that existed prior to the invoking update are accessible to the triggered SQL statements.

- INSERT operation

The BEFORE clause can verify that new data follows integrity constraint rules before performing an INSERT. Column values passed as parameters are visible to the triggered SQL statements but the inserted rows are not. In the trigger, you can set default column values or derived column values before performing an INSERT.

The AFTER clause can perform operations on newly inserted data. For example, after insertion of a sales order, the total order can be computed to see if a customer is eligible for a discount.

Column values are passed as parameters and inserted rows are visible to the triggered SQL statements.

- DELETE operation

The BEFORE clause can perform operations on data about to be deleted. Column values passed as parameters and inserted rows that are about to be deleted are visible to the triggered SQL statements.

The AFTER clause can be used to confirm the deletion of data. Column values passed as parameters are visible to the triggered SQL statements. The deleted rows are visible to the triggering SQL statement.

A.24.3 INSERT | UPDATE | DELETE clause

The INSERT | UPDATE | DELETE clause indicates the trigger action when a user action (INSERT, UPDATE, DELETE) is attempted.

Statements related to processing a trigger occur first before commits and autocommits from the invoking DML (INSERT, UPDATE, DELETE) statements on tables. If a trigger body or a procedure called within the trigger body attempts to execute a COMMIT or ROLLBACK, a solidDB server returns an appropriate run-time error.

INSERT specifies that the trigger is activated by an INSERT on the table. Loading n rows of data is considered as n inserts.

Note:

There may be some performance impact if you try to load the data with triggers enabled. Depending on your business need, you may want to disable the triggers before loading and enable them after loading. For details, see A.5, "ALTER TRIGGER," on page 162.

DELETE specifies that the trigger is activated by a DELETE on the table.

UPDATE specifies that the trigger is activated by an UPDATE on the table. The following rules apply when using the UPDATE clause:

- Within the REFERENCES clause of a trigger, a column may be referenced (aliased) no more than once in the BEFORE sub-clause and once in the AFTER sub-clause. Also, if the column is referenced in both the BEFORE and AFTER sub-clauses, the column's alias must be different in each sub-clause.
- The solidDB server allows for recursive update to the same table and does not prohibit recursive updates to the same row.

The solidDB server does not detect situations where the actions of different triggers cause the same data to be updated. For example, assume there are two update triggers (one that is a BEFORE trigger and one that is an AFTER trigger) on different columns, Col1 and Col2, of table Table1. When an update is attempted on all the columns of Table1, the two triggers are activated. Both triggers call stored procedures which update the same column, Col3, of a second table, Table2. The first trigger updates Table2.Col3 to 10 and the second trigger updates Table2.Col3 to 20.

Likewise, a solidDB server does not detect situations where the result of an UPDATE which activates a trigger conflicts with the actions of the trigger itself. For example, consider the following SQL statement:

```
UPDATE t1 SET c1 = 20 WHERE c3 = 10;
```

If the trigger is activated by this UPDATE then calls a procedure that contains the following SQL statement, the procedure overwrites the result of the UPDATE that activated the trigger:

```
UPDATE t1 SET c1 = 17 WHERE c1 = 20;
```

Note: The above example can lead to recursive trigger execution, which you should try to avoid.

A.24.4 table_name

The *table_name* is the name of the table on which the trigger is created. solidDB server allows you to drop a table that has dependent triggers defined on it. When you drop a table all dependent objects including triggers are dropped. However, you might still get runtime errors. For example, assume you create two tables A and B. If a procedure SP-B inserts data into table A, and table A is then dropped, a user will receive a run-time error if table B has a trigger which invokes SP-B.

A.24.5 trigger_body

The *trigger_body* contains the statement(s) to be executed when a trigger fires. The *trigger_body* definition is identical to the stored procedure definition. Read A.15, "CREATE PROCEDURE," on page 180 for details on creating a stored procedure body.

Note: It is valid syntactically to create a trigger with an empty body, although it is not useful.

A trigger body may also invoke any procedure registered with a solidDB server. solidDB procedure invocation rules follow standard procedure invocation practices.

You must explicitly check for business logic errors and raise an error.

A.24.6 REFERENCING clause

The REFERENCING clause is optional when creating a trigger on an INSERT/UPDATE/DELETE operation. It provides a way to reference the current column identifiers in the case of INSERT and DELETE operations, and both the old column identifier and the new updated column identifier by aliasing the column(s) on which an UPDATE operation occurs.

You must specify the OLD or NEW *column_identifier* to access it. A solidDB server does not provide access to the *column_identifier* unless you define it using the REFERENCING subclause.

{OLD | NEW} column_name AS col_identifier

The {OLD | NEW} *column_name AS col_identifier* subclause of the REFERENCING clause allows you to reference the values of columns both before and after an UPDATE operation. It produces a set of old and new column values which can be passed to a stored procedure; once passed, the procedure contains logic (for example, domain constraint checking) used to determine these parameter values.

Use the OLD AS clause to alias the table's old identifier as it exists before the UPDATE. Use the NEW AS clause to alias the table's new identifier as it exists after the UPDATE.

If you reference both the old and new values of the same column, you must use different *column_identifiers*.

Each column that is referenced as NEW or OLD should have a separate REFERENCING subclause.

The statement atomicity in a trigger is such that operations made in a trigger are visible to the subsequent SQL statements inside the trigger. For example, if you execute an INSERT statement in a trigger and then also perform a select in the same trigger, then the inserted row is visible.

In the case of AFTER trigger, an inserted row or an updated row is visible in the AFTER insert trigger, but a deleted row cannot be seen for a select performed within the trigger. In the case of a BEFORE trigger, an inserted or updated row is invisible within the trigger and a deleted row is visible. In the case of an UPDATE, the pre-update values are available in a BEFORE trigger.

The table below summarizes the statement atomicity in a trigger, indicating whether the row is visible to the SELECT statement in the trigger body.

Table 35. Statement atomicity in a trigger

Operation	BEFORE TRIGGER	AFTER TRIGGER
INSERT	row is invisible	row is visible
UPDATE	previous value is visible	new value is visible
DELETE	row is visible	row is invisible

A.24.7 Triggers usage notes and restrictions

- To use the stored procedure that a trigger calls, provide the catalog, schema/owner and name of the table on which the trigger is defined and specify whether to enable or disable the triggers on the table.
- To create a trigger on a table, you must have DBA authority or be the owner of the table on which the trigger is being defined.
- You can define, by default, up to one trigger for each combination of table, action (INSERT, UPDATE, DELETE) and time (BEFORE and AFTER). This means there can be a maximum of six triggers per table.

Note: The triggers are applied to each row. This means that if there are ten inserts, a trigger is executed ten times.

- You cannot define triggers on a view (even if the view is based on a single table).
- You cannot alter a table that has a trigger defined on it when the dependent columns are affected.
- You cannot create a trigger on a system table.
- You cannot execute triggers that reference dropped or altered objects. To prevent this error:
 - Recreate any referenced object that you drop.
 - Restore any referenced object you changed back to its original state (known by the trigger).
- You can use reserved words in trigger statements if they are enclosed in double quotation marks. For example, the following CREATE TRIGGER statement references a column named "data" which is a reserved word.

```
"CREATE TRIGGER TRIG1 ON TMPT BEFORE INSERT
REFERENCING NEW "DATA" AS NEW_DATA
BEGIN
END"
```


Related information:

3.3.3, “Triggers and procedures,” on page 64

3.3.4, “Triggers and transactions,” on page 67

A.25 CREATE USER

```
CREATE USER user_name IDENTIFIED {BY password | EXTERNALLY}
```

Usage

The CREATE USER statement creates a new user with a given password. You must have administrator privileges to execute the statement.

BY *password*

BY *password* defines that the user is authenticated internally: the password is stored in the SYS_USERS table in scrambled format.

EXTERNALLY

EXTERNALLY defines that the user is authenticated by services provided by the operating system. When using external authentication, the *username* must match the user name of an operating system user account.

Additionally, to use external authentication, you must enable the use of IBM Global Security Kit (GSKit). For more details, see Enabling encryption with IBM Global Security Kit (GSKit).

Note: When using external authentication on Windows systems, your environment might require you to provide the user name in the two-part format of *domain_name\username* or *username@domain_name*. You can specify the *domain_name* using the **General.DefaultDomainName** parameter. The solidDB server uses the value of the **General.DefaultDomainName** parameter to resolve the two-part user ID as expected by the Windows operating system.

Alternatively, you can leave the **General.DefaultDomainName** parameter empty (default) and provide the domain name as part of the user ID of each externally authenticated user.

When specifying the domain name as part of the user name, the following formats are supported:

```
domain_name\username  
username@domain_name
```

When using the CREATE USER *user_name* EXTERNALLY statement, the *user_name* string with \ or @ character must be given in double quotation marks.

For example:

```
CREATE USER "chicago_prod\solid1" IDENTIFIED EXTERNALLY  
CREATE USER "solid1@chicago_prod" IDENTIFIED EXTERNALLY
```

Examples

The following example creates a new, internally authenticated user named HOBBS with the password CALVIN.

```
CREATE USER HOBBS IDENTIFIED BY CALVIN;
```

The following example creates a new, externally authenticated user named solid1.


```
CREATE USER solid1 IDENTIFIED EXTERNALLY;
```

Related reference:

A.6, “ALTER USER,” on page 162

A.47, “DROP USER,” on page 229

A.26 CREATE VIEW

```
CREATE VIEW viewed_table_name [( column_identifier
    [,column_identifier ]... )]
AS query-specification
```

Usage

The CREATE VIEW statement creates a *view*. A view is like a virtual table that provides a different way of looking at the data in one or more tables; views are named specifications of a result table. The specification is a SELECT statement that is run whenever the view is referenced in an SQL statement. A view has columns and rows just like a table. All views can be used just like tables for data retrieval.

Example

```
CREATE VIEW TEST_VIEW
    (VIEW_I, VIEW_C, VIEW_ID)
AS SELECT I, C, ID FROM TEST;
```

Related reference:

A.48, “DROP VIEW,” on page 229

A.27 DELETE

```
DELETE FROM table_name [WHERE search_condition]
DELETE FROM table_name WHERE CURRENT OF cursor_name
```

Usage

There are two types of DELETE statements:

- The *searched* DELETE statement, DELETE FROM *table_name* [WHERE *search_condition*], deletes all rows in a given table or the rows specified with A.79.5, “search_condition,” on page 308.
- The *positioned* DELETE statement, DELETE FROM *table_name* WHERE CURRENT OF *cursor_name*, deletes the current row of the cursor.

table_name can be a table name or an alias.

Tip: If you want to delete all rows in a table, consider using the TRUNCATE statement.

Example of searched DELETE

```
DELETE FROM TEST WHERE ID = 5;
DELETE FROM TEST;
```

Example of positioned DELETE

```
DELETE FROM TEST WHERE CURRENT OF MY_CURSOR;
```

Related reference:

A.45, “DROP TABLE,” on page 228

A.74, “TRUNCATE TABLE,” on page 301

A.28 DESCRIBE

```
DESC[RIBE] [type option] {[context_specifier.]object} [format option]
```

where

```
type option := CATALOG | SCHEMA | TABLE | VIEW | PROCEDURE | FUNCTION | TRIGGER
```

```
context_specifier := catalog | [catalog.]schema | [[catalog.]schema.]table
```

```
object := catalog | schema | table | view | procedure | function | trigger
```

```
format option := [NICE (default) | RAW ]
```

Usage

The DESCRIBE statement describes a database object, such as catalog, schema, table, view, procedure, or function: it prints the column definitions for the specified table or view, or the specifications for the specified function or procedure.

The result of DESCRIBE depends of the output format, NICE or RAW.

- NICE format is brief and simplified. It can be used to make quick ad-hoc queries over the metadata of the database.
- RAW format is more specific and more verbose. It outputs information that would normally be acquired by using solidDB Data Dictionary (**solidd**) or complex SQL queries.

Both NICE and RAW formats are not available for all types.

- DESCRIBE CATALOG my_catalog prints the table name, the creator, the owner and the list of schemas of the specified catalog in NICE format.
- DESCRIBE SCHEMA my_schema prints the schema name, the catalog name to which it belongs, the creator and the owner in NICE format.
- DESCRIBE TABLE my_table prints the table name, the storage type used, list of parent tables, list of child tables, columns, column types, column precisions, column nullability, whether column is primary key and whether column is sec key in NICE format. In RAW format, the create statement is printed along with create statements of associated indexes and constraints.
- DESCRIBE VIEW my_view prints the columns of the view in NICE format and the create statement in RAW format.
- DESCRIBE PROCEDURE my_procedure prints the create statement in the RAW format.
- DESCRIBE FUNCTION my_function prints the create statement in the RAW format.
- DESCRIBE TRIGGER my_trigger prints the create statement in the RAW format.

Note: In many cases, there is only single output format available. In those cases, the only output available is produced despite the output format specified by the user.

Examples

The following example outputs information on a table named PERSON in NICE format.

```
DESCRIBE PERSON;
Table name : PERSON
Table type : In-memory
Memory usage: 7935 KB (total), 7925 KB (active), 6192 KB (rows), 1733 KB (indexes).

Parent tables : CITY
Child tables  : OWNER, IN_RELATION
```

COLUMN	TYPE	PRECISION	NULLABLE	PRIMARY KEY	SECONDARY KEY
FIRSTNAME	VARCHAR	30	YES	NO	NULL
LASTNAME	VARCHAR	60	NO	YES	NULL
HOME_CITY	INT	NULL	YES	NO	'CITY'

1 rows fetched.

The following example outputs information about a table named PERSON in RAW format.

```
DESCRIBE PERSON RAW;
```

```
CREATE TABLE "DBA"."DBA"."PERSON" (
  "FIRSTNAME" VARCHAR(30),
  "LASTNAME" VARCHAR(60) NOT NULL,
  "HOME_CITY" INTEGER,
  FOREIGN KEY("HOME_CITY")
    REFERENCES "DBA"."DBA"."CITY"("CITY"),
  PRIMARY KEY ("LASTNAME")
);
```

1 rows fetched.

See also

A.58, “LIST,” on page 245

A.29 DROP CACHE SEGMENT

DROP CACHE SEGMENT *name*

Usage

The DROP CACHE SEGMENT statement removes cache segments from the database. You can only drop segments that do not have tables or rows assigned to it. If the segment name does not exist or the segment has tables/rows assigned to it, the operation fails with error 13471. Before dropping a cache segment, all table assignments to the segment must be dropped.

Parameters

name Identifies the cache segment name.

Example

To drop the cache segment called HISTORY_CACHE:

1. Drop all table or row associations for the segment.
2. Use the following statement to remove the segment:

```
DROP CACHE SEGMENT HISTORY_CACHE
```

A.30 DROP CATALOG

DROP CATALOG *catalog_name* [CASCADE | RESTRICT]

Usage

The DROP CATALOG statement drops the specified catalog from the database.

If you do not specify either RESTRICT or CASCADE, all the database objects in the catalog must have been dropped prior to using this statement.

The RESTRICT keyword means that the catalog is dropped only if all database objects in the catalog have been dropped before.

The CASCADE keyword means that if the catalog contains database objects (such as tables), those will automatically be dropped. If you use the CASCADE keyword and the objects in other catalogs reference an object in the catalog being dropped, the references will automatically be resolved by dropping those referencing objects or updating them to eliminate the reference.

Only the creator of the database or users having SYS_ADMIN_ROLE have privileges to create or drop a catalog. Even the creator of a catalog cannot drop that catalog without SYS_ADMIN_ROLE privileges.

Examples

```
DROP CATALOG C1;  
DROP CATALOG C2 CASCADE;  
DROP CATALOG C3 RESTRICT;
```

Related reference:

A.10, “CREATE CATALOG,” on page 170

A.31 DROP EVENT

```
DROP EVENT event_name  
DROP EVENT [[catalog_name.]schema_name.]event_name
```

Usage

The DROP EVENT statement removes user-defined events from the database. System events cannot be dropped.

Examples

```
DROP EVENT event_test;
```

The following example uses the catalog, schema, and event name to specify the event to be dropped.

```
DROP EVENT HR_database.smith_schema.event1;
```

Related reference:

A.11, “CREATE EVENT,” on page 172

A.32 DROP FUNCTION

```
DROP FUNCTION [catalog-name[.schema].]function_name
```

Usage

The DROP FUNCTION statement removes the specified user-defined stored SQL function or external function from the database.

Related reference:

A.12, “CREATE FUNCTION,” on page 173

A.13, “CREATE FUNCTION (external),” on page 175

A.33 DROP INDEX

DITA

```
DROP INDEX index_name
DROP INDEX[[catalog_name.]schema_name.]index_name
```

Usage

The DROP INDEX statement removes the specified index from the database.

Primary key indexes cannot be removed.

Examples

```
DROP INDEX test_index;
-- Using catalog, schema, and index name
DROP INDEX bank_accounts.bankteller.first_name_index;
```

A.34 DROP MASTER

```
DROP MASTER master_name
```

Usage

The DROP MASTER statement drops the master database definitions from a replica database. After the drop, the replica cannot synchronize with the master database. The DROP master statement can only be used in the replica database.

If *master_name* is a reserved word, it must be enclosed in double quotation marks.

To use the DROP MASTER statement, autocommit must be set off.

Note: The preferred way to stop using the master database is to unregister the replica. The DROP MASTER statement should be used only when the MESSAGE APPEND UNREGISTER REPLICA statement cannot be executed. For details on unregistering a replica, see A.60.1, “MESSAGE APPEND,” on page 249.

Return values

Table 36. DROP MASTER return values

Error code	Description
13047	No privilege for operation
25007	Master <i>master_name</i> not found
25019	Database is not a replica database
25056	Autocommit not allowed
25065	Unfinished message <i>message_name</i> found for master <i>master_name</i>

Examples

```
DROP MASTER "MASTER";
DROP MASTER MY_MASTER;
```

A.35 DROP PROCEDURE

```
DROP PROCEDURE procedure_name
DROP PROCEDURE [[catalog_name.]schema_name.]procedure_name
```

Usage

The DROP PROCEDURE statement removes the specified procedure from the database.

Examples

```
DROP PROCEDURE PROCTEST;

DROP PROCEDURE telecomm_database.technician1.add_new_IP_address;
```

A.36 DROP PUBLICATION

```
DROP PUBLICATION publication_name
```

Usage

The DROP PUBLICATION statement drops a publication definition in the master database. All subscriptions to the dropped publication are dropped automatically as well.

Usage in master

When issued in the master database, the publication definition is removed from the master and replicas will not be able to refresh from it.

Usage in replica

There is no need to define the publications in the replicas, thus the DROP PUBLICATION statement is typically not issued on the replica. If you have created a publication on the replica, the DROP PUBLICATION removes the publication definition from the replica.

Return values

Table 37. DROP PUBLICATION return values

Error Code	Description
25010	Publication <i>publication_name</i> not found.
13111	Ambiguous entity name <i>name</i>

Example

```
DROP PUBLICATION customers_by_area;
```

Related reference:

A.17, "CREATE [OR REPLACE] PUBLICATION," on page 197

A.37 DROP PUBLICATION REGISTRATION

```
DROP PUBLICATION publication_name REGISTRATION
```

Usage

The DROP PUBLICATION REGISTRATION statement drops a registration for a publication in the replica database. The publication definition remains on the master database, but a user will be unable to refresh from the publication. All subscriptions to the registered publication are dropped automatically as well.

The DROP PUBLICATION REGISTRATION statement can only be issued on the replica database.

Return values

Table 38. DROP PUBLICATION REGISTRATION Return Values

Error Code	Description
13047	No privilege for operation
25019	Database is not a replica database
25025	Node name not defined
25071	Not registered to publication <i>publication_name</i>

Example

```
DROP PUBLICATION customers_by_area REGISTRATION;
```

A.38 DROP REMOTE SERVER

```
DROP REMOTE SERVER
```

Usage

The DROP REMOTE SERVER statement removes the back-end login information from the SYS_SERVER system table. The login data is used for SQL passthrough purposes in Universal Cache.

See also

A.18, “CREATE [OR REPLACE] REMOTE SERVER,” on page 201

A.3, “ALTER REMOTE SERVER,” on page 157

A.39 DROP REPLICA

```
DROP REPLICA replica_name
```

Usage

The DROP REPLICA statement drops a replica database from the master database. After this operation, the dropped replica cannot synchronize with the master database. The DROP REPLICA statement can only be issued on the master.

If *replica_name* is a reserved word, it must be enclosed in double quotation marks.

To use the DROP REPLICA statement, autocommit must be set off.

Note: The preferred way for stopping to use the replica is to unregister the replica. The DROP REPLICA statement is used only when the MESSAGE APPEND UNREGISTER REPLICA statement cannot be executed. For details on unregistering a replica, see A.60.1, “MESSAGE APPEND,” on page 249.

Return values

Table 39. DROP REPLICA return values

Error code	Description
13047	No privilege for operation
25009	Replica <i>replica_name</i> not found
25020	Database is not a master database
25056	Autocommit not allowed
25064	Unfinished message <i>message_name</i> found for replica <i>replica_name</i>

Examples

```
DROP REPLICA salesman_smith ;
DROP REPLICA "REPLICA";
```

A.40 DROP ROLE

```
DROP ROLE role_name
```

Usage

The DROP ROLE statement removes the specified role from the database.

You cannot drop the solidDB system roles (SYS_ADMIN_ROLE and so on).

Example

```
DROP ROLE GUEST_USERS;
```

A.41 DROP SCHEMA

```
DROP SCHEMA schema_name [CASCADE | RESTRICT]
DROP SCHEMA [catalog_name.] schema_name [CASCADE | RESTRICT]
```

Usage

The DROP SCHEMA statement drops the specified schema from the database.

If you do not specify the RESTRICT or CASCADE keyword, all the objects associated with the specified *schema_name* must have been dropped prior to using this statement.

The RESTRICT keyword means that the schema is dropped only if all objects associated with the specified *schema_name* have been dropped prior to using this statement.

The CASCADE keyword means that all the database objects (such as tables) within the specified schema will be dropped automatically. If you use the CASCADE keyword and the objects in other schemas reference an object in the schema being dropped, those references will automatically be resolved by dropping those referencing objects or updating them to eliminate the reference.

Examples

```
DROP SCHEMA finance;
DROP SCHEMA finance CASCADE;
DROP SCHEMA finance RESTRICT;
DROP SCHEMA forecasting_db.securities_schema CASCADE;
```

A.42 DROP SEQUENCE

```
DROP SEQUENCE sequence_name
DROP SEQUENCE [[catalog_name.]schema_name.]sequence_name
```

Usage

The DROP SEQUENCE statement removes the specified sequence from the database.

Example

```
DROP SEQUENCE SEQ1;
-- Using catalog, schema, and sequence name
DROP SEQUENCE bank_db.checking_acct_schema.account_num_seq;
```

A.43 DROP SUBSCRIPTION

In replica:

```
DROP SUBSCRIPTION publication_name [{(parameter_list) | ALL}]
    [COMMITBLOCK number_of_rows] [OPTIMISTIC | PESSIMISTIC]
```

In master:

```
DROP SUBSCRIPTION publication_name [{(parameter_list) | ALL}]
    REPLICA replica_name
```

Usage

The DROP SUBSCRIPTION statement is used to delete data that is no longer needed in a replica database by dropping the subscription that was used to retrieve the data from the master database. To issue the DROP SUBSCRIPTION statement, autocommit must be set off.

When issued on a master, the DROP SUBSCRIPTION statement removes a subscription for the specified replica.

When issued on a replica, the DROP SUBSCRIPTION statement removes a subscription from that replica.

By default, the data of a subscription is deleted in one transaction. If the amount of data is large, for example, tens of thousands of rows, it is recommended that the COMMITBLOCK be defined. When using the COMMITBLOCK option the data is deleted in more than one transaction. This ensures good performance for the operation.

In a replica, you can define that the DROP SUBSCRIPTION statement uses table-level pessimistic locking when it is initially executed. If the PESSIMISTIC mode is specified, all other concurrent access to the table affected is blocked until the drop has completed. If the optimistic mode is used, the DROP SUBSCRIPTION can fail due to a concurrency conflict.

A subscription can be dropped also from the master database. In this case, the replica name is included in the command. Names of all replicates that have been registered in the master database can be found in the SYS_SYNC_REPLICAS table. This operation deletes only the internal information about the subscription for this replica. The actual data in the replica is kept.

Dropping a subscription from the master is useful when a replica is no longer using the subscription and the replica has not dropped the subscription itself. Dropping old subscriptions releases old history data from the database. This history data is automatically deleted from the master databases after dropping the subscription.

If a replica's subscription has been dropped in the master database, the replica will receive the full data in the next refresh.

If a subscription is dropped in this case, DROP SUBSCRIPTION also drops the publication registration if the last subscription to the publication was dropped. Otherwise, registration must be dropped explicitly using the DROP PUBLICATION REGISTRATION statement or MESSAGE APPEND UNREGISTER PUBLICATION.

You can define the DROP SUBSCRIPTION statement to use table-level pessimistic locking when it is initially executed. If the PESSIMISTIC mode is specified, all other concurrent access to the tables affected is blocked until the import has completed. If the optimistic mode is used, the DROP SUBSCRIPTION can fail due to a concurrency conflict.

When a transaction acquires an exclusive lock to a table, the **General.TableLockWaitTimeout** parameter setting determines the transaction's wait period until the exclusive or shared lock is released.

Return values

Table 40. DROP SUBSCRIPTION return values

Error Code	Description
13047	No privileges for operation
25004	Dynamic parameters are not supported
25009	Replica <i>replica_name</i> not found
25010	Publication <i>publication_name</i> not found
25019	Database is not a replica database
25020	Database is not a master database
25041	Subscription to publication <i>publication_name</i> not found
25056	Autocommit not allowed

Example

```
Drop subscription from a master database
DROP SUBSCRIPTION customers_by_area('south')
FROM REPLICA salesman_joe
```

A.44 DROP SYNC BOOKMARK

```
DROP SYNC BOOKMARK bookmark_name
```

Usage

The DROP SYNC BOOKMARK statement drops a bookmark defined on a master database. Bookmarks are typically used when exporting data to a file. After a file is successfully imported to a replica from the master database, it is recommended that you drop the bookmark that you used to export the data to a file.

To drop a bookmark, you must have administrative privileges (SYS_ADMIN_ROLE or SYS_SYNC_ADMIN_ROLE). The DROP SYNC BOOKMARK statement cannot be used in a replica database.

If a bookmark remains, all subsequent changes to data on the master including deletes and updates are tracked on the master database for each bookmark to facilitate incremental refreshes.

If you do not drop bookmarks, the history information takes up disk space and unwanted disk I/O is incurred for each bookmark registered in the master database. This may result in performance degradation.

Note: Bookmarks should only be dropped after the exported data is imported into all intended replicas and after all the replicas have synchronized at least once. You should drop a bookmark only when you no longer have replicas to import and those replicas have refreshed once from the publication after the import.

When dropping bookmarks, solidDB uses the following rules to delete history records:

- Find the oldest REFRESH delivered to any replica on that table.
- Find the oldest bookmark.
- Determine which is older, the oldest REFRESH or the oldest bookmark.
- Delete all rows from history up to what is older, the oldest REFRESH or oldest bookmark.

Return values

Table 41. DROP SYNC BOOKMARK return values

Error Code	Description
25067	Synchronizer bookmark <i>bookmark_name</i> not found
13047	No privilege for operation

Examples

```
DROP SYNC BOOKMARK new_database;  
DROP SYNC BOOKMARK database_after_dataload;
```

A.45 DROP TABLE

```
DROP TABLE base_table_name [CASCADE [CONSTRAINTS]]  
DROP TABLE [[catalog_name.]schema_name.]table_name [CASCADE  
[CONSTRAINTS]]
```

Purpose

The DROP TABLE statement removes the specified table from the database.

[CASCADE [CONSTRAINTS]] means that all referential integrity constraints that refer to primary and unique keys in the dropped table are removed. This means that all the foreign keys specifications are removed from tables referencing the table being dropped. However, no column values are changed in the referencing tables.

Note: By default, objects are dropped with a RESTRICT type drop behavior, that is, as if the RESTRICT referential action was defined in all the related foreign keys.

However, there are exceptions such as the following:

- If your table has a synchronization history table, that synchronization history table will be dropped automatically.
- If a table has indexes on it, you do not need to drop the indexes first; they will be dropped automatically when the table is dropped.

Examples

```
DROP TABLE table1;  
-- Using catalog, schema, and table name  
DROP TABLE domains_db.demand_schema.bad_address_table;  
--remove foreign key constraints in referencing tables  
DROP TABLE table2 CASCADE CONSTRAINTS;
```

A.46 DROP TRIGGER

```
DROP TRIGGER [[catalog_name.]schema_name.]trigger_name
```

Usage

The DROP TRIGGER statement drops (or deletes) a trigger defined on a table from the system catalog. You must be the owner of a table or have DBA authority to delete a trigger from a table.

Examples: Dropping triggers

```
DROP TRIGGER update_acct_balance;  
-- Using schema and trigger name  
DROP TRIGGER savings_accounts.update_acct_balance;  
-- Using catalog, schema, and trigger name  
DROP TRIGGER accounts.savings_accounts.update_acct_balance;
```

Example: Dropping and re-creating a trigger

```
DROP TRIGGER TRIGGER_BI;  
COMMIT WORK;  
  
"CREATE TRIGGER TRIGGER_BI ON TRIGGER_TEST  
BEFORE INSERT
```

```

        REFERENCING NEW BI AS NEW_BI
BEGIN
    EXEC SQL PREPARE BI INSERT INTO TRIGGER_OUTPUT VALUES(
        'BI_NEW', TRIG_NAME(0), TRIG_SCHEMA(0));
    EXEC SQL EXECUTE BI;
    SET NEW_BI = 'TRIGGER_BI_NEW';
END";
COMMIT WORK;

INSERT INTO TRIGGER_TEST(XX) VALUES ('XX');
COMMIT WORK;

SELECT * FROM TRIGGER_TEST;
SELECT * FROM TRIGGER_OUTPUT;
COMMIT WORK;

```

A.47 DROP USER

```
DROP USER user_name
```

Purpose

The DROP USER statement removes the specified user from the database. All the objects associated with the specified *user_name* must be dropped prior to using this statement; the DROP USER statement is not a cascaded operation.

Example

```
DROP USER HOBBS;
```

A.48 DROP VIEW

```
DROP VIEW view_name
DROP VIEW [[catalog_name.]schema_name.]view_name
```

Usage

The DROP VIEW statement removes the specified view from the database.

Examples

```

DROP VIEW sum_of_acct_balances;
-- Using schema and view name
DROP VIEW acct_manager_schema.sum_of_acct_balances;
-- Using catalog, schema, and view name
DROP VIEW account_db.acct_manager_schema.sum_of_acct_balances;

```

A.49 EXPLAIN PLAN FOR

```
EXPLAIN PLAN FOR sql_statement
```

Usage

The EXPLAIN PLAN FOR statement shows the execution plan that the SQL optimizer has selected for a given SQL statement.

An execution plan is a series of primitive operations, and an ordering of these operations, that solidDB performs to execute the statement. Each operation in the execution plan is called a unit.

Table 42. EXPLAIN PLAN FOR units

Unit	Description
JOIN UNIT*	Join unit joins two or more tables. The join can be done by using loop join or merge join.
TABLE UNIT	The table unit is used to fetch the data rows from a table or index.
ORDER UNIT	Order unit is used to order rows for grouping or to satisfy ORDER BY. The ordering can be done in memory or using an external disk sorter.
GROUP UNIT	Group unit is used to do grouping and aggregate calculation (SUM, MIN, and so on).
UNION UNIT*	Union unit performs the UNION operation. The unit can be done by using loop join or merge join.
INTERSECT UNIT*	Intersect unit performs the INTERSECT operation. The unit can be done by using loop join or merge join.
EXCEPT UNIT*	Except unit performs the EXCEPT operation. The unit can be done by using loop join or merge join.

*This unit is generated also for queries that reference only a single table. In that case no join is executed in the unit; it simply passes the rows without manipulating them.

The table returned by the EXPLAIN PLAN FOR statement contains the following columns.

Table 43. EXPLAIN PLAN FOR table columns

Column Name	Description
ID	The output row number, used only to guarantee that the rows are unique.
UNIT_ID	This is the internal unit id in the SQL interpreter. Each unit has a different id. The unit id is a sparse sequence of numbers, because the SQL interpreter generates unit ids also for those units that are removed during the optimization phase. If more than one row has the same unit id it means that those rows belong to the same unit. For formatting reasons the info from one unit may be divided into several different rows.
PAR_ID	Parent unit id for the unit. The parent id number refers to the id in the UNIT_ID column.
JOIN_PATH	For join, union, intersect, and except units there is a join path which specifies which tables are joined in the unit and the join order for tables. The join path number refers to the unit id in the UNIT_ID column. It means that the input to the unit comes from that unit. The order in which the tables are joined is the order in which the join path is listed. The first listed table is the outermost table in a loop join.
UNIT_TYPE	Unit type is the execution graph unit type.

Table 43. EXPLAIN PLAN FOR table columns (continued)

Column Name	Description
INFO	Info column is reserved for additional information. It may contain, for example, index usage, the database table name and constraints used in solidDB to select rows. Note that the constraints listed here may not match those constraints given in the SQL statement.

Example

```
EXPLAIN PLAN FOR select * from tables;
```

A.50 EXPORT SUBSCRIPTION

```
EXPORT SUBSCRIPTION publication_name [(publication_parameters)]
  TO 'filename'
  USING BOOKMARK bookmark_name;
  [WITH [NO] DATA];
```

where

publication_name and *bookmark_name* are identifiers that must exist in the database.

filename represents a literal value enclosed in single quotation marks. You can export several publications to a single file by specifying the same file name.

WITH DATA option creates a replica with data, for example, when data is exported to an existing database that does not contain master data and requires a subset of data.

WITH NO DATA option creates a replica with only the metadata, for example, when a subscription is imported to a database that already contains the data (for example, using a backup copy of an existing master).

By default, the export file is created using the WITH DATA option and includes all rows. If there is more than one publication specified, the exported file can have a combination of WITH DATA and WITH NO DATA.

Usage

The EXPORT SUBSCRIPTION statement exports a version of data from a master database to a file. You can then use the IMPORT statement to import the data in the file into a replica database. The EXPORT SUBSCRIPTION statement is not available in a replica database.

There are several uses for the EXPORT SUBSCRIPTION statement. For example:

- Create a large replica database (greater than 2 GB) from an existing master. This procedure requires that you export a subscription with or without data to a file first, then import the subscription to the replica.
- Export specific versions of the data to a replica. For performance reasons, you could choose to export the data instead of using the MESSAGE APPEND REFRESH statement to send the data to a replica.
- Export metadata information without the actual row data.

You could create a replica that already contains existing data and only needs the schema and version information associated with a publication.

Unlike the MESSAGE APPEND REFRESH statement where a REFRESH is requested from a replica, you request an export directly from the master database. The export output is saved to a user-specified file rather than a solidDB message.

The EXPORT SUBSCRIPTION statement exports publication from the master database using REFRESH synchronization. If the publication uses input parameters, they are included in the REFRESH. The statement is based on a given bookmark, which means that export data is consistent up to this bookmark. When you export data, the EXPORT SUBSCRIPTION statement includes all rows as in a full publication up to the bookmark. However, since export is based on a given bookmark, the subsequent REFRESHes are incremental.

If a bookmark is created in a master for the purpose of exporting and importing data, the bookmark must exist when:

- The EXPORT SUBSCRIPTION statement is executed on the master database. If the bookmark does not exist at this point, error message 25067 is generated, indicating the bookmark cannot be found.
- The IMPORT statement is executed on all intended replicas and replicas receive their first set of data ("REFRESH").

During a file import, a connection to a master database is not needed and the existence of the bookmark is not checked. However, if the bookmark does not exist at the time a replica receives its first REFRESH, the REFRESH fails with error message 25067 and the import data is unusable. The remedy is to create a new bookmark on the master, re-export the data, and re-import the data.

For more details, see *Exporting and importing subscriptions* in the *IBM solidDB Advanced Replication User Guide*.

Usage rules

The following rules apply to the EXPORT SUBSCRIPTION statement:

- Only one file per subscription is allowed for export. You can use the same file name to include multiple subscriptions on the same file.
- The file size of an export file depends on the underlying operating system. If the respective platform (such as SUN, or HP) allows more than 2 GB, you can write files greater than 2 GB. This means that replica (recipient) should also have a compatible platform and file system. Otherwise, the replica is not able to accept the export file. If both the operating system on a master and replica support file sizes greater than 2 GB, export files greater than 2 GB are permitted.
- The export file can contain more than one subscription. Subscriptions can be exported using either the WITH DATA or WITH NO DATA options. An exported file with multiple subscriptions can have a combination of WITH DATA and WITH NO DATA included.
- When a subscription is exported to a file using the WITH NO DATA option, only metadata (that is, schema and version information corresponding to that publication) is exported to the file.
- Autocommit must be set off when using the EXPORT SUBSCRIPTION statement.

Return values

Table 44. EXPORT SUBSCRIPTION return values

Error Code	Description
25056	Autocommit not allowed
25067	Bookmark is not found.
25068	Export file <i>file_name</i> open failure.
25010	Publication <i>name</i> not found.

Example

```
EXPORT SUBSCRIPTION FINANCE_PUBLICATION(2004) TO 'FINANCE.EXP'  
USING BOOKMARK BOOKMARK_FOR_FINANCE_DB WITH NO DATA;
```

Related reference:

A.17, “CREATE [OR REPLACE] PUBLICATION,” on page 197

A.22, “CREATE SYNC BOOKMARK,” on page 206

A.51 EXPORT SUBSCRIPTION TO REPLICA

```
EXPORT SUBSCRIPTION publication_name [(publication_parameters)]  
TO REPLICA replica_node_name  
USING BOOKMARK bookmark_name  
[COMMITBLOCK number_of_rows]
```

where

publication_name and *bookmark_name* are identifiers that must exist in the database.

COMMIT BLOCK keywords specify how many rows of the exported data are committed in the replica database in one transaction.

Usage

The EXPORT SUBSCRIPTION TO REPLICA statement exports large volumes of data specified by a publication from master database to a replica database. After the EXPORT operation has completed, the replica can use MESSAGE APPEND REFRESH statement to refresh the data of the subscription in an incremental manner. The EXPORT SUBSCRIPTION TO REPLICA statement is not available in a replica database.

Because the EXPORT SUBSCRIPTION TO REPLICA statement does not use the disk-based advanced replication MESSAGEs to send data from master to replica, it provides a significantly more efficient way to send large volumes of data from master to replica as the usage of disk during the operation is minimized.

Publication data is exported from the master database as a REFRESH with a set of input parameter values (if they are used in the publication).

The EXPORT SUBSCRIPTION TO REPLICA statement is based on a given bookmark, which means that export data is consistent up to this bookmark. When you export data, the EXPORT SUBSCRIPTION statement includes all rows as in a

full publication up to the bookmark. However, since export is based on a given bookmark, the subsequent REFRESHes are incremental.

If a bookmark is created in a master for the purpose of exporting data, then the bookmark must exist when the EXPORT SUBSCRIPTION statement is executed on the master database. If the bookmark does not exist at this point, error message 25067 is generated, indicating the bookmark cannot be found.

Specifying a commit block when a large number of rows are to be exported improves the performance of the operation. However, it is recommended that the replica database is not used by other applications when export operation with commit block is active.

Return values

Table 45. EXPORT SUBSCRIPTION TO REPLICA return values

Error Code	Description
25056	Autocommit not allowed
25067	Bookmark is not found.
25010	Publication <i>name</i> not found.

Example

```
EXPORT SUBSCRIPTION FINANCE_PUBLICATION(2004) TO REPLICA replica_1
USING BOOKMARK BOOKMARK_FOR_FINANCE_DB COMMITBLOCK 10000;
```

Related reference:

A.17, "CREATE [OR REPLACE] PUBLICATION," on page 197

A.22, "CREATE SYNC BOOKMARK," on page 206

A.52 GRANT

```
GRANT {ALL | grant_privilege [, grant_privilege]...}
      ON table_name
TO {PUBLIC | user_name [, user_name]... |
   role_name [, role_name]... }
[WITH GRANT OPTION]

GRANT role_name TO user_name

grant_privilege ::= DELETE | INSERT | SELECT |
                 UPDATE [( column_identifier [, column_identifier]... )] |
                 REFERENCES [( column_identifier [, column_identifier]... )]

GRANT EXECUTE ON function_name
  TO {PUBLIC | user_name [, user_name]... | role_name [, role_name]... }

GRANT EXECUTE ON procedure_name
  TO {PUBLIC | user_name [, user_name]... | role_name [, role_name]... }

GRANT {SELECT | INSERT} ON event_name
  TO {PUBLIC | user_name [, user_name]... | role_name [, role_name]... }

GRANT {SELECT | UPDATE} ON sequence_name
  TO {PUBLIC | user_name [, user_name]... | role_name [, role_name]... }
```

Usage

The GRANT statement is used to:

- grant privileges to the specified user or role
- grant privileges to the specified user by giving the user the privileges of the specified role

When you grant a role to a user, the role may be one that you have created, or it may be a system-defined role, such as SYS_SYNC_ADMIN_ROLE or SYS_ADMIN_ROLE.

If you use the optional WITH GRANT OPTION, the user who receives the privilege can grant the privilege to other users.

Examples

```
GRANT GUEST_USERS TO CALVIN;  
GRANT INSERT, DELETE ON TEST TO GUEST_USERS;
```

Related reference:

A.64, “REVOKE,” on page 273

Related information:

4.2, “Managing user privileges and roles,” on page 90

A.53 GRANT PASSTHROUGH

```
GRANT PASSTHROUGH {READ | WRITE}  
TO {PUBLIC | user_name [, user_name]... | role_name [, role_name]... }  
[WITH GRANT OPTION]
```

Supported in

This command requires the use of the SQL passthrough feature.

Usage

The GRANT PASSTHROUGH statement grants access rights for SQL passthrough.

Examples

```
GRANT PASSTHROUGH READ TO cdcuser1, cdcuser2  
GRANT PASSTHROUGH WRITE TO cdcuser1, cdcuser2
```

See also

A.65, “REVOKE PASSTHROUGH,” on page 274

A.54 GRANT REFRESH

```
GRANT REFRESH ON publication_name TO {PUBLIC |  
user_name,  
[ user_name ] ... | role_name , [ role_name ] ...}
```

Usage

This GRANT REFRESH statement grants access rights on a publication to a user or role defined in the master database. The GRANT REFRESH statement applies only to advanced replication configurations and it can only be executed on the master database.

Note: The keyword REFRESH can be used to be equivalent to the keyword SUBSCRIBE; however, the keyword SUBSCRIBE has been deprecated in the GRANT statement.

Return values

Table 46. GRANT REFRESH return values

Error Code	Description
13137	Illegal grant/revoke mode
13048	No grant option privilege
25010	Publication <i>name</i> not found

Examples

```
GRANT REFRESH ON customers_by_area TO salesman_jones;  
GRANT REFRESH ON customers_by_area TO all_salesmen;
```

A.55 HINT

```
--(* vendor (SOLID), product (Engine), option(hint)  
--hint *)--
```

```
hint::=  
[MERGE JOIN |  
TRIPLE MERGE JOIN |  
LOOP JOIN |  
JOIN ORDER FIXED |  
INTERNAL SORT |  
EXTERNAL SORT |  
INDEX [REVERSE] table_name.index_name |  
PRIMARY KEY [REVERSE] table_name |  
FULL SCAN table_name |  
[NO] SORT BEFORE GROUP BY |  
UNION FOR OR |  
OR FOR OR |  
LOOP FOR OR]
```

Usage

Hints are used to specify directives to the SQL optimizer for determining the query execution plan that is to be used. Hints are specified through embedded pseudo comments within query statements. The optimizer detects these directives or hints and bases its query execution plan accordingly. Optimizer hints allow applications to be optimized under various conditions to the data, query type, and the database. They not only provide solutions to performance problems occasionally encountered with queries, but shift control of response times from the system to the user.

Hints are available for:

- Selecting merge or nested loop join
- Using a fixed join order as given in the from list
- Selecting internal or external sort
- Selecting a particular index
- Selecting a table scan over an index scan
- Selecting sorting before or after grouping

You can place hints in SQL statements as static strings, just after a SELECT, UPDATE, or DELETE keyword. Hints are not allowed after the INSERT keyword.

Pseudo comment identifier

The pseudo comment prefix is followed by identifying information. You must specify the vendor as SOLID, product as Engine, and the option, which is the pseudo comment class name, as a valid hint.

The terminator may be on its own line, or it may be at the end of the last line of the hint. For example, either of the following is acceptable:

```
--(* vendor (Solid), product (Engine), option(hint)
--hint
-- *)--
```

or

```
--(* vendor (Solid), product (Engine), option(hint)
--hint *)--
```

Spacing is sensitive. In the pseudo comment prefix `--(* and postfix *)--`, there can be no space between the parenthesis and the asterisk. There must be a space prior to the `*)--` terminator, that is, prior to the asterisk (see the examples above). No space is required prior to the opening parenthesis in `--(*`. The terminator `*)--` cannot be on a line by itself without being after the comment delimiter `--`.

Table name resolution in optimizer hints

Table name resolution in optimizer hints is the same as in any table name in an SQL statement. This means that if there is a table alias name in the query, you must use the alias, not the table name, in the optimizer hints. For example:

```
SELECT
  --(* vendor(SOLID), product(Engine), option(hint)
  -- FULL SCAN emp_alias *)--
  emp_alias.emp_id, employee_name, dependent_name
FROM employee_table AS emp_alias LEFT OUTER JOIN dependent_table
AS dep_alias
  ON (dep_alias.emp_id = emp_alias.emp_id)
ORDER BY emp_alias.emp_id;
```

If you specify the table name when you should have specified the alias name, you will get the following error message:

```
102: Unused optimizer hint.
```

If you are not using an alias and you are using a table that is in another schema and/or another catalog, then make sure that in the hint you precede the table name with the schema and/or catalog name. For example:

```

SELECT
  --(* vendor(SOLID), product(Engine), option(hint)
  -- FULL SCAN sally_schema.employee_table *)--
  emp_id, employee_name
FROM sally_schema.employee_table;

```

Using hints with ODBC or JDBC

If you are using hints and you compose a query as a string and then submit that string using ODBC or JDBC, you must ensure that appropriate newline characters are embedded within that string to mark the end of the comments. Otherwise, you will get a syntax error. If you do not include any newlines, then all of the statement after the start of the first comment will look like a comment. For example, suppose that your code looks like the following:

```
strcpy(s, "SELECT --(* hint... *)-- col_name FROM table;");
```

Everything after the first "--" looks like a comment, and therefore your statement looks incomplete. You must do something like the following:

```
strcpy(s, "SELECT --(* hint... *)-- \n col_name FROM table;");
```

The embedded newline `\n` character terminates the comment. A useful technique for debugging is to print out the strings to make sure that they look correct. They should look like:

```
SELECT --(* hint ... *)--
column_name FROM table_name...;
```

or

```
SELECT --(* hint ... *)--
column_name FROM table_name...;
```

Subselects and multiple hints

Each subselect requires its own hint; for example, the following are valid uses of hints syntax:

```

INSERT INTO ... SELECT hint FROM ...
UPDATE hint TABLE ... WHERE column = (SELECT hint ... FROM ...)
DELETE hint TABLE ... WHERE column = (SELECT hint ... FROM ...)

```

When specifying multiple hints in one pseudo comment, you must separate the hints with commas as shown in the following examples:

Example 1:

```

SELECT
  --(* vendor(SOLID), product(Engine), option(hint)
  --MERGE JOIN
  --JOIN ORDER FIXED *)--
  *
FROM TAB1 A, TAB2 B;
WHERE A.INTF = B.INTF;

```

Example 2:

```

SELECT
  --(* vendor(SOLID), product(Engine), option(hint)
  --INDEX TAB1.INDEX1
  --INDEX TAB1.INDEX1 FULL SCAN TAB2 *)--
  *
FROM TAB1, TAB2
WHERE TAB1.INTF = TAB2.INTF;

```

Enabling and disabling hints

Hints are enabled and disabled using the **SQL.EnableHints** parameter. By default, hints are enabled (**SQL.EnableHints=yes**).

Examples

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
--MERGE JOIN
--JOIN ORDER FIXED
-- *)--
col1, col2
FROM TAB1 A, TAB2 B;
WHERE A.INTF = B.INTF;
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
--INDEX TAB1.INDEX1
--INDEX TAB1.INDEX1 FULL SCAN TAB2
-- *)--
*
FROM TAB1, TAB2
WHERE TAB1.INTF = TAB2.INTF;
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- INDEX TAB1.IDX1 *)--
* FROM TAB1 WHERE I > 100
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- INDEX MyCatalog.mySchema.TAB1.IDX1 *)--
* FROM TAB1 WHERE I > 100
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- JOIN ORDER FIXED *)--
* FROM TAB1, TAB2 WHERE TAB1.I >= TAB2.I
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- LOOP JOIN *)--
* FROM TAB1, TAB2 WHERE TAB1.I >= TAB2.I
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- INDEX REVERSE MyCatalog.mySchema.TAB1.IDX1 *)--
* FROM TAB1 WHERE I > 100
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- SORT BEFORE GROUP BY *)--
AVG(I) FROM TAB1 WHERE I > 10 GROUP BY I2
```

```
SELECT
--(* vendor(SOLID), product(Engine), option(hint)
-- INTERNAL SORT *)--
* FROM TAB1 WHERE I > 10 ORDER BY I2
```

A.55.1 Hints supported by solidDB

Table 47. Supported hints

Hint	Definition
MERGE JOIN	<p>Directs the optimizer to choose the merge join access plan in a SELECT query for all tables listed in the FROM clause. The MERGE JOIN option is used when two tables are approximately equal in size and the data is distributed equally. It is faster than a LOOP JOIN when an equal amount of rows are joined. For joining data, MERGE JOIN supports a maximum of three tables. The joining table is ordered by joining columns and combining the results of the columns.</p> <p>You can use this hint when the data is sorted by a join key and the nested loop join performance is not adequate. The optimizer selects the merge join only where there is an equal predicate between tables (for example, "table1.col1 = table2.col1"). Otherwise, the optimizer selects LOOP JOIN even if the MERGE JOIN hint is specified.</p> <p>When data is not sorted before performing the merge operation, the solidDB query executor sorts the data.</p> <p>Keep in mind that the merge join with a sort is more resource intensive than the merge join without the sort.</p>
TRIPLE MERGE JOIN	<p>TRIPLE MERGE JOIN is a variant of MERGE JOIN. It has three table sources which are merged on equal basis instead of the two in MERGE JOIN. The TRIPLE MERGE JOIN hint instructs the SQL interpreter to use the triple merge join algorithm whenever possible. The triple merge join algorithm can only be used in situations where in all three table sources there is one single field that should be equal in all the resulting rows after evaluating the WHERE condition.</p>
LOOP JOIN	<p>Directs the Optimizer to pick the nested loop join in a select query for all tables listed in the FROM clause. By default, the Optimizer does not pick the nested loop join.</p> <p>The LOOP JOIN loops through both inner and outer tables to find matches between columns in the inner and outer tables. For better performance, the joining columns should be indexed.</p> <p>Using the loop join when tables are small and fit in memory may offer greater efficiency than using other join algorithms.</p>
JOIN ORDER FIXED	<p>Specifies that the optimizer use tables in a join in the order listed in the FROM clause of the query. This means that the optimizer does not attempt to rearrange the join order and does not try to find alternate access paths to complete the join.</p> <p>Tip: You should test the hint by running the EXPLAIN PLAN output to ensure that the generated plan is optimal for the given query.</p>

Table 47. Supported hints (continued)

Hint	Definition
INTERNAL SORT	<p>Specifies that the query executor use the internal sorter. Use this hint if the expected resultset is small (hundreds of rows as opposed to thousands of rows); for example, if you are performing some aggregates, ORDER BY with small resultsets, or GROUP BY with small resultsets, and so on.</p> <p>This hint avoids the use of the more expensive external sorter.</p>
EXTERNAL SORT	<p>Specifies that the query executor use the external sorter. Use this hint when the expected resultset is large and does not fit in memory; for example, if the expected result set has thousands of rows.</p>
INDEX [REVERSE] <i>table_name.index_name</i>	<p>Forces a given index scan for a given table. In this case, the optimizer does not proceed to evaluate if there are any other indexes that can be used to build the access plan or whether a table scan is better for the given query. Tip: You should test the hint by running the EXPLAIN PLAN output to ensure that the generated plan is optimal for the given query.</p> <p>The optional keyword REVERSE returns the rows in the reverse order. In this case, the query executor begins with the last page of the index and starts returning the rows in the descending (reverse) key order of the index.</p> <p>In <i>tablename.indexname</i>, the <i>tablename</i> is a fully qualified table name which can include the <i>catalogname</i> and <i>schemaname</i>.</p>
PRIMARY KEY [REVERSE] <i>table_name</i>	<p>Forces a primary key scan for a given table.</p> <p>The optional keyword REVERSE returns the rows in the reverse order.</p> <p>If the primary key is not available for the given table, you will receive a runtime error.</p>
FULL SCAN <i>table_name</i>	<p>Forces a table scan for a given table. In this case, the optimizer does not proceed to evaluate if there are any other indexes that can be used to build the access plan or whether a table scan is better for the given query. Tip: You should test the hint by running the EXPLAIN PLAN output to ensure that the generated plan is optimal for the given query.</p>
[NO] SORT BEFORE GROUP BY	<p>Indicates whether the SORT operation occurs before the resultset is grouped by the GROUP BY columns.</p> <p>If the grouped items are few (hundreds of rows), use NO SORT BEFORE. On the other hand, if the grouped items are large (thousands of rows), use SORT BEFORE.</p>

Table 47. Supported hints (continued)

Hint	Definition
UNION FOR OR	<p>The UNION FOR OR hint instructs the SQL interpreter to replace an OR condition of style $A = 1 \text{ OR } A = 2$ with a construction of the following type:</p> <pre>SELECT ... WHERE A = 1 UNION ALL SELECT ... WHERE A = 2</pre> <p>In most cases the SQL interpreter performs the replacement automatically; the UNION FOR OR hint ensures the UNION-type execution is used always. Note: Conditions of type $A = 1 \text{ OR } B = 2$ can also be handled, but this may be problematic since the conditions are not mutually exclusive. Because of this, the construction for $A = 1 \text{ OR } B = 2$ is the following:</p> <pre>SELECT ... WHERE A = 1 UNION ALL SELECT ... WHERE B = 2 AND UtoT NOT (A = 1)</pre> <p>where UtoT stands for UNKNOWN TO TRUE.</p> <p>The UtoT operator is needed for handling cases with NULL values. Without the UtoT operator, a row which has values $A = \text{NULL}$ and $B = 2$ would not appear correctly in the UNION variant.</p>
OR FOR OR	<p>The OR FOR OR hint is the opposite for UNION FOR OR. It prevents the interpreter from using the UNION-type solution.</p>
LOOP FOR OR	<p>The LOOP FOR OR hint is an alternative query execution plan that falls between UNION FOR OR and OR FOR OR. With LOOP FOR OR the OR values are passed individually to the data table level, but conditions like $A = 1 \text{ OR } B = 2$ cannot be handled (see description of UNION FOR OR for details on $A = 1 \text{ OR } B = 2$).</p>

A.56 IMPORT

```
IMPORT 'file_name' [COMMITBLOCK number_of_rows]
[{'OPTIMISTIC' | PESSIMISTIC}]
```

Usage

This IMPORT statement imports data to a replica database from a data file created by the EXPORT SUBSCRIPTION statement. The IMPORT statement can only be issued on the replica database.

The *file_name* represents a literal value enclosed in single quotation marks. The import command can accept a single filename only. Therefore, all the data to be imported to a replica must be contained in one file.

The COMMITBLOCK option indicates the number of rows processed before the data is committed. The *number_of_rows* is an integer value used with the optional

COMMITBLOCK clause to indicate the commitblock size. The use of the COMMITBLOCK option improves the performance of the import and releases the internal transaction resources frequently.

The optimal value for the COMMITBLOCK size varies depending on various resources at the server. A good example is a COMMITBLOCK size of 1000 for 10,000 rows. If you do not specify the COMMITBLOCK option, the IMPORT command uses all rows in the publication as one transaction. This might work well for a small number of rows, but is problematic for thousands and millions of rows.

You can define the import to use table-level pessimistic locking when it is initially executed. If the PESSIMISTIC mode is specified, all other concurrent access to the table affected is blocked until the import has completed. Otherwise, if the optimistic mode is used, the IMPORT may fail due to a concurrency conflict.

When a transaction acquires an exclusive lock to a table, the **General.TableLockWaitTimeout** parameter setting determines the transaction's wait period until the exclusive or shared lock is released.

Imported data is not valid in a replica until it is refreshed once after the import. At the time a replica makes its first REFRESH, the bookmark used to export the file must exist in the master database. If it does not exist, the REFRESH statement fails. This means that you are required to create a new bookmark on the master database, re-export the data, and re-import the data on the replica.

The following rules apply when using the IMPORT statement:

- Only one file per subscription is allowed for import.
- The file size of an export file is dependent upon the underlying operating system. If a respective platform (such as SUN, or HP) allows more than 2 GB, you can write files greater than 2 GB. This means that a replica (recipient) should also have a compatible platform and file system. Otherwise, the replica is not able to accept the export file. If both the operating system on a master and replica support file sizes greater than 2 GB, export files greater than 2 GB are permitted.
- Back up replica databases before using the IMPORT command. If a COMMITBLOCK option is used and fails, the imported data is only partially committed; you will need to restore the replica with a backup file.
- To use the IMPORT statement, autocommit must be set OFF.

Return values

Table 48. IMPORT return values

Error Code	Description
25007	Master <i>master_name</i> not found.
25019	Database is not a replica database.
25069	Import file <i>file_name</i> open failure.
13XXX	Table level error

Table 48. *IMPORT* return values (continued)

Error Code	Description
13124	User id <i>num</i> not found This message is generated, for example, if the user has been dropped.
10006	Concurrency conflict (simultaneous other operation)
13047	No privilege for operation
13056	Insert not allowed for pseudo column
21XXX	Communication error
25024	Master not defined
25026	Not a valid master user
25031	Transaction is active, operation failed
25036	Publication <i>publication_name</i> not found or publication version mismatch
25040	User id <i>user_id</i> is not found While executing a message reply an attempt to map a master user to a local replica id failed.
25041	Subscription to publication <i>publication_name</i> not found
25048	Publication <i>publication_name</i> request info not found
25054	Table <i>table_name</i> is not set for synchronization history
25056	Autocommit not allowed
25060	Column <i>column_name</i> does not exist on publication <i>publication_name</i> resultset in table <i>table_name</i>

Example

```
IMPORT 'FINANCE.EXP';
```

A.57 INSERT

```
INSERT INTO table_name insert_columns_and_source
```

```
insert_columns_and_source::=  
from_subquery  
| from_constructor  
| from_default
```

```
from_subquery ::=  
[insert_column_name_list] query expression
```

```
insert_column_name_list ::=  
([column name [, column name]... ])
```

```

from constructor ::=
[insert_column_name_list] VALUES row_constructor[, row_constructor]... ]

row_constructor ::= ([insert_item[, insert_item]...])

insert_item ::= insert_value | DEFAULT | NULL

from default ::= DEFAULT VALUES

```

Usage

The INSERT statement inserts rows into a table.

There are several variations of the INSERT statement. In the simplest instance, a value is provided for each column of the new row in the order specified at the time the table was defined (or altered). In the preferable form of the INSERT statement, the columns are specified as part of the statement and they do not need to be in any specific order as long as the orders of the column list matches the order of the value list.

insert_value can be a literal, a scalar function, or a variable in a procedure.

Examples

```

INSERT INTO TEST (C, ID) VALUES (0.22, 5);
INSERT INTO TEST VALUES (0.35, 9);

```

Multirow inserts can also be done. For example, to insert three rows in one statement, you can use the following command:

```

INSERT INTO employees VALUES
(10021, 'Peter', 'Humlaut'),
(10543, 'John', 'Wilson'),
(10556, 'Bunba', '01o');

```

You can insert default values by using the DEFAULT VALUES statement as shown in the second example below. An equivalent form is "INSERT INTO TEST() VALUES()". You can also assign a specific value for one column and use the default value for another column. These methods as shown in the examples below:

```

INSERT INTO TEST () VALUES ();
INSERT INTO TEST DEFAULT VALUES;
INSERT INTO TEST (C, ID) VALUES (0.35, DEFAULT);
INSERT INTO TEST (C, ID) SELECT A, B FROM INPUT_TO_TEST;

```

A.58 LIST

```
LIST <statement>
```

```
where
statement :=
```

```
CATALOGS |
USERS |
ROLES |
```

```
[catalog.]SCHEMAS |
[catalog.]MASTERS |
[catalog.]REPLICAS |
[catalog.]SUBSCRIPTIONS |
[catalog.]PUBLICATIONS |
[[catalog.]schema.]TABLES |
[[catalog.]schema.]VIEWS |
```

```

[[catalog.]schema.]PROCEDURES |
[[catalog.]schema.]FUNCTIONS |
[[catalog.]schema.]EVENTS |
[[catalog.]schema.]SEQUENCES |
[[[catalog.]schema.]table.]INDEXES |
[[[catalog.]schema.]table.]TRIGGERS

```

Usage

The LIST statement can be used to view the existing database objects in a solidDB database. The LIST statement prints a list of specific database object types.

The object type is given as an argument. Any object name can be replaced with a wildcard expression (for example, foo%b_r). The wildcard '%' extends the scope to every possible object on that particular level.

The LIST statement can be used in companion with the DESCRIBE command as follows:

```

LIST my_catalog.my_schema.PROCEDURES
DESCRIBE my_catalog.my_schema.certain_procedure

```

Note: LIST PROCEDURES lists all user-defined SQL and external procedures. System procedures are not listed. LIST FUNCTIONS lists all user-defined SQL and external functions.

Examples

```
LIST my_schema.TABLES;
```

```

Catalog: my_catalog
Schema: my_schema
TABLES:
-----
CITY
PERSON

```

1 rows fetched.

```
LIST my_catalog.%.TABLES;
```

```

Catalog: my_catalog
Schema: my_schema
TABLES:
-----
PERSON
IN_RELATION
CITY
CAR
Schema: another_schema
TABLES:
-----
HORSENAME

```

1 rows fetched.

See also

A.28, "DESCRIBE," on page 218

A.59 LOCK TABLE

```
LOCK lock-definition [lock-definition] [wait-option]  
lock-definition ::= TABLE tablename [,tablename]  
IN { SHARED | [LONG] EXCLUSIVE } MODE  
wait-option ::= NOWAIT | WAIT <#seconds>
```

where

- **tablename:** The name of the table to lock. You can also specify the catalog and schema of the table by qualifying the table name. You may only lock tables, not views.
- **SHARED:** Shared mode allows others to perform read and write operations on the table. DDL operations are not allowed. Also, shared mode prohibits others from issuing an EXCLUSIVE lock on the same table.
- **EXCLUSIVE:** If a D-table uses pessimistic locking, then an exclusive lock prevents other user from accessing the table in any way (for example, inserting or deleting data, DDL operations, acquiring a lock), except for SELECT statements. If the case of M-tables (always pessimistic) and optimistic D-tables, an exclusive lock allows other users to perform SELECT and SELECT FOR UPDATE statements on the locked table but prohibits any other activity (inserting or deleting data, DDL operations, acquiring a lock) on that table.
- **LONG:** By default, locks are released at the end of a transaction. If the LONG option is specified, then the lock is not released when the locking transaction commits. If the locking transaction aborts or is rolled back, then all locks, including LONG locks, are released. You must unlock long locks explicitly using the UNLOCK command. LONG duration locks are allowed only in EXCLUSIVE mode. LONG shared locks are not supported.
- **NOWAIT:** Specifies that control is returned to you immediately even if any specified table is locked by another user. If requested lock is not granted, an error is returned.
- **WAIT:** Specifies a timeout (in seconds) for how long system should wait to get requested locks. If requested lock is not granted within that time, an error is returned.

Note: The WAIT option is effective on disk-based tables only.

Usage

The LOCK and UNLOCK commands allow you to manually lock and unlock tables. Putting a lock on a table limits access to that object. The LONG option allows you to extend the duration of a manual exclusive lock past the end of the current transaction; in other words, you can keep the table exclusively locked through a series of multiple transactions.

Manual locking is not needed very often. The server's automatic locking is usually sufficient. For a detailed discussion of locking in general, and the server's automatic locking in particular, see 5.2, "Concurrency control and locking," on page 111.

Explicit locking of tables is primarily intended to help database administrators execute maintenance operations in a database without being disturbed by other users. For example, manual locking is typically used in advanced replication setups when making schema changes. For more details, see *Upgrading the schema of a distributed system* in the *IBM solidDB Advanced Replication User Guide*.

Table locks can be either SHARED or EXCLUSIVE.

An EXCLUSIVE lock on a table prohibits any other user or connection from changing the table or any records within the table. If you have an exclusive lock on a table, then other users/connections cannot do any of the following on that table until your exclusive lock is released:

- INSERT, UPDATE, DELETE
- ALTER TABLE
- DROP TABLE
- LOCK TABLE (in shared or exclusive mode)

Furthermore, if a D-table uses pessimistic locking, then an exclusive lock also prevents others users/connections from doing the following:

- SELECT FOR UPDATE

Exclusive locks do not prevent other users from SELECTing records from that table. Most other database servers behave differently – they do not allow SELECTs on a table that is locked exclusively.

A SHARED lock is less restrictive than an exclusive lock. If you have a shared lock on a table, then other users/connections cannot do any of the following on that table until your shared lock is released:

- ALTER TABLE
- DROP TABLE
- LOCK TABLE (in exclusive mode)

If you use a shared lock on a table, other users/connections may insert, update, delete, and select from the table.

Shared locks on a table are somewhat different from shared locks on a record. If you have a shared lock on a record, then no other user may change data in the record. If you have a shared lock on a table, then other users may still change data in that table.

More than one user at a time may have shared locks on a table. If you have a shared lock on the table, other users may also get shared locks on the table. However, no user may get an exclusive lock on a table when another user has a shared lock (or exclusive lock) on that table.

The LOCK command takes effect at the time it is executed. If you do not use the LONG option, then the lock will be released at the end of the transaction. If you use the LONG option, the table lock lasts until you explicitly unlock the table. The table lock will also be released if you roll back the transaction in which the lock was placed – LONG locks only persist across transactions if you commit the transaction in which you placed the LONG lock.

The LOCK/UNLOCK TABLE commands apply only to tables. There is no command to manually lock or unlock individual records within a table.

You can lock more than one table and specify different modes within one LOCK command. If the lock command fails, then none of the tables are locked. If the lock command was successful, then all requested locks are granted.

If the user does not specify a wait option (NOWAIT or WAIT seconds), then the default wait time is used. That is the same as the deadlock detection timeout. The WAIT option is effective on disk-based tables only.

To use the LOCK TABLE command to issue a lock on a table, you must have insert, delete or update privileges on that table. There is no GRANT command to give other users LOCK and UNLOCK privileges on a table.

Return values

Error code	Description
10014	Resource is locked.
13047	No privilege for operation.
13011	Table <tablename> is not found.

Examples

```
LOCK TABLE emp IN SHARED MODE;
LOCK TABLE emp IN SHARED MODE TABLE dept IN EXCLUSIVE MODE;
LOCK TABLE emp,dept IN SHARED MODE NOWAIT;
LOCK TABLE emp IN LONG EXCLUSIVE MODE;
```

Related reference:

A.75, “UNLOCK TABLE,” on page 302

Related information:

“SET SYNC MODE” on page 290

A.60 MESSAGE

The MESSAGE statements are used to create and execute the synchronization messages that are sent between a replica and the master database in an advanced replication setup.

A.60.1 MESSAGE APPEND

```
MESSAGE unique_message_name APPEND
[
  PROPAGATE TRANSACTIONS
  [ { IGNORE_ERRORS | LOG_ERRORS | FAIL_ERRORS } ]
  [WHERE { property_name {=|<|<=|>|=|<>} 'value_string' | ALL } ]
]
[ { REFRESH | SUBSCRIBE }
publication_name[(publication_parameters)]
timeout[(timeout_in_seconds)]
[FULL]
]
[REGISTER PUBLICATION publication_name]
[UNREGISTER PUBLICATION publication_name]
[REGISTER REPLICA]
[UNREGISTER REPLICA]
[SYNC_CONFIG ('sync_config_arg')]

```

Usage

The MESSAGE APPEND statement appends tasks to a message that has been created in the replica database with the MESSAGE BEGIN statement. The tasks include the following:

- Propagate transactions to the master database
- Refresh a publication from the master database
- Register or unregister a publication for replica subscription
- Register or unregister a replica database to the master
- Download master user information (list of user names and passwords) from the master database

The MESSAGE APPEND statement applies only to advanced replication configurations and it can only be issued on the replica databases.

The PROPAGATE TRANSACTIONS task may contain a WHERE clause that is used to propagate only those transactions where a transaction property defined with the SAVE PROPERTY statement meets specific criteria.

Using the keyword ALL overrides any default propagation condition set earlier with the following statement:

```
SAVE DEFAULT PROPAGATE PROPERTY WHERE property_name {=|<|<=|>|=|<>} 'value'
```

This enables you to propagate transactions that do not contain any properties.

The REGISTER REPLICA task adds a new replica database to the list of replicas in the master database. Replicas must be registered with the master database before any other synchronization functions can be performed in the replica database.

Synchronizing each master database to the replica in a multimaster environment requires registration of a replica to each master database by setting up catalogs. One replica catalog can register only to one master catalog. This statement performs the actual registration once catalogs are created in a synchronization environment. For synchronization to the replica, a new catalog for each master database is required.

Note:

- A single-master environment does not require the use of catalogs. By default, when catalogs are not used, registration of the replica occurs automatically with a base catalog that maps to a master base catalog, whose name is given when the database is created.
- A single replica *node* may have multiple masters, but only if the node has a separate replica *catalog* for each master catalog.
- A single replica *catalog* may not have multiple masters.

The UNREGISTER REPLICA option removes an existing replica database from the list of replicas in the master database.

The REFRESH task may contain arguments to the publication (if used in the publication). The parameters must be literals; you cannot use stored procedure variables, for example. Using keyword FULL with REFRESH forces fetching full data to the replica. The publication requested must be registered. The keywords REFRESH and SUBSCRIBE are synonymous; however, SUBSCRIBE is deprecated in the MESSAGE APPEND statement.

The REGISTER PUBLICATION task registers a publication in the replica so that it can be refreshed from. Users can only refresh from publications that are registered. In this way, publication parameters are validated, preventing users from accidentally subscribing to unwanted subscriptions or requesting ad hoc subscriptions. All tables that the registered publication refers to must exist in the replica.

The UNREGISTER PUBLICATION option removes an existing registered publication from the list of registered publications in the master database.

The input argument of the SYNC_CONFIG task defines the search pattern of the user names that are returned from the master database to the replica. SQL wildcards (such as the symbol %) that follow the conventions of the LIKE keyword are used in this argument with a *match_string*, which is a character string. For details on using the LIKE keyword, see A.81, “Wildcard characters,” on page 311.

Return values

Table 49. MESSAGE APPEND return values

Error Code	Description
13133	Not a valid license for this product
25004	Dynamic parameters are not supported
25005	Message <i>message_name</i> is already active.
25006	Message <i>message_name</i> not active
25015	Syntax error: <i>error_message</i> , line <i>line_number</i>
25018	Illegal message state. An appending message in the replica must be placed between the MESSAGE BEGIN and MESSAGE END statements.
25024	Master not defined
25025	Node name not defined
25026	Not a valid master user
25028	Message <i>message_name</i> can include only one system subscription
25035	Message <i>message_name</i> is in use. A user is currently creating or forwarding this message.
25044	SYNC_CONFIG system publication takes only character arguments
25056	Autocommit not allowed
25071	Not registered to publication <i>publication_name</i>
25072	Already registered to publication <i>publication_name</i>

Examples

```
MESSAGE MyMsg0001 APPEND PROPAGATE TRANSACTIONS;  
MESSAGE MyMsg0001 APPEND REFRESH PUB_CUSTOMERS_BY_AREA('SOUTH');  
MESSAGE MyMsg0001 APPEND REGISTER REPLICA;  
MESSAGE MyMsg0001 APPEND SYNC_CONFIG ('S%');  
MESSAGE MyMsg0001 APPEND REGISTER PUBLICATION pub1_customer;
```

A.60.2 MESSAGE BEGIN

```
MESSAGE unique_message_name BEGIN [TO master_node_name]
```

Usage

The MESSAGE BEGIN statement is used to explicitly begin each message that is sent from a replica to the master database. The MESSAGE BEGIN statement is applicable only in advanced replication configurations. It can only be issued on the replica databases.

Each message must have a name that is unique within a replica. To construct unique message names, you may use the GET_UNIQUE_STRING() function, which is documented in B.1, “String functions,” on page 313. After a message has been processed, that message name may be reused. However, if the message fails for any reason, the master will keep a copy of the failed message, and if you try to reuse the message name before you delete the failed message, then the name will not be unique. You may want to use a new message name even in situations where you might be able to re-use an existing name. Note that it is possible for two replicas of the same master to have the same message name.

When registering a replica to a master catalog, other than the master system catalog, you must provide the master node name in the MESSAGE BEGIN command. The master node name is used to resolve the correct catalog at the master database. Note that specifying a master node name only applies when using the REGISTER REPLICA statement. Later messages are automatically sent to the correct master node.

If you use the optional "TO *master_node_name*" clause, then you must put double quotation marks around the *master_node_name*.

Note:

When working with messages, be sure the autocommit mode is always switched off.

Return values from replica

Table 50. MESSAGE BEGIN return values from replica

Error Code	Description
25005	Message <i>message_name</i> is already active. A message of the specified name was created and appears to still be active. The message is automatically deleted when the reply of the message has been successfully executed in the replica.
25035	Message <i>message_name</i> is in use. A user is currently creating or forwarding this message.

Table 50. MESSAGE BEGIN return values from replica (continued)

Error Code	Description
25056	Autocommit not allowed

Return values from master

Table 51. MESSAGE BEGIN return values from master

Error Code	Description
25019	Database is not a replica database.
25025	Node name not defined.
25056	Autocommit not allowed

Examples

```
MESSAGE MyMsg0001 BEGIN ;
MESSAGE MyMsg0002 BEGIN TO "BerkeleyMaster";
```

A.60.3 MESSAGE DELETE

```
MESSAGE message_name [FROM REPLICA replica_name] DELETE
```

Usage

The MESSAGE DELETE statement deletes a message from the database to recover from an error. For example, the execution of a message can be terminated because of an error. The MESSAGE DELETE statement is applicable only in advanced replication configurations. It can be issued both on the master and replica databases.

When issued on the master database, the node name of the replica database that forwarded the message must be specified with the FROM REPLICA *replica_name* clause.

Note: The current transaction and all subsequent transactions that were propagated to the master in this message are permanently lost when the message is deleted.

When deleting messages, the autocommit mode must always be switched off.

To use this statement, you must have SYS_SYNC_ADMIN_ROLE access.

Note: As an alternative, the MESSAGE DELETE CURRENT TRANSACTION command can provide better recovery because it lets you delete only the offending transaction.

Return values

Table 52. MESSAGE DELETE return values from replica

Error code	Description
25005	Message <i>message_name</i> is already active
25013	Message <i>message_name</i> not found
25035	Message <i>message_name</i> is in use. A user is currently creating or forwarding this message.
25056	Autocommit not allowed

Table 53. MESSAGE DELETE return values from master

Error code	Description
13047	No privilege for operation
25009	Replica <i>replica_name</i> not found
25013	Message <i>message_name</i> not found
25020	Database is not a master database
25035	Message <i>message_name</i> is in use. A user is currently executing this message.
25056	Autocommit not allowed

Examples

```
MESSAGE MyMsg0000 DELETE;  
MESSAGE MyMsg0001 FROM REPLICA bills_laptop DELETE;
```

Related reference:

A.60.3, "MESSAGE DELETE," on page 253

A.60.4 MESSAGE DELETE CURRENT TRANSACTION

```
MESSAGE message_name FROM REPLICA replica_name  
DELETE CURRENT TRANSACTION
```

Usage

The MESSAGE DELETE CURRENT TRANSACTION statement deletes the current transaction from a given message in the master database.

The MESSAGE DELETE CURRENT TRANSACTION statement is applicable only in advanced replication configurations. It can only be issued on the master database.

To use this statement, you must have SYS_SYNC_ADMIN_ROLE privilege.

If a DBMS level error such as a duplicate insert occurs during the execution, the execution of a message stops. This kind of error can be resolved by deleting the offending transaction from the message. Once deleted with the MESSAGE FROM REPLICA DELETE CURRENT TRANSACTION, an administrator can proceed with the synchronization process.

When deleting the current transaction, the autocommit mode must be switched off.

This statement is used only when the message is in an error state; if used otherwise, an error message is returned. This statement is a transactional operation and must be committed before message execution can continue. To restart the message after the deletion is committed, use the following statement:

```
MESSAGE msgname FROM REPLICA replicaname EXECUTE
```

The deletion is completed before the MESSAGE FROM REPLICA EXECUTE statement is executed; that is, the statement starts the message from replica, but waits until the active statement is completed before actually executing the message. Thus the statement performs asynchronous message execution.

Note: Delete a transaction only as a last resort; normally transactions should be written to prevent unresolved conflicts in a master database. MESSAGE FROM REPLICA DELETE CURRENT TRANSACTION is intended for use in the development phase, when unresolved conflicts occur more frequently.

Use caution when deleting a transaction. Because subsequent transactions can be dependent on the results of a deleted transaction, the result might be more transaction errors.

Return values

Table 54. MESSAGE DELETE CURRENT TRANSACTION return values

Error code	Description
13047	No privilege for operation
25009	Replica <i>replica_name</i> not found
25013	Message name <i>message_name</i> not found
25018	Illegal message state. An attempt was made to delete a transaction from a message that is not in error.
25056	Autocommit not allowed

Examples

```
MESSAGE somefailures FROM REPLICA laptop1 DELETE
CURRENT TRANSACTION;
COMMIT WORK;
MESSAGE somefailures FROM REPLICA laptop1 EXECUTE;
COMMIT WORK;
```

Related reference:

A.60.3, "MESSAGE DELETE," on page 253

A.60.9, "MESSAGE FROM REPLICA EXECUTE," on page 264

A.60.5 MESSAGE END

MESSAGE *unique_message_name* END

Usage

The MESSAGE END statement closes a message in a replica database. Closing a message means that you can no longer append anything to it. A message must be "wrapped up" and made persistent before it can be sent to the master database. You must commit the transaction make the message persistent.

The MESSAGE END statement is applicable only in advanced replication configurations. It can be issued only on replica databases.

Note: When working with messages, the autocommit mode must be switched off.

Return values

Table 55. MESSAGE END return values from replica

Error Code	Description
13133	Not a valid license for this product
25005	Message <i>message_name</i> is already active
25013	Message <i>message_name</i> not found
25018	Illegal message state. The MESSAGE BEGIN statement must exist to begin a transaction and the MESSAGE END statement can be executed only once per message.
25026	Not a valid master user
25035	Message <i>message_name</i> is in use A user is currently creating or forwarding this message.
25056	Autocommit not allowed

Table 56. MESSAGE END return values from master

Error Code	Description
25019	Database is not a replica database
25056	Autocommit not allowed

Examples

```
MESSAGE MyMsg001 END;  
COMMIT WORK;
```

The following example shows a complete message that propagates transactions and refreshes from publication PUB_CUSTOMERS_BY_AREA.


```

MESSAGE MyMsg001 BEGIN;
MESSAGE MyMsg001 APPEND PROPAGATE TRANSACTIONS;
MESSAGE MyMsg001 APPEND REFRESH PUB_CUSTOMERS_BY_AREA('△SOUTH');
MESSAGE MyMsg001 END;
COMMIT WORK;

```

A.60.6 MESSAGE EXECUTE

```
MESSAGE message_name EXECUTE [{OPTIMISTIC | PESSIMISTIC}]
```

Usage

The MESSAGE EXECUTE statement allows a message to be re-executed on the replica, if the execution of a reply message fails in a replica. The execution of a reply message can fail, for example, if the database server detects a concurrency conflict between a REFRESH and an ongoing user transaction. The MESSAGE EXECUTE statement is applicable only in advanced replication configurations. It can be issued only on replica databases. On the master database, use the MESSAGE FROM REPLICA EXECUTE statement.

In the PESSIMISTIC mode, all other concurrent access to the table affected is blocked until the synchronization message has completed. If you anticipate concurrency conflicts to happen often and the re-execution of the message fails because of a concurrency conflict, you can execute the message using the PESSIMISTIC option for table-level locking; this ensures the message execution is successful.

In the OPTIMISTIC mode, the MESSAGE EXECUTE statement can fail due to a concurrency conflict.

When a transaction acquires an exclusive lock to a table, the **General.TableLockWaitTimeout** parameter determines the wait period of the transaction after which the exclusive or shared lock is released.

Note: When working with messages, the autocommit mode must be switched off.

Result set

The MESSAGE EXECUTE statements returns a result set. The returned result set is the same as with command MESSAGE GET REPLY.

The columns of the result set are as follows:

Table 57. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set

Column Name	Description
Partno	Message part number

Table 57. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set (continued)

Column Name	Description
Type	The type of result set row. Possible types are: 0: Message part start 1: This type is not in use 2: The message was a propagation message and the status of that operation is stored in the return message 3: Task 4: Subscription task 5: Type of refresh (FULL or INCREMENTAL) 6: MESSAGE DELETE status
Masterid	Master ID
Msgid	Message ID
Errcode	Message error code. Zero if successful.
Errstr	Message error string. NULL is successful.
Insertcount	Number of inserted rows to replica. Type=3: Total number of insert Type=4: Row inserts restored from replica history to replica base table Type=5: Insert operations received from master
Deletecount	Type = 3: Total number of deletes Type = 4: Row deletes restored from replica base table Type = 5: Delete operations received from master
Bytecount	Size of message in bytes. Indicated in result received from command MESSAGE END. Otherwise 0.
Info	Information of the current task. Type = 0: then Message name Type = 3: Publication name Type = 4: Table name Type = 5: FULL/INCREMENTAL

Return values

Table 58. MESSAGE EXECUTE return values

Error code	Description
13XXX	Table level error
10006	Concurrency conflict (simultaneous other operation)
13047	No privilege for operation
13056	Insert not allowed for pseudo column
25005	Message <i>message_name</i> is already active
25013	Message name <i>message_name</i> not found
25018	Illegal message state
25024	Master not defined
25026	Not a valid master user
25031	Transaction is active, operation failed
25035	Message <i>message_name</i> is in use A user is currently creating or forwarding this message.
25040	User id <i>user_id</i> is not found While executing a message reply an attempt to map a master user to a local replica id failed.
25041	Subscription to publication <i>publication_name</i> not found
25048	Publication <i>publication_name</i> request info not found
25056	Autocommit not allowed

Example

```
MESSAGE MyMsg0002 EXECUTE;
```

Related reference:

A.60.9, "MESSAGE FROM REPLICA EXECUTE," on page 264

A.60.7 MESSAGE FORWARD

```
MESSAGE unique_message_name FORWARD  
[TO {'connect_string' | node_name | "node_name"} ]  
[TIMEOUT {number_of_seconds | FOREVER} ]  
[COMMITBLOCK block_size_in_rows]  
[{'OPTIMISTIC' | PESSIMISTIC}]
```

Usage

After a message has been completed and made persistent with the MESSAGE END statement, it can be sent to the master database using the MESSAGE FORWARD

statement. The MESSAGE FORWARD statement is applicable only in advanced replication configurations. It can be issued only on the replica databases.

It is only necessary to specify the recipient of the message with keyword TO when a new replica is being registered with the master database; that is, when the first message from a replica to the master server is sent.

The *connect_string* must be a valid connect string, such as:

```
tcp [host_computer_name] server_port_number
```

In the context of a MESSAGE FORWARD command, a connect string must be delimited in single quotation marks.

The *node_name* (without quotation marks) is a valid alphanumeric sequence that is not a reserved word. The "*node_name*" (in double quotation marks) is used if the node name is a reserved word; in this case, the double quotation marks ensure that the node name is treated as a delimited identifier. For example, since the word "master" is a reserved word, the word is placed in double quotation marks when it is used as a node name:

```
-- On master
SET SYNC NODE "master";
--On replica
MESSAGE refresh_severe_bugs2 FORWARD TO "master" TIMEOUT FOREVER;
```

Each sent message has a reply message. The TIMEOUT property defines how long the replica server will wait for the reply message.

If a TIMEOUT is not defined, the message is forwarded to the master and the replica does not fetch the reply. In this case the reply can be retrieved with a separate MESSAGE GET REPLY call.

If the reply of the sent message contains REFRESHes of large publications, the size of the REFRESH commit block, that is, the number of rows that are committed in one transaction, can be defined using the COMMITBLOCK property. This has a positive impact on the performance of the replica database. It is recommended that there are no on-line users accessing the database when the COMMITBLOCK property is being used.

As part of the MESSAGE FORWARD operation, you can specify table-level pessimistic locking when the reply message is initially executed in the replica. If the PESSIMISTIC mode is specified, all other concurrent access to the table affected is blocked until the synchronization message has completed. If the optimistic mode is used, the MESSAGE FORWARD operation can fail due to a concurrency conflict.

When a transaction acquires an exclusive lock to a table, the **General.TableLockWaitTimeout** parameter setting determines the wait period of the transaction after which the exclusive or shared lock is released.

If a forwarded message fails in delivery due to a communication error, you must explicitly use the MESSAGE FORWARD to resend the message. Once re-sent, MESSAGE FORWARD re-executes the message.

Note:

When working with the messages, the autocommit mode must be switched off.

Result set

If the MESSAGE FORWARD also retrieves the reply, the statement returns a result set. The result set returned is the same as the one returned with the statements MESSAGE GET REPLY and MESSAGE.

The columns of the result set are as follows:

Table 59. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set

Column Name	Description
Partno	Message part number
Type	The type of result set row. Possible types are: 0: Message part start 1: This type is not in use 2: The message was a propagation message and the status of that operation is stored in the return message 3: Task 4: Subscription task 5: Type of refresh (FULL or INCREMENTAL) 6: MESSAGE DELETE status
Masterid	Master ID
Msgid	Message ID
Errcode	Message error code. Zero if successful.
Errstr	Message error string. NULL is successful.
Insertcount	Number of inserted rows to replica. Type=3: Total number of insert Type=4: Row inserts restored from replica history to replica base table Type=5: Insert operations received from master
Deletecount	Type = 3: Total number of deletes Type = 4: Row deletes restored from replica base table Type = 5: Delete operations received from master
Bytecount	Size of message in bytes. Indicated in result received from command MESSAGE END. Otherwise 0.

Table 59. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set (continued)

Column Name	Description
Info	Information of the current task. Type = 0: then Message name Type = 3: Publication name Type = 4: Table name Type = 5: FULL/INCREMENTAL

Return values

Table 60. MESSAGE FORWARD return values from replica

Error Code	Description
13XXX	Table level error
21XXX	Communication error
10006	Concurrency conflict (simultaneous other operation)
13047	No privilege for operation
13056	Insert not allowed for pseudo column
25005	Message <i>message_name</i> is already active
25013	Message name <i>message_name</i> not found
25018	Illegal message state In the replica, the message can only be executed using the MESSAGE FORWARD statement if the message is ended and the ending transaction is committed.

Table 60. MESSAGE FORWARD return values from replica (continued)

Error Code	Description
25024	<p>Master not defined</p> <p>This message is produced if double quotation marks, rather than single quotation marks, are used around the <i>connect_string</i> in a MESSAGE FORWARD statement.</p> <p>For example, if the master node is given the node name "master" (which is a reserved word and therefore should be delimited by double quotation marks), and if that node's connect string is:</p> <p>tcp localhost 1315</p> <p>then the MESSAGE statements shown below are correct:</p> <pre>--On the replica ... --double quotation marks MESSAGE msg1 BEGIN TO "master"; ... --single quotation marks MESSAGE msg2 FORWARD TO 'tcp localhost 1315';</pre> <p>Note that the MESSAGE BEGIN statement defines (within the replica server) what the node name of the master is. The MESSAGE FORWARD statement may contain the connect string to the server.</p>
25026	Not a valid master user
25031	Transaction is active, operation failed
25035	<p>Message <i>message_name</i> is in use.</p> <p>A user is currently creating or forwarding this message.</p>
25040	<p>User id <i>user_id</i> is not found.</p> <p>While executing a message reply an attempt to map a master user to a local replica id failed.</p>
25041	Subscription to publication <i>publication_name</i> not found
25048	Publication <i>publication_name</i> request info not found
25052	Failed to set node name to <i>node_name</i> .
25054	Table <i>table_name</i> is not set for synchronization history
25055	<p>Connect information is allowed only when not registered</p> <p>The connect info in MESSAGE <i>message_name</i> FORWARD TO <i>connect_info options</i> is allowed only if the replica has not yet been registered to the master database.</p>
25056	Autocommit not allowed
25057	The replica database has already been registered to a master database

Table 60. MESSAGE FORWARD return values from replica (continued)

Error Code	Description
25060	Column <i>column_name</i> does not exist on publication <i>publication_name</i> resultset in table <i>table_name</i>

Table 61. MESSAGE FORWARD return values from master

Error code	Description
13XXX	Table level error
13124	User id <i>num</i> not found This message is generated, for example, if the user has been dropped.
25016	Message not found, replica ID <i>replica_id</i> , message ID <i>message_id</i>
25056	Autocommit not allowed

Examples

Forward message, wait for the reply for 60 seconds:

```
MESSAGE MyMsg001 FORWARD TIMEOUT 60 ;
```

Forward message to a master server that runs on the "mastermachine.acme.com" machine. Do not wait for the reply message.

```
MESSAGE MyRegistrationMsg FORWARD TO  
'tcp mastermachine.acme.com 1313';
```

Forward message, wait for the reply for 5 minutes (300 seconds) and commit the data of the refreshed publications to replica database in transactions of maximum of 1000 rows.

```
MESSAGE MyMsg001 FORWARD TIMEOUT 300 COMMITBLOCK 1000 ;
```

A.60.8 MESSAGE FROM REPLICA DELETE

```
MESSAGE msgid FROM REPLICA replicaname DELETE;  
MESSAGE msgid FROM REPLICA replicaname DELETE CURRENT TRANSACTION;
```

This command can only be executed on the master.

A.60.9 MESSAGE FROM REPLICA EXECUTE

```
MESSAGE message_name FROM REPLICA replica_name EXECUTE
```

Usage

The MESSAGE FROM REPLICA EXECUTE statement executes a failed message on the master database.

The MESSAGE FROM REPLICA EXECUTE statement is applicable only in advanced replication configurations. It can only be issued on the master database. On replica databases, use the MESSAGE EXECUTE statement.

The execution of a message stops, for example, if a DBMS level error such as a duplicate insert occurs during the execution. Also the execution of a message can stop if an error is raised from a procedure by putting the SYS_ROLLBACK parameter to the transactions bulletin board. This kind of error is recoverable by fixing the reason for the error, for example, by removing the duplicate row from the database, and then executing the message.

When the transaction in error is deleted with MESSAGE DELETE CURRENT TRANSACTION, the deletion is completed before the MESSAGE FROM REPLICA EXECUTE command is executed; that is, the statement starts the message from replica, but waits until the active statement is completed before actually executing the message. Thus the command performs asynchronous message execution.

Note: When working with the messages, the autocommit mode must always be switched off.

Return values

Table 62. MESSAGE FROM REPLICA EXECUTE return values

Error code	Description
13047	No privilege for operation
25009	Replica <i>replica_name</i> not found
25013	Message name <i>message_name</i> not found
25018	Illegal message state. An attempt was made to delete a transaction from a message that is not in error.
25056	Autocommit not allowed

Example

```
MESSAGE MyMsg0002 FROM REPLICA bills_laptop EXECUTE;
```

Related reference:

A.60.6, "MESSAGE EXECUTE," on page 257

A.60.4, "MESSAGE DELETE CURRENT TRANSACTION," on page 254

A.60.10 MESSAGE FROM REPLICA RESTART

```
MESSAGE msgid FROM REPLICA replicaname RESTART <err-options>;
```

Where <err-options> can be IGNORE_ERRORS or LOG_ERRORS or FAIL_ERRORS

This command can only be executed on the master.

This command allows you to re-execute a failed transaction that has been stored in the system tables and that can be retrieved using the SYNC_FAILED_MESSAGES view.

A.60.11 MESSAGE GET REPLY

```
MESSAGE unique_message_name GET REPLY  
[TIMEOUT {FOREVER | seconds}]  
[COMMITBLOCK block_size_in_rows]  
[NO EXECUTE]  
[{OPTIMISTIC | PESSIMISTIC}]
```

Usage

The MESSAGE GET REPLY statement requests a message from the master database, if a reply to a sent message has not been received by the MESSAGE FORWARD statement. The MESSAGE GET REPLY statement is applicable only in advanced replication configurations. It can only be issued on the replica databases.

TIMEOUT defines how long the replica waits for the reply from the master. For example, if the master database takes a long time to start due to, for example, a server restart, the message reply to replica can time out. To recover, reissue the MESSAGE GET REPLY command and set the timeout to a larger value or FOREVER.

If the reply message contains REFRESHes of large publications, the size of the REFRESH commit block, that is, the number of rows that are committed in one transaction, can be limited using the COMMITBLOCK property. This has a positive impact on the performance of the replica database. It is recommended that there are no online users in the database when the COMMITBLOCK property is in use.

If the execution of a reply message with the COMMITBLOCK property fails in the replica database, it cannot be re-executed. The failed message must be deleted from the replica database and refreshed from the master database.

If NO EXECUTE is specified, when the reply message is available at the master, it is only read and stored for later execution. Otherwise, the reply message is downloaded from the master and executed in the same statement. Using NO EXECUTE reduces bottlenecks in communication lines by allowing reply messages for later execution in different transactions.

You can define the reply message to use table-level pessimistic locking when it is initially executed. If the PESSIMISTIC mode is specified, all other concurrent access to the table affected is blocked until the synchronization message has completed. If the optimistic mode is used, the MESSAGE GET REPLY operation may fail due to a concurrency conflict.

When a transaction acquires an exclusive lock to a table, the **General.TableLockWaitTimeout** parameter setting determines the wait period of the transaction after which the exclusive or shared lock is released.

If a reply message fails in delivery due to a communication error (without COMMITBLOCK), you must explicitly use the MESSAGE GET REPLY to resend the message. Once resent, MESSAGE GET REPLY re-executes the message.

Note: When working with the messages, the autocommit mode must be switched off.

Result set

MESSAGE GET REPLY returns a result set table. The returned result set is the same as with command MESSAGE EXECUTE.

The columns of the result set are as follows:

Table 63. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set

Column Name	Description
Partno	Message part number
Type	The type of result set row. Possible types are: 0: Message part start 1: This type is not in use 2: The message was a propagation message and the status of that operation is stored in the return message 3: Task 4: Subscription task 5: Type of refresh (FULL or INCREMENTAL) 6: MESSAGE DELETE status
Masterid	Master ID
Msgid	Message ID
Errcode	Message error code. Zero if successful.
Errstr	Message error string. NULL is successful.
Insertcount	Number of inserted rows to replica. Type=3: Total number of insert Type=4: Row inserts restored from replica history to replica base table Type=5: Insert operations received from master
Deletecount	Type = 3: Total number of deletes Type = 4: Row deletes restored from replica base table Type = 5: Delete operations received from master
Bytecount	Size of message in bytes. Indicated in result received from command MESSAGE END. Otherwise 0.

Table 63. MESSAGE EXECUTE, MESSAGE FORWARD, and MESSAGE GET REPLY result set (continued)

Column Name	Description
Info	Information of the current task. Type = 0: then Message name Type = 3: Publication name Type = 4: Table name Type = 5: FULL/INCREMENTAL

Return values

Fatal errors in transaction propagation abort the message and return an error code to the replica. To propagate the aborted message, you need to correct the fatal errors and restart the message with command MESSAGE FROM REPLICA EXECUTE.

If a REFRESH fails in the master, an error message about the failed REFRESH is added to the result set. Other parts of the message are executed normally. The failed REFRESH must be REFRESHed from the master in a separate synchronization message.

If a REFRESH (that is, the execution of the reply message) fails in the replica, the message is still available in the replica database and can be restarted with the MESSAGE *msg_name* EXECUTE command.

Table 64. MESSAGE GET REPLY return values from replica

Error code	Description
13XXX	Table level error
13124	User id <i>num</i> not found This message is generated, for example, if the user has been dropped.
10006	Concurrency conflict (simultaneous other operation)
13047	No privilege for operation
13056	Insert not allowed for pseudo column
21XXX	Communication error
25005	Message <i>message_name</i> is already active
25013	Message name <i>message_name</i> not found
25018	Illegal message state In the replica, the message can only be executed using the MESSAGE GET REPLY statement if the message is forwarded to the master.

Table 64. MESSAGE GET REPLY return values from replica (continued)

Error code	Description
25024	Master not defined
25026	Not a valid master user
25031	Transaction is active, operation failed
25035	Message <i>message_name</i> is in use. A user is currently creating or forwarding this message.
25036	Publication <i>publication_name</i> not found or publication version mismatch
25040	User id <i>user_id</i> is not found While executing a message reply, an attempt to map a master user to a local replica id failed.
25041	Subscription to publication <i>publication_name</i> not found
25048	Publication <i>publication_name</i> request info not found
25054	Table <i>table_name</i> is not set for synchronization history
25056	Autocommit not allowed
25057	Already registered to master <i>master_name</i>
25060	Column <i>column_name</i> does not exist on publication <i>publication_name</i> resultset in table <i>table_name</i>

Table 65. MESSAGE GET REPLY return values from master

Error code	Description
13XXX	Table level error
13124	User id <i>num</i> not found This message is generated, for example, if the user has been dropped.
25012	Message reply timed out
25016	Message not found, replica id <i>replica-id</i> , message id <i>message-id</i>
25043	Reply message is too long (<i>size_of_messages</i> bytes). Maximum is set to <i>max_message_size</i> bytes.
25056	Autocommit not allowed

Examples

```
MESSAGE MyMessage001 GET REPLY TIMEOUT 120
MESSAGE MyMessage001 GET REPLY TIMEOUT 300 COMMITBLOCK 1000
```

Related reference:

A.60.9, “MESSAGE FROM REPLICA EXECUTE,” on page 264

A.62, “REFRESH”

A.61 POST EVENT

```
post_statement ::= POST EVENT event_name [( parameters) ]
                [UNIQUE | DATA UNIQUE]
```

The POST EVENT statement is allowed only inside stored procedures; it can be used for posting system events and user-defined events. For more details, see A.15.6, “post_statement,” on page 185.

Related reference:

A.15.6, “post_statement,” on page 185

Related information:

3.5, “Events,” on page 80

A.62 REFRESH

```
REFRESH publication [parameters] [FULL]
[OPTIMISTIC|PESSIMISTIC]
[COMMITBLOCK number_of_rows]
[TIMEOUT {DEFAULT | FOREVER | timeout_ms}]
```

Usage

The REFRESH statement is a storageless refresh command. It conserves memory by streaming the associated data. It also saves I/O bandwidth because no messages are written to disk. Each command blocks until it is successfully executed.

The optional property OPTIMISTIC|PESSIMISTIC define the way the replica table is being locked.

- The OPTIMISTIC mode (the default value) defines that the concurrency control method depends on the table type and the isolation level.
 - With disk-based tables (D-tables) in the OPTIMISTIC mode, the REFRESH always succeeds.
 - With in-memory tables (M-tables), row-level locking is used.
 - With D-tables in the PESSIMISTIC mode, if a lock cannot be obtained, PESSIMISTIC fails and returns an error.
- PESSIMISTIC defines that the table is exclusively locked, regardless of the table type and isolation level chosen, for the time of refresh. If the lock cannot be obtained, the refresh request fails and returns an error.

The COMMITBLOCK property defines the size of the REFRESH commit block, that is, the number of rows that are committed in one transaction. The COMMITBLOCK keyword is useful if the reply to the REFRESH request contains REFRESHes of large publications. Defining the commit block size has a positive impact on the performance of the replica database. It is recommended that there are no online users accessing the database when the COMMITBLOCK property is being used.

- If COMMITBLOCK is not used, the execution of REFRESH is a part of the current transaction.

The effect of REFRESH can be revoked by issuing the ROLLBACK command. To make the effect of REFRESH durable, COMMIT WORK has to be issued.

REFRESH is idempotent in the sense that it can be issued repeatedly, over the rollbacks and commits, and the effects are (in the quiescent state of the database) always the same.

- If the COMMITBLOCK clause is used, each transfer part (of the specified size) is committed in Replica implicitly. The ROLLBACK statement removes the effect of the latest transfer part only. COMMIT WORK commits the last transfer part.

The TIMEOUT property defines how long the replica server waits for the reply message. If TIMEOUT is not defined, FOREVER is used.

Return values

Table 66. REFRESH return values

Error Code	Description
13133	Not a valid license for this product
25004	Dynamic parameters are not supported
25015	Syntax error: <i>error_message</i> , line <i>line_number</i>
25024	Master not defined
25025	Node name not defined
25026	Not a valid master user
25044	SYNC_CONFIG system publication takes only character arguments
25056	Autocommit not allowed
25071	Not registered to publication <i>publication_name</i>
25072	Already registered to publication <i>publication_name</i>
13XXX	Table level error
21XXX	Communication error
10006	Concurrency conflict (simultaneous other operation)
13047	No privilege for operation
13056	Insert not allowed for pseudo column
25005	Message <i>message_name</i> is already active
25018	Illegal message state In the replica, the message can be executed only using the MESSAGE FORWARD statement if the message is ended and the ending transaction is committed.

Table 66. REFRESH return values (continued)

Error Code	Description
25024	<p>Master not defined</p> <p>This message is produced if double quotation marks, rather than single quotation marks, are used around the <i>connect_string</i> in a MESSAGE FORWARD statement.</p> <p>For example, if the master node is given the node name "master" (which is a reserved word and therefore should be delimited by double quotation marks), and if connect string for the is tcp localhost 1315, the following MESSAGE statements shown are correct:</p> <pre>--On the replica ... --double quotation marks MESSAGE msg1 BEGIN TO "master"; ... --single quotation marks MESSAGE msg2 FORWARD TO 'tcp localhost 1315';</pre> <p>The MESSAGE BEGIN statement defines (within the replica server) what the node name of the master is. The MESSAGE FORWARD statement may contain the connect string to the server.</p>
25026	Not a valid master user
25031	Transaction is active, operation failed
25035	<p>Message <i>message_name</i> is in use.</p> <p>A user is currently creating or forwarding this message.</p>
25040	<p>User ID <i>user_id</i> is not found.</p> <p>While executing a message reply an attempt to map a master user to a local replica id failed.</p>
25041	Subscription to publication <i>publication_name</i> not found
25048	Publication <i>publication_name</i> request info not found
25052	Failed to set node name to <i>node_name</i> .
25054	Table <i>table_name</i> is not set for synchronization history
25055	<p>Connect information is allowed only when not registered</p> <p>The connect info in MESSAGE <i>message_name</i> FORWARD TO <i>connect_info options</i> is allowed only if the replica has not yet been registered to the master database.</p>
25056	Autocommit not allowed
25057	The replica database is registered to a master database
25060	Column <i>column_name</i> does not exist on publication <i>publication_name</i> result set in table <i>table_name</i>
13XXX	Table level error

Table 66. REFRESH return values (continued)

Error Code	Description
13124	User ID <i>num</i> not found This message is generated, for example, if the user has been dropped.
25056	Autocommit not allowed

Examples

The following example shows a synchronous, messageless refresh:

```
REFRESH publ_states;
PESSIMISTIC;
COMMITBLOCK 1000;
COMMIT WORK;
```

A.63 REGISTER EVENT

```
wait_register_statement ::= REGISTER EVENT event_name
```

The REGISTER EVENT statement tells the server that you would like to be notified of all future occurrences of the specified event, even if you are not yet waiting for it.

The REGISTER EVENT statement is allowed only inside stored procedures. For more details, see A.15.7, “wait_register_statement,” on page 185.

Related reference:

A.15.7, “wait_register_statement,” on page 185

Related information:

3.5, “Events,” on page 80

A.64 REVOKE

```
REVOKE { role_name [, role_name ]... }
      FROM {PUBLIC | user_name [, user_name ]... }

REVOKE
  {ALL | revoke_privilege [, revoke_privilege ]... } ON table-name
  FROM {PUBLIC | user_name [, user_name ]... | role_name [, role_name ]... }

revoke-privilege ::= DELETE | INSERT | SELECT |
  UPDATE [( column_identifier [, column_identifier ]... )] |
  REFERENCES

REVOKE EXECUTE ON procedure_name
  FROM {PUBLIC | user_name [, user_name ]... | role_name [, role_name ]... }

REVOKE {SELECT | INSERT} ON event_name FROM
  {PUBLIC | user_name [, user_name ]... | role_name [, role_name ]... }

REVOKE {SELECT | INSERT} ON sequence_name
  FROM {PUBLIC | user_name [, user_name ]... | role_name [, role_name ]... }
```

Usage

The REVOKE statement removes a role from users or privileges from users and roles.

Note: The keywords CASCADE and RESTRICT are not supported in REVOKE statements.

Examples

The following command removes the role GUEST_USERS from the user HOBBS:

```
REVOKE GUEST_USERS FROM HOBBS;
```

The following command removes the privilege to insert on event TEST from the role GUEST_USERS:

```
REVOKE INSERT ON TEST FROM GUEST_USERS;
```

Related reference:

A.52, “GRANT,” on page 234

Related information:

4.2, “Managing user privileges and roles,” on page 90

A.65 REVOKE PASSTHROUGH

```
REVOKE PASSTHROUGH  
FROM {PUBLIC | user_name [, user_name]... | role_name [, role_name]... }
```

Supported in

This command requires the use of the SQL passthrough feature.

Usage

The REVOKE PASSTHROUGH statement revokes access rights for SQL passthrough.

Examples

```
REVOKE PASSTHROUGH READ FROM cdcuser1, cdcuser2  
REVOKE PASSTHROUGH WRITE FROM cdcuser1, cdcuser2
```

See also

A.53, “GRANT PASSTHROUGH,” on page 235

A.66 REVOKE REFRESH

```
REVOKE { REFRESH | SUBSCRIBE } ON publication_name FROM {PUBLIC |  
    user_name, [ user_name ] ... |  
    role_name , [ role_name ] ... }
```

Usage

The REVOKE REFRESH statement revokes access rights to a publication from a user or role defined in the master database. The REVOKE REFRESH statement is applicable only in advanced replication configurations. It can only be issued on the master database.

The keywords REFRESH and SUBSCRIBE are synonymous. However, SUBSCRIBE is deprecated in the REVOKE statement.

Return values

Table 67. REVOKE REFRESH return values

Error Code	Description
13137	Illegal grant/revoke mode
13048	No grant option privilege
25010	Publication <i>name</i> not found

Examples

```
REVOKE REFRESH ON customers_by_area FROM joe_smith;  
REVOKE REFRESH ON customers_by_area FROM all_salesmen;
```

A.67 ROLLBACK WORK

```
ROLLBACK [WORK]
```

Usage

The ROLLBACK or ROLLBACK WORK statement discards the changes made in the database by the current transaction. It also terminates the transaction.

Related reference:

A.8, “COMMIT WORK,” on page 169

A.68 SAVE

```
SAVE [NO CHECK] [ { IGNORE_ERRORS | LOG_ERRORS | FAIL_ERRORS } ]  
[ { AUTOSAVE | AUTOSAVEONLY } ] sql_statement
```

Usage

The SAVE statement saves statements in the replica for propagation to the master. The SAVE statement is applicable only in advanced replication configurations. It can only be issued on the replica databases.

The statements of a transaction that need to be propagated to the master database must be explicitly saved to the transaction queue of the replica database. Adding a SAVE statement before the transaction statements does this.

Only master users are allowed to save statements. When the saved statements are executed on the master, they must be executed using the appropriate access rights of a user on the master. The saved statements are executed in the master database using the access rights of the master user that was active in the replica when the statement was saved. If a user in the replica was mapped to a user in the master, the SAVE statement uses the access rights of the user in the master.

The default behavior for error handling with transaction propagation is that a failed transaction halts execution of the message; this aborts the

currently-executing transaction and prevents execution of any subsequent transactions that are in that same message. However, you can choose a different error-handling behavior.

NO CHECK means that the statement is not prepared in the replica. This option is useful if the command would not make sense on the replica. For example, if the SQL command calls a stored procedure that exists on the master but not on the replica, then you do not want the replica to try to prepare the statement. If you use this option, then the statement can not have parameter markers.

IGNORE_ERRORS means that if a statement fails while executing on the master, then the failed statement is ignored and the transaction is aborted. However, only the transaction, not the entire message, is aborted. The master continues executing the message, resuming with the first transaction after the failed one.

LOG_ERRORS means that if a statement failed while executing on the master, then the failed statement is ignored and the current transaction is aborted. The failed transaction's statements are saved in `SYS_SYNC_RECEIVED_STMTS` system table for later execution or investigation. The failed transactions can be examined using `SYS_SYNC_FAILED_MESSAGES` system view and they can be re-executed from there using `MESSAGE <msg_id> FROM REPLICA <replica_name> RESTART -statement`.

Note: As with the `IGNORE_ERROR` option, aborting the transaction does not abort the entire message. The master continues executing the message, resuming with the first transaction after the failed one.

FAIL_ERRORS means that if a statement fails, the master stops executing the message. This is the default behavior.

AUTOSAVE means that the statement is executed in the master and automatically saved for further propagation if the master is also a replica to some other master (that is, a middle-tier node).

AUTOSAVEONLY means that the statement is NOT executed in the master but instead is automatically saved for further propagation if the master is also a replica to some other master (that is, a middle-tier node).

Return values

Table 68. SAVE return values

Error Code	Description
25001	Internal error Master database has exceeded the database size limit required to save the statement.
25003	Cannot save SAVE statements
25070	Statement can be saved only for one master in transaction.

Examples

```
SAVE INSERT INTO mytbl (col1, col2) VALUES ('calvin', 'hobbes')
SAVE CALL SP_UPDATE_MYTAB('calvin_1', 'hobbes')
SAVE CALL SP_DELETE_MYTAB('calvin')
SAVE NO CHECK IGNORE_ERRORS insert into mytab values(1,2)
```

A.69 SAVE PROPERTY

```
SAVE PROPERTY property_name VALUE {'value_string' | NONE}

SAVE DEFAULT PROPERTY property_name VALUE {'value_string' | NONE}

SAVE DEFAULT PROPAGATE PROPERTY {WHERE name {=<|<=>|>=>|<>} 'value' | NONE}
```

Usage

The SAVE PROPERTY statements set properties in the replica for a transaction that is saved for propagation to the master. The SAVE PROPERTY statements are applicable only in advanced replication configurations. They can only be issued on the replica databases. The property values can be read in the master database.

The SAVE PROPERTY statement assigns properties to the current active transaction.

The statements of the transaction in the master database can access these properties by calling the GET_PARAM() function. Properties are only available in the replica database that apply to the following command:

```
MESSAGE APPEND unique_message_name PROPAGATE TRANSACTIONS
WHERE property > 'value_string'
```

When the transaction is executed in the master database, the saved properties are placed on the parameter bulletin board of the transaction. If the saved property already exists, the new value overwrites the previous one.

The SAVE DEFAULT PROPERTY statement defines default properties that are saved to all transactions of the current connection.

The SAVE DEFAULT PROPAGATE PROPERTY WHERE statement saves default transaction propagation criteria. This can be used for example to set the propagation priority of transactions created in the current connection.

The SAVE DEFAULT PROPAGATE PROPERTY WHERE statement can be used on connection level to append all MESSAGE *unique_message_name* APPEND PROPAGATE TRANSACTIONS statements to have the default WHERE statement. If the WHERE statement is entered also in the PROPAGATE statement, it will override the statement set with the DEFAULT PROPAGATE PROPERTY.

A property or a default property can be removed by re-saving the property with value string NONE.

Return values

Table 69. SAVE PROPERTY return values

Error Code	Description
13086	Invalid data type in a parameter.

Result set

The SAVE PROPERTY statements do not return a result set.

Examples

```
SAVE PROPERTY conflict_rule VALUE 'override'  
SAVE DEFAULT PROPERTY userid VALUE 'scott'  
SAVE DEFAULT PROPERTY userid VALUE NONE  
SAVE DEFAULT PROPAGATE PROPERTY WHERE priority > '2'
```

Related reference:

B.6.2, “PUT_PARAM() function,” on page 321

A.70 SELECT

```
SELECT [ALL | DISTINCT] select-list  
  LEVEL  
  FROM table_reference_list  
  [WHERE search_condition]  
  [GROUP BY column_name [, column_name]... ]  
  [HAVING search_condition]  
  [hierarchical_condition]  
  [[UNION | INTERSECT | EXCEPT] [ALL] select_statement]...  
  [ORDER BY expression]  
  [ASC | DESC]  
  [LIMIT row_count [OFFSET skipped_rows] | LIMIT skipped_rows,row_count |  
  FETCH FIRST ROW ONLY | FETCH FIRST constant ROWS ONLY]  
  [FOR UPDATE]
```

Usage

The SELECT statement retrieves values from a database. You can select 0 or more records from one or more tables. A SELECT operation is typically called a *query*.

LEVEL

LEVEL is a pseudocolumn valid in the context of the hierarchical query only. If the result set is viewed as a tree of interreferenced rows, the LEVEL column produces the tree level number, assigning "1" to the top-level row.

select-list

See A.79.4, “*query_specification*,” on page 307 and A.79.7, “Pseudo columns in SELECT statement,” on page 309.

table_reference_list

See A.79.6, “*table_reference*,” on page 309.

search_condition

See A.79.5, “*search_condition*,” on page 308.

hierarchical_condition

```
::= START WITH search_condition CONNECT BY [PRIOR] search_condition
```

If your table contains hierarchical data, you can select rows in a hierarchical order using a hierarchical query clause. In a hierarchical query clause, START WITH specifies the root row(s) of the hierarchy and CONNECT BY specifies the relationship between parent rows and child rows of the hierarchy. The CONNECT BY condition cannot contain a subquery.

In a hierarchical query, one expression in the condition must be qualified with the PRIOR operator to refer to the parent row. PRIOR is a unary operator and has the same precedence as the unary + and - arithmetic operators. It evaluates the immediately following expression for the parent row of the current row in a hierarchical query. PRIOR is most commonly used when comparing column values with the equality operator. The PRIOR keyword can be on either side of the operator.

ORDER BY *expression*

ORDER SIBLINGS BY causes the rows at any level to be ordered accordingly.

See also A.79.3, “expression,” on page 306.

LIMIT

The non-standard clause LIMIT *row_count* OFFSET *skipped_rows* allows to mask out a portion of a result set with a sliding window having the size of *row_count* and positioned at the *skipped_rows*+1 row. A negative value of *skipped_rows* results in an error, while the negative value of *row_count* results in the whole result set produced.

Two formats are available: for example LIMIT 24 OFFSET 10 is equal to LIMIT 10, 24.

FETCH

The FETCH FIRST ROW ONLY clause returns only the first row of the table. FETCH FIRST *constant* ROWS ONLY returns the number of rows defined with *constant*.

FOR UPDATE

When a user accesses a row with the SELECT... FOR UPDATE statement, the row is locked with an update mode lock. This means that no other user can read or update the row, and ensures the current user can later update the row. See 5.2, “Concurrency control and locking,” on page 111 for more details.

Examples

```
SELECT ID FROM TEST;
SELECT DISTINCT ID, C FROM TEST WHERE ID = 5;
SELECT DISTINCT ID FROM TEST ORDER BY ID ASC;
SELECT NAME, ADDRESS FROM CUSTOMERS
UNION
SELECT NAME, DEP FROM PERSONNEL;
SELECT dept, count(*) FROM person
GROUP BY dept
ORDER BY dept
LIMIT 20 OFFSET 10
```

Example of LIMIT and FETCH

```
SELECT * FROM SYS_TABLES LIMIT 1;
SELECT * FROM SYS_TABLES FETCH FIRST ROW ONLY; -- same as above

SELECT * FROM SYS_TABLES WHERE TABLE_NAME LIKE 'A%' LIMIT 5;
SELECT * FROM SYS_TABLES WHERE TABLE_NAME LIKE 'A%' FETCH
FIRST 5 ROWS ONLY; -- same as above
```

Example of START WITH

```
SELECT last_name, employee_id, manager_id, LEVEL
FROM employees
START WITH employee_id = 100
CONNECT BY PRIOR employee_id = manager_id
ORDER SIBLINGS BY last_name;
```

LAST_NAME	EMPLOYEE_ID	MANAGER_ID
King	100	
Cambrault	148	100
Bates	172	148
Bloom	169	148
Fox	170	148
Kumar	173	148
Ozer	168	148
Smith	171	148
De Haan	102	100
Hunold	103	102
Austin	105	103
Ernst	104	103
Lorentz	107	103
Pataballa	106	103
Errazuriz	147	100
Ande	166	147
Banda	167	147

Example of LEVEL and ORDER SIBLINGS BY

```
SELECT last_name, employee_id, manager_id, LEVEL
FROM employees
START WITH last_name = 'King'
CONNECT BY PRIOR employee_id = manager_id
ORDER SIBLINGS BY last_name
ORDER BY LEVEL;
```

LAST_NAME	EMPLOYEE_ID	MANAGER_ID	LEVEL
King	100	NULL	1
Cambrault	148	100	2
De Haan	102	100	2
Bates	172	148	3
Bloom	169	148	3
Gates	104	148	3
Hunold	103	102	3
Hope	202	172	4
Smith	201	172	4

A.71 SET

SET commands apply to the user session (connection) in which they are executed. They do not affect other user sessions.

SET statements may be issued at any time; however, they do not all take effect immediately. The following statements take effect immediately:

- SET CATALOG
- SET IDLE TIMEOUT
- SET LOCK TIMEOUT
- SET OPTIMISTIC LOCK TIMEOUT
- SET SCHEMA
- SET STATEMENT MAXTIME

The following statements take effect after the next COMMIT WORK:

- SET DURABILITY
- SET ISOLATION LEVEL
- SET { READ ONLY | READ WRITE | WRITE }

SET statements are not subject to rollback; they remain in force even if the transaction they have been issued in has been aborted or rolled back. It is a good practice to issue them before any DDL/DML SQL statement in a transaction.

The settings continue in effect until the end of the session (connection) or until another SET command changes the settings, or in some cases until a higher-precedence command (for example, SET TRANSACTION) is executed.

Examples

```
SET CATALOG myCatalog;
SET DURABILITY STRICT;
SET IDLE TIMEOUT 30;
SET ISOLATION LEVEL REPEATABLE READ;
SET OPTIMISTIC LOCK TIMEOUT 30;
SET LOCK TIMEOUT 30;
SET LOCK TIMEOUT 500MS;
SET READ ONLY;
SET SCHEMA 'accounting_info';
SET SCHEMA 'john_smith';
SET STATEMENT MAXTIME 180;
```

See also

A.72, “SET TRANSACTION,” on page 295

A.71.1 Differences between SET and SET TRANSACTION

The SET and SET TRANSACTION commands are used for setting various transaction properties, such as isolation level, the read level, or SQL passthrough mode.

The SET commands set the properties for a session and they are often called session-level commands. The SET TRANSACTION commands set the properties for one transaction and they are called transaction-level commands.

The transaction-level commands follow different rules from the session-level commands. These differences are listed below.

- The transaction-level commands take effect in the transaction in which they are issued; the session-level commands take effect in the next transaction, that is, after the next COMMIT WORK.
- The transaction-level commands apply to only the current transaction; the session-level commands apply to all subsequent transactions — that is, until the end of the session (connection) or until another SET command changes them.
- The transaction-level commands must be executed at the beginning of a transaction, that is, before any DML or DDL statements. However, they may be executed after other SET statements. If this rule is violated, an error is returned. The session-level commands may be executed at any point in a transaction.
- The transaction-level commands take precedence over the session-level commands. However, the transaction-level commands apply only to the current

transaction. After the current transaction is committed or aborted, the settings will return to the value set by the most recent previous SET command (if any).
For example:

```
COMMIT WORK; -- Finish previous transaction;
SET ISOLATION LEVEL SERIALIZABLE;
COMMIT WORK; -- Isolation level is now SERIALIZABLE
...
COMMIT WORK;
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ;
-- Isolation level is now REPEATABLE READ because
-- transaction-level settings take precedence
-- over session-level settings.
COMMIT WORK;
-- Isolation level is now back to SERIALIZABLE, since the
-- transaction-level settings applied only to that
-- transaction.
```

The complete precedence hierarchy for isolation level and read level settings is below. Items closer to the top of the list have higher precedence.

1. SET TRANSACTION ... (transaction-level settings)
2. SET ... (session-level settings)
3. The server-level settings specified by the value in `solid.ini` configuration parameter (for example, **IsolationLevel** or **DurabilityLevel**)
4. The server's default setting (parameter factory value).

A.71.2 SET (read/write level)

```
SET {READ ONLY | READ WRITE | WRITE}
```

SET {READ ONLY | READ WRITE | WRITE} allows you to specify whether the connection be allowed only to read, read and write, or whether it be allowed to write only.

See also A.71.6, “SET ISOLATION LEVEL,” on page 283.

A.71.3 SET CATALOG

```
SET CATALOG catalog_name
```

SET CATALOG sets the current catalog context for a connection or session.

In advanced replication setups with multiple master databases, the SET CATALOG command defines explicitly the master database that is used for all advanced-replication-related operations. You must issue the SET CATALOG command before you issue any synchronization commands. The SET CATALOG command is required also when more than one catalog is defined in the database.

A.71.4 SET DELETE CAPTURE

```
SET DELETE CAPTURE {NONE | CHANGES}
```

Usage

The SET DELETE CAPTURE statement sets the data aging mode in solidDB Universal Cache. Data aging removes table rows from the front-end but preserves them in the back-end.

- SET DELETE CAPTURE NONE: data aging is enabled
- SET DELETE CAPTURE CHANGES: data aging is disabled

The SET DELETE CAPTURE statement takes effect immediately, starting from the next SQL statement, until it is reverted by a similar statement or SET TRANSACTION DELETE CAPTURE.

See also

A.72.3, “SET TRANSACTION DELETE CAPTURE,” on page 297

A.71.5 SET DURABILITY

```
SET DURABILITY { RELAXED | STRICT | DEFAULT }
```

SET DURABILITY sets the transaction durability level.

- STRICT means that as soon as a transaction is committed, the server writes the information to the transaction log file.
- RELAXED means that the server may not write the information as soon as the transaction is committed; instead, the server may wait, for example, until it is less busy, or until it can write multiple transactions in a single write operation. If you use relaxed durability and the server shuts down abnormally, you may lose a few of the most recent transactions.
- DEFAULT means that the durability level defined with the **Logging.DurabilityLevel** parameter is used.

For more details, see 5.3, “Choosing transaction durability level,” on page 121.

A.71.6 SET ISOLATION LEVEL

```
SET ISOLATION LEVEL {  
    READ COMMITTED |  
    REPEATABLE READ |  
    SERIALIZABLE }  
}
```

SET ISOLATION LEVEL allows you to specify the isolation level.

If the assigned workload server is Secondary, it can be changed programmatically to the Primary. At the session level, the following statements change the workload connection server to the Primary:

```
SET WRITE (nonstandard)  
SET ISOLATION LEVEL REPEATABLE READ  
SET ISOLATION LEVEL SERIALIZABLE
```

The statement takes effect immediately, if it is a first statement of a transaction, or from the next transaction otherwise.

If the above statement is not applicable, it returns SQL_SUCCESS, with no action performed. For example, such is the case when SET WRITE is applied to a standalone server. In that case the semantics of SET WRITE is equal to that of SET READ WRITE.

The effect of the SET WRITE statement may be reverted with the statement SET READ WRITE or ... READ ONLY (SQL:1999). Also, the isolation level statement has the same effect:

```
SET ISOLATION LEVEL READ COMMITTED
```

A.71.7 SET PASSTHROUGH

```
SET PASSTHROUGH {READ <passthrough level> [WRITE <passthrough level>]}  
| {WRITE <passthrough level> | [READ <passthrough level>]}  
| <passthrough level>
```

where

```
passthrough level ::= NONE | CONDITIONAL | FORCE | DEFAULT
```

Usage

The SET PASSTHROUGH statement sets the SQL passthrough mode in solidDB Universal Cache.

- NONE: SQL passthrough is not used; no commands are passed from the front-end to the back-end
- CONDITIONAL: SQL passthrough is activated by a missing table or a syntax error.
- FORCE: SQL passthrough is used to pass all statements from the front-end to the back-end
- DEFAULT: causes the current session default for SQL passthrough is used

The SET PASSTHROUGH statement takes effect immediately, starting from the next SQL statement, until it is reverted by a similar statement or SET TRANSACTION PASSTHROUGH.

See also

A.72.6, “SET TRANSACTION PASSTHROUGH,” on page 298

A.71.8 SET SAFENESS

```
SET SAFENESS {1SAFE | 2SAFE | DEFAULT}
```

SET SAFENESS determines whether the replication protocol is synchronous (2-safe) or asynchronous (1-safe).

- 1-safe: the transaction is first committed at Primary and then transmitted to Secondary
- 2-safe: the transaction is not committed before it has been acknowledged by Secondary (default).

SET SAFENESS sets the safeness level for the current session.

A.71.9 SET SCHEMA

```
SET SCHEMA {'schema_name' | USER | 'user_name'}
```

Usage

The SET SCHEMA statement changes the default schema.

USER means that the schema is changed to the current user name.

user_name must be a valid user name in the database.

As per ANSI SQL-92 standard, the *user_name* and *schema_name* can be enclosed in single quotation marks.

solidDB supports SQL89 style schemas. Schemas are used to help uniquely identify entities (tables, views) within a database. By using schemas, each user can create database objects without the risk of using names that already in use.

To identify an entity (such as a table) uniquely, you qualify it by specifying the catalog name and schema name using the syntax *[catalog_name.][schema_name.]database_object*. For example, *FinanceCatalog.AccountsReceivableSchema.CustomersTable* is a qualified table name.

The algorithm to resolve entity names *[catalog_name.][schema_name.]database_object* is as follows:

1. If *schema_name* is given, *database_object* is searched only from that schema.
2. If *schema_name* is not given:
 - a. First *database_object* is searched from default schema. Default schema is initially the same as user name, but can be changed with SET SCHEMA statement
 - b. Then *database_object* is searched from all schemas in the database. If more than one entity with same *database_object* and type (table, stored procedure, ...) is found, the following error is returned: 13111, Table Error: Ambiguous entity name *database_object*.

The SET SCHEMA statement affects only the default entity name resolution. It does not change any access rights to database entities. The SET SCHEMA statement sets the default schema name for unqualified names in statements that are prepared in the current session by an EXECDIRECT statement or a prepare statement.

Example

```
SET SCHEMA 'CUSTOMERS';
```

Related information:

4.6, “Managing database objects,” on page 105

In the solidDB server, you use catalogs and schemas to organize your data. The use of schemas in the solidDB server conforms to the SQL standard, but the use of catalogs is an extension to the SQL standard. In general, a catalog can be thought of as a logical database. A schema typically corresponds to a user.

4.6.3, “Uniquely identifying objects within catalogs and schemas,” on page 107

A.71.10 SET SEQUENCE

```
SET SEQUENCE <sequence_name> VALUE <value>
```

Usage

The SET SEQUENCE statement sets the sequence value dynamically.

The <value> is the sequence number as a numeric value. The sequence values are stored in BIGINT data type.

When the sequence value has been set, the next returned value is <value> + 1.

Examples

```
SELECT MYSEQ.NEXT;  
       MYSEQ.NEXT  
-----
```

```

      8
1 rows fetched.

SET SEQUENCE MYSEQ VALUE 5;
Command completed successfully, 0 rows affected.

SELECT MYSEQ.NEXT;
      MYSEQ.NEXT
      -----
      6
1 rows fetched.

```

See also

A.21, “CREATE SEQUENCE,” on page 204

A.71.11 SET SQL

```

SET SQL INFO {ON | OFF} [FILE {file_name | "{file_name}" | '{file_name'}}]
  [LEVEL info_level]
SET SQL SORTARRAYSIZE {array-size | DEFAULT}
SET SQL JOINPATHSPAN {path-span | DEFAULT}
SET SQL CONVERTORSTOUNIONS
  {YES [COUNT value] | NO | DEFAULT}

```

Usage

All the settings are read per user session (unlike the settings in the `solid.ini` file, which are automatically read each time solidDB is started).

SET SQL INFO

The SET SQL INFO statement controls the SQL Info facility which outputs trace information about each SQL statement processed by solidDB. You can use the SQL Info facility to provide information about an SQL statement to troubleshoot problems or tune queries.

The default file is a global `soltrace.out` shared by all users. If the file name is given, all future SET SQL INFO ON settings will use that file unless a new file is set. Use single quotation marks when specifying the file name. Without single quotation marks, the file name is converted to uppercase.

The output is appended to an existing file. The file is never truncated. After you no longer need the SQL INFO file, delete the file manually. If the file open fails, the info output is silently discarded.

The tracing level (*info_level*) is defined as an integer between 0 (no tracing) and 8 (solidDB info from every fetched row).

Table 70. SQL Info levels

Info level	Description
0	no output
1	table, index, and view info in SQL format
2	SQL execution graphs (technical support use only)
3	some SQL estimate info, solidDB selected key name

Table 70. SQL Info levels (continued)

Info level	Description
4	all SQL estimate info, solidDB selected key info
5	solidDB info also from discarded keys
6	solidDB table level info
7	SQL info from every fetched row
8	solidDB info from every fetched row

A good way to generate useful info output is to set info on with a new file name and then execute the SQL statement using EXPLAIN PLAN FOR syntax. This method gives all necessary estimator information but does not generate output from the fetches (which might generate a large output file).

Tip: The default behavior of the SQL Info facility is controlled with the **SQL.Info** and **SQL.InfoFileName** parameters.

SET SQL SORTARRAYSIZE

The SET SQL SORTARRAYSIZE statement sets the size of the array that SQL uses when ordering the result set of a query. The units are rows. For example, if you specify a value of 1000, the server creates an array large enough to sort 1000 rows.

SET SQL JOINPATHSPAN

This command is obsolete. The syntax is accepted, but the command has no effect.

SET SQL CONVERTORSTOUNIONS

The SET SQL CONVERTORSTOUNIONS statement controls whether a query that contains "OR" operations is converted into an equivalent query that uses "UNION" operations.

COUNT specifies the maximum number of OR operations that can be converted to UNION operations. The default value is 10, which is sufficient for most cases.

The following operations are logically equivalent:

```
select ... where x = 1 OR y = 1;
select ... where x = 1 UNION select... where y = 1;
```

By setting CONVERTORSTOUNIONS, you tell the optimizer that it may use equivalent UNION operations instead of OR operations if the UNIONS seem more efficient based on the volume and distribution of data.

You can also specify CONVERTORSTOUNIONS with the SQL.ConvertOrsToUnionsCount parameter.

Example

```
SET SQL INFO ON FILE 'sqlinfo.txt' LEVEL 5
```

A.71.12 SET STATEMENT MAXTIME

```
SET STATEMENT MAXTIME minutes
```

SET STATEMENT MAXTIME sets connection-specific maximum execution time in minutes. The setting is effective until a new maximum time is set. Zero time means no maximum time, which is also the default.

A.71.13 SET SYNC

The SET SYNC statements are used to set up and configure the databases used in an advanced replication system.

SET SYNC master_or_replica

```
SET SYNC master_or_replica yes_or_no
```

where:

```
master_or_replica ::= MASTER | REPLICA  
yes_or_no ::= YES | NO
```

Supported in: This command requires solidDB advanced replication.

Usage: When a database catalog is created and configured for synchronization use, you must use this command to specify whether the database is a master, replica, or both. Only a DBA or a user with SYS_SYNC_ADMIN_ROLE can set the database role.

The database catalog is a master database if there are replicas in the domain that refresh from publications from this database and/or propagate transactions to it. The database catalog is a replica catalog if it can refresh from publications that are in a master database. In multi-tier synchronization, intermediate level databases serve a dual role, as both master and replica databases.

Note that to use this command requires that you have already set the node name for the master or replica using the SET SYNC NODE command. For details, read “SET SYNC NODE” on page 291.

When you set the database for a dual role, you can use the statement once or twice. For example:

```
SET SYNC MASTER YES;  
SET SYNC REPLICA YES;
```

Note that when you set the database for dual roles, SET SYNC REPLICA YES does not override SET SYNC MASTER YES. Only the following explicit statement can override the status of the master database:

```
SET SYNC MASTER NO;
```

Once overridden, the current database is set as replica only.

Examples:

```
-- configure as replica  
SET SYNC REPLICA YES;  
-- configure as master  
SET SYNC MASTER YES;
```

Return values: For details on each error code, see the appendix titled Error Codes in *solidDB Administration Guide*.

Table 71. SET SYNC return values

Error code	Description
13047	No privilege for operation
13107	Illegal set operation
13133	Not a valid license for this product
25051	Unfinished messages found

SET SYNC CONNECT

```
SET SYNC CONNECT 'connect_string [,connect_string]' TO MASTER
master_name
SET SYNC CONNECT 'connect_string' TO REPLICA replica_name
```

Supported in: This command requires solidDB advanced replication.

Usage: This statement changes the network name associated with the database name. Use this statement in a replica (or master) whenever you have changed network names in databases that a replica (or master) connects to. Network names are defined in the Listen parameter of the `solid.ini` configuration file.

The second connect string in SET SYNC CONNECT ... TO MASTER facilitates transparent failover of a Replica server to a standby Master server, should the Primary Master server fail. The order of the connect strings is not significant. The connection is automatically maintained to the currently active Primary server.

Usage in master: Use this statement in a master to change the replica's network name.

Usage in replica: Use this statement in a replica to change the master's network name.

Example:

```
SET SYNC CONNECT 'tcp server.company.com 1313' TO MASTER hq_master;
```

Return values: For details on each error code, see the appendix titled Error Codes in *solidDB Administration Guide*.

Table 72. SET SYNC CONNECT return values

Error code	Description
13047	No privilege for operation
13107	Illegal set operation
21300	Illegal network protocol
25007	Master <i>master_name</i> not found
25019	Database is not a replica database

SET SYNC MODE

SET SYNC MODE { MAINTENANCE | NORMAL }

Supported in: This command requires solidDB advanced replication.

Usage: This command sets the current catalog's sync mode to either Maintenance mode or Normal mode.

This command applies only to catalogs that are involved in synchronization (i.e. are "master" catalogs or "replica" catalogs, or are both master and replica in a hierarchy with 3 or more levels).

This command applies only to the current catalog. If you want to set more than one catalog's sync mode to Maintenance, then you will have to switch to each catalog (by using the SET CATALOG command) and then issue the SET SYNC MODE MAINTENANCE command for that catalog.

While a catalog's sync mode is Maintenance, the following rules apply:

- The catalog will not send or receive synchronization messages and therefore will not engage in synchronization activities (for example, refresh or respond to a refresh request).
- DDL commands (for example, ALTER TABLE) will be allowed on tables that are referenced by publications.
- When the sync mode changes, the server will send the system event SYNC_MAINTENANCEMODE_BEGIN or SYNC_MAINTENANCEMODE_END.
- If the master catalog's publications are altered (dropped and re-created) by using the REPLACE option, then the publication's metadata (internal publication definition data) is refreshed automatically to each replica the next time that replica refreshes from the changed publication. (This is true whether or not the database was in Maintenance sync mode when the publication was REPLACed.)
- Each catalog has a read-only parameter named SYNC_MODE in the parameter bulletin board so that applications can check the catalog's mode. Values for that parameter are either 'MAINTENANCE' if the catalog is in maintenance sync mode or 'NORMAL' if the catalog is not in maintenance sync mode. The value is NULL if the catalog is not a master or a replica.
- The user must have DBA or synchronization administrations privileges to set sync mode to Maintenance or Normal.
- A user may have more than one catalog in Maintenance sync mode at a time.
- If the session that set the mode ON disconnects, then mode is set off.
- The normal synchronization history operations are disabled. For example, when a delete or update operation is done on a table that has synchronization history on, the synchronization history tables will not store the "original" rows (i.e. the rows before they were deleted or updated). Note, however, that deletes and updates apply to the synchronization history table, for example:

```
DELETE * FROM T WHERE c = 5
```

will delete rows from the history table as well as from the base table. The table below shows how various operations (INSERT, DELETE, and so on) apply to the synchronization history tables in master and replica when sync mode is set to Maintenance.

Table 73. How different operations apply to synchronization history tables

Operation	Master	Replica
INSERT	Rows are inserted to base table.	Rows are inserted to base table and marked as official.
UPDATE	Both base table and history is updated.	Both base table and history is updated. Tentative/official status is not updated so tentative rows remains tentative and official rows remains official.
DELETE	Rows are deleted from base table and from history.	Rows are deleted from base table and from history.
Add, alter, drop column	Same operation is done to history also.	Same operation is done to history also.
Altering table mode	History mode is not altered	History mode is not altered
Create index	Same index is created to history also	Same index is created to history also
Create triggers	Triggers are not created on history	Triggers are not created on history

Example:

```
SET SYNC MODE MAINTENANCE SET SYNC MODE NORMAL
```

Return values: For details on each error code, see the appendix titled Error Codes in *IBM solidDB Administrator Guide*.

Table 74. SET SYNC MODE return values

Error code	Description
13047	No privilege for operation.
13133	Not a valid license for this product.
25021	Database is not master or replica database. This operation only applies to master and replica databases.
25088	Catalog already in maintenance mode. You have set the mode on already.
25089	Not allowed to set maintenance mode off. Someone else has set the mode on so you can not set it off.
25090	Catalog is already in maintenance mode. Someone else has set the mode on so you can not set it on.
25091	Catalog is not in maintenance mode. You tried to set mode off and it is not currently on.

SET SYNC NODE

```
SET SYNC NODE {unique_node_name | NONE}
```

Supported in: This command requires solidDB advanced replication.

Usage: Assigning the node name is part of the registration process of a replica database. Each catalog of a solidDB environment must have a node name that is unique within the domain. One catalog can have only one node name. Two catalogs cannot have the same node name.

You can use the SET SYNC NODE *unique_node_name* option to rename a node name if:

- If the node is a replica database and it is not registered to a master

and/or

- If the node is a master database and there are no replicas registered in the master database

Following are examples for renaming a node name:

```
SET SYNC NODE A; -- Now the node name is A.
SET SYNC NODE B; -- Now the node name is B.
COMMIT WORK;
SET SYNC NODE C; -- Now the node name is C.
ROLLBACK WORK; -- Now the node name is rolled back to B.
SET SYNC NODE NONE; -- Now the node has no name.
COMMIT WORK;
```

The *unique_node_name* must conform to the rules that are used for naming other objects (such as tables) in the database. Do not put single quotation marks around the node name.

If you specify NONE, then this command will remove the current node name.

If you want to use a reserved word, such as "NONE", as a node name, then you must put the keyword in double quotation marks to ensure that it is treated as a delimited identifier. For example:

```
SET SYNC NODE "NONE"; -- Now the node name is "NONE"
```

You can verify the node name assignment with the following statement:

```
SELECT GET_PARAM('SYNC NODE')
```

The SET SYNC NODE NONE option removes the node name from the current catalog. This option is used when you are dropping a synchronized database and removing its registration.

Note:

When using the SET SYNC NODE NONE option, be sure the catalog associated with the node name is not defined as a master, replica, or both. To remove the node name, the catalog must be defined as SET SYNC MASTER NO and/or SET SYNC REPLICA NO. If you do try to set the node name to NONE on a master and/or replica catalog, solidDB returns error message 25082.

Usage in master: Use this statement in the master to set or remove the node name from the current catalog.

Usage in replica: Use this statement in the replica to set or remove the node name from the current catalog.

Example:

```
SET SYNC NODE SalesmanJones;
```

Return values: For details on each error code, see the appendix titled Error Codes in *solidDB Administration Guide*.

Table 75. SET SYNC NODE return values

Error code	Description
13047	No privilege for operation
13107	Illegal set operation
25059	After registration nodename cannot be changed
25082	Node name can not be removed if node is master or replica.

SET SYNC PARAMETER

```
SET SYNC PARAMETER parameter_name 'value_as_string';
SET SYNC PARAMETER parameter_name NONE;
```

Supported in: This command requires solidDB advanced replication.

Usage: This statement defines persistent catalog-level parameters that are visible via the parameter bulletin board to all transactions that are executed in that catalog. Each catalog has a different set of parameters.

If the parameter already exists, the new value overwrites the previous one. An existing parameter can be deleted by setting its value to NONE. All parameters are stored in the SYS_BULLETIN_BOARD system table.

These parameters are not propagated to the master.

In addition to system specific-parameters, you can also store in the system table a number of system parameters that configure the synchronization functionality. Available system parameters are listed at the end of the SQL reference.

Usage in master: Use the SET SYNC PARAMETER in the master for setting database parameters.

Usage in replica: Use the SET SYNC PARAMETER in replicas for setting database parameters.

Example:

```
SET SYNC PARAMETER db_type 'REPLICA'
SET SYNC PARAMETER db_type NONE
```

Return values: For details on each error code, see the appendix titled Error Codes in *solidDB Administration Guide*.

Table 76. SET SYNC PARAMETER Return Values

Error Code	Description
13086	Invalid data type in a parameter

See also: GET_PARAM

PUT_PARAM

SET SYNC PROPERTY

Syntax in Master:

```
SET SYNC PROPERTY <propertyname> = { 'value' | NONE } FOR REPLICA  
<replicaname>
```

Syntax in Replica:

```
SAVE SET SYNC PROPERTY <propertyname> = { 'value' | NONE }
```

Supported in: This command requires solidDB advanced replication.

Usage: This command allows you to specify a property name and value for a replica. Replicas that have properties may be grouped, and a group may be specified when using the START AFTER COMMIT statement. For example, you might have some replicas that are related to the bicycle industry and others that are related to the surfboard industry, and you may want to update each of those groups of replicas separately. You can use Property Names to group these replicas. All members of a group have the same property and have the same value for that property.

For more information, see the section titled "Replica Property Names" in *solidDB Advanced Replication Guide*.

Examples: *Master:*

```
SET SYNC PROPERTY color = 'red' FOR REPLICA replica1;  
SET SYNC PROPERTY color = NONE FOR REPLICA replica1;
```

Replica:

```
SAVE SET SYNC PROPERTY color = 'red';  
SAVE SET SYNC PROPERTY color = NONE;
```

SET SYNC USER

```
SET SYNC USER master_username IDENTIFIED BY password  
SET SYNC USER NONE
```

Supported in: This command requires solidDB advanced replication.

Usage: This statement is used to define the username and password for the registration process when the replica database is being registered in the master database. To use this command, you are required to have SYS_SYNC_ADMIN_ROLE access.

Note:

The SET SYNC USER statement is used for replica registration only. Aside from registration, all other synchronization operations require a valid master user ID in a replica database. If you want to designate a different master user for a replica, you must map the replica ID on the replica database with the master ID on the master database. For details, read the section titled "Mapping Replica User ID With Master User ID" in *solidDB Advanced Replication Guide*.

You define the registration username in the master database. The name you specify must have sufficient rights to execute the replica registration tasks. You can provide registration rights for a master user in the master database by designating the user with the SYS_SYNC_REGISTER_ROLE or the SYS_SYNC_ADMIN_ROLE using the GRANT *rolename* TO *user* statement.

After the registration has been successfully completed, you must reset the sync user to NONE; otherwise, if a master user saves statements, propagates messages, or refreshes from or registers to publications, the following error message is returned:

```
User definition not allowed for this operation.
```

Usage in master: This statement is not available in the master database.

Usage in replica: Use this statement in the replica to set the user name.

Example:

```
SET SYNC USER homer IDENTIFIED BY marge;  
SET SYNC USER NONE;
```

A.71.14 SET TIMEOUT

```
SET IDLE TIMEOUT { timeout_in_seconds |  
                  timeout_in_millisecondsMS | DEFAULT }  
SET LOCK TIMEOUT { timeout_in_seconds |  
                  timeout_in_millisecondsMS }  
SET OPTIMISTIC LOCK TIMEOUT { timeout_in_seconds |  
                              timeout_in_millisecondsMS }
```

- SET IDLE TIMEOUT sets the connection-specific maximum timeout in seconds. The setting is effective until a new timeout is given. If the timeout is set to DEFAULT, it means no maximum time.
- SET LOCK TIMEOUT sets the time in seconds that the engine waits for a lock to be released.

The SET LOCK TIMEOUT statement applies to tables used with the pessimistic concurrency control mode.

By default, lock timeout is 30 seconds. The maximum lock timeout is 1000 seconds; SET LOCK TIMEOUT of more than 1000 seconds fails.

SET LOCK TIMEOUT does not change the timeout for those tables for which the table-level timeout has been set with LOCK TABLE WAIT.

- SET OPTIMISTIC LOCK TIMEOUT sets the timeout for the row locks acquired with the SELECT FOR UPDATE statement in the optimistic concurrency control mode. The timeout affects only other SELECT FOR UPDATE statements competing for the same table rows. By default, the timeout is zero (no waiting).

By default, the timeout is defined in seconds. The lock timeout can be set at millisecond granularity by adding "MS" after the value, for example:

```
SET LOCK TIMEOUT 500MS;  
SET LOCK TIMEOUT 1500 MS;
```

Tip: Spacing of the "MS" is not significant, and you may use upper or lower case. Without the "MS", the lock timeout will be in seconds.

When the timeout interval is reached, solidDB terminates the timed-out statement. For more information, see "Setting lock timeout" on page 120.

A.72 SET TRANSACTION

The SET TRANSACTION commands have effect in the beginning of a transaction, and they affect the transaction until commit or abort. If the SET TRANSACTION statement is issued in the middle of a transaction, an error is returned.

The SET TRANSACTION command is based on ANSI SQL. However, the solidDB implementation has some differences from the ANSI definition. solidDB does not support combining two ANSI-defined transaction properties in one SET TRANSACTION statement, for example:

```
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE, READ WRITE;
```

However, solidDB supports multiple SET TRANSACTION statements in a single transaction, for example:

```
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE;  
SET TRANSACTION READ WRITE;
```

Examples

```
SET TRANSACTION DURABILITY RELAXED;  
SET TRANSACTION ISOLATION REPEATABLE READ;  
SET TRANSACTION READ WRITE;
```

See also

A.71, “SET,” on page 280

A.72.1 Differences between SET and SET TRANSACTION

The SET and SET TRANSACTION commands are used for setting various transaction properties, such as isolation level, the read level, or SQL passthrough mode.

The SET commands set the properties for a session and they are often called session-level commands. The SET TRANSACTION commands set the properties for one transaction and they are called transaction-level commands.

The transaction-level commands follow different rules from the session-level commands. These differences are listed below.

- The transaction-level commands take effect in the transaction in which they are issued; the session-level commands take effect in the next transaction, that is, after the next COMMIT WORK.
- The transaction-level commands apply to only the current transaction; the session-level commands apply to all subsequent transactions — that is, until the end of the session (connection) or until another SET command changes them.
- The transaction-level commands must be executed at the beginning of a transaction, that is, before any DML or DDL statements. However, they may be executed after other SET statements. If this rule is violated, an error is returned. The session-level commands may be executed at any point in a transaction.
- The transaction-level commands take precedence over the session-level commands. However, the transaction-level commands apply only to the current transaction. After the current transaction is committed or aborted, the settings will return to the value set by the most recent previous SET command (if any).
For example:

```
COMMIT WORK; -- Finish previous transaction;  
SET ISOLATION LEVEL SERIALIZABLE;  
COMMIT WORK; -- Isolation level is now SERIALIZABLE  
...  
COMMIT WORK;  
SET TRANSACTION ISOLATION LEVEL REPEATABLE READ;  
-- Isolation level is now REPEATABLE READ because  
-- transaction-level settings take precedence  
-- over session-level settings.
```



```
COMMIT WORK;
-- Isolation level is now back to SERIALIZABLE, since the
-- transaction-level settings applied only to that
-- transaction.
```

The complete precedence hierarchy for isolation level and read level settings is below. Items closer to the top of the list have higher precedence.

1. SET TRANSACTION ... (transaction-level settings)
2. SET ... (session-level settings)
3. The server-level settings specified by the value in `solid.ini` configuration parameter (for example, `IsolationLevel` or `DurabilityLevel`)
4. The server's default setting (parameter factory value).

A.72.2 SET TRANSACTION (read/write level)

```
SET TRANSACTION {READ ONLY | READ WRITE | WRITE}
```

The command `SET TRANSACTION { READ ONLY | READ WRITE | WRITE}` is based on ANSI SQL. It allows the user to specify whether the transaction is allowed to make any changes to data.

The transaction level read-write level setting can be used with dynamic load balancing in HotStandby configuration. If the assigned workload server is Secondary, it can be changed programmatically to the Primary for the time of one transaction. At the transaction level, the following statement changes the workload connection server to Primary for the time of one transaction:

```
SET TRANSACTION WRITE (nonstandard)
```

The affected transaction is the one that is started with the statement. After the transaction has been executed at the Primary, the workload connection server is reverted to the session default.

A.72.3 SET TRANSACTION DELETE CAPTURE

```
SET TRANSACTION DELETE CAPTURE {NONE | CHANGES | DEFAULT}
```

Usage

The `SET TRANSACTION DELETE CAPTURE` statement sets the data aging mode for the next transaction. Data aging removes table rows from the front-end but preserves them in the back-end.

- NONE: data aging is enabled
- CHANGES: data aging is disabled
- DEFAULT: the data capture mode returns to the mode set with a previous `SET DELETE CAPTURE` statement, or the server's default.

The `SET TRANSACTION DELETE CAPTURE` statement has effect in the beginning of a transaction, and it affects the transaction until commit or abort. If the statement is issued in the middle of a transaction, an error is returned.

See also

A.71.4, “`SET DELETE CAPTURE`,” on page 282

A.72.4 SET TRANSACTION DURABILITY

```
SET TRANSACTION DURABILITY {RELAXED | STRICT}
```

The command `SET TRANSACTION DURABILITY { RELAXED | STRICT }` controls whether the server uses *strict* or *relaxed* durability for transaction logging. This command is a solidDB extension to SQL; it is not part of the ANSI standard.

The `SET TRANSACTION DURABILITY` command must be given in the beginning of the transaction. The durability setting applies only to the current transaction, until the transaction is committed or aborted.

If the new transaction durability setting is `STRICT`, any previous committed transactions that have not yet been written to disk will be written at the time that the `SET TRANSACTION DURABILITY STRICT` transaction is committed. This means that the pending transactions are not written to disk as soon as the transaction durability level is changed to `STRICT`; the writes wait until the current transaction that was started with `SET TRANSACTION DURABILITY STRICT` is committed.

Note: There is no `DEFAULT` option in `SET TRANSACTION DURABILITY` to set the value to the whatever value the `Logging.DurabilityLevel` parameter has specified. Instead, after the transaction has been committed (or aborted), the server returns to the durability level that was set before the `SET TRANSACTION DURABILITY` statement was executed.

A.72.5 SET TRANSACTION ISOLATION LEVEL

```
SET TRANSACTION ISOLATION LEVEL {  
    READ COMMITTED |  
    REPEATABLE READ |  
    SERIALIZABLE}
```

The command `SET TRANSACTION ISOLATION` is based on ANSI SQL. It sets the transaction isolation level to one of the following:

- `READ COMMITTED`
- `REPEATABLE READ`
- `SERIALIZABLE`

For more details on the isolation levels, see *Choosing transaction isolation levels*.

A.72.6 SET TRANSACTION PASSTHROUGH

```
SET TRANSACTION PASSTHROUGH {READ <passthrough level> [WRITE <passthrough level>]}  
| {WRITE <passthrough level> | [READ <passthrough level>]}  
| <passthrough level>
```

where

```
passthrough level ::= NONE | CONDITIONAL | FORCE | DEFAULT
```

Usage

The `SET TRANSACTION PASSTHROUGH` statement sets the SQL passthrough mode for the next transaction in solidDB Universal Cache.

- `NONE`: SQL passthrough is not used; no commands are passed from the front-end to the back-end
- `CONDITIONAL`: SQL passthrough is activated by a missing table or a syntax error.

- FORCE: SQL passthrough is used to pass all statements from the front-end to the back-end
- DEFAULT: causes the current session default for SQL passthrough is used

The SET TRANSACTION PASSTHROUGH statement has effect in the beginning of a transaction, and it affects the transaction until commit or abort. If the statement is issued in the middle of a transaction, an error is returned.

See also

A.71.7, “SET PASSTHROUGH,” on page 284

A.72.7 SET TRANSACTION SAFENESS

```
SET TRANSACTION SAFENESS {1SAFE | 2SAFE | DEFAULT}
```

SET TRANSACTION SAFENESS determines whether the replication protocol is synchronous (2-safe) or asynchronous (1-safe).

- 1-safe: the transaction is first committed at Primary and then transmitted to Secondary
- 2-safe: the transaction is not committed before it has been acknowledged by Secondary (default).

SET TRANSACTION SAFENESS sets the safeness level for the current transaction.

A.73 START AFTER COMMIT

```
START AFTER COMMIT
  [FOR EACH REPLICA WHERE search_condition [RETRY retry_spec]]
  {UNIQUE | NONUNIQUE} stmt;
```

stmt ::= any SQL statement.
search_condition ::= *search_item* | *search_item* {AND|OR } *search_item*
search_item ::= {*search_test* | (*search_condition*)}
search_test ::= *comparison_test* | *like_test*
comparison_test ::= *property_name* { = | <> | > | >= | < | <= } *value*
property_name ::= name of a replica property
like_test ::= *property_name* [NOT] LIKE *value* [ESCAPE *value*]
value ::= *literal*
retry_spec ::= *seconds*,*count*

Usage

The START AFTER COMMIT statement specifies an SQL statement (such as a call to a stored procedure) that will be executed when the current transaction commits. If the transaction is rolled back, the specified SQL statement is not executed.

A statement started in the background using START AFTER COMMIT is executed in a separate transaction. That transaction is executed in autocommit mode, that is, it cannot be rolled back as soon as it has started.

The START AFTER COMMIT statement returns a result set with one INTEGER column. This integer is a unique "job" id and can be used to query the status of a statement that failed to start due to an invalid SQL statement, insufficient access rights, replica not available, and so on.

If you use the UNIQUE keyword before the <stmt>, that statement will be executed only if there is not already an identical statement executing or pending. Statements

are compared using simple string compare. For example call `foo(1)` is different from call `foo(2)`. The server also takes into account whether the statement already being executed (or pending for execution) is on the same replica or a different replica; only identical statements on the same replica are discarded.

Important: When duplicate statements are discarded by using the `UNIQUE` keyword, the most recent statements are the ones thrown out, and the oldest one is the one that keeps running. It is quite possible to create a situation where you do multiple updates, for example, and you trigger multiple `START AFTER COMMIT` operations, but only the oldest one executes and thus the newest updated data may not get sent to the replicas immediately.

`NONUNIQUE` means that duplicate statements can be executed simultaneously in the background.

`FOR EACH REPLICA` specifies that the statement is executed for each replica that fulfills the property conditions given in the `search_condition` part of the `WHERE` clause. Before executing the statement, a connection to the replica is established. If a procedure call is started, the procedure can get the "current" replica name using the keyword `DEFAULT`.

If `RETRY` is specified, the operation is re-executed after `N` seconds (defined by seconds in the `retry_spec`) if the replica is not reached on the first attempt. The count specifies how many times a retry is attempted.

Context of background statements

Statements started in the background are executed in the context of the user who issued the `START AFTER COMMIT` statement, and are executed in the catalog and schema in which the `START AFTER COMMIT` statement executed.

In the example below, `'CALL FOO'` is executed in the catalog `'katmandu'` and the schema `'steinbeck'`.

```
SET CATALOG katmandu;
SET SCHEMA steinbeck;
START AFTER COMMIT UNIQUE CALL FOO;
COMMIT WORK;
SET CATALOG irrelevant_catalog;
SET SCHEMA irrelevant_schema
```

Durability

Background statements are not durable. In other words, the execution of statements started with `START AFTER COMMIT` is not guaranteed.

Rollback

Background statements cannot be rolled back after they have been started. After a statement that has been started with `START AFTER COMMIT` has executed successfully, there is no way to roll it back.

The `START AFTER COMMIT` statement itself can be rolled back, and this will prevent the specified statement from executing. For example,

```
START AFTER COMMIT UNIQUE INSERT INTO MyTable VALUES (1);
ROLLBACK;
```

In the example above, the transaction rolls back and thus INSERT INTO MyTable VALUES (1) will not be executed.

Order of execution

Background statements are executed asynchronously and they do not have any guaranteed order even inside a transaction.

Examples

Start local procedure in the background.

```
START AFTER COMMIT NONUNIQUE CALL myproc;
```

Start the call if "CALL myproc" is not running in the background already.

```
START AFTER COMMIT UNIQUE call myproc;
```

Start procedure in the background using replicas which have property "color" = "blue".

```
START AFTER COMMIT FOR EACH REPLICA WHERE color='blue' UNIQUE CALL myproc;
```

The following statements are all considered different and therefore each is executed, despite the presence of the keyword UNIQUE. name is a unique property of each replica.

```
START AFTER COMMIT UNIQUE call myproc;  
START AFTER COMMIT FOR EACH REPLICA WHERE name='R1' UNIQUE call myproc;  
START AFTER COMMIT FOR EACH REPLICA WHERE name='R2' UNIQUE call myproc;  
START AFTER COMMIT FOR EACH REPLICA WHERE name='R3' UNIQUE call myproc;
```

But if the following statement is executed in the same transaction as the previous ones and the condition "color='blue'" matches some of the replicas R1, R2 or R3, the call is not executed for those replicas again.

```
START AFTER COMMIT FOR EACH REPLICA WHERE color='blue' UNIQUE call myproc;
```

Related information:

3, "SQL extensions," on page 23

The SQL support in solidDB is comparable to any advanced SQL-based system; solidDB offers the most commonly expected features and a set of useful extensions employing solidDB-specific (nonstandard) SQL syntax. Procedural SQL extensions such as stored procedures, stored functions, and triggers enable moving parts of the application logic into the database. These extensions help reduce network traffic, thus improving performance.

A.74 TRUNCATE TABLE

```
TRUNCATE TABLE tablename
```

Usage

The TRUNCATE statement removes all rows from a table. The TRUNCATE statement is semantically equivalent to a DELETE FROM *tablename* statement. However, it is much more efficient due to relaxed isolation.

Restrictions

During the execution of this statement, the defined isolation level is not maintained in concurrent transactions. The effect of removing the rows is seen

immediately in all concurrent transactions. Therefore, this statement is recommended for maintenance purposes only.

When the truncated table participates in a referential integrity constraint as a referenced table, the TRUNCATE statement fails if the referencing table is not empty, regardless of the referential action defined (RESTRICT, CASCADE, and so on). The restriction does not apply if the referencing table is the same as the referenced table (a tree-structured table).

If a collection of interrelated tables is to be truncated, the TRUNCATE statements should be issued in the order starting from the tables that are referencing tables only, up the way of the referencing chain, and ending up with the table that is a referenced table only.

In Universal Cache setups where solidDB is a source database, the TRUNCATE TABLE statement is not allowed on tables which are part of a subscription. If this rule is violated, an error is returned to the application.

A.75 UNLOCK TABLE

```
UNLOCK TABLE { ALL | tablename [,tablename]}
```

where

tablename is the name of the table to unlock. You can also specify the catalog and schema of the table by qualifying the table name.

ALL releases all table-level locks on all tables.

Usage

The UNLOCK TABLE statement unlock tables that you have locked manually by using the LOCK TABLE command with the LONG option. The LONG option allows you to hold a lock past the end of the transaction in which the lock was placed. Since there is no natural endpoint for the lock (other than the end of the transaction), you must explicitly release a LONG lock by using the UNLOCK command.

The UNLOCK TABLE command does not apply to the server's automatic locks, or to manual locks that were not locked with the LONG option. If a lock is automatic, or if it is manual and not LONG, the server will automatically release the lock at the end of the transaction in which the lock was placed. Thus there is no need to manually unlock those locks.

When the UNLOCK TABLE command is used, it does not take effect immediately; instead, the locks are released when the current transaction is committed.

Note:

If the current transaction (the one in which the UNLOCK TABLE command was executed) is not committed (for example, if it is rolled back), the tables are not unlocked; they will remain locked until another UNLOCK TABLE command is successfully executed and committed.

The LOCK/UNLOCK commands apply only to tables. There is no command to manually lock or unlock individual records.

If you have a table named ALL, you need to use the delimited identifier feature to specify the table name.

Return values

Table 77. UNLOCK TABLE return values

Error code	Description
10083	Table <table_name> not locked.
13011	Table <tablename> not found.

Examples

```
LOCK TABLE emp IN SHARED MODE;
LOCK TABLE emp IN SHARED MODE TABLE dept IN EXCLUSIVE MODE;
LOCK TABLE emp,dept IN SHARED MODE NOWAIT;

-- Get an exclusive lock that will persist past the end of the current
-- transaction. If you can't get an exclusive lock immediately, then
-- wait up to 60 seconds to get it.
LOCK TABLE emp, dept IN LONG EXCLUSIVE MODE WAIT 60;
-- Make the schema changes (or do whatever you needed the exclusive
-- lock for).
CALL DO_SCHEMA_CHANGES_1;
COMMIT WORK;
CALL DO_SCHEMA_CHANGES_2;
UNLOCK TABLE ALL; -- at the end of this transaction, release locks.
...
COMMIT WORK;
...
UNLOCK TABLE "ALL"; -- Unlock the table named "ALL".
```

Related reference:

A.59, "LOCK TABLE," on page 247

A.76 UNREGISTER EVENT

```
wait_register_statement ::= UNREGISTER EVENT event_name
```

The UNREGISTER EVENT statement unregisters the specified event. The UNREGISTER EVENT statement is allowed only inside stored procedures.

For more details, see A.15.7, "wait_register_statement," on page 185.

Related reference:

A.15.7, "wait_register_statement," on page 185

Related information:

3.5, "Events," on page 80

A.77 UPDATE

```
UPDATE table_name
  SET [table_name.]column_identifier = {expression | NULL}
  [, [table_name.]column_identifier = {expression | NULL}]...
  [WHERE search_condition]

UPDATE table_name
  SET [table_name.]column_identifier = {expression | NULL}
  [, [table_name.]column_identifier = {expression | NULL}]...
  WHERE CURRENT OF cursor_name
```

Usage

There are two types of UPDATE statements:

- The *searched* UPDATE statement, `UPDATE table_name ... [WHERE search_condition]`, modifies the values of one or more columns in one or more rows specified with A.79.5, “search_condition,” on page 308.
- The *positioned* UPDATE statement, `UPDATE table_name ... WHERE CURRENT OF cursor_name`, updates the current row of the cursor. The name of the cursor is defined using ODBC API function named `SQLSetCursorName`.

table_name can be a table name or an alias.

Example of searched UPDATE

```
UPDATE TEST SET C = 0.44
WHERE ID = 5
```

Example of positioned UPDATE

```
UPDATE TEST SET C = 0.33
WHERE CURRENT OF MYCURSOR
```

A.78 WAIT EVENT

```
wait_event_statement ::=
  WAIT EVENT
  [event_specification ...]
  END WAIT
```

```
event_specification ::=
  WHEN event_name [(parameters)]
  BEGIN statements
  END EVENT
```

The WAIT EVENT statement makes procedures to wait for an event to happen. You can use WAIT EVENT with both system-defined events and user-defined events.

For more details, see A.15.8, “wait_event_statement,” on page 186.

Related reference:

A.15.8, “wait_event_statement,” on page 186

Related information:

3.5, “Events,” on page 80

A.79 Common clauses

The solidDB SQL syntax reference uses various common clauses.

A.79.1 check_condition

Table 78. check_condition

check_condition	
<i>check_condition</i>	::= <i>check_item</i> <i>check_item</i> { AND OR } <i>check_item</i>
<i>check_item</i>	::= [NOT] { <i>check_test</i> (<i>check_condition</i>) }

Table 78. *check_condition* (continued)

check_condition	
<i>check_test</i>	::= <i>comparison_test</i> <i>between_test</i> <i>like_test</i> <i>null_test</i> <i>list_test</i>
<i>comparison_test</i>	::= <i>expression</i> { = <> < <= > >= } { <i>expression</i> <i>subquery</i> }
<i>between_test</i>	::= <i>column_identifier</i> [NOT] BETWEEN <i>expression</i> AND <i>expression</i>
<i>like_test</i>	::= <i>column_identifier</i> [NOT] LIKE <i>value</i> [ESCAPE <i>value</i>]
<i>null_test</i>	::= <i>column_identifier</i> IS [NOT] NULL
<i>list_test</i>	::= <i>expression</i> [NOT] IN ({ <i>value</i> [, <i>value</i>]... })

A.79.2 data_type

Table 79. *data_type*

Variable name	Data type
<i>data_type</i>	::= {BIGINT BINARY BLOB CHAR [<i>length</i>] CHARACTER LARGE OBJECT CHAR LARGE OBJECT CLOB DATE DECIMAL [(<i>precision</i> [, <i>scale</i>])] DOUBLE [PRECISION] FLOAT [(<i>precision</i>)] INTEGER LONG NATIONAL VARCHAR LONG VARBINARY LONG VARCHAR LONG WVARCHAR NCHAR LARGE OBJECT NUMERIC [(<i>precision</i> [, <i>scale</i>])] NATIONAL CHAR NATIONAL CHARACTER NATIONAL VARCHAR NCHAR NCHAR VARYING NCLOB NVARCHAR REAL SMALLINT TIME TIMESTAMP [(<i>timestamp precision</i>)] TINYINT VARBINARY VARCHAR [(<i>length</i>)] } WCHAR WVARCHAR [<i>length</i>]

A.79.3 expression

Table 80. Expression

Expression	
<i>expression</i>	<p>::= <i>expression_item</i> <i>expression_item</i> { + - * / } <i>expression_item</i></p> <p>Note: Spaces on each side of the operator are optional.</p>
<i>expression_item</i>	<p>::= [+ -] { <i>value</i> <i>column_identifier</i> <i>function</i> <i>case_expression</i> <i>cast_expression</i> (<i>expression</i>) }</p>
<i>value</i>	<p>::= <i>literal</i> USER <i>variable</i></p>
<i>function</i>	<p>::= <i>set_function</i> <i>null_function</i> <i>string_function</i> <i>numeric_function</i> <i>datetime_function</i> <i>system_function</i> <i>datatypeconversion_function</i></p> <p>Note: The string, numeric, datetime, and datatype conversion functions are scalar functions in which an operation denoted by a function name is followed by a pair of parenthesis enclosing zero or more specified arguments. Each scalar function returns a single value.</p>
<i>set_function</i>	<p>::= COUNT (*) { AVG MAX MIN SUM COUNT } ({ ALL DISTINCT } <i>expression</i>)</p>
<i>null_function</i>	<p>::= { NULLVAL_CHAR() NULLVAL_INT() }</p>
<i>datatypeconversion_function</i>	<p>::= CONVERT_CHAR(<i>value_exp</i>) CONVERT_DATE(<i>value_exp</i>) CONVERT_DECIMAL(<i>value_exp</i>) CONVERT_DOUBLE(<i>value_exp</i>) CONVERT_FLOAT(<i>value_exp</i>) CONVERT_INTEGER(<i>value_exp</i>) CONVERT_LONGVARCHAR(<i>value_exp</i>) CONVERT_NUMERIC(<i>value_exp</i>) CONVERT_REAL(<i>value_exp</i>) CONVERT_SMALLINT(<i>value_exp</i>) CONVERT_TIME(<i>value_exp</i>) CONVERT_TIMESTAMP(<i>value_exp</i>) CONVERT_TINYINT(<i>value_exp</i>) CONVERT_VARCHAR(<i>value_exp</i>)</p> <p>Note: These functions are used to implement the {fn CONVERT(<i>value</i>, <i>odbc_typename</i>)} escape clauses defined by ODBC. However, you should use CAST(<i>value</i> AS <i>sql_typename</i>) which is defined in SQL-92 and fully supported by solidDB.</p>
<i>case_expression</i>	<p>::= <i>case_abbreviation</i> <i>case_specification</i></p>

Table 80. Expression (continued)

Expression	
<i>case_abbreviation</i>	<pre>::= NULLIF(<i>value_exp</i>, <i>value_exp</i>) COALESCE(<i>value_exp</i> {, <i>value_exp</i> }...)</pre> <p>The NULLIF function returns NULL if the first parameter is equal to the second parameter; otherwise, it returns the first parameter. It is equivalent to IF (p1 = p2) THEN RETURN NULL ELSE RETURN p1; The NULLIF function is useful if you have a special value that serves as a flag to indicate NULL. You can use NULLIF to convert that special value to NULL. In other words, it behaves like IF (p1 = NullFlag) THEN RETURN NULL ELSE RETURN p1;</p> <p>COALESCE returns the first non-NULL argument. The list of arguments may be of almost any length. All arguments should be of the same (or compatible) data types.</p>
<i>case_specification</i>	<pre>::= CASE [<i>value_exp</i>] WHEN <i>value_exp</i> THEN {<i>value_exp</i> } [WHEN <i>value_exp</i> THEN { <i>value_exp</i> } ...] [ELSE { <i>value_exp</i> }] END</pre>
<i>cast_expression</i>	<pre>::= CAST (<i>value_exp</i> AS <i>-data-type</i>)</pre>
<i>row value constructor expression</i>	<p>A row value constructor (RVC) is an ordered sequence of values delimited by parentheses, for example:</p> <pre>(1, 4, 9)</pre> <pre>('Smith', 'Lisa')</pre> <p>A row construction is based on a series of elements/values, similar to a row of a table that is composed of a series of fields.</p> <p>For more information about row value constructors, see 2.5.5, “Row value constructors,” on page 19.</p>

A.79.4 query_specification

Table 81. query_specification

query_specification	
<i>query_specification</i>	<pre>::= SELECT [DISTINCT ALL] <i>select_list</i> <i>table_expression</i></pre>
<i>select_list</i>	<pre>::= * <i>select_sublist</i> [{, <i>select_sublist</i> } ...]</pre>
<i>select_sublist</i>	<pre>::= <i>derived_column</i> [<i>table_name</i> <i>table_identifier</i>].*</pre>
<i>derived_column</i>	<pre>::= <i>expression</i> [[AS] <i>column_alias</i>]]</pre>

Table 81. query_specification (continued)

query_specification	
table_expression	<pre> ::= FROM table_reference_list [WHERE search_condition] [GROUP BY column_name_list [[UNION INTERSECT EXCEPT] [ALL] [CORRESPONDING [BY (column_name_list)]] query_specification] [HAVING search_condition]</pre>

A.79.5 search_condition

Table 82. search_condition

search_condition	
search_condition	<pre> ::= search_item search_item { AND OR } search_item</pre>
search_item	<pre> ::= [NOT] { search_test (search_condition) }</pre>
search_test	<pre> ::= comparison_test between_test like_test null_test set_test quantified_test existence_test</pre>
comparison_test	<pre> ::= expression { = <> < <= > >= } { expression subquery }</pre> <p>Note: Spaces on each side of the operator are optional.</p>
between_test	<pre> ::= column_identifier [NOT] BETWEEN expression AND expression</pre>
like_test	<pre> ::= column_identifier [NOT] LIKE value [ESCAPE value]</pre>
null_test	<pre> ::= column_identifier IS [NOT] NULL</pre>
set_test	<pre> ::= expression [NOT] IN ({ value [, value]... subquery })</pre> <p>Note: If the IN clause contains only one item, parenthesis are not needed.</p> <pre> ::= expression [NOT] IN value</pre> <p>For example:</p> <pre> SELECT * FROM MY_TABLE1 WHERE A IN 2; SELECT * FROM MY_TABLE1 WHERE A IN (2,3,6);</pre>
quantified_test	<pre> ::= expression { = <> < <= > >= } [ALL ANY SOME] subquery</pre>
existence_test	<pre> ::= EXISTS subquery</pre>

A.79.6 table_reference

Table 83. table_reference

Table_reference	
<i>table_reference_list</i>	::= <i>table_reference</i> [, <i>table-reference</i> ...]
<i>table_reference</i>	::= <i>table_name</i> [[AS] <i>correlation_name</i>] <i>derived_table</i> [[AS] <i>correlation_name</i>] [(<i>derived_column_list</i>)] <i>joined_table</i>
<i>table_name</i>	::= <i>table_identifier</i> <i>schema_name.table_identifier</i>
<i>derived_table</i>	::= <i>subquery</i>
<i>derived_column_list</i>	::= <i>column_name_list</i>
<i>joined_table</i>	::= <i>cross_join</i> <i>qualified_join</i> (<i>joined_table</i>)
<i>cross_join</i>	::= <i>table_reference</i> CROSS JOIN <i>table_reference</i>
<i>qualified_join</i>	::= <i>table_reference</i> [NATURAL] [<i>join_type</i>] JOIN <i>table_reference</i> [<i>join_specification</i>]
<i>join_type</i>	::= INNER <i>outer_join_type</i> [OUTER] UNION
<i>outer_join_type</i>	::= LEFT RIGHT FULL
<i>join_specification</i>	::= <i>join_condition</i> <i>named_columns_join</i>
<i>join_condition</i>	::= ON <i>search_condition</i>
<i>named_columns_join</i>	::= USING (<i>column_name_list</i>)
<i>column_name_list</i>	::= <i>column_identifier</i> [{ , <i>column_identifier</i> } ...]

A.79.7 Pseudo columns in SELECT statement

The pseudo columns described in this section can be used in the select-list of a SELECT statement.

Table 84. Pseudo columns

Pseudo column	Type	Explanation
ROWVER	VARBINARY(10)	Version of the row in a table. Note: Because ROWVER refers to a single row, it can only be used with queries that return rows from a single table.
ROWID	VARBINARY(254)	Persistent id for a row in a table. Note: Because ROWID refers to a single row, it can only be used with queries that return rows from a single table.

Table 84. Pseudo columns (continued)

Pseudo column	Type	Explanation
ROWNUM	DECIMAL(16,2)	<p>Row number indicates the sequence in which a row was selected from a table or set of joined rows. The first row selected has a ROWNUM of 1, the second row has 2, and so on.</p> <p>Because ROWNUM is given to a row before the order by clause is evaluated, you cannot use ROWNUM to identify sorted rows.</p> <p>ROWNUM is chiefly useful for limiting the number of rows returned by a query for example, WHERE ROWNUM < 10.</p> <p>Note: You cannot use the following constructions:</p> <ul style="list-style-type: none"> • ROWNUM = n • ROWNUM ≠ n • ROWNUM IN (n) <p>For example, the following constructions are not supported:</p> <ul style="list-style-type: none"> • ROWNUM = 2 • ROWNUM IN (2,4,6) <p>Instead, use the constructions <, >, ≤, ≥, or BETWEEN.</p> <p>For example:</p> <pre>SELECT * FROM TABLE1 WHERE Y = 0 AND ROWNUM < 2; SELECT * FROM TABLE1 WHERE Y = 0 AND (ROWNUM > 1 AND ROWNUM < 3); SELECT * FROM TABLE1 WHERE Y = 0 AND ROWNUM BETWEEN 1 AND 3;</pre> <p>Additionally, comparing a ROWNUM condition to a result set of a subquery is not supported.</p> <p>For example, the following construction is not supported:</p> <pre>SELECT * FROM TABLE1 WHERE Y = 0 AND ROWNUM < (SELECT COUNT(*) FROM TABLE1 WHERE Y = 0);</pre> <p>ROWNUM can, however, be used within a subquery. For example, the following is supported:</p> <pre>SELECT * FROM (SELECT ROWNUM R, * FROM TABLE WHERE Y = 0) WHERE R = 5;</pre>

A.80 Date and time literals

Table 85. Date and time literals

Date/time literal	
<i>date_literal</i>	'YYYY-MM-DD'
<i>time_literal</i>	'HH:MM:SS'
<i>timestamp_literal</i>	'YYYY-MM-DD HH:MM:SS'

A.81 Wildcard characters

The following may be used as wildcard characters in certain expressions, such as LIKE '<string>'.

Table 86. Wildcard characters

Character	Explanation
_ (underscore)	The underscore character matches any single character. For example, 'J_NE' matches 'JANE' and 'JUNE'.
% (percent sign)	The percent sign character matches any group of 0 or more characters. For example 'ED%' matches 'EDWARD' and 'EDITOR'. As another example, '%ED%' matches 'EDWARD', 'TEDDY', and 'FRED'.

A.81.1 Using SQL wildcards

Exact match searches are conducted by specifying literal values, as in:

```
SELECT * FROM table1 WHERE name = 'SMITH';
```

The string 'SMITH' is a literal value.

Similar match searches are conducted by specifying an SQL wildcard that represents a character string that is similar to another character string. Logical expressions (such as those used in WHERE clauses and CHECK constraints) may use the "wildcard" characters and the keyword LIKE to match strings that are similar.

The underscore character (_) is a wildcard character that matches any single character. For example, the following query:

```
SELECT * FROM table1 WHERE first_name LIKE 'J_NE';
```

returns both JANE and JUNE (as well as any other four-character name where the first letter is J and the last two letters are NE).

The percent character (%) is a wildcard character that matches any occurrence of 0 or more characters. For example, the following query:

```
SELECT * FROM table1 WHERE first_name LIKE 'JOHN%';
```

could return JOHN, JOHNNY, JOHNATHAN, and so on.

The % wildcard is used most often at the end of strings, but it can be used anywhere. For example, the following search pattern:

```
LIKE '%JO%'
```

returns all people who have JO somewhere in their name, included but not limited to:

JOANNE, BILLY JO, and LONG JOHN SILVER

Multiple wildcards are allowed in a single string. For example, the string J_V_ matches JAVA and JIVE and any other four-character words or names that start

with J and have V as the third character. Note that because the underscore (_) only matches exactly one character, the string J_V_ does not match the string JOVIAL, which has more than four characters.

A.81.2 Wildcard characters as literals

A wildcard character may be used in one part of a string while the literal character % (percent) or underscore (_) may be used in another part of the same string.

To use a wildcard character as a literal, preface the wildcard character with an escape character and specify the escape character itself as part of the query. The syntax is the following:

```
... LIKE <string> {ESCAPE <escape character>}
```

For example, the expression below uses the backslash character (\) as the escape character:

```
LIKE 'MY\_EXPRESSION_' ESCAPE '\\';
```

The above expression matches the following:

- MY_EXPRESSION1
- MY_EXPRESSIONA
- MY_EXPRESSION_

The above expression does not match the following:

- MY#EXPRESSION1

Remember: As per ANSI standard, character strings must be delimited by single quotation marks.

For example:

```
...LIKE 'J_N_'; -- CORRECT  
...LIKE "J_N_"; --WRONG
```

Double quotation marks are used for delimited identifiers, not data. This is different from C and Java languages as they use double quotation marks to delimit strings as in "*C-language string*" and single quotation marks 'C' to delimit single characters.

Appendix B. Functions

A *function* is an operation that is denoted by a function name followed by a pair of parentheses enclosing the specification of arguments (there may be no arguments). In addition to built-in functions, solidDB supports both internal and external user-defined stored functions.

Built-in functions are provided within the database engine.

User-defined functions are registered to the database using the CREATE FUNCTION statement. There are two types of user-defined functions:

- User-defined *stored functions* can be written in the solidDB proprietary SQL procedure language.
- User-defined *external stored functions* can be written in C programming language. The external stored functions are loaded at runtime using standard dynamic library interfaces as provided by the operating system. External stored functions enable you to expand the server functionality by writing a C program.

Related reference:

A.12, “CREATE FUNCTION,” on page 173

A.13, “CREATE FUNCTION (external),” on page 175

B.1 String functions

Table 87. String functions

Function	Purpose
ASCII(<i>str</i>)	Returns the integer equivalent of string <i>str</i>
CHAR(<i>code</i>)	Returns the character equivalent of <i>code</i>
CONCAT(<i>str1</i> , <i>str2</i>)	Concatenates <i>str2</i> to <i>str1</i>
<i>str1</i> { + } <i>str2</i>	Concatenates <i>str2</i> to <i>str1</i> . For example: SELECT <i>str1</i> + <i>str2</i> , <i>col1</i> ... SELECT <i>str1</i> <i>str2</i> , <i>col1</i> ...
GET_UNIQUE_STRING(<i>str</i>)	This function generates a unique string, based on a “prefix” (the input string, which may be any string you choose) and a sequence number (which is created and used internally). If the input is NULL, the function still returns a string based on the unique sequence number.
INSERT(<i>str1</i> , <i>start</i> , <i>length</i> , <i>str2</i>)	Merges strings by deleting <i>length</i> characters from <i>str1</i> and inserting <i>str2</i>
LCASE(<i>str</i>)	Converts string <i>str</i> to lowercase
LEFT(<i>str</i> , <i>count</i>)	Returns leftmost <i>count</i> characters of string <i>str</i>
LENGTH(<i>str</i>)	Returns the number of characters in <i>str</i>

Table 87. String functions (continued)

Function	Purpose
LOCATE(<i>str1</i> , <i>str2</i> [, <i>start</i>])	Returns the starting position of <i>str1</i> within <i>str2</i> . If the optional argument, <i>start</i> , is specified, the search begins with the character position indicated by the value of <i>start</i> . If <i>string_exp1</i> is not found within <i>string_exp2</i> , the function returns 0. For both the return value and the input parameter <i>start</i> , string positions are numbered starting from 1 (not 0).
LTRIM(<i>str</i>)	Removes leading spaces of <i>str</i>
POSITION (<i>str1</i> IN <i>str2</i>)	Returns starting position of <i>str1</i> within <i>str2</i>
REPEAT(<i>str</i> , <i>count</i>)	Returns characters of <i>str</i> repeated <i>count</i> times
REPLACE(<i>str1</i> , <i>str2</i> , <i>str3</i>)	Replaces occurrences of <i>str2</i> in <i>str1</i> with <i>str3</i>
RIGHT(<i>str</i> , <i>count</i>)	Returns the rightmost count characters of string <i>str</i>
RTRIM(<i>str</i>)	Removes trailing spaces in <i>str</i>
SOUNDEX(<i>str</i>)	Calculate 4-character soundex (phonetic) code
SPACE(<i>count</i>)	Returns a string of count spaces
SUBSTRING(<i>str</i> , <i>start</i> , <i>length</i>)	Derives substring <i>length</i> bytes long from <i>str</i> beginning at <i>start</i> . For example, if <i>str</i> ="First Second Third", then SUBSTRING(<i>str</i> , 7, 6) would return "Second". String positions are numbered starting from 1 (not 0).
TO_CHAR(<i>expression</i> , <i>format-string</i>)	Returns a character representation of an input expression that has been formatted using a character template. The TO_CHAR scalar function is a synonym for the VARCHAR_FORMAT scalar function.
TRIM(<i>str</i>)	Removes leading and trailing spaces in <i>str</i>
UCASE(<i>str</i>)	Converts <i>str</i> to uppercase
VARCHAR_FORMAT(<i>expression</i> , <i>format-string</i>)	Returns a character representation of an input expression that has been formatted using a character template. The VARCHAR_FORMAT function is a synonym for the TO_CHAR function.

Related information:

A.81, "Wildcard characters," on page 311

B.2 Numeric functions

Table 88. Numeric functions

Function	Purpose
ABS(<i>numeric</i>)	Absolute value of <i>numeric</i>
ACOS(<i>float</i>)	Arccosine of <i>float</i> , where <i>float</i> is expressed in radians
ASIN(<i>float</i>)	Arcsine of <i>float</i> , where <i>float</i> is expressed in radians

Table 88. Numeric functions (continued)

Function	Purpose
ATAN(float)	Arctangent of <i>float</i> , where <i>float</i> is expressed in radians
ATAN2(float1, float2)	Arctangent of the <i>x</i> and <i>y</i> coordinates, specified by <i>float1</i> and <i>float2</i> , respectively, as an angle, expressed in radians
CEILING(numeric)	Smallest integer greater than or equal to <i>numeric</i>
COS(float)	Cosine of <i>float</i> , where <i>float</i> is expressed in radians
COT(float)	Cotangent of <i>float</i> , where <i>float</i> is expressed in radians
DEGREES(numeric)	Converts <i>numeric</i> radians to degrees
DIFFERENCE(str1, str2)	Return the value of phonetic difference: 0 - 4
EXP(float)	Exponential value of <i>float</i>
FLOOR(numeric)	Largest integer less than or equal to <i>numeric</i>
LOG(float)	Natural logarithm of <i>float</i>
LOG10(float)	Base 10 log of <i>float</i>
MOD(integer1, integer2)	Modulus of <i>integer1</i> divided by <i>integer2</i>
PI()	Pi as a floating point number
POWER(numeric, integer)	Value of <i>numeric</i> raised to the power of <i>integer</i>
RADIANS(numeric)	Converts from <i>numeric</i> degrees to radians
ROUND(numeric, integer)	<i>Numeric</i> rounded to <i>integer</i>
SIGN(numeric)	Sign of <i>numeric</i>
SIN(float)	Sine of <i>float</i> , where <i>float</i> is expressed in radians
SQRT(float)	Square root of <i>float</i>
TAN(float)	Tangent of <i>float</i> , where <i>float</i> is expressed in radians
TRUNCATE(numeric, integer)	<i>Numeric</i> truncated to <i>integer</i>

B.3 Date time functions

Table 89. Date time functions

Function	Purpose
ADD_MONTHS(<i>expression</i> , <i>numeric-expression</i>)	Returns a datetime value that represents <i>expression</i> plus a specified number of months. <i>expression</i> specifies the starting date. The expression must return a value of one of the following built-in data types: a DATE or a TIMESTAMP. <i>numeric-expression</i> returns a value of any built-in numeric data type. If the value is not of type INTEGER, it is implicitly cast to INTEGER before evaluating the function. The numeric-expression specifies the number of months to add to the starting date specified by <i>expression</i> . A negative numeric value is allowed.
CURDATE() CURRENT_DATE CURRENT DATE	Returns the current date
CURTIME() CURRENT_TIME CURRENT TIME	Returns the current time
CURRENT_TIMESTAMP CURRENT TIMESTAMP	Returns the current date and time as a timestamp
DATE(<i>expression</i>)	Returns a date from a value. The argument must be a DATE, TIMESTAMP, a positive number less than or equal to 3 652 059, or a valid string representation of a date or timestamp. <ul style="list-style-type: none"> • If the argument is a DATE, TIMESTAMP, or valid string representation of a date or timestamp, the result is the date part of the value. • If the argument is a string of length 7, it must represent a valid date in the form YYYYNNN, where YYYY are digits denoting a year, and NNN are digits between 001 and 366, denoting a day of that year. • If the argument is a number, the result is the date that is n-1 days after January 1, 0001, where n is the integral part of the number.
DAYNAME(<i>date</i>)	Returns a string with the day of the week
DAYOFMONTH(<i>date</i>)	Returns the day of the month as an integer between 1 and 31
DAYOFWEEK(<i>date</i>)	Returns the day of the week as an integer between 1 and 7, where 1 represents Sunday
DAYOFYEAR(<i>date</i>)	Returns the day of the year as an integer between 1 and 366
EXTRACT (<i>date field</i> FROM <i>date_exp</i>)	Isolates a single field of a datetime or a interval and converts it to a number.
HOUR(<i>time_exp</i>)	Returns the hour as an integer between 0 and 23
MINUTE(<i>time_exp</i>)	Returns the minute as an integer between 0 and 59
MONTH(<i>date</i>)	Returns the month as an integer between 1 and 12

Table 89. Date time functions (continued)

Function	Purpose
MONTHNAME(<i>date</i>)	Returns the month name as a string
NOW()	Returns the current date and time as a timestamp
QUARTER(<i>date</i>)	Returns the quarter as an integer between 1 and 4
SECOND(<i>time_exp</i>)	Returns the second as an integer between 0 and 59
SYSDATE	Returns the current date and time as a timestamp
TIMESTAMPADD(<i>interval</i> , <i>integer_exp</i> , <i>timestamp_exp</i>)	<p>Calculates a timestamp by adding <i>integer_exp</i> intervals of type <i>interval</i> to <i>timestamp_exp</i></p> <p>Keywords used to express valid TIMESTAMPADD interval values are:</p> <p>SQL_TSI_FRAC_SECOND</p> <p>SQL_TSI_SECOND</p> <p>SQL_TSI_MINUTE</p> <p>SQL_TSI_HOUR</p> <p>SQL_TSI_DAY</p> <p>SQL_TSI_WEEK</p> <p>SQL_TSI_MONTH</p> <p>SQL_TSI_QUARTER</p> <p>SQL_TSI_YEAR</p>
TIMESTAMPDIFF(<i>interval</i> , <i>timestamp-exp1</i> , <i>timestamp-exp2</i>)	<p>Returns the integer number of intervals by which <i>timestamp-exp2</i> is greater than <i>timestamp-exp1</i></p> <p>Keywords used to express valid TIMESTAMPDIFF interval values are:</p> <p>SQL_TSI_FRAC_SECOND</p> <p>SQL_TSI_SECOND</p> <p>SQL_TSI_MINUTE</p> <p>SQL_TSI_HOUR</p> <p>SQL_TSI_DAY</p> <p>SQL_TSI_WEEK</p> <p>SQL_TSI_MONTH</p> <p>SQL_TSI_QUARTER</p> <p>SQL_TSI_YEAR</p>

Table 89. Date time functions (continued)

Function	Purpose
TIMESTAMP_FORMAT(<i>input-string</i> , <i>format-string</i>) TO_DATE(<i>input-string</i> , <i>format-string</i>) TO_TIMESTAMP(<i>input-string</i> , <i>format-string</i>)	Returns a timestamp that is based on the interpretation of the input string using the specified format. If the timestamp precision of timestamp-expression is less than what is specified by the format, zero digits are padded onto the right of the specified digits. The following formats are supported: <ul style="list-style-type: none"> • YY Last two digits of the year (00-99) • YYYY 4-digit year (0000-9999) • RR behaves the same as YY • RRRR behaves the same as YYYY • MM Month (01-12) • DD Day of month (01-31) • HH24 Hour of the day (00-24) in 24-hour format • MI Minute (00-59) • SS Seconds (00-59) • FF or FFn Fractional seconds (0-999999999999). The number n is used to specify the number of digits to include in the returned value. Valid values for n are 1-12 with no leading zeros. Specifying FF is equivalent to specifying FF6. The TO_DATE, TIMESTAMP_FORMAT, and TO_TIMESTAMP functions are synonymous.
WEEK(<i>date</i>)	Returns the week of the year as an integer between 1 and 52
YEAR(<i>date</i>)	Returns the year as an integer

B.4 System functions

The system functions return information about the solidDB database.

Table 90. System functions

Function	Purpose
UIC()	Returns the connection id associated with the connection
CURRENT_USERID ()	Returns the current user id
LOGIN_USERID ()	Returns the login userid
CURRENT_CATALOG ()	Returns the current catalog
LOGIN_CATALOG ()	Returns the login catalog
CURRENT_SCHEMA ()	Returns the current schema
LOGIN_SCHEMA ()	Returns the login schema

B.5 Miscellaneous functions

Table 91. Miscellaneous functions

Function	Purpose
BITAND(<i>integer1</i> , <i>integer2</i>) BIT_AND(<i>integer1</i> , <i>integer2</i>)	Returns the result of the bitwise AND operation.
BITANDNOT(<i>expression1</i> , <i>expression2</i>) BIT_ANDNOT(<i>expression1</i> , <i>expression2</i>)	Clears any bit in the first argument that is in the second argument.
BITNOT(<i>expression</i>)	Performs a bitwise NOT operation.
BITOR(<i>expression1</i> , <i>expression2</i>)	Performs a bitwise OR operation.
BITXOR(<i>expression1</i> , <i>expression2</i>)	Performs a bitwise exclusive OR operation.
DATABASE_ENCRYPTION_LEVEL()	Returns the encryption level of the database <ul style="list-style-type: none">• 0 - no encryption• 1 - encrypted, the key is not protected (empty password)• 2 - encrypted, the key is protected by a separate startup password• 3 - encrypted, a custom encryption method is used
IFNULL(<i>exp</i> , <i>value</i>)	If <i>exp</i> is null, returns <i>value</i> ; if not, returns <i>exp</i> (if <i>value</i> is returned, it is converted to type of <i>exp</i>)
SLEEP(<i>milliseconds</i>)	The SLEEP function can only be called from a stored procedure or a trigger. It causes the stored procedure or trigger to "sleep" (temporarily suspend activity) for the specified number of milliseconds. Resolution is accurate to approximately 1 second (1000 milliseconds). The exact length of the sleep also depends upon how busy the computer is with other processes and threads. The value must be a literal, not a variable or expression.

B.6 Advanced replication functions

This sections describes the built-in functions that are specific to advanced replication configurations

B.6.1 GET_PARAM() function

```
get_param('param_name')
```

Usage

The GET_PARAM() function retrieves a parameter that was placed on the transaction bulletin board using the PUT_PARAM() function or with statements SAVE PROPERTY, SAVE DEFAULT PROPERTY, or SET SYNC PARAMETER. The parameter that is retrieved is specific to a catalog; each catalog has a different set of parameters.

If the parameter does not exist in the bulletin board, the function returns the VARCHAR value of the parameter or NULL.

Because get_param() is an SQL function, it can be used only in a procedure or as part of a SELECT statement.

The parameter name must be enclosed in quotation marks.

The GET_PARAM() function applies only to advanced replication configurations; it can be issued on both master and replica databases.

solidDB system parameters

solidDB system parameters are divided into the following categories:

Read-only system parameters

Read-only system parameters are maintained by solidDB and can be read only by using the following syntax:

```
GET_PARAM(parameter_name)
```

The lifecycle of parameters in this category is one transaction, that is, values of these parameters is always initialized at the beginning of the transaction.

Updatable transaction-level system parameters

Users can set and update transaction-level system parameters using the following function call inside the transaction:

```
PUT_PARAM(parameter_name, value)
```

Updatable system parameters are used by solidDB advanced replication to configure synchronization-related operations.

The lifecycle of parameters in this category is one transaction, that is, values of these parameters is always initialized at the beginning of the transaction.

Database catalog level system parameters

Database catalog level system parameters are set by using the following syntax:

```
SET SYNC PARAMETER(parameter_name value)
```

Parameters in this category are database catalog level parameters that are valid until changed or removed. They are specified as bulletin board parameters.

Return values

Table 92. GET_PARAM return values

Error Code	Description
13086	Invalid data type in a parameter.

Examples

```
SELECT put_param('myparam', '123abc');  
SELECT get_param('myparam');
```

Related reference:

B.6.2, “PUT_PARAM() function,” on page 321

A.69, “SAVE PROPERTY,” on page 277

Related information:

“SET SYNC PARAMETER” on page 293

B.6.2 PUT_PARAM() function

`PUT_PARAM(param_name, param_value)`

Use the `PUT_PARAM()` function to place a parameter on the bulletin board. If the parameter already exists, the new value overwrites the previous one.

With solidDB Intelligent Transaction, SQL statements or procedures of a transaction can communicate with each other by passing parameters to each other using a parameter bulletin board. The bulletin board is a storage of parameters that is visible to all statements of a transaction.

The `PUT_PARAM()` function can be used in both the master and replicas for setting parameters to the parameter bulletin board of the current transaction.

Parameters are specific to a catalog. Different replica and master catalogs have their own set of bulletin board parameters that are not visible to each other.

These parameters are not propagated to the master. You can use the `SAVE PROPERTY` statement to propagate properties from the replica to the master. For details, read A.69, "SAVE PROPERTY," on page 277.

Because `PUT_PARAM()` is an SQL function, it can be used only within a procedure or in an SQL statement.

Both the parameter name and value are of type `VARCHAR`.

Differences between "PUT_PARAM()" and "SAVE PROPERTY property_name VALUE property_value;"

You typically use `PUT_PARAM()` inside the (running) transaction to pass parameters between procedures. These parameter values disappear from the bulletin board when the transaction terminates (commits or rolls back).

You typically use the `SAVE PROPERTY` statement in the replica to set properties for the entire transaction. These properties can be used in the `WHERE` clause of the `PROPAGATE TRANSACTIONS` statement. When the transaction is executed in the master, the properties of the transaction are put to the parameter bulletin board of the transaction in the beginning of the transaction. Hence, they can be accessed by all procedures of the transaction by using the `GET_PARAM(param_name)` function.

Return values

Table 93. `PUT_PARAM()` return values

Error Code	Description
13086	Invalid data type in a parameter.

When executed successfully, `PUT_PARAM()` returns the new value of the assigned parameter.

Example

```
SELECT PUT_PARAM('myparam', '123abc');
```

Related reference:

B.6.1, "GET_PARAM() function," on page 319

A.69, "SAVE PROPERTY," on page 277

Related information:

"SET SYNC PARAMETER" on page 293

B.7 Trigger functions

The following system supported triggers stack functions are useful for analyzing and debugging purposes.

Note: The trigger stack refer to those triggers that are cached, regardless of whether they are executed or detected for execution. Trigger stack functions can be used in the application program like any other function.

The functions are:

- TRIG_COUNT()

This function returns the number of triggers in the trigger stack, including the current trigger. The return value is an integer.

- TRIG_NAME(n)

This function returns the nth trigger name in the trigger stack. The first trigger position or offset is zero.

- TRIG_SCHEMA(n)

This function returns the nth trigger schema name in the trigger stack. The first trigger position or offset is zero. The return value is a string.

Appendix C. Data types

The tables in this appendix list the supported data types by category.

Table 94. Abbreviations in data type tables

Abbreviation	Description
DEFLEN	The defined length of the column; for example, for CHAR(24) the precision and length is 24
DEFPREC	The defined precision; for example, for NUMERIC(10,3) it is 10
DEFSCALE	The defined scale; for example, for NUMERIC(10,3), it is 3
MAXLEN	The maximum length of column
N/A	Not applicable

Explanation of column values in the data type tables

Size (range)

The size (range) of a numeric column refers to the minimum and maximum values the column can store.

The size of character columns refers to the maximum length of data that can be stored in the column of that data type.

Precision

The precision of a numeric column refers to the maximum number of digits used by the data type of the column.

The precision of a non-numeric column refers to the defined length of the column.

Scale The scale of a numeric column refers to the maximum number of digits to the right of the decimal point. For the approximate floating point number columns, the scale is undefined because the number of digits to the right of the decimal point is not fixed.

Length

The length of a column is the maximum number of bytes returned to the application when data is transferred to its default C type.

For character data, the length does not include the null termination byte. The length of a column may differ from the number of bytes needed to store the data on the data source.

Display size

The display size of a column is the maximum number of bytes needed to display data in character form.

C.1 Character data types

Table 95. Character Data Types

Data Type	Size	Precision	Scale	Length	Display Size
CHARACTER CHAR	2 G - 1* (2147483647)	DEFLN	N/A	DEFLN	DEFLN
WCHAR NATIONAL CHARACTER NATIONAL CHAR NCHAR	2 G - 1* (2147483647)	DEFLN	N/A	DEFLN	DEFLN
VARCHAR CHARACTER VARYING CHAR VARYING	2 G - 1** (2147483647)	DEFLN	N/A	DEFLN	DEFLN
WVARCHAR NATIONAL VARCHAR NCHAR VARYING NVARCHAR	2 G - 1** (2147483647)	DEFLN	N/A	DEFLN	DEFLN
LONG VARCHAR CHARACTER LARGE OBJECT CHAR LARGE OBJECT CLOB	2 G - 1 (2147483647)	MAXLEN	N/A	MAXLEN	MAXLEN
LONG WVARCHAR LONG NATIONAL VARCHAR NCHAR LARGE OBJECT NCLOB	2 G - 1 (2147483647)	MAXLEN	N/A	MAXLEN	MAXLEN
<p>* default is 1</p> <p>** default is 254</p>					

C.2 Numeric data types

Table 96. Numeric data types

Data Type	Size	Precision	Scale	Length	Display Size
TINYINT	[-128, 255]	3	0	1 (bytes)	4 (signed) 3 (unsigned)
SMALLINT	[-32768, 65535]	5	0	2 (bytes)	6 (signed) 5 (unsigned)
INTEGER INT	$[-2^{31}, 2^{31} - 1]$	10	0	4 (bytes)	11 (signed) 10 (unsigned)
BIGINT	$[-2^{63}, 2^{63} - 1]$	19	0	8 (bytes)	20 (signed)
REAL	Positive numbers: 1.175494351e-38 to 1.7014117e+38 Negative numbers: -1.7014117e+38 to -1.175494351e-38 You can also use value zero (0) with this data type.	7	N/A	4 (bytes)	13
FLOAT	Positive numbers: 2.2250738585072014e-308 - 8.98846567431157854e+307 Negative numbers: -8.98846567431157854e+307 to -2.2250738585072014e-308 You can also use value zero (0) with this data type.	15	N/A	8 (bytes)	22
DOUBLE PRECISION	Positive numbers: 2.2250738585072014e-308 - 8.98846567431157854e+307 Negative numbers: -8.98846567431157854e+307 to -2.2250738585072014e-308 You can also use value zero (0) with this data type.	15	N/A	8 (bytes)	22

Table 96. Numeric data types (continued)

Data Type	Size	Precision	Scale	Length	Display Size
DECIMAL ¹	±1.0e254	DEFPREC Max 52 Default 52	DEFSCALE ¹ Default 0	2-27 (bytes)	Variable
NUMERIC ²	±1.0e254	DEFPREC Max 52 Default 52	DEFSCALE ² Default 2 or 0	2-27 (bytes)	Variable

¹ The default value of scale (DEFSCALE) for the DECIMAL data type depends on whether precision (DEFPREC) is specified:

- If precision is not specified, the DECIMAL values are represented as (exact) decimal floating point numbers of precision 52 and range ±1.0e254.
- If precision is specified, the default scale is 0.

Also, the precision of DECIMAL is always 52. The maximum value is 52.

DECIMAL = DECIMAL(52, floating point)
 DECIMAL(prec) = DECIMAL(52, 0)
 DECIMAL(prec, scale) = DECIMAL(52, scale)

² The default value of scale (DEFSCALE) for the NUMERIC data type depends on whether precision (DEFPREC) is specified:

- If precision is not specified, the default scale is 2.
- If precision is specified, the default scale is 0.

NUMERIC = NUMERIC(52, 2)
 NUMERIC(prec) = NUMERIC(prec, 0)
 NUMERIC(prec, scale) = NUMERIC(prec, scale)

Note: Although integer data types (TINYINT, SMALLINT, INT, and BIGINT) can be interpreted by the client program as either signed or unsigned, solidDB stores and orders them as signed integers. There is no way to tell the server to order the integer data types as though they were unsigned.

Note: BIGINT has approximately 19 significant digits. This means that you can lose the least significant digits when storing BIGINT into non-integer data types such as FLOAT (approximately 15 significant digits), SMALLFLOAT (approximately 7 significant digits), DECIMAL (16 significant digits).

C.3 Binary data types

Table 97. Binary Data Types

Data Type	Size	Precision	Scale	Length	Display Size
BINARY	2 G*	DEFLN	N/A	DEFLN	DEFLN x 2
VARBINARY	2 G**	DEFLN	N/A	DEFLN	DEFLN x 2
LONG VARBINARY BLOB	2 G	MAXLEN	N/A	MAXLEN	MAXLEN x 2

Table 97. Binary Data Types (continued)

Data Type	Size	Precision	Scale	Length	Display Size
* default is 1					
** default is 254					

Tip:

To insert values into BINARY, VARBINARY, and LONG VARBINARY fields, you may express the value as hexadecimal and use the CAST operator, for example:
`INSERT INTO table1 VALUES (CAST('FF00AA55' AS VARBINARY));`

Similarly, you may use CAST() expressions in WHERE clauses:

```
CREATE TABLE t1 (x VARBINARY);
INSERT INTO t1 (x) VALUES (CAST('000000A512' AS VARBINARY));
INSERT INTO t1 (x) VALUES (CAST('000000FF12' AS VARBINARY));
```

```
-- To compare the VARBINARY value(s) using LIKE, cast the
-- VARBINARY to VARCHAR.
SELECT * FROM t1 WHERE CAST(x AS VARCHAR) LIKE '000000A5%';
SELECT * FROM t1 WHERE CAST(x AS VARCHAR) LIKE '000000A5__';
```

```
-- NOTE: If you want to use "=" rather than "LIKE" then you
-- can cast either operand.
SELECT * FROM t1 WHERE CAST(x AS VARCHAR) = '000000A512';
SELECT * FROM t1 WHERE x = CAST('000000A512' AS VARBINARY);
```

WARNING: this kind of query cannot use indexed search for the LIKE predicate and results in poor query performance in many cases.

C.4 Date data type

Table 98. Date Data Type

Data Type	Size	Precision	Scale	Length	Display Size
DATE	N/A	10*	N/A	6**	10*
* the number of characters in the yyyy-mm-dd format					
** the size of the DATE_STRUCT structure					

C.5 Time data type

Table 99. Time data type

Data Type	Size	Precision	Scale	Length	Display Size
TIME	N/A	8*	N/A	6**	8*
* the number of characters in the hh:mm:ss format					
** the size of the TIME_STRUCT structure					

C.6 Timestamp data type

Table 100. Timestamp data type

Data Type	Size	Precision	Scale	Length	Display Size
TIMESTAMP	N/A	19*	9	16**	19/29***

* the number of characters in the 'yyyy-mm-dd hh:mm:ss.ffffff' format

** the size of the TIMESTAMP_STRUCT structure

*** size is 29 with a decimal fraction part

C.7 Smallest possible non-zero numbers

Table 101. Smallest possible non-zero numbers

Data Type	Value
DOUBLE PRECISION	2.2250738585072014e-308
REAL	1.175494351e-38

C.8 BLOBs and CLOBs

solidDB can store binary and character data up to 2147483647 (2G - 1) bytes long. When such data exceeds a certain length, the data is called a BLOB (Binary Large Object) or CLOB (Character Large Object), depending upon the data type that stores the information. CLOBs contain only "plain text" and can be stored in any of the following data types:

- CHAR, WCHAR
- VARCHAR, WVARCHAR
- LONG VARCHAR (mapped to standard type CLOB)
- LONG WVARCHAR (mapped to standard type NCLOB)

BLOBs can store any type of data that can be represented as a sequence of bytes, such as a digitized picture, video, audio, a formatted text document. They can also store plain text, but you have more flexibility if you store plain text in CLOBs.

BLOBs are stored in any of the following data types:

- BINARY
- VARBINARY
- LONG VARBINARY (mapped to standard type BLOB)

Since character data is a sequence of bytes, character data can be stored in BINARY fields, as well as in CHAR fields. CLOBs can be considered a subset of BLOBs.

Tip: The term *BLOB* is used to refer to both CLOBs and BLOBs.

For most non-BLOB data types, such as integer, float, or date, there is a rich set of valid operations that you can do on that data type. For example, you can add, subtract, multiply, divide, and do other operations with FLOAT values. Because a

BLOB is a sequence of bytes and the database server does not know the "meaning" of that sequence of bytes (it does not know whether the bytes represent a movie, a song, or the design of the space shuttle), the operations that you can do on BLOBs are very limited.

solidDB allows you to perform some string operations on CLOBs. For example, you can search for a particular substring (for example a person's name) inside a CLOB by using the LOCATE() function. Because such operations require a lot of the server's resources (memory and/or CPU time), solidDB allows you to limit the number of bytes of the CLOB that are processed. For example, you might specify that only the first 1 megabyte of each CLOB is searched when doing a string search. For more information, see the description of the **MaxBlobExpressionSize** configuration parameter in *IBM solidDB Administrator Guide*.

Although it is theoretically possible to store the entire BLOB "inside" a typical table, if the BLOB is large, then the server usually performs better if most or all of the BLOB is not stored in the table. In solidDB, if a BLOB is no more than N bytes long, then the BLOB is stored in the table. If the BLOB is longer than N bytes, then the first N bytes are stored in the table, and the rest of the blob is stored outside the table as disk blocks in the physical database file. The exact value of "N" depends in part upon the structure of the table, the disk page size that you specified when you created the database, and so on, but it is always at least 256. Data that is 256 bytes or shorter is always stored in the table.

If a data row size is larger than one third of the disk block size of the database file, you must store it partly as a BLOB.

The SYS_BLOBS system table is used as a directory for all BLOB data in the physical database file. One SYS_BLOB entry can accommodate 50 BLOB parts. If the BLOB size exceeds 50 parts, several SYS_BLOB entries per BLOB are needed.

The query below returns an estimate on the total size of BLOBs in the database.

```
select sum(totalsize) from sys_blobs
```

The estimate is not accurate, because the info is only maintained at checkpoints. After two empty checkpoints, this query should return an accurate response.

Appendix D. Reserved words

This section contains reserved words in several SQL standards: ODBC 3.0, X/Open and SQL Access Group SQL CAE specification, Database Language - SQL: ANSI X3H2 (SQL-92). Some words are used by solidDB SQL. This section also contains potential reserved words; these markings are enclosed in parenthesis. Avoid using any of these keywords for other purposes.

Some of the reserved words in this section can be used as identifiers (such as table name or column name) by surrounding the word in double quotation marks (""). Identifiers in double quotation marks are known as delimited identifiers and conform to the ANSI standard for SQL. In the following SQL statement example, the reserved word "NULL" is used as a table name identifier:

```
CREATE TABLE "NULL" (column_1 INTEGER);
```

Note: solidDB SQL allows some reserved words to be used as identifiers even if those words are not in double quotation marks. However, you should use double quotation marks around any reserved word that you want to use as an identifier; this will increase portability.

Table 102. Reserved words list

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
ABSOLUTE	•		•	
ACTION	•		•	
ADA	•			
ADD	•	•	•	•
ADMIN				•
AFTER			(•)	•
ALIAS			(•)	
ALL	•	•	•	•
ALLOCATE	•	•	•	
ALTER	•	•	•	•
AND	•	•	•	•
ANY	•	•	•	•
APPEND				•
ARE	•		•	
AS	•	•	•	•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
ASC	•	•	•	•
ASSERTION	•		•	
ASYNC			(*)	•
AT	•		•	
AUTHORIZATION	•		•	•
AVG	•	•	•	
BEFORE			(*)	•
BEGIN	•	•	•	•
BETWEEN	•	•	•	•
BINARY				•
BIT	•		•	
BIT_LENGTH	•		•	
BOOKMARK				•
BOOLEAN			(*)	
BOTH	•		•	
BREADTH			(*)	
BY	•	•	•	•
CALL			(*)	•
CASCADE	•	•	•	•
CASCADED*	•		•	•
CASE	•		•	•
CAST	•		•	•
CATALOG	•		•	•
CHAR	•	•	•	•
CHAR_LENGTH	•		•	
CHARACTER	•	•	•	•
CHARACTER_LENGTH	•		•	

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
CHECK	•	•	•	•
CLOSE	•	•	•	•
COALESCE	•			
COLLATE	•		•	
COLLATION	•		•	
COLUMN	•		•	•
COMMIT	•	•	•	•
COMMITBLOCK				•
COMMITTED				•
COMPLETION			(*)	
CONNECT	•	•	•	•
CONNECTION	•	•	•	
CONSTRAINT	•		•	•
CONSTRAINTS	•		•	
CONTINUE	•	•	•	
CONVERT	•		•	
CORRESPONDING	•		•	•
COUNT	•	•	•	
CREATE	•	•	•	•
CROSS	•		•	•
CURRENT	•	•		•
CURRENT_DATE	•		•	
CURRENT_TIME	•		•	
CURRENT_TIMESTAMP	•		•	
CURRENT_USER	•		•	
CURSOR	•	•	•	•
CYCLE			(*)	

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
DATA			(*)	•
DATE	•			
DAY	•			
DEALLOCATE				
DEC	•	•	•	•
DECIMAL	•	•	•	•
DECLARE	•	•	•	•
DEFAULT	•	•	•	•
DEFERRABLE	•		•	
DEFERRED	•		•	
DELETE	•	•	•	•
DENSE				•
DEPTH			(*)	
DESC	•	•	•	•
DESCRIBE	•	•	•	
DESCRIPTOR	•	•	•	
DIAGNOSTICS	•	•	•	
DICTIONARY			(*)	
DISCONNECT	•	•	•	
DISTINCT	•	•	•	•
DOMAIN	•		•	•
DOUBLE	•	•	•	•
DROP	•	•	•	•
EACH			(*)	
ELSE	•		•	•
ELSEIF			(*)	•
ENABLE				•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
END	•	•	•	•
END-EXEC	•		•	
EQUALS			(*)	
ESCAPE	•		•	•
EVENT				•
EXCEPT	•		•	•
EXCEPTION	•			
EXEC	•	•	•	•
EXECUTE	•	•	•	•
EXISTS	•	•	•	•
EXPLAIN				•
EXPORT				•
EXTERNAL	•		•	•
EXTRACT	•		•	•
FALSE	•		•	
FETCH	•	•	•	•
FIRST	•		•	
FIXED				•
FLOAT	•	•	•	•
FOR	•	•	•	•
FOREIGN	•	•	•	•
FOREVER				•
FORTRAN	•			
FORWARD				•
FOUND	•	•	•	
FROM	•	•	•	•
FROMFIXED				•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
FULL	•		•	•
GENERAL			(•)	
GET	•	•	•	•
GLOBAL	•		•	
GO	•		•	
GOTO	•	•	•	
GRANT	•	•	•	•
GROUP	•	•	•	•
HAVING	•	•	•	•
HINT				•
HOUR	•		•	
IDENTIFIED				•
IDENTITY	•		•	
IF			(•)	•
IGNORE	•		(•)	
IMMEDIATE	•	•	•	
IMPORT				•
IN	•	•	•	•
INCLUDE	•	•		
INDEX	•	•		•
INDICATOR	•		•	
INITIALLY	•		•	
INNER	•		•	•
INPUT	•		•	
INSENSITIVE	•		•	
INSERT	•	•	•	•
INT	•	•	•	•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
INTEGER	•	•	•	•
INTERNAL				•
INTERSECT	•		•	•
INTERVAL	•		•	
INTO	•	•	•	•
IS	•	•	•	•
ISOLATION	•		•	•
JAVA				•
JOIN	•		•	•
KEY	•	•	•	•
LANGUAGE	•		•	
LAST	•		•	
LEADING	•		•	
LEAVE			(•)	•
LEFT	•		•	•
LESS			(•)	
LEVEL	•		•	•
LIKE	•	•	•	•
LIMIT			(•)	
LOCAL	•		•	•
LOCK				•
LONG				•
LOOP			(•)	•
LOWER	•		•	
MAINMEMORY				•
MASTER				•
MATCH	•		•	

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
MAX	•	•	•	
MERGE				•
MESSAGE				•
MIN	•	•	•	
MINUTE	•		•	
MODIFY			(*)	•
MODULE	•		•	
MONTH	•		•	
NAMES	•		•	
NATIONAL	•		•	
NATURAL	•		•	•
NCHAR	•		•	
NEW			(*)	•
NEXT	•		•	•
NO	•		•	•
NONE	•		(*)	
NOT	•	•	•	•
NULL	•	•	•	•
NULLIF	•		•	•
NUMERIC	•	•	•	•
OBJECT			(*)	
OCTET_LENGTH	•		•	
OF	•	•	•	•
OFF				
OID			(*)	
OLD			(*)	•
ON	•	•	•	•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
ONLY	•		•	•
OPEN	•	•	•	
OPERATION			(*)	
OPERATORS			(*)	
OPTIMISTIC				•
OPTION	•			
OR				•
ORDER	•			
OTHERS				
OUTER				•
OUTPUT	•		•	
OVERLAPS	•		•	
PARAMETERS			(*)	
PARTIAL	•		•	
PASCAL	•			
PENDANT			(*)	
PESSIMISTIC				•
PLAN				•
PLI	•			
POSITION	•		•	
POST				•
PRECISION	•	•	•	•
PREORDER			(*)	
PREPARE	•			
PRESERVE	•			
PRIMARY	•	•	•	•
PRIOR	•		•	

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
PRIVATE			(*)	
PRIVILEGES	•		•	•
PROCEDURE	•		•	•
PROPAGATE				•
PROTECTED			(*)	
PUBLIC	•	•	•	•
PUBLICATION				•
READ			•	•
REAL		•	•	•
RECURSIVE			(*)	
REF			(*)	
REFERENCES	•	•	•	•
REFERENCING			(*)	•
REFRESH				•
REGISTER				•
RELATIVE	•		•	
RENAME				•
REPEATABLE				•
REPLACE			(*)	
REPLICA				•
REPLY				•
RESIGNAL			(*)	
RESTART				•
RESTRICT	•	•	•	•
RESULT				•
RETURN			(*)	•
RETURNS			(*)	•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
REVERSE				•
REVOKE	•	•	•	•
RIGHT	•		•	•
ROLE			(•)	•
ROLLBACK	•	•	•	•
ROUTINE			(•)	
ROW			(•)	
ROWID				•
ROWNUM				•
ROWSPERMESAGE				•
ROWVER				•
ROWS	•		•	
SAVEPOINT			(•)	•
SCAN				•
SCHEMA	•		•	•
SCROLL	•		•	
SEARCH			(•)	
SECOND	•		•	
SECTION	•	•	•	
SELECT	•	•	•	•
SENSITIVE			(•)	
SEQUENCE			(•)	•
SERIALIZABLE				•
SESSION	•		•	
SESSION_USER	•		•	
SET	•	•	•	•
SIGNAL			(•)	

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
SIMILAR			(*)	
SIZE	•		•	
SMALLINT	•	•	•	•
SOME	•		•	•
SORT				•
SPACE	•			
SQL	•	•	•	•
SQLCA	•	•		
SQLCODE	•		•	
SQLERROR	•	•	•	•
SQLEXCEPTION			(*)	
SQLSTATE	•			
SQLWARNING	•		(*)	
START				•
STRUCTURE			(*)	
SUBSCRIBE				•
SUBSCRIPTION				•
SUBSTRING	•		•	
SUM	•	•	•	
SYNC_CONFIG				•
SYSTEM	•			
SYSTEM_USER			•	
TABLE	•	•	•	•
TEMPORARY	•		•	
TEST			(*)	
THEN	•		•	•
THERE			(*)	

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
TIME	•		•	•
TIMEOUT				•
TIMESTAMP	•		•	•
TIMEZONE_HOUR	•		•	
TIMEZONE_MINUTE	•		•	
TINYINT				•
TO	•	•	•	•
TRAILING			•	
TRANSACTION	•		•	•
TRANSACTIONS				•
TRANSLATE	•		•	
TRANSLATION	•		•	
TRIGGER			(*)	•
TRIM	•		•	
TRUE	•		•	
TRUNCATE				•
TYPE			(*)	
UNDER			(*)	
UNION	•	•	•	•
UNIQUE	•	•	•	•
UNKNOWN	•		•	
UNREGISTER				•
UPDATE	•	•	•	•
UPPER	•		•	
USAGE	•		•	
USER	•	•	•	•
USING	•	•	•	•

Table 102. Reserved words list (continued)

Reserved word	ODBC	X/Open SQL	ANSI SQL-92	solidDB SQL
VALUE	•	•	•	•
VALUES	•	•	•	•
VARBINARY				•
VARCHAR	•	•	•	•
VARIABLE			(*)	
VARWCHAR				•
VARYING	•	•	•	
VIEW	•	•	•	•
VIRTUAL			(*)	
VISIBLE			(*)	
WAIT			(*)	•
WCHAR				•
WHEN	•		•	•
WHENEVER	•	•	•	
WHERE	•	•	•	•
WHILE			(*)	•
WITH	•	•	•	•
WITHOUT			(*)	
WORK	•	•	•	•
WRITE			•	•
WVARCHAR				•
YEAR	•		•	
ZONE			•	

*CASCADED: The word CASCADED is reserved in solidDB; however, the word is not currently used in any solidDB SQL statements.

Appendix E. Database system tables and system views

E.1 System tables

E.1.1 SQL_LANGUAGES

The SQL_LANGUAGES system table lists the SQL standards and SQL dialects that solidDB supports.

Table 103. SQL_LANGUAGES system table

Column name	Data type	Description
SOURCE	WVARCHAR	The organization that defined this specific SQL version.
SOURCE_YEAR	WVARCHAR	The year the relevant standard was approved.
CONFORMANCE	WVARCHAR	The conformance level at which conformance to the relevant standard.
INTEGRITY	WVARCHAR	Indicates whether the Integrity Enhancement Feature is supported.
IMPLEMENTATION	WVARCHAR	Identifies uniquely the vendor's SQL language; NULL if SOURCE is 'ISO'.
BINDING_STYLE	WVARCHAR	The binding style 'DIRECT', '*EMBED' or 'MODULE'.
PROGRAMMING_LANG	WVARCHAR	The host language used.

E.1.2 SYS_ATTAUTH

Table 104. SYS_ATTAUTH

Column name	Data type	Description
REL_ID	INTEGER	Table identifier.
UR_ID	INTEGER	User or role identifier.
ATTR_ID	INTEGER	Column identifier.
PRIV	INTEGER	Privilege info.
GRANT_ID	INTEGER	Grantor identifier.
GRANT_TIM	TIMESTAMP	Grant time.

E.1.3 SYS_AUDIT_TRAIL

The SYS_AUDIT_TRAIL system table contains auditing information. When auditing is enabled (`Sql.AuditTrailEnabled=yes`), information about the database activities are written into the SYS_AUDIT_TRAIL system table. Users with administrator rights can query the SYS_AUDIT_TRAIL system table with normal SQL syntax.

Table 105. SYS_AUDIT_TRAIL table definition

Column name	Data type	Description
CREATIME	TIMESTAMP	Timestamp when the audited statement was committed.
LOGIN_USER	WVARCHAR	User logged into the system to execute and commit the audited data.
MACHINE_ID	WVARCHAR	States where the audited statement was executed. If MACHINE_ID is not set, the value is NULL.
APP_INFO	WVARCHAR	States which application executed the statement. If the APP_INFO is not set, the value is NULL.
CURRENT_CATALOG	WVARCHAR	States which catalog was used when the transaction was committed.
CURRENT_SCHEMA	WVARCHAR	States which schema was used when the transaction was committed.
TYPE	VARCHAR	States what kind of an operation the audited change was. The possible values are described below.
SQLSTR	LONG WVARCHAR	States the executed SQL string that was audited.

The possible values for TYPE include:

```
STATUS
ALTER TABLE
ALTER USER (The password in the SQL statement is not recorded)
CREATE CATALOG
CREATE DOMAIN
CREATE EVENT
CREATE INDEX
CREATE PUBLICATION
CREATE ROLE
CREATE SCHEMA
CREATE SEQUENCE
CREATE TABLE
CREATE USER (The password in the SQL statement is not recorded)
CREATE VIEW
DELETE (Recorded if the DELETE statement affected the SYS_AUDIT_TRAIL system table)
DROP CATALOG
DROP DOMAIN
DROP EVENT
DROP INDEX
DROP PUBLICATION
DROP ROLE
DROP SCHEMA
DROP SEQUENCE
DROP TABLE
DROP USER
DROP VIEW
GRANT
GRANT ROLE
REVOKE
REVOKE ROLE
CREATE PROCEDURE
```

```

DROP PROCEDURE
CREATE TRIGGER
ALTER TRIGGER
DROP TRIGGER

```

E.1.4 SYS_BACKGROUNDJOB_INFO

If the body of a START AFTER COMMIT statement cannot be started, the reason is logged in the system table SYS_BACKGROUNDJOB_INFO. Only failed START AFTER COMMIT statements are logged in this table. If the statement (for example, a procedure call) starts successfully, no information is stored in this system table. Statements that start successfully but do not finish executing are not stored in this system table either.

The user can retrieve information from the table SYS_BACKGROUNDJOB_INFO by using either an SQL SELECT statement or by calling a system procedure SYS_GETBACKGROUNDJOB_INFO. See E.1.4, “SYS_BACKGROUNDJOB_INFO” for more details.

Also a system-defined event SYS_EVENT_SACFAILED is posted when a START AFTER COMMIT statement fails to start. See its description H.1, “Miscellaneous events,” on page 401 for more details. The application can wait for this event and use the jobid to retrieve the error message from the system table SYS_BACKGROUNDJOB_INFO.

The system table SYS_BACKGROUNDJOB_INFO can be emptied with the admin command:

```
ADMIN COMMAND 'cleanbgjobinfo';
```

Only a DBA can execute the 'cleanbgjobinfo' command.

Table 106. SYS_BACKGROUNDJOB_INFO

Column name	Data type	Description
ID	INTEGER	Job identifier.
STMT	WVARCHAR	The statement that could not be executed.
USER_ID	INTEGER	User or role identifier.
ERROR_CODE	INTEGER	The error that occurred when we tried to execute the statement.
ERROR_TEXT	WVARCHAR	A description of the error.

E.1.5 SYS_BLOBS

This table includes information about the blobs stored into the database. Furthermore, this table sees to it that the BLOB is physically saved on disk once only even if it is logically saved several times.

Table 107. SYS_BLOBS

Column name	Data type	Description
ID	BIGINT	Blob identifier.

Table 107. SYS_BLOBS (continued)

Column name	Data type	Description
STARTPOS	BIGINT	Byte offset from the beginning of the blob — the start position of the pages.
ENDSIZE	BIGINT	Byte offset of the end of the last page +1.
TOTALSIZE	BIGINT	Total size of the blob.
REFCOUNT	INTEGER	The number of references, that is, the number of existing instances of the same blob.
COMPLETE	INTEGER	Indicates whether the write to the blob is ready or not.
STARTCPNUM	INTEGER	Indicates on what checkpoint level the writing of the blob started.
NUMPAGES	INTEGER	The number of pages the blob consist of.
P01_ADDR	INTEGER	First page's byte offset from the beginning of the blob.
P01_ENDSIZE	BIGINT	Last byte of the first page + 1.
P[02...50]_ADDR	INTEGER	Byte offset of pages [2...50] from the beginning of the blob.
P[02...50]_ENDSIZE	BIGINT	Last byte of the pages [2...50] +1.

E.1.6 SYS_CACHESEGMENT_CONDITIONS

The SYS_CACHESEGMENT_CONDITIONS table contains information about the user-defined cache segments and the assignments of tables or data ranges to the segments.

Table 108. SYS_CACHESEGMENT_CONDITIONS table definition

Column name	Data type	Description
BASE_TABLE_ID	INTEGER	Internal identification number of the table with cache segment assignments
CACHE_SEGMENT_NAME	WVARCHAR	Name of the cache segment to which the table or data range is assigned
COLUMN_NO	INTEGER	Internal identification number of the column with cache segment assignments

Table 108. SYS_CACHESEGMENT_CONDITIONS table definition (continued)

Column name	Data type	Description
CONDITION	WVARCHAR	Type of range specification Possible values are: <ul style="list-style-type: none"> • < • > • = • <> • BETWEEN ... AND • OLDERTHAN ... DAYS • OLDERTHAN ... SECONDS • NEWERTHAN ... DAYS • NEWERTHAN ... SECONDS
VALUE1	WVARCHAR	First value in the range specification
VALUE2	WVARCHAR	Second value in the range specification if CONDITION is BETWEEN

E.1.7 SYS_CACHESEGMENTS

The SYS_CACHESEGMENTS table contains information about the user-defined cache segments that are created with the CREATE CACHE SEGMENT statement.

Table 109. SYS_CACHESEGMENTS table definition

Column name	Data type	Description
NAME	WVARCHAR	Name of the cache segment
MAX_PERCENTAGE	INTEGER	The size of the cache segment in percentages

E.1.8 SYS_CARDINAL

The data in this table is refreshed within every checkpoint, not at another time.

Table 110. SYS_CARDINAL

Column name	Data type	Description
REL_ID	INTEGER	The relation identifier as in SYS_TABLES.
CARDIN	INTEGER	The number of rows in the table.
SIZE	INTEGER	The size of the data in the table.
LAST_UPD	TIMESTAMP	The timestamp of the last update in the table.

E.1.9 SYS_CATALOGS

The SYS_CATALOGS lists available catalogs.

Table 111. SYS_CATALOGS

Column name	Data type	Description
ID	INTEGER	Catalog identifier.
NAME	WVARCHAR	Catalog name.
CREATIME	TIMESTAMP	Create date and time.
CREATOR	WVARCHAR	Creator name.

E.1.10 SYS_CHECKSTRINGS

The SYS_CHECKSTRINGS lists CHECK constraints of the tables.

Table 112. SYS_CHECKSTRINGS

Column name	Data type	Description
ID	INTEGER	Table identifier referring to SYS_TABLES.
CONSTRAINT_NAME	WVARCHAR	Name of the CHECK constraint (unique for the table) or an empty string for unnamed constraints (one string for all unnamed CHECK constraints. They are AND- concatenated).
CONSTRAINT	WVARCHAR	The constraint string itself. It is checked by the SQL interpreter while performing inserts/updates to the given table.

E.1.11 SYS_COLUMNS

This table lists all system table columns.

There are no owner or user viewing restrictions for viewing the system columns. This means that owners can view columns other than those they have created in this table and users with no access rights or with specific access rights can still view any system column in this table.

Table 113. SYS_COLUMNS

Column name	Data type	Description
ID	INTEGER	Unique column identifier.
REL_ID	INTEGER	The relation identifier as in SYS_TABLES.
COLUMN_NAME	WVARCHAR	The name of the column.
COLUMN_NUMBER	INTEGER	The number of the column in the table (in creation order).
DATA_TYPE	WVARCHAR	The data type of the column.

Table 113. SYS_COLUMNS (continued)

Column name	Data type	Description
SQL_DATA_TYPE_NUM	SMALLINT	ODBC compliant data type number.
DATA_TYPE_NUMBER	INTEGER	Internal data type number.
CHAR_MAX_LENGTH	INTEGER	Maximum length for a CHAR field.
NUMERIC_PRECISION	INTEGER	Numeric precision.
NUMERIC_PREC_RADIX	SMALLINT	Numeric precision radix.
NUMERIC_SCALE	SMALLINT	Numeric scale.
NULLABLE	CHAR	Are NULL values allowed (Yes, No).
NULLABLE_ODBC	SMALLINT	ODBC, are NULL values allowed (1,0).
FORMAT	WVARCHAR	Reserved for future use.
DEFAULT_VAL	WVARCHAR	Current default value (if set).
ATTR_TYPE	INTEGER	User defined (0) or internal (>0).
REMARKS	LONG WVARCHAR	Reserved for future use.

E.1.12 SYS_COLUMNS_AUX

If you insert a column with a default value to a table that has existing rows, the column default value is not appended to the existing rows. Instead, the default value defined in the column insert statement is written to the SYS_COLUMNS_AUX table. If an SQL query is targeted at a row that was inserted to the table before the column, the column value is read from the SYS_COLUMNS_AUX table unless the new column value on the row has been changed after it was inserted. Only the original default value is saved in the SYS_COLUMNS_AUX table.

Table 114. SYS_COLUMNS_AUX

Column name	Data type	Description
ID	INTEGER	Table identifier.
ORIGINAL_DEFAULT	WVARCHAR	The original default value.

E.1.13 SYS_DL_REPLICA_CONFIG

This table contains the diskless configurations in the master. This table is intended for updates only through the **sqlsetup** command. Users should not modify this table directly. Doing so can have adverse repercussions.

Table 115. SYS_DL_REPLICA_CONFIG

Column name	Data type	Description
CFG_NAME	WVARCHAR (254) PRIMARY KEY NOT NULL	The name of the diskless replica configuration.
INI_FILE	LONG WVARCHAR	The name of the replica configuration file. The solid.ini file contents are inserted into this column as a blob.
LIC_FILE	LONG WVARCHAR	The name of the replica license file. The solid.lic file contents are inserted into this column as a blob.
SCHEMA_FILE	LONG WVARCHAR	The name of the replica schema. The schema file contents are inserted into this column as a blob.

E.1.14 SYS_DL_REPLICA_DEFAULT

This table contains the diskless default configurations in the master. This table is intended for updates only through the **solidsetup** command. Users should not modify this table directly. Doing so can have adverse repercussions.

Table 116. SYS_DL_REPLICA_DEFAULT

Column name	Data type	Description
REPLICA_NAME	VARCHAR(254) NOT NULL PRIMARY KEY	The name of the replica.
INI_CFG	VARCHAR(254) REFERENCE SYS_DL_REPLICA_CONFIG(CFG_NAME)	The name of the replica configuration file.
LIC_CFG	VARCHAR(254) REFERENCE SYS_DL_REPLICA_CONFIG(CFG_NAME)	The name of the replica license file.
SCHEMA_CFG	VARCHAR(254) REFERENCE SYS_DL_REPLICA_CONFIG(CFG_NAME)	The name of the replica schema.

E.1.15 SYS_EVENTS

Table 117. SYS_EVENTS

Column name	Data type	Description
ID	INTEGER	Unique event identifier.
EVENT_NAME	WVARCHAR	The name of the event.
EVENT_PARAMCOUNT	INTEGER	Number of parameters.
EVENT_PARAMTYPES	LONG VARBINARY	Types of parameters.

Table 117. SYS_EVENTS (continued)

Column name	Data type	Description
EVENT_TEXT	WVARCHAR	The body of the event.
EVENT_SCHEMA	WVARCHAR	The owner of the event.
EVENT_CATALOG	WVARCHAR	The owner of the event.
CREATIME	TIMESTAMP	Creation time.
TYPE	INTEGER	Reserved for future use.

E.1.16 SYS_FEDT_DB_PARTITION

The SYS_FEDT_DB_PARTITION table describes the available Transaction Log Reader partitions.

Table 118. SYS_FEDT_DB_PARTITION table definition

Column name	Data type	Description
PARTITION_ID	CHAR(30)	This column contains the partition name.
DESCRIPTION	VARCHAR(255)	Not used
LOADACTIVE	INTEGER	Not used

E.1.17 SYS_FEDT_TABLE_PARTITION

The SYS_FEDT_TABLE_PARTITION table provides a mapping between tables and Log Reader partitions.

Table 119. SYS_FEDT_TABLE_PARTITION table definition

Column Name	Data Type	Description
TABLE_CATALOG	VARCHAR(255)	This column indicates the catalog name. Each table has an associated catalog name. A catalog is created in when the database is created, or afterwards by using a CREATE CATALOG statement. There is no factory value for a catalog name.
TABLE_SCHEMA	VARCHAR(255)	This column indicates the schema name. Each table has an associated schema name. By default, the schema name is the user ID (username). A schema may be also created by using the CREATE SCHEMA statement.

Table 119. SYS_FEDT_TABLE_PARTITION table definition (continued)

Column Name	Data Type	Description
PARTITION_ID	CHAR(30)	This column indicates the partition name. This column is a foreign key to the FEDT_TABLE_PARTITION table.
BE_TABLE_DATABASE	VARCHAR(255)	Not used
BE_TABLE_SCHEMA	VARCHAR(255)	Not used
BE_TABLE_NAME	CHAR(30)	Not used
RANGE_COL_NAME	CHAR(30)	Not used
LOWER_LIMIT	VARCHAR(255)	Not used
UPPER_LIMIT	VARCHAR(255)	Not used
SELECT_LIST	VARCHAR(255)	Not used
LOADORDER	INTEGER	Not used

E.1.18 SYS_FORKEYPARTS

Table 120. SYS_FORKEYPARTS

Column name	Data type	Description
KEY_CATALOG	INTEGER	Creator name or the owner of the key.
ID	INTEGER	Foreign key identifier.
KEYP_NO	INTEGER	Keypart number.
ATTR_NO	INTEGER	Column number.
ATTR_ID	INTEGER	Column identifier.
ATTR_TYPE	INTEGER	Column type.
CONST_VALUE	VARBINARY	Possible internal constant value; otherwise NULL.

E.1.19 SYS_FORKEYS

Table 121. SYS_FORKEYS

Column name	Data type	Description
ID	INTEGER	Foreign key identifier.
REF_REL_ID	INTEGER	Referenced table identifier.
CREATE_REL_ID	INTEGER	Creator table identifier.
REF_KEY_ID	INTEGER	Referenced key identifier.

Table 121. SYS_FORKEYS (continued)

Column name	Data type	Description
REF_TYPE	INTEGER	Reference type.
KEY_SCHEMA	WVARCHAR	Creator name.
KEY_CATALOG	WVARCHAR	Creator name or the owner of the key.
KEY_NREF	INTEGER	Number of referenced key parts.

E.1.20 SYS_HOTSTANDBY

Deprecated. Relevant to versions before 4.0.

E.1.21 SYS_INFO

Table 122. SYS_INFO

Column name	Data type	Description
PROPERTY	WVARCHAR	The name of the property.
VALUE_STR	WVARCHAR	Value as a string.
VALUE_INT	INTEGER	Value as an integer.

E.1.22 SYS_KEYPARTS

Table 123. SYS_KEYPARTS

Column name	Data type	Description
ID	INTEGER	This column is a foreign key reference to sys_keys.id, so that you can determine which key each keypart is part of.
REL_ID	INTEGER	The relation identifier as in SYS_TABLES.
KEYP_NO	INTEGER	Keypart identifier.
ATTR_ID	INTEGER	Column identifier.
ATTR_NO	INTEGER	The number of the column in the table (in creation order).
ATTR_TYPE	INTEGER	The type of the column.
CONST_VALUE	VARBINARY	Constant value or NULL.
ASCENDING	CHAR	Is the key ascending (Yes) or descending (No).

E.1.23 SYS_KEYS

All database tables must have one clustering key. This key defines the physical sorting order of the data. It has no capacity impact. If a primary key is defined, the primary key is used as the clustering key. If no primary key is defined, an entry with key_name "\$CLUSTKEY_xxxxx" will be automatically created in SYS_KEYS.

If there is a primary key definition for the table, there will be an entry in SYS_KEYS with a key_name like "\$PRIMARYKEY_xxxx" for this entry. The key_primary and key_clustering columns will have a value YES.

If there is no primary key definition for the table, there will be an entry in SYS_KEYS with a key_name like "\$CLUSTKEY_xxxxx". The key_primary column will have a value NO and key_clustering column will have a value YES.

Table 124. SYS_KEYS

Column name	Data type	Description
ID	INTEGER	Unique key identifier.
REL_ID	INTEGER	The relation identifier as in SYS_TABLES.
KEY_NAME	WVARCHAR	The name of the key.
KEY_UNIQUE	CHAR	Is the key unique (Yes, No).
KEY_NONUNIQUE_ODBC	SMALLINT	ODBC, is the key NOT unique (1, 0).
KEY_CLUSTERING	CHAR	Is the key a clustering key (Yes, No).
KEY_PRIMARY	CHAR	Is the key a primary key (Yes, No).
KEY_PREJOINED	CHAR	Reserved for future use.
KEY_SCHEMA	WVARCHAR	The owner of the key.
KEY_NREF	INTEGER	When creating a primary key, the server uses ALL fields of the table, even if the user specified N fields (the N fields specified by the user become the first N fields of the key). KEY_NREF = N, i.e. the number of fields specified by the user.

E.1.24 SYS_LOGPOS

The SYS_LOGPOS table contains information about the log intervals that the Transaction Log Reader needs for catchup.

Each row represents a log interval which is recorded with the approximate granularity defined with the parameter **LogReader.SaveLogposInterval**. A new row is added when the size of the current interval is reaching the value of **SaveLogposInterval**. The sum of the log interval sizes is approximately equal to the value of the parameter **LogReader.MaxLogSize**. A log interval row is deleted automatically when the distance of the represented log position from the latest recorded log interval head is bigger than **LogReader.MaxLogSize**.

The SYS_LOGPOS table is only used internally by the Log Reader while executing SYS_LOG table queries; the table does not contain any information to be used when designing applications that use the Log Reader.

If a query requests log records that are beyond the log intervals represented in SYS_LOGPOS, the query returns the error 'Log overflow'.

Table 125. SYS_LOGPOS table definition

Column name	Data type	Description
ID	INTEGER	Logical log interval ID (identifier of a log interval)
LOG_POS	VARCHAR(255)	Physical position of the head of the log interval
SIZE	INTEGER	Distance in bytes from the previous log interval head position (approximates the value defined with the LogReader.SaveLogposInterval parameter)

E.1.25 SYS_PROCEDURES

This system table lists procedures.

Specific users are restricted from viewing procedures. Owners are restricted to viewing procedures they have created. Users can only view procedures to which they have execute access to see the procedure definition. If users have no access rights, they are restricted from viewing all procedures. Note that execute access does not allow users to see procedure definitions. No restrictions apply to DBAs.

Table 126. SYS_PROCEDURES

Column name	Data type	Description
ID	INTEGER	Unique procedure identifier.
PROCEDURE_NAME	WVARCHAR	Procedure name.
PROCEDURE_TEXT	LONG WVARCHAR	Procedure body.
PROCEDURE_BIN	LONG VARBINARY	Compiled form of the procedure.
PROCEDURE_SCHEMA	WVARCHAR	The name of the schema containing PROCEDURE_NAME.
PROCEDURE_CATALOG	WVARCHAR	The name of the catalog containing PROCEDURE_NAME.
CREATETIME	TIMESTAMP	Creation time.
TYPE	INTEGER	Reserved for future use.

E.1.26 SYS_PROCEDURE_COLUMNS

The SYS_PROCEDURE_COLUMNS defines input parameters and result set columns.

Table 127. SYS_PROCEDURE_COLUMNS

Column name	Data type	Description
PROCEDURE_ID	INTEGER	Procedure identifier.
COLUMN_NAME	WVARCHAR	Procedure column name.
COLUMN_TYPE	SMALLINT	Procedure column type (SQL_PARAM_INPUT or SQL_RESULT_COL).
DATA_TYPE	SMALLINT	Column's SQL data type.
TYPE_NAME	WVARCHAR	Column's SQL data type name.
COLUMN_SIZE	INTEGER	Size of the procedure column.
BUFFER_LENGTH	INTEGER	Column size in bytes.
DECIMAL_DIGITS	SMALLINT	Decimal digits of the procedure column.
NUM_PREC_RADIX	SMALLINT	Radix for numeric data types (2, 10, or NULL if not applicable).
NULLABLE	SMALLINT	Whether the procedure column accepts a NULL value.
REMARKS	WVARCHAR	A description of the procedure column.
COLUMN_DEF	WVARCHAR	Column's default value. Always NULL, that is, no default value is specified.
SQL_DATA_TYPE	SMALLINT	SQL data type.
SQL_DATETIME_SUB	SMALLINT	Subtype code for datetime. Always NULL.
CHAR_OCTET_LENGTH	INTEGER	Maximum length in bytes of a character or binary data type column.
ORDINAL_POSITION	INTEGER	Ordinal position of the column.
IS_NULLABLE	WVARCHAR	Always "YES".

E.1.27 SYS_PROPERTIES

This table is for internal use of HSB only.

Table 128. SYS_PROPERTIES

Column name	Data type	Description
KEY	WVARCHAR	Property identifier.
VALUE	WVARCHAR	Value of a property.
MODTIME	TIMESTAMP	Creation time for the property.

E.1.28 SYS_RELAUTH

This table contains GRANT privileges issued for each table name and user name combination. When a database is created with no GRANT statements executed, this table is empty.

Table 129. SYS_RELAUTH

Column Name	Description
REL_ID	Table or object identifier.
UR_ID	User or role identifier.
PRIV	Information about privileges of a user or a role. Each privilege is related to someone (GRANT_ID) who has granted it.
GRANT_ID	Grantor identifier.
GRANT_TIM	Grant time.
GRANT_OPT	If set to "Yes", the user who receives the privilege may grant the privilege to other users. The possible values are "Yes" or "No".

E.1.29 SYS_SCHEMAS

The SYS_SCHEMAS lists available schemas.

Table 130. SYS_SCHEMAS

Column name	Data type	Description
ID	INTEGER	Schema identifier.
NAME	WVARCHAR	Schema name.
OWNER	WVARCHAR	Schema owner name.
CREATIME	TIMESTAMP	Create date and time.
SCHEMA_CATALOG	WVARCHAR	Schema catalog.

E.1.30 SYS_SEQUENCES

Table 131. SYS_SEQUENCES

Column name	Data type	Description
SEQUENCE_NAME	WVARCHAR	Sequence name.
ID	INTEGER	Unique identifier.
DENSE	CHAR	Is the sequence dense or sparse.
SEQUENCE_SCHEMA	WVARCHAR	The name of the schema containing SEQUENCE_NAME.

Table 131. SYS_SEQUENCES (continued)

Column name	Data type	Description
SEQUENCE_CATALOG	WVARCHAR	The name of the catalog containing SEQUENCE_NAME.
CREATIME	TIMESTAMP	Creation time.

E.1.31 SYS_SERVER

The SYS_SERVER table contains the back-end server login data for SQL passthrough use with the Universal Cache capability.

When the first subscription between solidDB and the back-end is created, the InfoSphere CDC for solidDB instance retrieves the backend data server login data from the backend InfoSphere CDC instance, and then stores it in a solidDB system table SYS_SERVER with a CREATE REMOTE SERVER statement.

Table 132. SYS_SERVER table definition

Column name	Data type	Description
NAME	VARCHAR	Name of the back-end server
DRIVER	VARCHAR	Not used
CONNECT	VARCHAR	Not used
UID	VARCHAR	Username to use when connecting to the back-end server
PWD	VARCHAR	Password to use when connecting to the back-end server Note: The password data is hidden.

E.1.32 SYS_SYNC_REPLICA_PROPERTIES

Table 133. SYS_SYNC_REPLICA_PROPERTIES

Column name	Data type	Description
ID	INTEGER	Replica ID.
NAME	VARCHAR	Property name.
VALUE	VARCHAR	Property value.

The primary key is on the ID and NAME fields.

E.1.33 SYS_SYNONYM

Table 134. SYS_SYNONYM

Column name	Data type	Description
TARGET_ID	INTEGER	Reserved for future use.
SYNON	INTEGER	Reserved for future use.

E.1.34 SYS_TABLEMODES

Table 135. SYS_TABLEMODES

Column name	Data type	Description
ID	INTEGER	Relation identifier.
MODE	WVARCHAR	Concurrency control mode (allowed values: OPTIMISTIC, PESSIMISTIC, MAINMEMORY, or MAINMEMORY PESSIMISTIC).
MODIFY_TIME	TIMESTAMP	Last modify time.
MODIFY_USER	WVARCHAR	Last user that modified.

SYS_TABLEMODES shows the mode only of tables for which the mode was explicitly set. SYS_TABLEMODES does not show the mode of tables that were left at the default mode. The default mode for D-tables is optimistic unless you set the `solid.ini` configuration parameter **General.Pessimistic** to 'Yes'. M-tables are always pessimistic.

To list the names and modes of tables that were explicitly set to optimistic or pessimistic, execute the command:

```
SELECT SYS_TABLEMODES.ID, SYS_TABLEMODES.MODE, SYS_TABLES.TABLE_NAME
FROM SYS_TABLEMODES, SYS_TABLES
WHERE SYS_TABLEMODES.ID = SYS_TABLES.ID
AND SYS_TABLES.TABLE_NAME = '<table_name>';
```

The output will look like:

```
   ID TABLE_NAME   MODE
   -- -
10054 TABLE2      OPTIMISTIC
10056 TABLE3      PESSIMISTIC
```

For more information about setting the concurrency control mode, see “Setting the concurrency (locking) mode to optimistic or pessimistic” on page 118.

E.1.35 SYS_TABLES

This table lists all the system tables.

There are no restrictions for viewing the system tables, which means even users with no access rights can view them. However, specific users are restricted from viewing the user table information. Owners are restricted to viewing user tables they have created and users can only view tables to which they have INSERT, UPDATE, DELETE, or SELECT access. Users are restricted from viewing any user tables if they have no access rights. No restrictions apply to DBAs.

Table 136. SYS_TABLES

Column name	Data type	Description
ID	INTEGER	Unique table identifier.
TABLE_NAME	WVARCHAR	The name of the table.

Table 136. SYS_TABLES (continued)

Column name	Data type	Description
TABLE_TYPE	WVARCHAR	The type of the table (BASE TABLE or VIEW).
TABLE_SCHEMA	WVARCHAR	The name of the schema containing TABLE_NAME.
TABLE_CATALOG	WVARCHAR	The name of the catalog containing TABLE_NAME.
CREATIME	TIMESTAMP	The creation time of the table.
CHECKSTRING	LONG WVARCHAR	Possible check option defined for the table.
REMARKS	LONG WVARCHAR	Reserved for future use.

E.1.36 SYS_TRIGGERS

This system table lists triggers.

Specific users are restricted from viewing triggers. Owners are restricted to viewing only those triggers that they have created. Normal users are restricted from viewing triggers. No restrictions apply to DBAs.

Table 137. SYS_TRIGGERS

Column name	Data type	Description
ID	INTEGER	Unique table identifier.
TRIGGER_NAME	WVARCHAR	Trigger name.
TRIGGER_TEXT	LONG WVARCHAR	Trigger body.
TRIGGER_BIN	LONG VARBINARY	Compiled form of the trigger.
TRIGGER_SCHEMA	WVARCHAR	The name of the schema containing TRIGGER_NAME.
TRIGGER_CATALOG	WVARCHAR	The name of the catalog containing TRIGGER_NAME.
TRIGGER_ENABLED	CHAR	If triggers are enabled "YES"; otherwise "NO".
CREATIME	TIMESTAMP	The creation time of the trigger.
TYPE	INTEGER	Reserved for future use.
REL_ID	INTEGER	The relation identifier.

E.1.37 SYS_TYPES

Table 138. SYS_TYPES

Column name	Data type	Description
TYPE_NAME	WVARCHAR	The name of the data type.
DATA_TYPE	SMALLINT	ODBC, data type number.
PRECISION	INTEGER	ODBC, the precision of the data type.
LITERAL_PREFIX	WVARCHAR	ODBC, possible prefix for literal values.
LITERAL_SUFFIX	WVARCHAR	ODBC, possible suffix for literal values.
CREATE_PARAMS	WVARCHAR	ODBC, the parameters needed to create a column of the data type.
NULLABLE	SMALLINT	ODBC, can the data type contain NULL values.
CASE_SENSITIVE	SMALLINT	ODBC, is the data type case sensitive.
SEARCHABLE	SMALLINT	ODBC, the supported search operations.
UNSIGNED_ATTRIBUTE	SMALLINT	ODBC, is the data type unsigned.
MONEY	SMALLINT	ODBC, whether the data is a money data type.
AUTO_INCREMENT	SMALLINT	ODBC, whether the data type is autoincrementing.
LOCAL_TYPE_NAME	WVARCHAR	ODBC, has the data type another implementation defined name.
MINIMUM_SCALE	SMALLINT	ODBC, the minimum scale of the data type.
MAXIMUM_SCALE	SMALLINT	ODBC, the maximum scale of the data type.

E.1.38 SYS_URole

The SYS_URole contains mapping of users to roles.

Table 139. SYS_URole

Column name	Data type	Description
U_ID	INTEGER	User identifier.
R_ID	INTEGER	Role identifier.

E.1.39 SYS_USERS

The SYS_USERS list information about users and roles.

Table 140. SYS_USERS

Column name	Data type	Description
ID	INTEGER	User or role identifier.
NAME	WVARCHAR	User or role name.
TYPE	WVARCHAR	User type, either USER or ROLE.
PRIV	INTEGER	Privilege information.
PASSW	VARBINARY	Password in encrypted format. If the user is authenticated externally, the value is NULL.
PRIORITY	INTEGER	Reserved for future use.
PRIVATE	INTEGER	Specifies whether user is private or public.
LOGIN_CATALOG	WVARCHAR	Reserved for future use.
AUTHENTICATION	INTEGER	Specifies whether user is authenticated internally (0) or externally (1).
CONNECTION_LIMIT	INTEGER	Defines the maximum number of connections for a given user name. The default value is 0, which means that there is no limit on concurrent connections for the specified user.

E.1.40 SYS_VIEWS

Table 141. SYS_VIEWS

Column name	Data type	Description
V_ID	INTEGER	Unique identifier for this view.
TEXT	LONG WVARCHAR	View definition.
CHECKSTRING	LONG WVARCHAR	Possible CHECK OPTION defined for the view.
REMARKS	LONG WVARCHAR	Reserved for future use.

E.2 System tables for data synchronization

solidDB contains a number of system tables that are used for implementing synchronization functionality. In general, these tables are for internal use only. However, you may need to know the contents of these tables when developing and troubleshooting a new application.

Note that the tables are presented in alphabetical order.

E.2.1 SYS_BULLETIN_BOARD

This table contains persistent parameters that are always available in the parameter bulletin board when transactions are executed in this database catalog.

Table 142. SYS_BULLETIN_BOARD

Column Name	Description
PARAM_NAME	Name of the persistent parameter.
PARAM_VALUE	Value of the parameter.
PARAM_CATALOG	Defines the master/replica catalog.

E.2.2 SYS_PUBLICATION_ARGS

This table contains the publication input arguments in this master database

Table 143. SYS_PUBLICATION_ARGS

Column Name	Description
PUBL_ID	Internal ID of the publication.
ARG_NUMBER	Sequence number of the argument.
NAME	Name of the argument.
TYPE	Type of the argument.
LENGTH_OR_PRECISION	Length or precision of the argument.
SCALE	Scale of the argument.

E.2.3 SYS_PUBLICATION_REPLICA_ARGS

This table contains the definition of the publication arguments in a replica database.

Table 144. SYS_PUBLICATION_REPLICA_ARGS

Column Name	Description
MASTER_ID	Internal ID of the master from which the data is refreshed.
PUBL_ID	Internal ID of the publication.
ARG_NUMBER	Sequence number of the argument.
NAME	Name of the argument.
LENGTH_OR_PRECISION	Length or precision of the argument.
SCALE	Scale of the argument.

E.2.4 SYS_PUBLICATION_REPLICA_STMTARGS

This table contains the mapping between the publication arguments and the statements in the replica.

Table 145. SYS_PUBLICATION_REPLICA_STMTARGS

Column Name	Description
MASTER_ID	Internal ID of the master from which the data is refreshed.
PUBL_ID	Internal ID of the publication.
STMT_NUMBER	Sequence number of the statement.
STMT_ARG_NUMBER	Sequence number of the statement argument.
PUBL_ARG_NUMBER	Sequence number of the publication argument.

E.2.5 SYS_PUBLICATION_REPLICA_STMTS

This table contains the definition of the publication statements in a replica database.

Table 146. SYS_PUBLICATION_REPLICA_STMTS

Column Name	Description
MASTER_ID	Internal ID of the master from which the data is refreshed.
PUBL_ID	Internal ID of the publication.
STMT_NUMBER	Sequence number of the statement.
REPLICA_CATALOG	Name of the target catalog in the replica database.
REPLICA_SCHEMA	Name of the target schema in the replica database.
REPLICA_TABLE	Name of the target table in the replica database.
TABLE_ALIAS	Alias name of the target table.
REPLICA_FROM_STR	SQL FROM tables as string.
WHERE_STR	SQL WHERE arguments as string.
LEVEL	Level of this SQL statement in this publication hierarchy.

E.2.6 SYS_PUBLICATION_STMTARGS

This table contains mapping between the publication arguments and the statements in the master database.

Table 147. SYS_PUBLICATION_STMTARGS

Column Name	Description
PUBL_ID	Internal ID of the publication.

Table 147. SYS_PUBLICATION_STMTARGS (continued)

Column Name	Description
STMT_NUMBER	Sequence number of the statement.
STMT_ARG_NUMBER	Sequence number of the statement argument.
PUBL_ARG_NUMBER	Sequence number of the publication argument.

E.2.7 SYS_PUBLICATION_STMTS

This table contains the publication statements in the master database.

Table 148. SYS_PUBLICATION_STMTS

Column Name	Description
PUBL_ID	Internal ID of the publication.
MASTER_SCHEMA	Name of the publication schema in the master database.
MASTER_TABLE	Name of the table in the master database.
REPLICA_SCHEMA	Name of the schema in the replica database.
REPLICA_TABLE	Name of the table in the replica database.
TABLE_ALIAS	The alias name of the target table.
MASTER_SELECT_STR	SQL SELECT INTO columns as string.
REPLICA_SELECT_STR	SQL SELECT INTO columns as string.
MASTER_FROM_STR	SQL SELECT FROM tables as string.
REPLICA_FROM_STR	SQL SELECT FROM tables as string.
WHERE_STR	SQL WHERE arguments as a string.
DELETEFLAG_STR	For internal use.
LEVEL	Level of this SQL statement in the publication hierarchy.

E.2.8 SYS_PUBLICATIONS

This table contains the publications that have been defined in this master database.

Table 149. SYS_PUBLICATIONS

Column Name	Description
ID	Internal ID of the publication.
NAME	Name of the publication.
CREATOR	User ID of the creator of the publication.

Table 149. SYS_PUBLICATIONS (continued)

Column Name	Description
CREATETIME	Date and time when the publication was created.
ARGCOUNT	Number of input arguments for this publication.
STMTCOUNT	Number of statement contained in this publication.
TIMEOUT	N/A.
TEXT	Contents of the CREATE PUBLICATION statement.
PUBL_CATALOG	Defines the master catalog.

E.2.9 SYS_PUBLICATIONS_REPLICA

This table contains publications that are being used in this replica database.

Table 150. SYS_PUBLICATIONS_REPLICA

Column Name	Description
MASTER_ID	Internal ID of the master from which the data is refreshed.
ID	Internal ID of the publication.
NAME	Name of the publication.
CREATOR	User ID of the creator of the publication.
ARGCOUNT	Number of input arguments for this publication.
STMTCOUNT	Number of statements contained by this publication.

E.2.10 SYS_SYNC_BOOKMARKS

This table contains bookmarks that are being used in a master database.

Table 151. SYS_SYNC_BOOKMARKS

Column Name	Description
BM_ID	Internal ID of the bookmark.
BM_CATALOG	Reserved for future use.
BM_NAME	Name of the bookmark.
BM_VERSION	Internal version information of the bookmark in the master.
BM_CREATOR	User ID of the creator of the bookmark.
BM_CREATETIME	Create time of the bookmark.

E.2.11 SYS_SYNC_HISTORY_COLUMNS

If you turn on synchronization history for a table, you may turn it on for all columns, or only for a subset of columns. If you turn it on for a subset of columns, then the SYS_SYNC_HISTORY_COLUMNS table records which columns you are keeping synchronization history information for. There is one row in SYS_SYNC_HISTORY_COLUMNS for each column that you keep synchronization history for.

Table 152. SYS_SYNC_HISTORY_COLUMNS

Column Name	Description
REL_ID	The ID of the table to keep sync history for.
COLUMN_NUMBER	The ordinal number of the column in that table that we keep sync history for. (For example, if we keep sync history for the second column in the table, then this field will hold the number 2.

E.2.12 SYS_SYNC_INFO

This table contains synchronization information, one row for each node.

Table 153. SYS_SYNC_INFO

Column Name	Description
NODE_NAME	Master or replica node.
NODE_CATALOG	Catalog where node belongs.
IS_MASTER	IF YES, this node is a master.
IS_REPLICA	If YES, this node is a replica.
CREATIME	Node create data and time.
CREATOR	Node creator user name.

E.2.13 SYS_SYNC_MASTER_MSGINFO

This table contains information about the currently active message in the master database.

Data in this table is used to control the synchronization process between the replica and master database. This table also contains information that is useful for troubleshooting purposes. If the execution of a message halts in the master database due to an error, you can query this table to obtain the reason for the problem, as well as the transaction and statement that caused the error.

Table 154. SYS_SYNC_MASTER_MSGINFO

Column Name	Description
STATE	Current state of the message. The following values are possible: <ul style="list-style-type: none"> • 0 = DELETED N/A (internal non-persistent state) • 1 = ERROR - Error has occurred during message processing; the reason for the error was recorded in the error-columns of the row. • 10 = RECEIVED - master has received a message from the replica • 11 = SAVED - message has been saved in the master database and is being processed • 12 = READY - master has processed the message • 13 = SENT - N/A (internal non-persistent state)
REPLICA_ID	ID of the replica database from which the message was sent.
MASTER_ID	ID of the database to which the master is sent.
MSG_ID	Internal ID of the message.
MSG_NAME	Name of the message given by the user.
MSG_TIME	Create time of the message.
MSG_BYTE_COUNT	Size of the message in bytes.
CREATE_UID	ID of the user who created the message.
FORWARD_UID	ID of the user who forwarded the message.
ERROR_CODE	Code of the error that caused the termination of the message execution. You can determine the transaction and statement that caused the error from the TRX_ID and STMT_ID information.
ERROR_TEXT	Description of the error that caused the termination of the message execution.
TRX_ID	Sequence number of the transaction that caused the error.
STMT_ID	Sequence number of the statement of a transaction that caused an error.
ORD_ID_COUNT	N/A (internal use only).
ORD_ID	N/A (internal use only).
FLAGS	NULL or 0 = Normal message. 1 = Message is deleted when reply is sent to replica.
FAILED_MSG_ID	This is an INTEGER column which is part of the primary key. The value is zero for normal messages. The value is msg_id if LOG_ERRORS option is ON and any errors exists.

E.2.14 SYS_SYNC_MASTER_RECEIVED_BLOB_REFS

The received BLOBs are stored in this table on the master. The implementation sees to it that the BLOB is physically saved on disk once only even if it is logically saved several times.

Table 155. SYS_SYNC_MASTER_RECEIVED_BLOB_REFS

Column Name	Description
REPLICA_ID	Internal ID of the replica database from which the message was received.
MSG_ID	Internal ID of the message.
BLOB_NUM	The number that identifies the BLOB.
DATA	A reference to the BLOB.

E.2.15 SYS_SYNC_MASTER_RECEIVED_MSGPARTS

This table contains parts of the messages that were received in the master database from a replica database, but not yet processed in the master database.

Table 156. SYS_SYNC_MASTER_RECEIVED_MSGPARTS

Column Name	Description
REPLICA_ID	Internal ID of the replica database from which the message was received.
MSG_ID	Internal ID of the message.
PART_NUMBER	Sequence number of the message part.
DATA_LENGTH	Length of the data in the message part.
DATA	Data of the message part.

E.2.16 SYS_SYNC_MASTER_RECEIVED_MSGS

This table contains messages that were received in the master database from a replica database, but are not yet processed in the master database.

Table 157. SYS_SYNC_MASTER_RECEIVED_MSGS

Column Name	Description
REPLICA_ID	Internal ID of the replica database from which the message has been received.
MSG_ID	Internal ID of the message.
CREATIME	Create time of the message.
CREATOR	User ID of the user who created the message.

E.2.17 SYS_SYNC_MASTER_STORED_BLOB_REFS

The BLOBs to be sent are stored in this table on the master. The implementation sees to it that the BLOB is physically saved on disk once only even if it is logically saved several times.

Table 158. SYS_SYNC_MASTER_STORED_BLOB_REFS

Column Name	Description
REPLICA_ID	Internal ID of the replica database to which the message will be sent.
MSG_ID	Internal ID of the message.
BLOB_NUM	The number that identifies the BLOB.
DATA	A reference to the BLOB.

E.2.18 SYS_SYNC_MASTER_STORED_MSGPARTS

This table contains parts of the message result sets that were created in the master database, but not yet sent to the replica database.

Table 159. SYS_SYNC_MASTER_STORED_MSGPARTS

Column Name	Description
REPLICA_ID	Internal ID of the replica database to which the message will be sent.
MSG_ID	Internal ID of the message.
ORDER_ID	Sequence number of the result set.
RESULT_SET_ID	Internal ID of the result set.
RESULT_SET_TYPE	Type of the result set.
PART_NUMBER	Sequence number of the message part in the result set.
DATA_LENGTH	Length of the data in the message part in the result set.
DATA	Data of the message part.

E.2.19 SYS_SYNC_MASTER_STORED_MSGS

This table contains messages that were created in the master database, but not yet sent to the replica database.

Table 160. SYS_SYNC_MASTER_STORED_MSGS

Column Name	Description
REPLICA_ID	Internal ID of the replica database to which the message will be sent.
MSG_ID	Internal ID of the message.

Table 160. SYS_SYNC_MASTER_STORED_MSGS (continued)

Column Name	Description
CREATIME	Create time of the message.
CREATOR	User ID of the user who created the message.

E.2.20 SYS_SYNC_MASTER_SUBSC_REQ

This table contains the list of requested subscriptions waiting to be executed in the master.

Table 161. SYS_SYNC_MASTER_SUBSC_REQ

Column Name	Description
REPLICA_ID	Internal ID of the replica from which the statement has arrived.
MSG_ID	Internal ID of the message in which the statement has arrived.
ORD_ID	Sequence number of the subscription.
TRX_ID	Internal ID of the transaction to which the subscription belongs.
STMT_ID	Internal ID of the statement in the subscription.
REQUEST_ID	N/A.
PUBL_ID	Internal ID of the subscribed/refreshed publication.
VERSION	Internal version information of the subscription in the master.
REPLICA_VERSION	Internal version information of the subscription in the replica.
FULLSUBSC	Indicates if the subscription is full or incremental.

E.2.21 SYS_SYNC_MASTER_VERSIONS

This table contains the list of subscriptions (that have been subscribed) to replica databases from the master database.

Table 162. SYS_SYNC_MASTER_VERSIONS

Column Name	Description
REPLICA_ID	Internal ID of the replica database.
REQUEST_ID	Sequence number of the subscription.
VERS_TIME	Create time of the subscription.
PUBL_ID	ID of the publication.
TABNAME	Name of the table of the publication.
TABSCHEMA	Name of the schema of the table.

Table 162. SYS_SYNC_MASTER_VERSIONS (continued)

Column Name	Description
PARAM_CRC	N/A (for internal use only).
PARAM	Parameters of the publication in binary format.
VERSION	Version of the data that has been requested from the replica database.

E.2.22 SYS_SYNC_MASTERS

This table contains the list of master databases accessed by the replica.

Table 163. SYS_SYNC_MASTERS

Column Name	Description
NAME	Given name of the master database.
ID	Internal ID of the master database.
REMOTE_NAME	N/A.
REPLICA_NAME	Given name of the replica database.
REPLICA_ID	Surrogate identifier for the replica database.
REPLICA_CATALOG	Defines the replica catalog which is registered to this master.
CONNECT	Connect string of the master database.
CREATOR	ID of the user who set the database as a master.
ISDEFAULT	Reserved for future use.

E.2.23 SYS_SYNC_RECEIVED_BLOB_ARGS

This table is on the master. The BLOB parameters are saved in this table when the message from the replica is extracted. The rows only exist until the transaction in the message has been executed.

Table 164. SYS_SYNC_RECEIVED_BLOB_ARGS

Column Name	Description
REPLICA	Internal ID of the replica from which the BLOB parameters have arrived.
MSG	Internal ID of the message.
ORD_ID	Sequence number of the BLOB part.
TRX_ID	The transaction ID identifies the transaction.
ID	Internal ID of the user.

Table 164. SYS_SYNC_RECEIVED_BLOB_ARGS (continued)

Column Name	Description
ARGNO	Number of the parameter.
ARG_VALUE	Value of the parameter in binary format.

E.2.24 SYS_SYNC_RECEIVED_STMTS

This table contains the propagated statements that have been received in the master database.

Table 165. SYS_SYNC_RECEIVED_STMTS

Column Name	Description
REPLICA	Internal ID of the replica from which the statement has arrived.
MSG	Internal ID of the message in which the statement has arrived.
ORD_ID	N/A.
TXN_ID	Internal ID of the transaction to which the statement belongs.
ID	Sequence number of the statement within the transaction.
CLASS	Type of the constant.
STRING	the SQL statement as a string.
ARG_COUNT	Number of parameters bound to the statement.
ARG_TYPES	Types of the parameters bound to the statement.
ARG_VALUES	Values of the parameters in binary format.
USER_ID	ID of the user who has saved the statement.
REQUEST_ID	N/A.
FLAGS	This indicates the error-handling mode (for example, IGNORE_ERRORS, LOG_ERRORS, and so on).
ERRCODE	This has the error code if a statement failed while executing on the master.
ERR_STR	This has a description of the error that occurred if a statement failed while executing on the master.

E.2.25 SYS_SYNC_REPLICA_MSGINFO

This table contains information about currently active messages in the replica database.

Data in this table is used to control the synchronization process between the replica and master database. This table also contains information that is useful for

troubleshooting purposes. If the execution of a message halts in the replica database due to an error, you can query this table to obtain the reason for the problem, as well as the transaction and statement that caused the error.

Table 166. SYS_SYNC_REPLICA_MSGINFO

Column Name	Description
STATE	Current state of the message. The following values are possible: <ul style="list-style-type: none"> • 0 = DELETED N/A (internal non-persistent state) • 1 = ERROR - Internal error has occurred during message processing; the reason for the error was recorded in the error-columns of the row. • 20 = R_INIT - N/A (internal non-persistent state) • 21 = R_INITEND - N/A (internal non-persistent state) • 22 = R_SAVED - Replica has saved an outgoing message • 23 = R_SENT - Replica has sent a message to the master • 24 = R_RECEIVED - Replica has received a reply message from the master • 25 = R_EXECUTE - The reply message in a replica is ready for execution • 26 = R_EXECUTE_NOTIFYMASTER - Replica has received a reply, but not yet confirmed it with the master
MASTER_ID	ID of the master database to which the message is sent.
MASTER_NAME	Name of the master database to which the message is sent.
MSG_ID	Internal ID of the message.
MSG_NAME	Name of the message given by the user.
MSG_TIME	Create time of the message.
MSG_BYTE_COUNT	Size of the message in bytes.
CREATE_UID	ID of the user who created the message.
FORWARD_UID	ID of the user who sent the message.
ERROR_CODE	Code of the error that caused the message execution to terminate.
ERROR_TEXT	Description of the error that caused the message execution to terminate.
FLAGS	NULL or 0 = Normal message. 1 = Message is deleted when a reply is received from master. 3 = Message is a registration message.

E.2.26 SYS_SYNC_REPLICA_RECEIVED_BLOB_REFS

The received BLOBs are stored in this table. The implementation sees to it that the BLOB is physically saved on disk once only even if it is logically saved several times.

Table 167. SYS_SYNC_REPLICA_RECEIVED_BLOB_REFS

Column Name	Description
MASTER_ID	Internal ID of the master database from which the message has been received.
MSG_ID	Internal ID of the message.
BLOB_NUM	The number that identifies the BLOB.
DATA	A reference to the BLOB.

E.2.27 SYS_SYNC_REPLICA_RECEIVED_MSGPARTS

This table contains parts of the reply messages that have been received into the replica database from the master database, but are not yet processed in the replica database.

Table 168. SYS_SYNC_REPLICA_RECEIVED_MSGPARTS

Column Name	Description
MASTER_ID	Internal ID of the master database from which the message has been received.
MSG_ID	Internal ID of the message.
PART_NUMBER	Sequence number of the message part.
DATA_LENGTH	Length of the data in the message part.
RESULT_SET_TYPE	Type of the result set.
DATA	Data of the message part.

E.2.28 SYS_SYNC_REPLICA_RECEIVED_MSGS

This table contains reply messages that were received in the replica database from the master database, but not yet processed in the replica database.

Table 169. SYS_SYNC_REPLICA_RECEIVED_MSGS

Column Name	Description
MASTER_ID	Internal ID of the master database from which the message has been received.
MSG_ID	Internal ID of the message.
CREATIME	Create time of the message.
CREATOR	User ID of the user who created the message.

E.2.29 SYS_SYNC_REPLICA_STORED_BLOB_REFS

The BLOBs in the flow message are stored in this table. The implementation sees to it that the BLOB is physically saved on disk once only even if it is logically saved several times.

Table 170. SYS_SYNC_REPLICA_STORED_BLOB_REFS

Column Name	Description
MASTER_ID	Internal ID of the master database from which the message has been received.
MSG_ID	Internal ID of the message.
BLOB_NUM	The number that identifies the BLOB.
DATA	A reference to the BLOB.

E.2.30 SYS_SYNC_REPLICA_STORED_MSGS

This table contains messages that were created in the replica database, but not yet sent to the master database.

Table 171. SYS_SYNC_REPLICA_STORED_MSGS

Column Name	Description
MASTER_ID	Internal ID of the master database to which the message will be sent.
MSG_ID	Internal ID of the message.
CREATIME	Create time of the message.
CREATOR	User ID of the user who has created the message.

E.2.31 SYS_SYNC_REPLICA_STORED_MSGPARTS

This table contains parts of the messages that were created in the replica database, but not yet sent to the master database.

Table 172. SYS_SYNC_REPLICA_STORED_MSGPARTS

Column Name	Description
MASTER_ID	Internal ID of the master database to which the message will be sent.
MSG_ID	Internal ID of the message.
PART_NUMBER	Sequence number of the message part.
DATA_LENGTH	Length of the data in the message part.
DATA	Data of the message part.

E.2.32 SYS_SYNC_REPLICA_VERSIONS

This table contains the list of subscriptions (that have been subscribed) to this replica database from the master database.

Table 173. SYS_SYNC_REPLICA_VERSIONS

Column Name	Description
BOOKMARK_ID	Internal ID of the bookmark in the subscription.
REQUEST_ID	Internal ID of the publication request in the subscription.
VERS_TIME	Create time of the subscription.
PUBL_ID	ID of the subscribed publication.
MASTER_ID	ID of the master database from which the publication has been subscribed.
PARAM_CRC	Internal use only.
PARAM	Parameters of the subscription.
VERSION	Version number of subscribed publication in the master database.
LOCAL_VERSION	Version number of subscribed publication in the replica database.
PUBL_NAME	Name of the publication.
REPLY_ID	ID of the publication reply.

E.2.33 SYS_SYNC_REPLICAS

This table contains the list of replica databases registered with the master.

Table 174. SYS_SYNC_REPLICAS

Column Name	Description
NAME	Given name of the replica database.
ID	Internal ID of the replica database.
MASTER_NAME	N/A.
MASTER_CATALOG	Defines the catalog where the replica is registered
CONNECT	This contains the connect string (for example, 'tcp MyWorkstation 1315') of the replica.

E.2.34 SYS_SYNC_SAVED_BLOB_ARGS

If the user saves a transaction with a BLOB parameter at the replica, a reference to the BLOB is saved in the SYS_SYNC_SAVED_BLOB_ARGS table. The reference points to the SYS_SYNC_REPLICA_STORED_BLOB_REFS table. The rows only exist until the sent message has been prepared.

Table 175. SYS_SYNC_SAVED_BLOB_ARGS

Column Name	Description
MASTER	ID of the master database to which the parameters are sent.
TRX_ID	The transaction ID identifies the transaction.
ID	Internal ID of the user.
ARGNO	Number of the parameter.
ARG_VALUE	Value of the parameter in binary format.

E.2.35 SYS_SYNC_SAVED_STMTS

This table contains statements that have been saved in replica database for later propagation.

Table 176. SYS_SYNC_SAVED_STMTS

Column Name	Description
MASTER	Internal ID of the master database to which the statement will be propagated.
TRX_ID	Internal ID of the transaction to which the statement belongs.
ID	Sequence number of the statement within the transaction.
CLASS	Type of the constant.
STRING	The SQL statement as a string.
ARG_COUNT	Number of parameters bound to the statement.
ARG_TYPES	Types of parameters bound to the statement.
ARG_VALUES	Values of the parameters in binary format.
USER_ID	ID of the user who has saved the statement.
REQUEST_ID	N/A.
FLAGS	This indicates the error-handling mode (for example, IGNORE_ERRORS, LOG_ERRORS, and so on).

E.2.36 SYS_SYNC_TRX_PROPERTIES

When you save transactions, you can assign properties for them. These properties can later be used to select transactions for propagation. The properties are saved in the SYS_SYNC_TRX_PROPERTIES table.

Table 177. SYS_SYNC_TRX_PROPERTIES

Column Name	Description
TRX_ID	The transaction ID identifies the transaction.
NAME	The transaction property name (for example, COLOR).
VALUE_STR	The transaction property value (for example, RED).

E.2.37 SYS_SYNC_USERMAPS

This table maps replica user ids to master users in the SYS_SYNC_USERS table.

Table 178. SYS_SYNC_USERMAPS

Column Name	Description
REPLICA_UID	Replica user ID mapped to master user.
MASTER_ID	Master ID.
REPLICA_USERNAME	Replica user name.
MASTER_USERNAME	Master user name.
PASSW	Encrypted password for master user name.

E.2.38 SYS_SYNC_USERS

This table contains a list of users that have access to the synchronization functions of the replica database. These functions include saving transactions and creating synchronization messages.

In a replica the data of this table is downloaded from the master in a message with the command:

```
MESSAGE unique-message-name APPEND SYNC_CONFIG
['sync-config-arg']
```

Table 179. SYS_SYNC_USERS

Column Name	Description
MASTER_ID	Internal ID of the master database.
ID	Internal ID of the user.
NAME	User name.
PASSW	Encrypted password of the user.

E.3 System views

solidDB supports views as specified in the X/Open SQL Standard.

E.3.1 COLUMNS

The COLUMNS system view identifies the columns which are accessible to the current user.

Table 180. COLUMNS

Column name	Data type	Description
TABLE_CATALOG	WVARCHAR	The name of the catalog containing TABLE_NAME.
TABLE_SCHEMA	WVARCHAR	The name of the schema containing TABLE_NAME.
TABLE_NAME	WVARCHAR	The name of the table or view.
COLUMN_NAME	WVARCHAR	The name of the column of the specified table or view.
DATA_TYPE	WVARCHAR	The data type of the column.
SQL_DATA_TYPE_NUM	SMALLINT	ODBC compliant data type number.
CHAR_MAX_LENGTH	INTEGER	Maximum length for a character data type column; for others NULL.
NUMERIC_PRECISION	INTEGER	The number of digits of mantissa precision of the column, if DATA_TYPE is approximate numeric data type, NUMERIC_PREC_RADIX indicates the units of measurement; for other numeric types contains the total number of decimal digits allowed in the column; for character data types NULL.
NUMERIC_PREC_RADIX	SMALLINT	The radix of numeric precision if DATA_TYPE is one of the approximate numeric data types; otherwise NULL.
NUMERIC_SCALE	SMALLINT	Total number of significant digits to the right of the decimal point; for INTEGER and SMALLINT 0; for others NULL.
NULLABLE	CHAR	If column is known to be not nullable 'NO'; otherwise 'YES'.
NULLABLE_ODBC	SMALLINT	ODBC, if column is known to be not nullable '0'; otherwise '1'.
REMARKS	LONG WVARCHAR	Reserved for future use.

E.3.2 SERVER_INFO

The SERVER_INFO system view provides attributes of the current database system or server.

Table 181. SERVER_INFO

Column name	Data type	Description
SERVER_ATTRIBUTE	WVARCHAR	Identifies an attribute of the server.
ATTRIBUTE_VALUE	WVARCHAR	The value of the attribute.

E.3.3 TABLES

The TABLES system view identifies the tables accessible to the current user.

Table 182. TABLES

Column name	Data type	Description
TABLE_CATALOG	WVARCHAR	The name of the catalog containing TABLE_NAME.
TABLE_SCHEMA	WVARCHAR	The name of the schema containing TABLE_NAME.
TABLE_NAME	WVARCHAR	The name of the table or view.
TABLE_TYPE	WVARCHAR	The type of the table.
REMARKS	LONG WVARCHAR	Reserved for future use.

E.3.4 USERS

The USERS system view identifies users and roles.

Table 183. USERS

Column name	Data type	Description
ID	INTEGER	User or role identifier.
NAME	WVARCHAR	User or role name.
TYPE	WVARCHAR	User type, either USER or ROLE.
PRIV	INTEGER	Privilege information.
PRIORITY	INTEGER	Reserved for future use.
PRIVATE	INTEGER	Specifies whether user is private or public.

E.4 Synchronization-related views

solidDB provides four views that show information about synchronization messages between masters and replicas. One pair of views (SYNC_FAILED_MESSAGES and SYNC_FAILED_MASTER_MESSAGES) shows failed messages. The other pair (SYNC_ACTIVE_MESSAGES and SYNC_ACTIVE_MASTER_MESSAGES) shows active messages.

E.4.1 SYNC_FAILED_MESSAGES

This table is on the master and holds information about messages received from the replica. It is possible to view all necessary information about failed messages using one simple view:

```
SELECT * FROM SYNC_FAILED_MESSAGES.
```

This returns the following columns:

Table 184. SYNC_FAILED_MESSAGES

Column name	Data type	Description
REPLICA_NAME	WVARCHAR	Given node name of the replica from which the message was sent.
MESSAGE_NAME	WVARCHAR	Name of the message given by the user.
TRANSACTION_ID	BINARY	Internal ID of the replica transaction that has failed.
STATEMENT_ID	INTEGER	Sequence number of the statement within the transaction.
STATEMENT_STRING	WVARCHAR	SQL statement as a string.
ERROR_CODE	INTEGER	Code of the error that caused the termination of the message execution.
ERROR_MESSAGE	VARCHAR	Description of the error.

All users have access to this view; no particular privileges are required.

E.4.2 SYNC_FAILED_MASTER_MESSAGES

This table is on the replica and holds information about messages sent to the master. It is possible to view all necessary information about failed messages using one simple view:

```
SELECT * FROM SYNC_FAILED_MASTER_MESSAGES.
```

This returns the following columns:

Table 185. SYNC_FAILED_MASTER_MESSAGES

Column name	Data type	Description
MASTER_NAME	WVARCHAR	Given node name of the master.
MESSAGE_NAME	WVARCHAR	Name of the message given by user.
ERROR_CODE	INTEGER	Code of the error that caused the termination of the message execution.
ERROR_MESSAGE	VARCHAR	Description of the error.

All users have access to this view; no particular privileges are required.

E.4.3 SYNC_ACTIVE_MESSAGES

This table is on the master and holds information about messages received from the replica. This returns the following columns:

Table 186. SYNC_ACTIVE_MESSAGES

Column name	Data type	Description
REPLICA_NAME	WVARCHAR	Given node name of the replica.
MESSAGE_NAME	WVARCHAR	Name of the message given by user.
MESSAGE STATE	VARCHAR	Current state of the message as a string. See details in system table SYS_SYNC_MASTER_MSGINFO.

All users have access to this view; no particular privileges are required.

E.4.4 SYNC_ACTIVE_MASTER_MESSAGES

This table is on the replica and holds information about messages sent to the master. It is possible to view all necessary information about failed messages using one simple view:

```
SELECT * FROM SYNC_FAILED_MASTER_MESSAGES.
```

This returns the following columns:

Table 187. SYNC_ACTIVE_MASTER_MESSAGES

Column name	Data type	Description
MASTER_NAME	WVARCHAR	Given node name of the master.
MESSAGE_NAME	WVARCHAR	Name of the message given by user.
MESSAGE STATE	VARCHAR	Current state of the message as a string. See details in system table SYS_SYNC_REPLICA_MSGINFO.

All users have access to this view; no particular privileges are required.

Appendix F. Database virtual tables

Database virtual tables are tables that are dynamically generated result sets.

F.1 SYS_LOG

The SYS_LOG table is a virtual table: when the Log Reader receives an SQL request for the SYS_LOG table, the appropriate result set is generated dynamically from the internal log structures.

Table 188. SYS_LOG table definition

Column name	Data type	Description
RECID	INTEGER	Record identifier See section “Record identifier descriptions” on page 388 for a detailed description.
RECNAME	VARCHAR	Descriptive name for log record. This is a more human readable format for the RECID field.
TRXID	INTEGER	Transaction identifier
STMTRXID	INTEGER	Statement identifier in the transaction
RELID	INTEGER	Table identifier as stored into the SYS_TABLES table.
FLAGS	INTEGER	Flags field, this field has the following meaning: <ul style="list-style-type: none">• 0 NOP No data is available, in this case all other column values are NULL.• Bit 0 set DATA Data is available.• Bit 1 set SHUTDOWN Server shutdown started.• Bit 6 set CAPTURE_OFF If set in the log record as DBE_LOGREADER_LOG_REC_TRX_START, then operations in this transaction should not be propagated.
LOGADDR	VARBINARY	This is the log address of current log record. The LOGADDR can be used as a starting point for the read from SYS_LOG. The length of the log address is 20 bytes. The log addresses can be compared using binary compare to see which log address comes first.
DATA	VARBINARY	The actual user data in the table. See sections “DATA column format for row data” on page 388 and “DATA column format for DDL data” on page 389 below for detailed descriptions.
TEXTDATA	LONG VARCHAR	Not used, always NULL. (Intended for user readable text format for DATA.)

Record identifier descriptions

Table 189. Record identifier descriptions

RECNAME	RECID	Description
DBE_LOGREADER_LOG_REC_EMPTY	0	Empty record, no data available.
DBE_LOGREADER_LOG_REC_INSERT	1	Row insert.
DBE_LOGREADER_LOG_REC_UPDATE	3	After-image for update. If primary key is changed, it is not shown as an update. Instead it shows as insert and delete.
DBE_LOGREADER_LOG_REC_DELETE_FULL	4	Row delete where all column values stored into log.
DBE_LOGREADER_LOG_REC_UPDATE_BEFOREIMAGE	5	Before-image for update. If primary key is changed, it is not shown as an update. Instead it shows as insert and delete.
DBE_LOGREADER_LOG_REC_TRX_START	7	Transaction start. See FLAGS field how to detect unpropagated deletes.
DBE_LOGREADER_LOG_REC_COMMIT	12	Transaction commit.
DBE_LOGREADER_LOG_REC_DDL	13	DDL operation.
DBE_LOGREADER_LOG_REC_SQL	6	SQL string of DDL operation.

DATA column format for row data

This section describes the DATA column format for rows where the value of the RECNAME field is the following:

- DBE_LOGREADER_LOG_REC_INSERT
- DBE_LOGREADER_LOG_REC_UPDATE
- DBE_LOGREADER_LOG_REC_UPDATE_BEFOREIMAGE
- DBE_LOGREADER_LOG_REC_DELETE_FULL

DATA column is formatted so that for each user column in a table, the DATA column contains a <length><data> pair in binary format. Columns are ordered in the order they are defined in the table.

The format for integer values is so called Most Significant Byte (MSB) first format. This means, for example, that bytes in a four byte hexadecimal integer 0x12345678 is stored as bytes 12 34 56 78.

The <length> field is always a four byte integer in the MSB first format.

NULL values are stored as <length> field of -1.

BLOB values are stored as <length> field of -2, followed by 8 byte integer in MSB first format in <data> field.

The 8 byte integer is a unique BLOB ID. Whenever a new BLOB values is stored or old value is changed, it will be allocated a unique BLOB ID.

Table 190. DATA type descriptions for row data

Name	Description	SQL data type
BINARY UNICODE CHAR	Stored as raw data bytes, the same format as stored into column.	CHAR VARCHAR LONG VARCHAR WCHAR WVARCHAR LONG WVARCHAR BINARY VARBINARY LONG VARBINARY
DOUBLE FLOAT	Stored as 8 byte IEEE floating point number in MSB first format.	FLOAT REAL DOUBLE PRECISION
INTEGER	Stored as 4 byte integer in MSB first format.	BIT TINYINT
BIGINT	Stored as 8 byte integer in MSB first format.	BIGINT
DATE	Stored as 11 byte raw data format. The bytes have the following meaning: <ul style="list-style-type: none"> • 0-1 Year, 2 bytes in MSB first format • 2 Month, 1 byte • 3 Day of month, 1 byte • 4 Hours, 1 byte • 5 Minutes, 1 byte • 6 Seconds, 1 byte • 7 Fractions of second, 4 bytes in MSB first format 	DATE TIME TIMESTAMP
DFLOAT	Decimal floating point in string format.	NUMERIC DECIMAL

DATA column format for DDL data

For rows where RECNAME field has value DBE_LOGREADER_LOG_REC_DDL, the DATA column format is the following:

- DATA column is formatted so that it contains a <length><logrecid><length><info> values.

For rows where RECNAME field has value DBE_LOGREADER_LOG_REC_SQL, the DATA column format is the following:

- DATA column is formatted so that it contains a <length><info> values. The value of <info> field is the SQL string.

The format for integer values is so called Most Significant Byte (MSB) first format. It means that bytes in a four byte hexadecimal integer 0x12345678 is stored as bytes 12 34 56 78.

The <length> field is always a four byte integer in MSB-first format.

The possible <logrecid> field descriptions and <info> values are described in the table below.

Table 191. DATA type descriptions for DDL data

LOGRECID	Description	INFO
45	Create table.	Fully qualified table name
17	Drop table.	Table name.
47	Rename table.	Fully qualified table name.
22	Alter table	Table name.
73	Truncate table.	Table name.
16	Create index.	Index name.
18	Drop index.	Index name.
46	Create view.	Fully qualified view name.
20	Drop view.	View name.
28	Create sequence.	Sequence name.
30	Drop sequence.	Sequence name.
27	Create counter.	Counter name.
29	Drop counter.	Counter name.

Appendix G. System stored procedures

System stored procedures are stored procedures that are provided with the solidDB to help simplify tasks such as administering synchronization.

G.1 System stored procedures for advanced replication

The system procedures include procedures specific to advanced replication setups. You can use the system procedures to simplify commonly needed synchronization tasks. To execute synchronization system procedures, you must have `SYS_ADMIN_ROLE` or `SYS_SYNC_ADMIN_ROLE` privileges.

G.1.1 SYNC_SETUP_CATALOG

The `SYNC_SETUP_CATALOG` system stored procedure creates a catalog, assigns it a node name, and sets the role of the catalog to be master, replica, or both.

```
CALL SYNC_SETUP_CATALOG (  
    catalog_name,  -- WVARCHAR  
    node_name,    -- WVARCHAR  
    is_master,    -- INTEGER  
    is_replica    -- INTEGER  
)
```

EXECUTES ON: master or replica.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

If the `catalog_name` parameter is NULL, the current catalog is assigned the specified node name and role(s).

For `is_master` and `is_replica`, a value of 0 means "no"; any other value means "yes". At least one of these should be non-zero. Because a single catalog can be both a replica and a master, you can set both `is_master` and `is_replica` to non-zero values.

Table 192. `SYNC_SETUP_CATALOG` error codes

RC	Text	Description
13047	No privilege for operation	
13110	NULL not allowed	Only the catalog name can be NULL; all other parameters must be non-NULL.
13133	Not a valid license for this product.	
25031	Transaction is active, operation failed.	The user has made some changes that have not yet been committed.
25052	Failed to set node name to <code>node_name</code> .	The <code>node_name</code> is invalid.
25059	After registration nodename cannot be changed.	Catalog has a name already and has one or more replicas.

G.1.2 SYNC_REGISTER_REPLICA

The SYNC_REGISTER_REPLICA system stored procedure creates a new catalog and registers the replica with the specified master.

```
CALL SYNC_REGISTER_REPLICA (  
    replica_node_name,    -- WVARCHAR  
    replica_catalog_name, -- WVARCHAR  
    master_network_name, -- WVARCHAR  
    master_node_name,    -- WVARCHAR  
    user_id,             -- WVARCHAR  
    password              -- WVARCHAR  
)
```

EXECUTES ON: replica.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

The *master_network_name* is the connect string of the master database server.

If the specified catalog does not exist, it is created automatically.

If the *replica_catalog_name* is NULL, the current catalog is used. Also, the *master_node_name* can be NULL. No other parameter may be NULL.

If registration fails, both master and replica end are reset back to their original status. If any of the parameters have illegal values, an error is returned.

If there are any open transactions that have modified data, an error is returned.

This system procedure does not return a resultset.

Table 193. SYNC_REGISTER_REPLICA error codes

RC	Text	Description
13047	No privilege for operation	
13110	NULL not allowed	Only the catalog name and master node name can be NULL; all other parameters must be non-NULL.
13133	Not a valid license for this product.	
21xxx	Communication error	Was not able to connect to master. For more details about 21xxx errors, see <i>solidDB communication errors</i> in the <i>IBM solidDB Administrator Guide</i> .
25005	Message is already active.	
25031	Transaction is active, operation failed.	The user has made some changes that have not yet been committed.
25035	Message is in use.	
25051	Unfinished messages found.	
25052	Failed to set node name to <i>node_name</i> .	The <i>node_name</i> is invalid.

Table 193. SYNC_REGISTER_REPLICA error codes (continued)

RC	Text	Description
25056	Autocommit not allowed.	You must run this stored procedure with autocommit off.
25057	The replica database has already been registered to a master database.	
25059	After registration nodename cannot be changed.	

G.1.3 SYNC_UNREGISTER_REPLICA

The SYNC_UNREGISTER_REPLICA system stored procedure unregisters the specified replica catalog from the master and optionally drops the replica catalog.

```
CALL SYNC_UNREGISTER_REPLICA (
    replica_catalog_name, -- WVARCHAR
    drop_catalog,        -- INTEGER
    force                 -- INTEGER
)
```

EXECUTES ON: replica.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

If the *drop_catalog* parameter has a non-zero value, the replica catalog is dropped. Any possibly hanging messages for the replica are deleted in both ends of the system.

If the *replica_catalog_name* is NULL, the current catalog is used.

If *force* is non-zero, the master accepts unregistration even if messages for this replica exist in the master. In that case, those messages are deleted.

If the user has any uncommitted changes (that is, open transactions), the call will fail with an error.

This system procedure does not return a resultset.

Table 194. SYNC_UNREGISTER_REPLICA error codes

RC	Text	Description
13047	No privilege for operation	
13110	NULL not allowed	Catalog name cannot be NULL if drop_catalog is non-zero.
13133	Not a valid license for this product.	
21xxx	Communication error	Was not able to connect to master. For more details about 21xxx errors, see <i>solidDB communication errors</i> in the <i>IBM solidDB Administrator Guide</i> .

Table 194. SYNC_UNREGISTER_REPLICA error codes (continued)

RC	Text	Description
25005	Message is already active.	
25019	Database is not a replica database.	
25020	Database is not a master database.	
25023	Replica not registered.	
25031	Transaction is active, operation failed.	The user has made some changes that have not yet been committed.
25035	Message is in use.	
25051	Unfinished messages found.	
25056	Autocommit not allowed.	You must run this stored procedure with autocommit off.
25079	Catalog is master and there are registered replicas. Catalog is not dropped.	
25093	A master database for this replica exists, operation failed.	This message is returned when the user either tries to drop a replica catalog which is registered to a master.

G.1.4 SYNC_REGISTER_PUBLICATION

The SYNC_REGISTER_PUBLICATION system stored procedure registers a publication from the master database.

```
CALL SYNC_REGISTER_PUBLICATION (
    replica_catalog_name, -- WVARCHAR
    publication_name      -- WVARCHAR
)
```

EXECUTES ON: replica.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

If *replica_catalog_name* is NULL, the current catalog is used.

If there are uncommitted changes, the call will fail with an error.

This system procedure does not return a resultset.

Table 195. SYNC_REGISTER_PUBLICATION error codes

RC	Text	Description
13047	No privilege for operation	
13110	NULL not allowed	Only the catalog name can be NULL; all other parameters must be non-NULL.
13133	Not a valid license for this product.	

Table 195. SYNC_REGISTER_PUBLICATION error codes (continued)

RC	Text	Description
21xxx	Communication error	Was not able to connect to master. For more details about 21xxx errors, see <i>solidDB communication errors</i> in the <i>IBM solidDB Administrator Guide</i> .
25005	Message is already active.	
25010	Publication not found	
25019	Database is not a replica database	
25020	Database is not a master database.	
25023	Replica not registered.	
25035	Message is in use.	
25056	Autocommit not allowed.	You must run this stored procedure with autocommit off.
25072	Already registered to publication.	

G.1.5 SYNC_UNREGISTER_PUBLICATION

The SYNC_UNREGISTER_PUBLICATION system stored procedure unregisters a publication.

```
CALL SYNC_UNREGISTER_PUBLICATION (
    replica_catalog_name,  -- WVARCHAR
    publication_name,      -- WVARCHAR
    drop_data              -- INTEGER
)
```

EXECUTES ON: replica.

If *drop_data* flag is set to a non-zero value, all subscriptions to the publication are automatically dropped.

If *replica_catalog_name* is NULL, then the current catalog is used.

If there are uncommitted changes, the call will fail with an error.

This system procedure does not return a resultset.

Table 196. SYNC_UNREGISTER_PUBLICATION error codes

RC	Text	Description
13047	No privilege for operation	
13110	NULL not allowed	Only the catalog name can be NULL; all other parameters must be non-NULL.
13133	Not a valid license for this product.	

Table 196. SYNC_UNREGISTER_PUBLICATION error codes (continued)

RC	Text	Description
21xxx	Communication error	Was not able to connect to master. For more details about 21xxx errors, see <i>solidDB communication errors</i> in the <i>IBM solidDB Administrator Guide</i> .
25005	Message is already active.	
25010	Publication not found.	
25019	Database is not a replica database.	
25020	Database is not a master database.	
25023	Replica not registered.	
25031	Transaction is active, operation failed.	User has made some changes that are not yet committed.
25035	Message is in use.	
25056	Autocommit not allowed.	You must run this stored procedure with autocommit off.
25071	Not registered to publication.	

G.1.6 SYNC_SHOW_SUBSCRIPTIONS

The SYNC_SHOW_SUBSCRIPTIONS system stored procedure shows which subscriptions (publication name and parameters as string representation) of a publication are active in a master database.

```
CALL SYNC_SHOW_SUBSCRIPTIONS (
publication_name -- WVARCHAR
)
```

EXECUTES ON: replica.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

Tip: Same functionality is available in both master and replica catalogs. Use SYNC_SHOW_SUBSCRIPTIONS in the replica catalog and SYNC_SHOW_REPLICA_SUBSCRIPTIONS in the master catalog.

The result set of this procedure call is:

Table 197. SYNC_SHOW_SUBSCRIPTIONS result set

Column Name	Data Type	Description
SUBSCRIPTION	WVARCHAR	Publication name and parameters as a string
SUBSCRIPTION_TIME	TIMESTAMP	Time of last subscription.

Table 198. SYNC_SHOW_SUBSCRIPTIONS error codes

RC	Text	Description
13047	No privilege for operation	
13133	Not a valid license for this product.	
25009	Replica not found.	
25010	Publication not found	
25019	Database is not a replica database	
25020	Database is not a master database.	
25023	Replica not registered.	
25071	Not registered to publication.	

Related reference:

G.1.7, “SYNC_SHOW_REPLICA_SUBSCRIPTIONS”

The SYNC_SHOW_REPLICA_SUBSCRIPTIONS system stored procedure shows which subscriptions (publication name and parameters as string representation) of a publication are active in a specified replica database.

G.1.7 SYNC_SHOW_REPLICA_SUBSCRIPTIONS

The SYNC_SHOW_REPLICA_SUBSCRIPTIONS system stored procedure shows which subscriptions (publication name and parameters as string representation) of a publication are active in a specified replica database.

```
CALL SYNC_SHOW_REPLICA_SUBSCRIPTIONS (
    replica_name,          -- WVARCHAR
    publication_name      -- WVARCHAR
)
```

EXECUTES ON: master.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

Tip: Same functionality is available in both master and replica catalogs. Use SYNC_SHOW_SUBSCRIPTIONS in the replica catalog and SYNC_SHOW_REPLICA_SUBSCRIPTIONS in the master catalog.

If the publication name is NULL, then subscriptions to all publications are listed.

The resultset of this procedure call is:

Table 199. SYNC_SHOW_REPLICA_SUBSCRIPTIONS result set

Column Name	Data Type	Description
REPLICA_NAME	WVARCHAR	Replica name.
SUBSCRIPTION	WVARCHAR	Publication name and parameters as a string

Table 199. SYNC_SHOW_REPLICA_SUBSCRIPTIONS result set (continued)

Column Name	Data Type	Description
SUBSCRIPTION_TIME	TIMESTAMP	Time of last subscription.

Table 200. SYNC_SHOW_REPLICA_SUBSCRIPTIONS error codes

RC	Text	Description
13047	No privilege for operation	
13133	Not a valid license for this product.	
25009	Replica not found.	
25010	Publication not found	
25019	Database is not a replica database	
25020	Database is not a master database.	
25023	Replica not registered.	
25071	Not registered to publication.	

Related reference:

G.1.6, "SYNC_SHOW_SUBSCRIPTIONS," on page 396

The SYNC_SHOW_SUBSCRIPTIONS system stored procedure shows which subscriptions (publication name and parameters as string representation) of a publication are active in a master database.

G.1.8 SYNC_DELETE_MESSAGES

The SYNC_DELETE_MESSAGES system stored procedure deletes hanging messages in the replica database. There might be several hanging messages if a replica application creates messages and does not check and handle errors properly.

```
CALL SYNC_DELETE_MESSAGES (
    replica_catalog_name, -- WVARCHAR
)
```

EXECUTES ON: replica.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

If *replica_catalog_name* is NULL, the current catalog is used.

This procedure does not return a resultset.

Table 201. SYNC_DELETE_MESSAGES error codes

RC	Text	Description
13047	No privilege for operation	

Table 201. SYNC_DELETE_MESSAGES error codes (continued)

RC	Text	Description
13133	Not a valid license for this product.	
25005	Message is already active.	
25009	Replica not found.	
25019	Database is not a replica database	
25020	Database is not a master database.	
25035	Message is in use.	

Related reference:

G.1.9, "SYNC_DELETE_REPLICA_MESSAGES"

The SYNC_DELETE_REPLICA_MESSAGES system stored procedure deletes the messages of a specified replica in the master database. There might be several hanging messages if a replica application creates messages and does not check and handle errors properly.

G.1.9 SYNC_DELETE_REPLICA_MESSAGES

The SYNC_DELETE_REPLICA_MESSAGES system stored procedure deletes the messages of a specified replica in the master database. There might be several hanging messages if a replica application creates messages and does not check and handle errors properly.

```
CALL SYNC_DELETE_REPLICA_MESSAGES (
    master_catalog_name -- WVARCHAR,
    replica_name        -- WVARCHAR
)
```

EXECUTES ON: master.

You must have administrator or synchronization administrator access rights to execute this stored procedure.

The *master_catalog_name* parameter specifies the catalog in the master database from which the messages of the specified replica are searched. If *master_catalog_name* is set to NULL, the current catalog is used.

This procedure does not return a resultset.

Table 202. SYNC_DELETE_REPLICA_MESSAGES error codes

RC	Text	Description
13047	No privilege for operation	
13133	Not a valid license for this product.	
25005	Message is already active.	
25009	Replica not found.	

Table 202. SYNC_DELETE_REPLICA_MESSAGES error codes (continued)

RC	Text	Description
25019	Database is not a replica database	
25020	Database is not a master database.	
25035	Message is in use.	

Related reference:

G.1.8, “SYNC_DELETE_MESSAGES,” on page 398

The SYNC_DELETE_MESSAGES system stored procedure deletes hanging messages in the replica database. There might be several hanging messages if a replica application creates messages and does not check and handle errors properly.

G.1.10 SYNC_GET_CATALOG_SYNC_NODE_NAME_AND_MODE

The SYNC_GET_CATALOG_SYNC_NODE_NAME_AND_MODE system stored procedure is used internally for advanced replication housekeeping.

G.2 Miscellaneous stored procedures

G.2.1 SYS_GETBACKGROUNDJOB_INFO

The SYS_GETBACKGROUNDJOB_INFO system stored procedure retrieves information from the SYS_BACKGROUNDJOB_INFO table.

```
CREATE PROCEDURE SYS_GETBACKGROUNDJOB_INFO(
    jobid INTEGER)
RETURNS(
    ID INTEGER,
    STMT WVARCHAR,
    USER_ID INTEGER,
    ERROR_CODE INTEGER,
    ERROR_TEXT INTEGER)
```

The SYS_GETBACKGROUNDJOB_INFO procedure returns the row that matches the given *jobid*. The *jobid* is the job ID of the START AFTER COMMIT statement that was executed. The job ID is returned by the server when the START AFTER COMMIT statement is executed.

G.2.2 SYS_HSBSTATE_WAIT

The SYS_HSBSTATE_WAIT system stored procedure is used internally for HotStandby housekeeping.

G.2.3 SYS_TF1_WAIT

The SYS_TF1_WAIT system stored procedure is used internally for HotStandby housekeeping.

Appendix H. System events

This section describes the solidDB system events. These events are provided with the solidDB to allow programs to be notified when certain actions occur. You can use these events to monitor the progress of activities such as synchronization between master and replica databases.

The system events follow most of the same rules as user-defined events.

Because these events are pre-defined, you do not create them. Furthermore, you cannot post system events; you can only wait on system events.

Most system events have the same five parameters:

- `ename`: The event name.
- `postsrvtime`: The time that the server posted the event.
- `uid`: The user ID (if applicable).
- `numdatainfo`: Miscellaneous numeric data — the exact meaning depends upon the event. For example, the event `SYS_EVENT_BACKUP` is posted both when a backup is started and when a backup is completed. The value in the `numdatainfo` parameter indicates which case applies — i.e. whether the backup has just started or has just completed. This parameter may be NULL if there is no numeric data.
- `textdata`: Miscellaneous text data — the exact meaning depends upon the event. This parameter may be NULL if there is no numeric data.

Note:

- For a description of events related to HotStandby, see *IBM solidDB High Availability User Guide*.
- For a description of events related to advanced replication, see *IBM solidDB Advanced Replication User Guide*.

H.1 Miscellaneous events

The miscellaneous events are mostly related to the server's internal scheduling and housekeeping, such as backups, checkpoints, and merges.

Although you cannot post these events, you can cause events indirectly, for example, when requesting a backup, or when turning on Maintenance Mode in advanced replication setups. You can monitor these events as necessary.

Table 203. Miscellaneous events

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_BACKUP	<p>The system has started or completed a backup operation. The "state" parameter (NUMDATAINFO) indicates:</p> <p>0: backup completed.</p> <p>1: backup started.</p> <p>Note that the server also posts a second event (SYS_EVENT_MESSAGES) when it starts or completes a backup.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_BACKUPREQ	<p>A backup operation has been requested (but has not yet started).</p> <p>If the user application's callback function returns non-zero, backup is not performed.</p> <p>This event can be caught by the user <i>only</i> if the user is using shared memory access or linked library access.</p> <p>None of the parameters are used.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_CHECKPOINT	<p>The system has started or completed a checkpoint operation.</p> <p>If the system started a checkpoint, then the "state" parameter (NUMDATAINFO) is 1, and the message (TEXTDATA) parameter is "started".</p> <p>If the system completed a checkpoint, then the "state" parameter (NUMDATAINFO) is 0, and the message (TEXTDATA) parameter is "completed".</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_CHECKPOINTREQ	<p>A checkpoint operation has been requested (but has not yet started). Checkpoints are typically executed each time a certain number of log writes has completed.</p> <p>If the user application's callback function returns non-zero, then the merge is not performed.</p> <p>This event can be caught by the user <i>only</i> if the user is using shared memory access or linked library access.</p> <p>None of the parameters are used.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_ERROR	<p>Some type of server error has occurred. The message parameter (TEXTDATA) contains the error text. See H.2, "Errors that cause SYS_EVENT_ERROR," on page 408 for a list of server errors that can cause this event to be posted.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>

Table 203. Miscellaneous events (continued)

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_IDLE	<p>The system is idle.</p> <p>Some tasks have a priority of "idle" and are only run when the system is not running any other tasks. Because very low priority tasks may be running in an "idle" system, the system is not necessarily truly idle in the sense of not doing anything.</p> <p>This event can be caught by the user <i>only</i> if the user is using shared memory access or linked library access.</p> <p>None of the parameters are used.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_IMDB_MEMORY	<p>The system has encountered an event related to in-memory database memory limits.</p> <p>The NUMDATAINFO parameter indicates the current memory allocation in Kilobytes.</p> <p>The TEXTDATA parameter can have one of the following values:</p> <ul style="list-style-type: none"> • IMDB_LIMIT_ABOVE - The amount of available virtual memory is above the limit specified by using the MME.ImdbMemoryLimit parameter • IMDB_LIMIT_BELOW - The amount of available virtual memory is below the limit specified by using the MME.ImdbMemoryLimit parameter • IMDB_LOW_LEVEL_ABOVE - The amount of available virtual memory is above the limit specified by using the MME.ImdbMemoryLowPercentage parameter • IMDB_LOW_LEVEL_BELOW - The amount of available virtual memory is below the limit specified by using the MME.ImdbMemoryLowPercentage parameter • IMDB_WARNING_LEVEL_ABOVE - The amount of available virtual memory is above the limit specified by using the MME.ImdbMemoryLowPercentage parameter • IMDB_WARNING_LEVEL_BELOW - The amount of available virtual memory is below the limit specified by using the MME.ImdbMemoryLowPercentage parameter 	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>

Table 203. Miscellaneous events (continued)

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_ILL_LOGIN	There has been an illegal login attempt. The username (TEXTDATA) and userid (NUMDATAINFO) indicate the user who tried to log in.	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYNC_MAINTENANCEMODE_BEGIN	When the sync mode changes from NORMAL to MAINTENANCE, the server will send this system event. The node_name is the name of the node in which maintenance mode started. A single solidDB server can have multiple "nodes" (catalogs.) For more details about sync mode, see "SET SYNC MODE" on page 290.	node_name WVARCHAR.
SYNC_MAINTENANCEMODE_END	When the sync mode changes from MAINTENANCE to NORMAL, the server will send this system event. The node_name is the name of the node in which maintenance mode started. A single solidDB server can have multiple "nodes" (catalogs.) For more details about sync mode, see "SET SYNC MODE" on page 290.	node_name WVARCHAR
SYS_EVENT_MERGE	An event associated with the "merge" operation (merging data from the Bonsai Tree to the main storage tree) has occurred. The parameter STATE (NUMDATAINFO) gives more details: 0: stop the merge 1: start the merge 2: merge progressing 3: merge accelerated.	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_MERGEREQ	A merge operation has been requested (but has not yet started). If the user application's callback function returns non-zero, the merge is not performed. This event can be caught by the user <i>only</i> if the user is using shared memory access or linked library access. None of the parameters are used.	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR

Table 203. Miscellaneous events (continued)

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_MESSAGES	<p>This event is posted when the server has a message (error message or warning message) to log to <code>so1error.out</code> or <code>so1msg.out</code>. In this case, the <code>TEXTDATA</code> contains the message text and <code>NUMDATAINFO</code> the code. If the message to be written is an error, <i>both</i> <code>SYS_EVENT_ERROR</code> and <code>SYS_EVENT_MESSAGES</code> will be posted. If the message is only a warning, only <code>SYS_EVENT_MESSAGES</code> is posted. For a list of the warnings that can cause <code>SYS_EVENT_MESSAGES</code>, see H.3, "Conditions or warnings that cause <code>SYS_EVENT_MESSAGES</code>," on page 409.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, MESSAGE WVARCHAR</p>
SYS_EVENT_NOTIFY	<p>Event sent with admin command 'notify'.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_PARAMETER	<p>This event is posted if a configuration parameter is changed with the command ADMIN COMMAND 'parameter...';</p> <p>The parameter MESSAGE (<code>TEXTDATA</code>) contains the section name and the parameter name.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>

Table 203. Miscellaneous events (continued)

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_PROCESS_MEMORY	<p>The system has encountered an event related to process size memory limits.</p> <p>The NUMDATAINFO parameter indicates the current memory allocation in Kilobytes.</p> <p>The TEXTDATA parameter can have one of the following values:</p> <ul style="list-style-type: none"> • PROCESS_LIMIT_ABOVE - The amount of available virtual memory is above the limit specified by using the Srv.ProcessMemoryLimit parameter • PROCESS_LIMIT_BELOW - The amount of available virtual memory is below the limit specified by using the Srv.ProcessMemoryLimit parameter • PROCESS_LOW_LEVEL_ABOVE - The amount of available virtual memory is above the limit specified by using the Srv.ProcessMemoryLowPercentage parameter • PROCESS_LOW_LEVEL_BELOW - The amount of available virtual memory is below the limit specified by using the Srv.ProcessMemoryLowPercentage parameter • PROCESS_WARNING_LEVEL_ABOVE - The amount of available virtual memory is above the limit specified by using the Srv.ProcessMemoryWarningPercentage parameter • PROCESS_WARNING_LEVEL_BELOW - The amount of available virtual memory is below the limit specified by using the Srv.ProcessMemoryWarningPercentage parameter 	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_ROWS2MERGE	<p>This event indicates that there are rows that need to be merged from the Bonsai Tree to the main storage tree. The rows parameter (NUMDATAINFO) indicates the number of non-merged rows in the Bonsai Tree.</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>
SYS_EVENT_SACFAILED	<p>This event is posted when a START AFTER COMMIT (SAC) fails. The application can wait for this event and use the job ID (which is in the NUMDATAINFO field) to retrieve the error message from the system table SYS_BACKGROUNDJOB_INFO. (The job ID in NUMDATAINFO matches the job ID that is returned when the START AFTER COMMIT statement is executed.)</p>	<p>ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR</p>

Table 203. Miscellaneous events (continued)

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_SHUTDOWNREQ	<p>A shutdown request has been received. If the user application's callback function returns non-zero, then shutdown is not performed.</p> <p>This event can be caught by the user <i>only</i> if the user is using shared memory access or linked library access.</p> <p>None of the parameters are used.</p>	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_STATE_MONITOR	<p>This event is posted when monitoring settings are changed.</p> <p>State (NUMDATAINFO) is one of the following:</p> <p>0: monitoring off.</p> <p>1: monitoring on.</p> <p>UID is the user ID of the user for whom monitoring was turned on or off.</p>	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_STATE_OPEN	<p>This event is posted when the "state" of the database is changed. The parameter STATE (NUMDATAINFO) indicates the new state:</p> <p>0: Closed. No new connections allowed.</p> <p>1: Opened: New connections allowed.</p>	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_STATE_SHUTDOWN	<p>This event is posted when a server shutdown is started. The NUMDATAINFO and TEXTDATA parameters provide no information.</p>	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_STATE_TRACE	<p>Server trace is turned on or off with the ADMIN COMMAND 'trace' command.</p> <p>The parameter STATE (NUMDATAINFO) indicates the new trace state:</p> <p>0: tracing off.</p> <p>1: tracing on.</p>	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_TMCMD	<p>This event is posted when an "AT" command (that is, a timed command) is executed. The message parameter (TEXTDATA) contains the command.</p>	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR

Table 203. Miscellaneous events (continued)

EVENT NAME	EVENT DESCRIPTION	PARAMETERS
SYS_EVENT_TRX_TIMEOUT	This event is not used.	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR
SYS_EVENT_USERS	The parameter REASON (NUMDATAINFO) contains the reason for the event: 0: User connected. 1: User disconnected. 2: User disconnected abnormally. 4: User disconnected because of timeout.	ENAME WVARCHAR, POSTSRVTIME TIMESTAMP, UID INTEGER, NUMDATAINFO INTEGER, TEXTDATA WVARCHAR

H.2 Errors that cause SYS_EVENT_ERROR

This section lists the errors that cause the server to post the event SYS_EVENT_ERROR.

The numbers in the **Error code** column match the solidDB error code numbers, as listed in *Error codes* in the *IBM solidDB Administrator Guide*. These values get passed in the NUMDATAINFO event parameter.

Table 204. Errors that cause SYS_EVENT_ERROR

Error code	Error description
30104	Shutdown aborted; denied by user callback
30208	Merge not started; denied by user callback
30284	Checkpoint not started; denied by user callback
30302	Backup start failed. Shutdown is in progress
30302	Backup start failed. Backup is already active
30303	Backup aborted
30304	Backup failed. <i>error_description</i>
30305	Backup not started; denied by user callback
30306	Backup not started; Backup is not supported on diskless server.
30307	Backup not started, index check failed. Errors written to file <i>ssdebug.log</i> .

Table 204. Errors that cause SYS_EVENT_ERROR (continued)

Error code	Error description
30360	AT command failed. <reason>
30403	Log file write failure.
30454	Failed to save configuration file <i>file_name</i>
30573	Network backup failed. <i>reason</i>
30640	<Server RPC error message>

H.3 Conditions or warnings that cause SYS_EVENT_MESSAGES

This section lists the warning messages that cause the server to post the event SYS_EVENT_MESSAGES

The numbers in the **Error code** column match the solidDB error code numbers, as listed in *Error codes* in the *IBM solidDB Administrator Guide*. These values get passed in the NUMDATAINFO event parameter.

Table 205. Warnings that cause SYS_EVENT_MESSAGES

Error code	Error description
30010	User <i>username</i> failed to connect, version mismatch. Client version <i>version</i> , server version <i>version</i> .
30011	User <i>username</i> failed to connect, collation version mismatch.
30012	User <i>username</i> failed to connect, there are too many connected clients.
30020	Server is in fatal state, no new connections are allowed
30020	User <i>username</i> failed to connect, database character set is utf8, unsupported by client.
30282	Checkpoint creation not started because shutdown is in progress.
30283	Checkpoint creation not started because it's disabled.
30300	Backup completed successfully. The server also posts a second event (SYS_EVENT_BACKUP) when it starts or completes a backup.
30301	Backup started to <i>directory_path</i> . The server also posts a second event (SYS_EVENT_BACKUP) when it starts or completes a backup.
30359	Server noticed time inconsistency during at-command execution. If the system time has been changed, restart the server.
30361	Illegal at command <i>command</i> ignored.

Table 205. Warnings that cause SYS_EVENT_MESSAGES (continued)

Error code	Error description
30362	Illegal immediate at command <i>command</i> ignored.
30405	Unable to open message log file <i>file_name</i>
30800	<p>Unable to reserve requested <i>number</i> memory blocks for external sorter.</p> <p>Only <i>number</i> memory blocks were available. SQL: <i>SQL_statement</i></p>
30801	<p>Unable to reserve requested <number> memory blocks for external sorter.</p> <p>Only <i>number</i> memory blocks were available.</p>

Index

Special characters

_ (underscore) 311
- (minus) 306
/ (slash) 306
* (asterisk) 306
% 311
>> (greater than) 308
>>= (greater than or equal to) 308
< (less than) 308
<> (not equal to) 308
<= (less than or equal to) 308
|| (concatenation operator) 313
+ (plus) 306, 313
= (equal to) 308

A

ABS (function) 314
access rights
 publications 235, 274
 registration user 294
ACCESS RIGHTS (remote stored procedures) 48
ACOS (function) 314
ADD CONSTRAINT 104
ADD_MONTHS (function) 316
ADMIN COMMAND
 abort 142
 assertexit 142
 backgroundjob 143
 backup 143
 backuplist 143
 checkpointing 143
 cleanbgjobinfo 143
 close 143
 commands 141
 describe 143
 errorcode 143
 errorexit 143
 filespec 144
 getreadonlyflag 144
 help 144
 hotstandby 144
 indexusage 144
 info 145
 logreader 146
 makecp 146
 memory 146
 messages 146
 monitor 146
 netbackup 147
 netbackuplist 147
 netstat 147
 notify 147
 open 147
 parameter 148
 passthrough status 148
 perfmon 149
 perfmon diff 149
 perfmon timers 150
 pid 150

ADMIN COMMAND (*continued*)

 proctrace 150
 protocols 150
 runmerge 150
 save parameters 150
 shutdown 150
 sqllist 151
 startmerge 151
 status 151
 throwout 152
 tid 152
 trace 152
 userid 153
 userlist 154, 155
 usertrace 156
 version 156
ADMIN EVENT statement 156
ALL (keyword)
 PROPAGATE TRANSACTIONS 249
ALTER REMOTE SERVER statement 157
ALTER TABLE SET HISTORY COLUMNS 159
ALTER TABLE SET NOSYNCHISTORY statement
 description 160
ALTER TABLE SET SYNCHISTORY
 description 160
ALTER TABLE statement 157
ALTER TRIGGER statement 162
ALTER USER statement 162
AND (operator) 30, 304
APPEND (keyword) 249
AS 17
 SELECT statement 17
ASCending 98
ASCII (function) 313
ASIN (function) 314
ATAN (function) 315
ATAN2 (function) 315
autocommit 181
AVG (function) 306

B

backup
 SYS_EVENT_BACKUP 402
batch inserts and updates 138
bcktime ADMIN COMMAND 145
BEGIN 180
BIGINT data type 325
binary data 326
BINARY data type
 CAST function 327
BIT_AND function (bit-wise AND operator) 319
BITANDNOT() function 319
BLOBs (Binary Large Objects) 14, 328
 CAST function 327
BLOBs and CLOBs 328
bookmarks
 dropping 206, 227
Bulletin board parameters 319

C

- CALL statement 165
 - example with EXECDIRECT 192
 - invoking procedures 24
- candidate key 100
- CASCADE 95, 219, 224
 - keyword in REVOKE statements 273
- CASCADED (reserved word) 331
- CASE 17, 307
- CAST (function) 17, 306
 - binary values 327
- catalogs
 - creating 108, 170
 - deleting 108
 - description 106
- CEILING (function) 315
- CHAR (function) 313
- CHAR data type 324
- CHAR LARGE OBJECT data type 324
- CHAR VARYING data type 324
- CHARACTER data type 324
- character data types 324, 325
- CHARACTER LARGE OBJECT data type 324
- CHARACTER VARYING data type 324
- CHECK 104
- checkpoint
 - 'makecp' command 146
- checkpoints
 - SYS_EVENT_CHECKPOINT 402
- client-server architecture
 - description 4
 - multi-user capability 4
- CLOB data type 324, 328
- clustering 97
 - key 356
- COALESCE 307
- columns 1, 7
 - adding to a table 96
 - deleting from a table 96
- COLUMNS system view 382
- commit block
 - defining refresh size 259, 266
- COMMIT WORK statement 5, 21, 169
 - stored procedures 45
- COMMITBLOCK (keyword)
 - DROP SUBSCRIPTION 225
 - MESSAGE FORWARD 259
 - MESSAGE GET REPLY 266
 - REFRESH 270
- committing work
 - after altering table 96
 - after altering users and roles 93
- comparison operators 30
- CONCAT (function) 313
- concatenated indexes 137
- CONCURRENCY 111
- concurrency (locking) mode
 - optimistic or pessimistic 118
- concurrency control
 - mixed 119
 - mode
 - displaying 361
 - MAINMEMORY 361
 - MAINMEMORY PESSIMISTIC 361
 - OPTIMISTIC 361
 - PESSIMISTIC 361
 - PESSIMISTIC vs. OPTIMISTIC 112
 - concurrency control (*continued*)
 - optimistic 112
 - pessimistic 112
 - purpose 111
 - setting 118, 119
- configuring
 - synchronization 293
- connect string
 - changing to master name 289
- ConnectStrForMaster (parameter) 168
- constraints
 - foreign key 207
- control structures, stored procedures 32
- CONVERT_CHAR 306
- CONVERT_DATE 306
- CONVERT_DECIMAL 306
- CONVERT_DOUBLE 306
- CONVERT_FLOAT 306
- CONVERT_INTEGER 306
- CONVERT_LONGVARCHAR 306
- CONVERT_NUMERIC 306
- CONVERT_REAL 306
- CONVERT_SMALLINT 306
- CONVERT_TIME 306
- CONVERT_TIMESTAMP 306
- CONVERT_TINYINT 306
- CONVERT_VARCHAR 306
- CONVERTORSTOUNIONS 286
- COS (function) 315
- COT (function) 315
- COUNT (function) 306
- cptime ADMIN COMMAND 145
- CREATE [OR REPLACE] REMOTE SERVER statement 201
- CREATE CATALOG statement 108, 170
- CREATE EVENT statement 80, 172
- CREATE INDEX statement 179
- CREATE PROCEDURE (external) statement 193
- CREATE PROCEDURE statement 180
 - declare section 27
 - parameter section 24
- CREATE PUBLICATION statement
 - description 197
- CREATE ROLE statement 202
- CREATE SCHEMA statement 203
- CREATE SEQUENCE statement 79
- CREATE SYNC BOOKMARK statement
 - description 206
- CREATE TABLE statement 207
- CREATE TRIGGER statement 210
- CREATE USER statement 216
- CREATE VIEW statement 217
- creating
 - publications 197
- CURDATE (function) 316
- CURRENT_CATALOG (system function) 318
- CURRENT_SCHEMA (system function) 318
- CURRENT_USERID (system function) 318
- CURSORNAME 180, 188
 - example usage 188, 191
- cursors
 - default management in stored procedures 45
 - parameter markers 39
 - stored procedures 46
 - closing 39
 - dropping 39
 - executing 38
 - fetching 39

cursors (*continued*)
 stored procedures (*continued*)
 handling 37
 preparing 38
CURTIME (function) 316

D

data
 exporting to file 231
 importing from a file 242
 returning in a stored procedure 36
data management
 solidDB SQL 111
data types 9, 305
 SQL 89
database
 column 1, 7
 creation time 145
 free space in 145
 relational 1
 row 1, 7
 table 1, 7
database objects 105
DATE (function) 316
date and time literals 310
DATE data type
 description 327
date time functions 316
DAYNAME (function) 316
DAYOFMONTH (function) 316
DAYOFWEEK (function) 316
DAYOFYEAR (function) 316
dbconfigsize ADMIN COMMAND 145
dbcreatetime ADMIN COMMAND 145
dbfreesize ADMIN COMMAND 145
dbpagesize ADMIN COMMAND 145
dbsize ADMIN COMMAND 145
DECIMAL data type 326
DEFAULT 47
deferred procedure calls 50
DEGREES (function) 315
DELETE statement 217
deleting
 failed messages 254
 messages 253
DESCending 98
DESCRIBE statement 218
diagnosing poor performance
 diagnosis 140
 solutions 140
 symptoms 140
DIFFERENCE (function) 315
DOUBLE data type 325
DOUBLE PRECISION data type 328
DROP BOOKMARK statement 206
DROP CATALOG statement 219
DROP CONSTRAINT statement 104
DROP EVENT statement 80, 220
DROP INDEX statement 221
DROP MASTER statement 221
DROP PROCEDURE statement 222
DROP PUBLICATION REGISTRATION statement 222
DROP PUBLICATION statement 222
DROP REMOTE SERVER statement 223
DROP REPLICA statement 223
DROP ROLE statement 224

DROP SCHEMA statement 224
DROP SEQUENCE statement 225
DROP SUBSCRIPTION statement 225
DROP SYNC BOOKMARK statement 227
DROP TABLE statement 228
DROP TRIGGER statement 228
DROP USER statement 229
DROP VIEW statement 229
dropping
 bookmarks 206, 227
 master database 221
 publications 222
 replica databases 223
 subscriptions 225
duplicate inserts 264

E

EnableHints (parameter) 139
END 180
END LOOP 187
ending messages 256
error handling
 stored procedures 41
errors
 DBMS 254, 264
 fatal errors, synchronization errors 266
 SYS_EVENT_ERROR 408
escape character 312
escape sequencefn 29
evaluating application performance 133
EVENT
 dropping an event 220
 posting an event 180
 registering for an event 180
 unregistering for an event 180
 waiting on an event 180
events
 ADMIN EVENT command 156
 code example 80
 using 80
 waiting on 137
EXCLUSIVE (lock mode) 116
exclusive locks 115
EXECDIRECT 188
 example usage 192
 using an SQL statement in a VARCHAR variable 192
executing
 failed messages 264
 messages 257
EXP (function) 315
EXPLAIN PLAN FOR statement 124, 140, 229
EXPORT SUBSCRIPTION
 description 231
expressions 306
 stored procedures 30
EXTRACT FROM 316

F

fatal errors (advanced replication) 266
FLOAT data type 325
FLOOR (function) 315
fn
 usage in {fn func_name} 28, 37
FOR EACH REPLICA 50

- FOREIGN KEY (constraint) 105
- foreign keys 100, 101, 207
- free space in database 145
- FULL (keyword) 249
- full table scan 136
- functions
 - AVG 306
 - COUNT 306
 - MAX 306
 - MIN 306
 - miscellaneous 319
 - scalar 28
 - SET_PARAM() 319
 - SQL functions 321
 - stack viewing in stored procedures 46
 - SUM 306
 - triggers 322

G

- GET_PARAM()
 - description 319
- GET_UNIQUE_STRING 188, 313
 - example usage 191, 192
- GRANT EXECUTE ON statement 46
- GRANT PASSTHROUGH statement 235
- GRANT REFRESH ON
 - description 235
- GRANT statement 234

H

- HINT statement 236
- HOUR (function) 316

I

- IF statement
 - description 32
- IF-THEN construct
 - description 32
- IF-THEN-ELSE construct
 - description 32
- IF-THEN-ELSEIF construct
 - description 33
- IFNULL (system function) 319
- imdbsize ADMIN COMMAND 145
- IMPORT statement
 - description 242
- incremental publication
 - specifying 160
- indexes 135
 - columns 137
 - concatenated 137
 - creating 99
 - unique index 100
 - deleting 100
 - foreign key 101
 - managing 97
 - multi-column 137
 - primary key indexes 97
 - secondary key indexes 98
- INSERT (string function) 313
- INSERT statement 244
 - multirow 244
 - using default values 244

- INT data type 325
- INTEGER data type 325
- Intelligent Transaction
 - parameter bulletin board 277
 - using saved properties 277
- IS NULL (operator)
 - description 31

L

- large replicas
 - creating 231
- LCASE (function) 313
- LEFT (function) 313
- LENGTH (function) 313
- LIKE 305, 308, 311
 - in START AFTER COMMIT 299
- LIST statement 245
- listing users 155
- LOCATE (function) 314
- LOCK TABLE statement 247
- Lock timeout
 - setting 120
 - setting, optimistic tables 121
- locking
 - description 118
 - mode, displaying 361
 - optimistic 112
 - pessimistic 112
- locking modes
 - mixed 119
- locks
 - duration 117
 - exclusive 115
 - EXCLUSIVE LOCK 115
 - modes
 - EXCLUSIVE 116
 - SHARED 115
 - UPDATE 116
 - shared 115
 - SHARED LOCK 115
 - update 115
 - UPDATE LOCK 115
- LOG (function) 315
- LOG10 (function) 315
- logical conditions
 - description 32
- logical database 170
- logical operators
 - AND 30
 - description 30
 - IS NULL 31
 - NOT 30, 35
 - OR 30
- LOGIN_CATALOG (system function) 318
- LOGIN_SCHEMA (system function) 318
- LOGIN_USERID (system function) 318
- logsize ADMIN COMMAND 145
- LONG NATIONAL VARCHAR data type 324
- LONG VARBINARY data type 326
 - using CAST to enter values 327
- LONG VARCHAR data type 324
- LONG WVARCHAR data type 324
- LOOP 187
- loops in stored procedures 34
- LTRIM (function) 314

M

- MAINTENANCE
 - set sync mode maintenance 290
- Maintenance Mode 290
- master database
 - changing network name 289
 - dropping 221
 - granting access to publications 235
 - MESSAGE APPEND statement 249
 - parameters 293
 - propagating transactions to 249
 - properties in 277
 - requesting reply messages 266
 - revoking access to publications 274
 - setting node name 291
 - user information 249
- master users
 - downloading list 249
- MAX (function) 306
- MaxStartStatements (parameter) 131
- maxusers ADMIN COMMAND 145
- memtotal ADMIN COMMAND 145
- MESSAGE APPEND PROPAGATE TRANSACTIONS
 - description 249
- MESSAGE APPEND PROPAGATE WHERE
 - using properties 277
- MESSAGE APPEND REFRESH 249
 - description 249
- MESSAGE APPEND REGISTER PUBLICATION
 - description 249
- MESSAGE APPEND REGISTER REPLICA
 - description 249
- MESSAGE APPEND SUBSCRIBE 249
- MESSAGE APPEND SYNC_CONFIG
 - description 249
- MESSAGE APPEND UNREGISTER PUBLICATION
 - description 249
- MESSAGE APPEND UNREGISTER REPLICA
 - description 249
- MESSAGE BEGIN
 - statement 252
- MESSAGE DELETE
 - description 253
- MESSAGE END
 - description 256
- MESSAGE EXECUTE
 - description 257
- MESSAGE FORWARD
 - description 259
- MESSAGE FROM REPLICA DELETE 264
 - description 254
- MESSAGE FROM REPLICA EXECUTE
 - description 264
- MESSAGE FROM REPLICA RESTART 265
- MESSAGE GET REPLY
 - description 266
- messages
 - beginning 252
 - deleting 253
 - ending 256
 - error messages, failed messages, reply messages 254, 264
 - executing 257
 - re-executing 257
 - requesting replies from the master database 266
 - saving 256
 - sending 259

- metadata
 - exporting 231
- MIN (function) 306
- MINUTE (function) 316
- MOD (function) 315
- monitorstate ADMIN COMMAND 145
- MONTH (function) 316
- MONTHNAME (function) 317
- multi-column indexes 137

N

- name ADMIN COMMAND 145
- named parameters 24
- NATIONAL CHAR data type 324
- NATIONAL CHARACTERdata type 324
- NATIONAL VARCHAR data type 324
- NCHAR data type 324
- NCHAR LARGE OBJECT data type 324
- NCHAR VARYING data type 324
- NCLOB data type 324
- node
 - setting 291
- node-def 47
- NONUNIQUE 50
- NORMAL
 - set sync mode normal 290
- NOT (operator) 30, 304
- NOT NULL 17
- NOTUNIQUE 299
- NOW (function) 317
- NULL 15
- NULLIF 307
- nulls
 - handling 36
- numcursors ADMIN COMMAND 145
- NUMERIC data type 326
- numeric functions 314
- numlocks ADMIN COMMAND 145
- nummerges ADMIN COMMAND 145
- numtransactions ADMIN COMMAND 145
- numusers ADMIN COMMAND 145
- NVARCHAR data type 324

O

- openstate ADMIN COMMAND 145
- optimistic locking 112
- optimizer hints
 - using 139
- optimizing
 - batch inserts and updates 138
- OR (operator) 30, 304

P

- parameter bulletin board 319
 - defining database-level parameters 293
 - description 321
 - Intelligent Transaction 277
- parameter modes 181
 - Input parameters 181
 - Input/output parameters 181
 - Output parameters 181
- parameters
 - database-level 319

- parameters (*continued*)
 - defining persistent database-level 293
 - deleting 293
 - EnableHints 139
 - get_param() 319
 - GET_PARAM() 319
 - MaxStartStatements 131
 - placing on bulletin board 321
 - put_param() 319
 - PUT_PARAM() 321
 - read-only 319
 - retrieving from bulletin board 319
 - SimpleSQLOpt 135
 - triggers 65
 - updatable 319
 - password
 - reset 91
 - passwords
 - administrator 93
 - changing 91, 93
 - entering 91
 - reset 93
 - percent sign character 311
 - performance 133
 - diagnosing problems 140
 - indexes 135
 - observing 123
 - single-table SQL queries 135
 - using indexes to improve 135
 - pessimistic locking 112
 - PI (function) 315
 - POSITION (function) 314
 - POWER (function) 315
 - PRECISION data type 325
 - primary keys 96, 100
 - indexes 97
 - primarystarttime ADMIN COMMAND 145
 - privileges
 - managing 90
 - stored procedures 46
 - PROC_COUNT function
 - stored procedure stack 46
 - PROC_NAME (N) function
 - stored procedure stack 46
 - PROC_SCHEMA (N) function
 - stored procedure 46
 - procedures
 - stored procedures 24
 - processsize ADMIN COMMAND 145
 - proctrace 130
 - propagating transactions
 - SAVE command 275
 - setting default properties 277
 - setting priority 277
 - terminated messages 266
 - usage 249
 - properties
 - assigning 277
 - saving as default 277
 - saving default transaction propagation criteria 277
 - pseudo columns 309
 - psize ADMIN COMMAND 145
 - publications
 - creating 197
 - dropping 222
 - granting access 235
 - refreshing 249
 - publications (*continued*)
 - revoking access 274
 - Push Synchronization 50
 - Example 59
 - PUT_PARAM()
 - description 321
- ## Q
- QUARTER (function) 317
- ## R
- RADIANS (function) 315
 - re-executing
 - messages 257
 - READ COMMITTED 280
 - REAL data type 325, 328
 - recovery
 - and transaction logging 5
 - DBMS level error 254, 264
 - referenced table 100
 - REFERENCES (keyword) 207, 234, 273
 - referencing table 100
 - referential actions
 - Cascade 104
 - No action 103
 - Restrict 104
 - Set default 104
 - Set null 104
 - referential integrity
 - constraints 103
 - defining 207
 - dynamic constraint management 104
 - overview 100
 - transient tables 207
 - REFRESH statement 270
 - defining commit block 259
 - refreshes
 - handling failure in the master database 266
 - handling failure in the replica database 266
 - publications 249
 - REGISTER EVENT statement 273
 - registering
 - replica databases 249
 - setting replica node names 291
 - registrating databases
 - registration user 294
 - relational databases 1
 - remote stored procedures 47
 - REPEAT (function) 314
 - REPEATABLE READ 280
 - REPLACE (function) 314
 - replica databases
 - deleting messages 253
 - dropping 223
 - properties in 277
 - refreshing from publications 249
 - registering 249, 291, 294
 - requesting reply messages from the master database 266
 - saving transactions 275
 - setting parameters in 293
 - unregistering 249
 - Replica Property Names 50
 - reply messages
 - requesting from the master database 266

- reply messages (*continued*)
 - setting timeout 259
- RESTRICT keyword 95, 219, 224
 - REVOKE statements 273
- RETURN keyword 36
- REVOKE (privilege from role or user) statement 273
- REVOKE (role from user) statement 273
- REVOKE PASSTHROUGH statement 274
- REVOKE REFRESH 274
- REVOKE REFRESH ON
 - description 274
- REVOKE SUBSCRIBE 274
- RIGHT (function) 314
- roles
 - _SYSTEM 91, 202
 - PUBLIC 91, 202
 - SYS_ADMIN_ROLE 91, 202
 - SYS_CONSOLE_ROLE 91, 202
 - SYS_SYNC_ADMIN_ROLE 91, 202
 - SYS_SYNC_REGISTER_ROLE 91, 202
- ROLLBACK (statement) 5
 - stored procedures 45
- ROLLBACK WORK statement 275
- ROUND (function) 315
- row 1, 7
- row value constructors 19
- ROWID 135
- ROWNUM 135, 309, 341
- RTRIM (function) 314
- RVC (row value constructors) 19

S

- SAVE
 - description 275
- SAVE DEFAULT PROPAGATE PROPERTY WHERE
 - description 277
- SAVE DEFAULT PROPERTY
 - description 277
- SAVE PROPERTY statement 277
- saving
 - messages 256
- scalar functions 28
 - description 28, 306
- schemas
 - creating 108
 - deleting 109
 - description 107, 203
- SECOND 317
- secondary key
 - indexes 98
- secondarystarttime ADMIN COMMAND 145
- SELECT statement 278
- sending
 - messages 259
- sequences 79
- SERIALIZABLE 280
- sernum ADMIN COMMAND 145
- SERVER_INFO system view 382
- SET
 - differences between SET and SET TRANSACTION 281, 296
- SET CATALOG catalog_name 280
- SET CATALOG statement 107
- SET DELETE CAPTURE statement 282
- SET DURABILITY 121, 280
- SET HISTORY COLUMNS
 - description 160
- SET IDLE TIMEOUT 280
- SET ISOLATION LEVEL 280
- SET LOCK TIMEOUT 280
- SET NOSYNCHISTORY
 - description 160
- SET OPTIMISTIC LOCK TIMEOUT 280
- SET PASSTHROUGH statement 284
- SET READ-ONLY 280
- SET READ-WRITE 280
- SET SAFENESS 280
- SET SCHEMA 280
- SET SCHEMA statement 107, 284
- SET SCHEMA USER statement 284
- SET SEQUENCE 79
- SET SQL statement 286
- SET statement 280
 - in stored procedures 28
- SET STATEMENT MAXTIME 280
- SET SYNC CONNECT 168
 - description 289
- SET SYNC MODE statement 290
- SET SYNC NODE
 - description 291
- SET SYNC PARAMETER
 - description 293
- SET SYNC USER IDENTIFIED BY
 - description 294
- SET SYNCHISTORY 159
 - description 160
- set theory 9
- SET TRANSACTION
 - differences between SET and SET TRANSACTION 281, 296
- SET TRANSACTION DELETE CAPTURE statement 297
- SET TRANSACTION DURABILITY 121
- SET TRANSACTION PASSTHROUGH statement 298
- SET TRANSACTION statement 295
- SET TRANSACTION WRITE 295
- SET vs. SET TRANSACTION 281, 296
- SET WRITE 280
- SHARED (lock mode) 115
- shared locks 115
- SIGN (function) 315
- SimpleSQLOpt (parameter) 135
- SIN 315
- SLEEP 319
- SMALLINT data type 325
- solidDB
 - data management 111
- solidDB SQL
 - data management 111
 - data types 89
 - extensions 89
 - functions 89
 - using for database administration 89
- solidDB SQL syntax
 - compliance 89
- soltrace.out 129
- SOUNDEX (function) 314
- SPACE (function) 314
- space ADMIN COMMAND 145
- SQL
 - getting started 7
 - subqueries 12
 - using in stored procedures 46

- SQL functions
 - GET_PARAM() 319
 - PUT_PARAM() 321
- SQL Info facility 123
- SQL optimization 135
- SQL scripts 90
 - sample.sql 94
 - users.sql 90
- SQL statements
 - examples of 94
 - tuning 133
 - using 89
- SQL wildcards 311
- SQL_LANGUAGES system table 345
- SQL_TSI_DAY 317
- SQL_TSI_FRAC_SECOND 317
- SQL_TSI_HOUR 317
- SQL_TSI_MINUTE 317
- SQL_TSI_MONTH 317
- SQL_TSI_QUARTER 317
- SQL_TSI_SECOND 317
- SQL_TSI_WEEK 317
- SQL_TSI_YEAR 317
- SQL-92 89
- SQL-99 89
- SQLERRNUM (variable)
 - error code 42
- SQLERROR (variable)
 - error string 42
- SQLERROR OF cursorname (variable) 42
- SQLERRSTR (variable)
 - error string 42
- SQLROWCOUNT (variable)
 - row count 42
- SQLSUCCESS (variable)
 - stored procedure 41
- SQRT (function) 315
- SSC_TASK_BACKGROUND 131
- START AFTER COMMIT statement 299
 - analyzing failures in 131
 - tuning performance of 131
- stored procedures
 - assigning values to variables 28
 - autocommit 181
 - CREATE PROCEDURE statement 24
 - cursors 46
 - declaring local variables 27
 - default cursor management 45
 - default values 24
 - description 23
 - error handling 41
 - exiting 36
 - input parameters 24
 - input/output parameters 24
 - loops 34
 - nesting procedures 43
 - output parameters 24
 - parameter markers in cursors 39
 - positioned updates and deletes 44
 - privileges 46
 - procedure body 28
 - procedure stack viewing 46
 - remote 47
 - tracing facilities for 129
 - transactions 45
 - triggers 64
 - using events 80
 - stored procedures (*continued*)
 - using parameters 24
 - using SQL 46
 - using SQL in 37
- string functions 313
- strings
 - zero-length 36
- SUBSCRIBERefresh 249
- subscriptions
 - defining commit block 266
 - dropping 225
 - exporting 231
 - importing 242
- SUBSTRING (function) 314
- SUM (function) 306
- Sync Pull Notify 50
 - Example 59
- SYNC_CONFIG 249
- SYNC_DELETE_MESSAGES stored procedure 398
- SYNC_DELETE_REPLICA_MESSAGES stored procedure 399
- SYNC_GET_CATALOG_SYNC_NODE_NAME_AND_MODE stored procedure 400
- SYNC_MAINTENANCEMODE_BEGIN (event) 290, 404
- SYNC_MAINTENANCEMODE_END (event) 290, 404
- SYNC_REGISTER_PUBLICATION stored procedure 394
- SYNC_REGISTER_REPLICA stored procedure 392
- SYNC_SETUP_CATALOG stored procedure 391
- SYNC_SHOW_REPLICA_SUBSCRIPTIONS stored procedure 397
- SYNC_SHOW_SUBSCRIPTIONS stored procedure 396
- SYNC_UNREGISTER_PUBLICATION stored procedure 395
- SYNC_UNREGISTER_REPLICA stored procedure 393
- SYNHISTORY 159
- synchronization
 - messages 321
- SYS Attauth system table 345
- SYS_BACKGROUNDJOB_INFO 131
- SYS_BACKGROUNDJOB_INFO system table 347
- SYS_BLOBS system table 347
- SYS_BULLETIN_BOARD system table 365
- SYS_CARDINAL system table 349
- SYS_CATALOGS system table 350
- SYS_CHECKSTRINGS system table 350
- SYS_COLUMNS system table 350
- SYS_COLUMNS_AUX system table 351
- SYS_DL_REPLICA_CONFIG system table 351
- SYS_DL_REPLICA_DEFAULT system table 352
- SYS_EVENT_BACKUP 402
- SYS_EVENT_BACKUPREQ 402
- SYS_EVENT_CHECKPOINT (event) 402
- SYS_EVENT_CHECKPOINTREQ 402
- SYS_EVENT_ERROR 402, 408, 409
- SYS_EVENT_IDLE 403
- SYS_EVENT_ILL_LOGIN 404
- SYS_EVENT_IMDB_MEMORY 403
- SYS_EVENT_MERGE 404
- SYS_EVENT_MERGEREQ 404
- SYS_EVENT_MESSAGES 405
- SYS_EVENT_NOTIFY 405
- SYS_EVENT_PARAMETER 405
- SYS_EVENT_PROCESS_MEMORY 406
- SYS_EVENT_ROWS2MERGE 406
- SYS_EVENT_SACFAILED 131, 406
- SYS_EVENT_SHUTDOWNREQ 407
- SYS_EVENT_STATE_MONITOR 407
- SYS_EVENT_STATE_OPEN 407
- SYS_EVENT_STATE_SHUTDOWN 407

SYS_EVENT_STATE_TRACE 407
 SYS_EVENT_TMCMD 407
 SYS_EVENT_TRX_TIMEOUT 408
 SYS_EVENT_USERS 408
 SYS_EVENTS system table 352
 SYS_FORKEYPARTS system table 354
 SYS_FORKEYS system table 354
 SYS_GETBACKGROUNDJOB_INFO 131
 SYS_GETBACKGROUNDJOB_INFO stored procedure 400
 SYS_HOTSTANDBY system table 355
 SYS_HSBDDSTATE_WAIT stored procedure 400
 SYS_KEYPARTS system table 355
 SYS_KEYS system table 356
 SYS_PROCEDURE_COLUMNS system table 357
 SYS_PROCEDURES system table 357
 SYS_PROPERTIES system table 358
 SYS_PUBLICATION_ARGS system table 365
 SYS_PUBLICATION_REPLICA_ARGS system table 365
 SYS_PUBLICATION_REPLICA_STMTARGS system table 366
 SYS_PUBLICATION_REPLICA_STMTS system table 366
 SYS_PUBLICATION_STMTARGS system table 366
 SYS_PUBLICATION_STMTS system table 367
 SYS_PUBLICATIONS system table 367
 SYS_PUBLICATIONS_REPLICA system table 368
 SYS_RELAUTH system table 359
 SYS_SCHEMAS system table 359
 SYS_SEQUENCES system table 359
 SYS_SYNC_BOOKMARKS system table 368
 SYS_SYNC_HISTORY_COLUMNS system table 369
 SYS_SYNC_INFO system table 369
 SYS_SYNC_MASTER_MSGINFO system table 369
 SYS_SYNC_MASTER_RECEIVED_BLOB_REFS system table 371
 SYS_SYNC_MASTER_RECEIVED_MSGPARTS system table 371
 SYS_SYNC_MASTER_RECEIVED_MSGS system table 371
 SYS_SYNC_MASTER_STORED_BLOB_REFS system table 372
 SYS_SYNC_MASTER_STORED_MSGPARTS system table 372
 SYS_SYNC_MASTER_STORED_MSGS system table 372
 SYS_SYNC_MASTER_SUBSC_REQ system table 373
 SYS_SYNC_MASTER_VERSIONS system table 373
 SYS_SYNC_MASTERS system table 374
 SYS_SYNC_RECEIVED_BLOB_ARGS system table 374
 SYS_SYNC_RECEIVED_STMTS system table 375
 SYS_SYNC_REPLICA_MSGINFO system table 375
 SYS_SYNC_REPLICA_PROPERTIES system table 360
 SYS_SYNC_REPLICA_RECEIVED_BLOB_REFS system table 376
 SYS_SYNC_REPLICA_RECEIVED_MSGPARTS system table 377
 SYS_SYNC_REPLICA_RECEIVED_MSGS system table 377
 SYS_SYNC_REPLICA_STORED_BLOB_REFS system table 378
 SYS_SYNC_REPLICA_STORED_MSGPARTS system table 378
 SYS_SYNC_REPLICA_STORED_MSGS system table 378
 SYS_SYNC_REPLICA_VERSIONS system table 379
 SYS_SYNC_REPLICAS system table 379
 SYS_SYNC_SAVED_BLOB_ARGS system table 379
 SYS_SYNC_SAVED_STMTS system table 380
 SYS_SYNC_TRX_PROPERTIES system table 380
 SYS_SYNC_USERMAPS system table 381
 SYS_SYNC_USERS system table 381
 SYS_SYNONYM system table 360
 SYS_TABLEMODES system table 361
 SYS_TABLES system table 361
 SYS_TF1_WAIT stored procedure 401
 SYS_TRIGGERS system table 362

SYS_TYPES system table 363
 SYS_URole system table 363
 SYS_USERS system table 364
 SYS_VIEWS system table 364
 system functions 318
 system parameters 319
 system tables 345
 description 94
 granting access 94
 viewing 94
 system views 381

T

table locks 115
 tables 1, 7
 adding columns to 96
 aliases 12
 committing work after altering 96
 creating 95
 deleting columns from 96
 managing 94
 removing 95
 TABLES system view 383
 TAN 315
 THEN
 keyword in CASE statement 307
 TIME data type 327
 timeout
 setting for reply messages 259
 TIMEOUT (keyword)
 MESSAGE FORWARD 259
 MESSAGE GET REPLY 259
 TIMESTAMP data type 328
 TIMESTAMP_FORMAT (function) 318
 TIMESTAMPADD 317
 TIMESTAMPDIF 317
 TINYINT data type 325
 TO (keyword)
 MESSAGE FORWARD 259
 TO_CHAR (function) 314
 TO_DATE (function) 318
 tracestate ADMIN COMMAND 145
 Tracing facilities for stored procedures and triggers 129
 transaction bulletin board 319
 transaction durability level
 improving performance with 121
 setting 121
 Transaction durability level
 choosing 121
 transactions 21
 assigning properties 277
 COMMIT WORK 5
 defining 111
 description 5, 117
 Intelligent Transactions 249, 275
 logs
 overview 5
 propagating 249
 propagation 249
 read-only 111
 read-write 111
 ROLLBACK 5
 saving 275
 saving default properties 277
 setting default properties for propagation 277
 setting propagation priority 277

- transactions (*continued*)
 - stored procedures 45
 - transaction log 5
 - using triggers in 67
- triggers
 - code example 76
 - comments and restrictions 215
 - creating 64
 - error handling 73
 - functions for analyzing and debugging 322
 - how they work 63
 - nested triggers 73
 - obtaining information 75
 - parameter settings 75
 - parameters and variables 65
 - privileges and security 74
 - procedures 64
 - recursive triggers 73
 - setting default or derived columns 65
 - tracing facilities for 129
 - transactions 67
 - using 63
- TRIM (function) 314
- TRUNCATE (function) 315
- TRUNCATE TABLE statement 301
- tuning
 - SQL statements 133
 - SQL statements and applications 133

U

- UCASE (function) 314
- UIC (system function) 318
- underline 311
- underscore 311
- UNIQUE 50, 105, 299
- unique constraint 96
- UNLOCK TABLE statement 302
- unregistering
 - replica databases 249
- UPDATE (lock mode) 116
- UPDATE (searched) statement 303
- update locks 115
- uptime ADMIN COMMAND 145
- user privileges 90
 - granting 92
 - granting administrator privileges 93
 - revoking 92
- user roles 90
 - administrator 91, 93, 202
 - changing password 91
 - creating 92
 - deleting 92
 - giving a user a role 92
 - granting privileges to 92
 - reserved role names 90
 - revoking privileges from 92
 - revoking the role of a user 93
 - system console role 91, 202
- userlist ADMIN COMMAND 154, 155
- usernames
 - reserved names 90
- users
 - creating 91
 - deleting 91
- users and roles
 - committing work after altering 93

- USERS system view 383
- usertrace 129

V

- VARBINARY data type 326
- VARCHAR data type 324
- VARCHAR_FORMAT (function) 314
- variables
 - triggers 65
- Variables
 - assigning in stored procedures 28
 - SQLERRNUM 42
 - SQLERROR 42
 - SQLERROR OF cursorname 42
 - SQLERRSTR 42
 - SQLROWCOUNT 42
 - SQLSUCCESS 41

W

- WCHAR data type 324
- WEEK (function) 318
- WHEN
 - in case_specification 307
 - keyword in event specification 172
- WHERE (keyword)
 - PROPAGATE TRANSACTIONS 249
- WHILE-LOOP statement
 - description 34
- wildcard characters 311
- WRITETRACE 129
- WVARCHAR data type 324

Y

- YEAR (function) 318

Z

- Zero-length strings 36

Notices

© Copyright Oy IBM Finland Ab 1993, 2013.

All rights reserved.

No portion of this product may be used in any way except as expressly authorized in writing by IBM.

This product is protected by U.S. patents 6144941, 7136912, 6970876, 7139775, 6978396, 7266702, 7406489, 7502796, and 7587429.

This product is assigned the U.S. Export Control Classification Number ECCN=5D992b.

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 Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

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.

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_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, ibm.com[®], Solid, solidDB, InfoSphere, DB2, Informix[®], and WebSphere[®] 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" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.



SC27-3841-05

