

IBM solidDB
IBM solidDB Universal Cache
Version 7.0

Programmer Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 303.

First edition

This edition applies to version 7, release 0 of IBM solidDB (product number 5724-V17) and IBM solidDB Universal Cache (product number 5724-W91) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Oy International Business Machines Ab Ltd. 1993, 2011

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

About this manual	xi
------------------------------------	-----------

Typographic conventions	xi
Syntax notation conventions.	xii

1 Introduction to solidDB APIs 1

1.1 solidDB ODBC Driver	1
1.1.1 Using solidDB ODBC Driver functions	2
1.1.2 ODBC API basic application steps	2
1.1.3 Format of the solidDB connect string	3
1.1.4 Client-side solid.ini configuration file.	5
1.1.5 ODBC non-standard behavior	5
1.2 solidDB JDBC Driver	6
1.3 solidDB Application Programming Interface (SA API)	7
1.4 solidDB Server Control API (SSC API)	7
1.5 Building client applications	7
1.5.1 What is a client?	7
1.5.2 How is the query passed to the server?.	7
1.5.3 How are the results passed back to the client?	9
1.5.4 Statement cache	9

2 Using solidDB ODBC API 11

2.1 Installing solidDB ODBC Driver	11
2.2 Using the ODBC driver library.	12
2.3 solidDB ODBC Driver 3.51 Features Support	12
2.4 Overview of usage on Windows operating systems.	13
2.5 Calling functions	14
2.6 Connecting to a data source.	16
2.6.1 Network name and connect string syntax	16
2.6.2 Using logical data source names.	19
2.6.3 Empty data source name	20
2.6.4 Configuring the solidDB ODBC Data Source for Windows	20
2.6.5 Using solidDB ODBC Driver with unixODBC.	21
2.6.6 Retrieving user login information	23
2.7 ODBC handle validation	24
2.8 Executing transactions	24
2.9 Retrieving information about the data source's catalog	26
2.10 Using ODBC extensions to SQL	27
2.10.1 Procedures	27
2.10.2 Hints	28
2.10.3 Additional ODBC extension functions	29
2.11 solidDB Extensions for ODBC API	29
2.12 Using cursors	32
2.12.1 Assigning storage for rowsets (binding).	32
2.12.2 Cursor support	33

2.13 Using bookmarks	39
2.14 Error text format	39
2.14.1 Processing error messages	40
2.15 Terminating transactions and connections.	41
2.16 Constructing an application	41
2.17 Testing and debugging an application	50

3 Using solidDB JDBC Driver 51

3.1 What is solidDB JDBC Driver	51
3.2 Getting started with solidDB JDBC Driver	51
3.2.1 Registering solidDB JDBC Driver	53
3.2.2 Connecting to the database	53
3.3 Special notes about solidDB and JDBC	56
3.4 JDBC driver interfaces and methods	56
3.5 solidDB JDBC Driver extensions	62
3.5.1 WebSphere compatibility	63
3.5.2 Connection timeout in JDBC	64
3.5.3 Non-standard JDBC connection properties	64
3.6 JDBC 2.0 optional package API support.	68
3.6.1 JDBC connection pooling	68
3.6.2 solidDB Connected RowSet Class: SolidJDBCRowSet	76
3.6.3 Java Naming and Directory Interface (JNDI)	78
3.7 Code examples	78
3.8 solidDB JDBC Driver type conversion matrix	89

4 Using solidDB SA 91

4.1 What is solidDB SA?	91
4.2 Getting started with solidDB SA	92
4.3 Writing data by using solidDB SA without SQL	93
4.4 Reading data by using solidDB SA without SQL	95
4.5 Running SQL Statements by Using solidDB SA	97
4.6 Transactions and autocommit mode	97
4.7 Handling database errors	98
4.8 Special notes about solidDB SA	99
4.9 solidDB SA Function Reference	101
4.9.1 SaArrayFlush	103
4.9.2 SaArrayInsert	104
4.9.3 SaColSearchCreate	104
4.9.4 SaColSearchFree.	105
4.9.5 SaColSearchNext	105
4.9.6 SaConnect.	106
4.9.7 SaCursorAscending	106
4.9.8 SaCursorAtleast	107
4.9.9 SaCursorAtmost.	107
4.9.10 SaCursorBegin	108
4.9.11 SaCursorClearConstr.	108
4.9.12 SaCursorColData	109
4.9.13 SaCursorColDate	110
4.9.14 SaCursorColDateFormat.	111
4.9.15 SaCursorColDfloat	111
4.9.16 SaCursorColDouble	112
4.9.17 SaCursorColDynData	113
4.9.18 SaCursorColDynStr	114
4.9.19 SaCursorColFloat	115

4.9.20 SaCursorColInt	116
4.9.21 SaCursorColLong	117
4.9.22 SaCursorColNullFlag	117
4.9.23 SaCursorColStr	118
4.9.24 SaCursorColTime	119
4.9.25 SaCursorColTimestamp	120
4.9.26 SaCursorCreate	120
4.9.27 SaCursorDelete	121
4.9.28 SaCursorDescending	121
4.9.29 SaCursorEnd	122
4.9.30 SaCursorEqual	122
4.9.31 SaCursorErrorInfo	123
4.9.32 SaCursorFree	123
4.9.33 SaCursorInsert	124
4.9.34 SaCursorLike	124
4.9.35 SaCursorNext	125
4.9.36 SaCursorOpen	125
4.9.37 SaCursorOrderByVector	125
4.9.38 SaCursorPrev	126
4.9.39 SaCursorReSearch	127
4.9.40 SaCursorSearch	127
4.9.41 SaCursorSearchByRowid	128
4.9.42 SaCursorSearchReset	128
4.9.43 SaCursorSetLockMode	130
4.9.44 SaCursorSetPosition	131
4.9.45 SaCursorSetRowsPerMessage	132
4.9.46 SaCursorUpdate	132
4.9.47 SaDateCreate	133
4.9.48 SaDateFree	133
4.9.49 SaDateSetAsciiiz	133
4.9.50 SaDateSetTimet	134
4.9.51 SaDateToAsciiiz	135
4.9.52 SaDateToTimet	136
4.9.53 SaDefineChSet	136
4.9.54 SaDfloatCmp	137
4.9.55 SaDfloatDiff	137
4.9.56 SaDfloatOverflow	138
4.9.57 SaDfloatProd	138
4.9.58 SaDfloatQuot	139
4.9.59 SaDfloatSetAsciiiz	139
4.9.60 SaDfloatSum	140
4.9.61 SaDfloatToAsciiiz	140
4.9.62 SaDfloatUnderflow	141
4.9.63 SaDisconnect	141
4.9.64 SaDynDataAppend	141
4.9.65 SaDynDataChLen	142
4.9.66 SaDynDataClear	143
4.9.67 SaDynDataCreate	143
4.9.68 SaDynDataFree	144
4.9.69 SaDynDataGetData	144
4.9.70 SaDynDataGetLen	145
4.9.71 SaDynDataMove	145
4.9.72 SaDynDataMoveRef	146
4.9.73 SaDynStrAppend	147
4.9.74 SaDynStrCreate	148
4.9.75 SaDynStrFree	148
4.9.76 SaDynStrMove	148
4.9.77 SaErrorInfo	149
4.9.78 SaGlobalInit	150
4.9.79 SaSetDateFormat	150
4.9.80 SaSetSortBufSize	151

4.9.81 SaSetSortMaxFiles	151
4.9.82 SaSetTimeFormat	152
4.9.83 SaSetTimestampFormat	153
4.9.84 SaSQLExecDirect	153
4.9.85 SaTransBegin	154
4.9.86 SaTransCommit	154
4.9.87 SaTransRollback	155
4.9.88 SaUserId	155

5 Using Unicode 157

5.1 What is Unicode?	158
5.2 Designing Unicode databases	159
5.3 Using solidDB tools with Unicode	161
5.4 Compatibility between Unicode and partial Unicode databases	162
5.4.1 Converting partial Unicode databases to Unicode	162
5.5 Developing applications for Unicode	164
5.5.1 ODBC applications and Unicode databases	165
5.5.2 JDBC applications and Unicode databases	167

6 Using Transaction Log Reader . . . 169

6.1 Considerations for developing applications with Log Reader	169
6.2 Configuring the Log Reader	171
6.3 Reading log data with the Log Reader	172
6.4 Partitioning and filtering log records	173
6.4.1 Creating and deleting partitions	173
6.4.2 Using partition filters	173
6.5 Setting transaction batches	174

Appendix A. solidDB supported ODBC functions 175

Appendix B. solidDB ODBC Driver 3.5.1 attributes support 185

Appendix C. SQLSTATE error codes 193

Appendix D. Minimum SQL grammar requirements for ODBC 215

D.1 SQL statements	215
D.1.1 Control statements (logical condition)	216
D.2 Data type support	217
D.3 Parameter data types	218
D.4 Literals in ODBC	219
D.5 List of reserved keywords	220

Appendix E. Data types 223

E.1 SQL data types	223
E.2 C data types	223
E.3 Data type identifiers	224
E.4 SQL data types	224
E.5 C data types	228
E.6 Numeric literals	233
E.7 Overriding default precision and scale for numeric data types	235
E.8 Data type identifiers and descriptors	236
E.9 Decimal digits	237

E.10 Transfer octet length	238
E.11 Constraints of the gregorian calendar	240
E.12 Converting data from SQL to C data types	241
E.12.1 Data conversion tables from SQL to C	243
E.12.2 SQL to C data conversion examples	253
E.13 Converting data from C to SQL data types	254
E.13.1 Data conversion tables from C to SQL	256
E.13.2 C to SQL data conversion examples	267

Appendix F. Scalar functions. 269

F.1 ODBC and SQL-92 scalar functions	269
F.2 String functions	269
F.3 Numeric functions	273
F.4 Time and date functions	276
F.5 System functions	281
F.6 Explicit data type conversion	282
F.7 SQL-92 CAST function	283

Appendix G. Timeout controls 285

G.1 Client timeouts	285
G.2 Server timeouts	288
G.3 HotStandby timeouts	291

Appendix H. Client-side configuration parameters 293

H.1 Setting client-side parameters through the solid.ini configuration file	293
H.2 Client section	294
H.3 Com section	295
H.4 Data Sources section	295
H.5 SharedMemoryAccess section	296
H.6 TransparentFailover section	296

Index 297

Notices 303

Figures

1. ODBC driver setup 21
2. ODBC data source administrator 23

Tables

1. Typographic conventions	xi	54. SaColSearchCreate Parameters	105
2. Syntax notation conventions	xii	55. SaColSearchNext Parameters	105
3. Connect string options	3	56. SaColSearchNext Return Value.	106
4. Connect string options	17	57. SaConnect Parameters	106
5. Additional ODBC Extension Functions	29	58. SaConnect Return Value	106
6. solidDB-specific ODBC functions to ODBC API	30	59. SaCursorAscending parameters	107
7. A Sample Resultset	35	60. SaCursorAtleast Parameters.	107
8. A sample resultset	36	61. SaCursorAtmost Parameters	108
9. A sample resultset	37	62. SaCursorBegin Parameters	108
10. A Sample Resultset	37	63. SaCursorClearConstr Parameters	108
11. A Sample Resultset	38	64. SaCursorColData Parameters	110
12. Errors in a Data Source	39	65. SaCursorColDate Parameters	110
13. Sample Error Messages	40	66. SaCursorColDateFormat parameters	111
14. SQLSTATE values	40	67. SaCursorColDfloat Parameters	112
15. Differences to the Standard CallableStatement Interface	57	68. SaCursorColDouble Parameters	112
16. Differences to the Standard Connection Interface	58	69. SaCursorColDynData Parameters	114
17. Differences to the Standard PreparedStatement Interface	59	70. SaCursorColDynStr Parameters	115
18. Differences to the Standard ResultSet Interface	60	71. SaCursorColFloat Parameters	115
19. Differences to the Standard Statement Interface	61	72. SaCursorColInt Parameters	116
20. Differences to the Standard ResultSet Interface	62	73. SaCursorColLong Parameters	117
21. Constructor	68	74. SaCursorColNullFlag Parameters	118
22. Constructor	69	75. SaCursorColStr Parameters	119
23. setDescription.	69	76. SaCursorColTime parameters	119
24. getDescription	69	77. SaCursorColTimestamp parameters	120
25. setURL	70	78. SaCursorCreate Parameters	121
26. getURL	70	79. Return Value.	121
27. setUser	70	80. SaCursorDelete parameters	121
28. getUser	71	81. SaCursorDescending parameters	122
29. setPassword	71	82. SaCursorEnd parameters.	122
30. getPassword	71	83. SaCursorEqual parameters	122
31. setConnectionURL	72	84. SaCursorErrorInfo parameters	123
32. getConnectionURL	72	85. SaCursorFree parameters.	123
33. getLoginTimeout.	72	86. SaCursorInsert parameters	124
34. getLogWriter	73	87. SaCursorLike parameters	124
35. getPooledConnection	73	88. SaCursorNext parameters	125
36. getPooledConnection	73	89. SaCursorOpen parameters	125
37. setLoginTimeout	74	90. SaCursorOrderByVector parameters	126
38. setLogWriter	74	91. SaCursorPrev parameters	127
39. addConnectionEventListener.	75	92. SaCursorReSearch Parameters	127
40. close	75	93. SaCursorSearch parameters	127
41. getConnection.	75	94. SaCursorSearchByRowid parameters.	128
42. removeConnectionEventListener	76	95. SaCursorSearchReset Parameters	130
43. Java data type to SQL data type conversion	90	96. SaCursorSetLockMode Parameters	131
44. Insert operation steps	93	97. SaCursorSetPosition parameters	131
45. Update and delete operation steps	94	98. SaCursorSetRowsPerMessage parameters	132
46. Query Operation Steps.	96	99. SaCursorUpdate parameters	132
47. solidDB SA Function Return Codes	98	100. SaDateCreate Return Values.	133
48. Supported SQL Datatype.	100	101. SaDateFree parameters	133
49. solidDB SA Parameter Usage Types	102	102. SaDateSetAsciiZ Parameters.	134
50. Return Usage Types for Pointers	103	103. SaDateSetTimet parameters	135
51. SaArrayFlush Parameters	103	104. SaDateToAsciiZ parameters	135
52. SaArrayInsert Parameters	104	105. SaDateToTimet parameters	136
53. SaColSearchCreate Parameters	105	106. SaDefineChSet parameters	136
		107. SaDfloatCmp parameters.	137
		108. SaDfloatDiff parameters	137
		109. SaDfloatOverflow parameters	138

110.	SaDfloatProd Parameters	138	160.	ODBC Functions' Return Parameter	237
111.	SaDfloatQuot parameters	139	161.	SQL data type decimal digits	238
112.	SaDfloatSetAsciz parameters	139	162.	Descriptor field corresponding to decimal digits	238
113.	SaDfloatSum parameters	140	163.	ODBC Functions' Return parameter Decimal Attributes	239
114.	SaDfloatToAsciz parameters	140	164.	Transfer Octet Lengths	239
115.	SaDfloatUnderflow parameters.	141	165.	Constraints of the Gregorian Calendar	240
116.	SaDisconnect Parameters	141	166.	C Data Type — SQL_C_datatype where Datatype Is:	242
117.	SaDynDataAppend parameters	142	167.	Character SQL Data to ODBC C Data Types	244
118.	SaDynDataChLen Parameters	142	168.	SQL Data to ODBC C Data Types	247
119.	SaDynDataClear Parameters	143	169.	Binary SQL Data to ODBC C Data Types	249
120.	SaDynDataCreate Return Value	144	170.	Date SQL Data to ODBC C Data Types	250
121.	SaDynDataFree parameters	144	171.	Time SQL Data to ODBC C Data Types	251
122.	SaDynDataGetData parameters	144	172.	Timestamp SQL Data to ODBC C Data Types	252
123.	SaDynDataGetData Parameters	145	173.	SQL to C Data Conversion Examples	253
124.	SaDynDataMove Parameters	146	174.	SQL Data Type — SQL_datatype where Datatype Is:	255
125.	SaDynDataMoveRef Parameters	147	175.	C Character Data to ODBC SQL Data Types	257
126.	SaDynStrAppend parameters	147	176.	Numeric C Data to ODBC SQL Data Types	261
127.	SaDynStrCreate Return Value	148	177.	Bit C Data to ODBC SQL Data Types	262
128.	SaDynStrFree parameters	148	178.	Binary C Data to ODBC SQL Data Types	263
129.	SaDynStrMove Parameters	149	179.	Date C Data to ODBC SQL Data Types	264
130.	SaErrorInfo Parameters	149	180.	Time C Data to ODBC SQL Data Types	265
131.	SaSetDateFormat Parameters	150	181.	Timestamp C Data to ODBC SQL Data Ttypes	266
132.	SaSetSortBufSize Parameters	151	182.	C Data to SQL Data	267
133.	SaSetSortMaxFiles parameters	152	183.	String Function Arguments	270
134.	SaSetTimeFormat Parameters	152	184.	List of String Functions	270
135.	SaSetTimestampFormat parameters	153	185.	Numeric Function Arguments	273
136.	SaSQLExecDirect Parameters	153	186.	List of Numeric Functions	274
137.	SaTransBegin parameters.	154	187.	Time and Data Arguments	276
138.	SaTransCommit Parameters	155	188.	List of Time and Date Functions	276
139.	SaTransRollback Parameters.	155	189.	System Function Arguments	281
140.	SaUserId Parameters	155	190.	List of System Functions	281
141.	Command line options for solidDB tools for partial Unicode and Unicode databases	162	191.	Login timeouts	285
142.	solidDB supported ODBC functions	175	192.	Connection timeout	286
143.	001 Environment Level	185	193.	Query Timeout	288
144.	002 Connection Level	186	194.	SQL statement execution timeouts	289
145.	03 Statement Level.	187	195.	Lock wait timeout	289
146.	04 Column Attributes	190	196.	Optimistic lock wait timeout	290
147.	Error code class values	193	197.	Table lock wait timeout	290
148.	SQLSTATE codes	193	198.	Transaction Idle Timeout.	290
149.	Control Statements.	216	199.	connection idle timeout	291
150.	Determining Data Ttype for Several Types of Parameters	218	200.	Connect timeout	291
151.	List of Reserved Keywords	220	201.	Ping timeout.	292
152.	Common SQL Data Type Names, Ranges, and Limits	225	202.	Transparent connection timeout	292
153.	Data types SQLGetTypeInfo returns (1)	227	203.	Client parameters	294
154.	Data Types SQLGetTypeInfo Returns (2)	227	204.	Client-side communication parameters	295
155.	Data Types SQLGetTypeInfo Returns (3)	227	205.	Data Sources parameters.	295
156.	C vs ODBC Naming Correspondence	229	206.	Shared memory access parameters (client-side)	296
157.	Conversions Involving Numeric Literals	233	207.	TransparentFailover parameters	296
158.	Override Default Precision and Scale Values for Numeric Data Type	235			
159.	Concise Type Identifier, Verbose Identifier, and Type Subcode for Each Datetime	236			

About this manual

This guide contains information about using IBM® solidDB® through the different Application Programming Interfaces, with or without the shared memory access (SMA), linked library access (LLA), or HotStandby.

solidDB ODBC Driver, solidDB Light Client, and solidDB JDBC Driver help your client application access solidDB.

- The solidDB ODBC Driver conforms to the Microsoft ODBC 3.51 API standard.
- The solidDB Light Client is a lightweight version of the solidDB ODBC API and is intended for environments where the footprint of the client application must be very small.
- The solidDB JDBC Driver is a solidDB implementation of the JDBC 2.0 standard.

This guide assumes general knowledge of relational databases and SQL. It also assumes familiarity with solidDB. If you will use the ODBC driver, this manual assumes a working knowledge of the C programming language. If you will use the JDBC driver, this manual assumes a working knowledge of the Java programming language.

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.

Table 1. Typographic conventions (continued)

Format	Used for
<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.
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 Introduction to solidDB APIs

solidDB supports ODBC and JDBC interfaces that enable application developers to write applications with C or Java languages calling ODBC functions or JDBC methods and writing or generating the SQL string at the application level. solidDB also provides two proprietary interfaces, solidDB Application Programming Interface (SA API) and solidDB Server Control API (SSC API). These allow, for example, C programs to directly call functions inside the database server. These proprietary interfaces are provided with the linked library access (LLA) and shared memory access (SMA) libraries.

Other programming environments

Multiple ways exist to raise the abstraction level from the ODBC/JDBC level. It can be done either by enabling database access from various (usually higher level programming or scripting languages, such as Visual Basic, Perl, and PHP) or enabling database access directly through application level objects that are able to load or save themselves without the application programmer having to be aware of database connections, transactions, or even SQL Strings.

Database access from higher level programming is usually based on some middleware component translating the higher level language calls to regular ODBC or JDBC calls. In these conditions, the middleware component is seen as an application from the database perspective. Usually the middleware components hide the difference between database brands. However, as solidDB is a relatively new product, not all middleware vendors explicitly list it among the supported database products. In those cases, there is usually an option to have a generic ODBC database or generic JDBC database that works with solidDB drivers.

Certain programming environments do not have a direct counterpart in solidDB applications, such as Embedded SQL or Java-based stored procedures. Applications designed to run on these programming environments must be redesigned to fit solidDB.

1.1 solidDB ODBC Driver

The native solidDB ODBC Driver conforms to the Microsoft ODBC 3.5.1 API standard.

The solidDB ODBC Driver is distributed in the form of a library (a Dynamic Link Library (DLL) on Windows). The solidDB installation package includes two ODBC drivers: one for Unicode and one for ASCII. The Unicode version is a superset of the ASCII version; you may use it with either Unicode or ASCII character sets.

The solidDB implementation of the ODBC API supports a rich set of database access operations sufficient for creating robust database applications, including:

- Allocating and deallocating handles
- Getting and setting attributes
- Opening and closing database connections
- Accessing descriptors
- Executing SQL statements

- Accessing schema metadata
- Controlling transactions
- Accessing diagnostic information

Depending on the application's request, the solidDB ODBC Driver can automatically commit each SQL statement or wait for an explicit commit or rollback request. When the driver performs a commit or rollback operation, the driver resets all statement requests associated with the connection.

1.1.1 Using solidDB ODBC Driver functions

Users on all platforms can access ODBC Driver supported functions with solidDB ODBC API.

The solidDB ODBC API is the native call level interface (CLI) for solidDB databases. It is distributed in the form of a library (a Dynamic Link Library (DLL) on Windows). The solidDB ODBC API is compliant with ANSI X3H2 SQL CLI standard.

solidDB's implementation of ODBC API supports a rich set of database access operations sufficient for creating robust database applications, including:

- Allocating and deallocating handles
- Getting and setting attributes
- Opening and closing database connections
- Accessing descriptors
- Executing SQL statements
- Accessing schema metadata
- Controlling transactions
- Accessing diagnostic information

Depending on the application's request, the solidDB ODBC Driver can automatically commit each SQL statement or wait for an explicit commit or rollback request. When the driver performs a commit or rollback operation, the driver resets all statement requests associated with the connection.

1.1.2 ODBC API basic application steps

A client database application calls the solidDB ODBC API directly (or through the ODBC Driver Manager) to perform all interactions with a database. For example, to insert, delete, update, or select records, you make a series of calls to functions in the ODBC API.

An application using ODBC API performs the following tasks:

1. The application allocates memory and creates handles, and establishes a connection to the database.
 - a. The application allocates memory for an environment handle (henv) and a connection handle (hdbc); both are required to establish a database connection.

An application may request multiple connections for one or more data sources. Each connection is considered a separate transaction space. In other words, a COMMIT or ROLLBACK on one connection will not commit or rollback any statements executed through any other connection.

- b. The SQLConnect() call establishes the database connection, specifying the server name (a connect string or a data source name), user id, and password.
 - c. The application then allocates memory for a statement handle.
 2. The application executes the statement. This requires a series of function calls.
 - a. The application calls either SQLExecDirect(), which both prepares and executes an SQL statement, or SQLPrepare() and SQLExecute(), which allows statements to be executed multiple times.
 - b. If the statement was a SELECT, the result columns must be bound to variables in the application so that the application can see the returned data. The SQLBindCol() function will bind the application's variables to the columns of the result set. The rows can then be fetched using SQLFetch() repeatedly. SELECT statements must be committed as soon as processing of the resultset is done.

If the statement was an UPDATE, DELETE, or INSERT, then the application needs to check if the execution succeeded and call SQLEndTran() to commit the transaction.
 3. Finally the application closes the connection and frees any handles.
 - a. The application frees the statement handle.
 - b. The application closes the connection.
 - c. The application frees the connection and environment handles (hdbc and henv).

Note that step 2 (executing SQL statements) may be done repeatedly, depending upon how many SQL statements need to be executed.

Read 2, "Using solidDB ODBC API," on page 11 for more information about using these API calls.

1.1.3 Format of the solidDB connect string

The client processes use the solidDB connect string (network name) to specify which server it will connect to.

A default connect string can be defined with the client-side **Com.Connect** configuration parameter. The connect string can also be supplied, for example, at connection time or when configuring data sources with an ODBC driver manager.

The same format of the connect string applies to the **Com.Connect** parameter as well as to the connect string used by solidDB tools or ODBC applications.

The format of a connect string is the following:

protocol_name [*options*] [*host_computer_name*] *server_name*

where

- *options* can be any combination of the following:

Table 3. Connect string options

Option	Description	Protocol
-4	Specifies that client connects using IPv4 protocol only.	TCP/IP
-6	Specifies that client connects using IPv6 protocol only. In Windows environments, this option is mandatory if IPv6 protocol is used.	TCP/IP

Table 3. Connect string options (continued)

Option	Description	Protocol
-i <i>source_address</i>	Specifies an explicit connecting socket source address for cases where the system default source IP address binding does not meet application needs. <i>source_address</i> can be an IP address or a host name.	TCP/IP
-z	Enables data compression for this connection	All
-c <i>milliseconds</i>	Specifies the login timeout (the default is operating-system-specific). A login request fails after the specified time has elapsed.	TCP/IP
-r <i>milliseconds</i>	Specifies the connection (or read) timeout. A network request fails when no response is received during the time specified. The value 0 (default) sets the timeout to infinite (operating system default timeout applies).	TCP/IP
-o <i>filename</i>	Turns on the Network trace facility and defines the name of the trace output file See <i>Network trace facility</i> in the <i>IBM solidDB Administrator Guide</i> for details.	All
-p <i>level</i>	Pings the server at the given level (0-5). Clients can always use the solidDB Ping facility at level 1 (0 is no operation/default). Levels 2, 3, 4 or 5 may only be used if the server is set to use the Ping facility at least at the same level. See <i>Ping facility</i> in the <i>IBM solidDB Administrator Guide</i> for details.	All
-t	Turns on the Network trace facility See <i>Network trace facility</i> in the <i>IBM solidDB Administrator Guide</i> for details.	All

- *host_computer_name* is needed with TCP/IP and Named Pipes protocols, if the client and server are running on different machines.
- *server_name* depends on the communication protocol:
 - In TCP/IP protocol, *server_name* is a service port number, such as '2315'.
 - In other protocols, *server_name* is a name, such as 'soliddb' or 'chicago_office'.
 For details on the syntax in different communication protocols, see *Communication protocols* in the *IBM solidDB Administrator Guide*.

Note:

- The *protocol_name* and the *server_name* must match the ones that the server is using in its network listening name.
- If given at the connection time, the connect string must be enclosed in quotes.
- All components of the connect string are case insensitive.

Examples

```
[Com]
Connect=tcp -z -c1000 1315

[Com]
Connect=nmpipe host22 SOLID
so1sql "tcp localhost 1315"
so1sql "tcp 192.168.255.1 1315"
rc = SQLConnect(hdbc, "upipe SOLID", (SWORD)SQL_NTS, "dba", 3, "dba", 3);
rc = SQLDriverConnect(hdbc,
                      (SQLHWND)NULL,
                      (SQLCHAR*)"DSN=tcp localhost 1964;UID=dba;PWD=dba",
                      38,
                      out_string,
                      255,
                      &out_length,
                      SQL_DRIVER_NOPROMPT);
```


1.1.4 Client-side solid.ini configuration file

The solidDB ODBC Driver gets its client configuration information from the client-side `solid.ini` file. The client-side configuration file must be located in the working directory of the application.

In most cases, only solidDB server-side parameters are used when programming for the solidDB. However, occasionally there is a need to use client-side parameters. For example, you might want to create an application that defines no data source, but takes the data source from the connect string defined in the client-side configuration file.

Note: In solidDB documentation, references to `solid.ini` file are usually for the server-side `solid.ini` file.

When the solidDB ODBC Driver is started, it attempts to open the configuration file `solid.ini`. If the file does not exist, solidDB will use the factory values for the parameters. If the file exists, but a value for a particular parameter is not set in the `solid.ini` file, solidDB will use a factory value for that parameter. The factory values may depend on the operating system you are using.

By default, the driver looks for the `solid.ini` file in the current working directory, which is normally the directory from which you started the client. When searching for the file, the following precedence (from high to low) is used:

- location specified by the SOLIDDIR environment variable (if this environment variable is set)
- current working directory

Client-side parameters

This section describes the most important solidDB client-side parameters.

- **Com.Connect**

The **Connect** parameter in the [Com] section defines the default network name (connect string) for a client to connect to when it communicates with a server. Since the client should talk to the same network name as the server is listening to, the value of the **Com.Connect** parameter on the client should match the value of the **Com.Listen** parameter on the server.

- **Com.Trace** and **Com.TraceFile**

If you set the **Com.Trace** parameter default setting from No to **Yes**, solidDB starts logging trace information about network messages for the established network connection to the default trace file or to the file specified in the **Com.TraceFile** parameter.

1.1.5 ODBC non-standard behavior

This section describes the non-standard behavior and limitations of solidDB ODBC driver.

Error information

Regardless of the version set by the client, the driver returns error information based on the ODBC 3.0 specification.

Error in SQLPutData using SQL_NULL_DATA as parameter length

If you try to insert or update one or more data items where one of the items has SQL_NULL_DATA as the length specifier, no data will be inserted. The column value will become NULL.

SQLAllocHandle can return incomplete error information

If you call SQLAllocHandle with an invalid handle type, for example, `SQLAllocHandle(-5, hdbc, &hstmt);`

the function will return SQL_ERROR but not Error State "HY092" or message "Invalid Attribute/Option Identifier".

MSAccess - linking the table with certain column types

After linking the table with data types WCHAR, WVARCHAR, and LONG WVARCHAR, when a user inserts a particular record and then inserts/updates/deletes another record, the driver shows '#deleted' for the previous newly added/updated record.

ADO - OpenSchema methods

The following OpenSchema methods are not supported through ADO:

- adSchemaCatalogs
- adSchemaColumnPrivileges
- adSchemaConstraintColumnUsage
- adSchemaConstraintTableUsage
- adSchemaTableConstraint
- adSchemaForeignKeys
- adSchemaTablePrivileges
- adSchemaViews
- adSchemaViewTableUsage

The above mentioned OpenSchema methods are not supported by ADO with any ODBC Driver. This is a limitation of the Microsoft OLE DB Provider for ODBC. This is not specific to the solidDB ODBC Driver.

1.2 solidDB JDBC Driver

The JDBC 2.0 Driver provides support for JDBC 2.0.

For solidDB JDBC Driver, Java Development Kit (JDK) 1.4.2 or newer is supported.

solidDB JDBC Driver allows you to develop your application with a Java tool that accesses the database using JDBC. The JDBC API, the core API for JDK 1.2, defines Java classes to represent database connections, SQL statements, result sets, database metadata, and so on. It allows you to issue SQL statements and process the results. JDBC is the primary API for database access in Java.

In order to use JDBC, you have to install the solidDB JDBC Driver. Usage of JDBC drivers varies depending on your Java development environment.

Instructions and samples for using the solidDB JDBC Driver are located in the /jdbc subdirectory in the solidDB installation package and in 3, "Using solidDB JDBC Driver," on page 51.

1.3 solidDB Application Programming Interface (SA API)

solidDB SA is a C-language client library to connect solidDB database products. This library is used internally in solidDB products and provides access to data in solidDB database tables. The library contains 90 functions providing low-level mechanisms for connecting the database and running cursor-based operations.

1.4 solidDB Server Control API (SSC API)

The solidDB Server Control API (SSC API) is proprietary API that contains a set of functions that provide a simple and efficient means to control the tasking system of a solidDB.

The SSA API is provided with the shared memory access (SMA) and linked library access (LLA) libraries. For more details, see *solidDB Server Control API (SSC API)* in the *IBM solidDB Shared Memory Access and Linked Library Access User Guide*.

1.5 Building client applications

This section provides an overview of how to create a client application that will work with solidDB. The information in this section applies primarily to C-language programs that use the ODBC driver.

1.5.1 What is a client?

A *client application*, or *client* for short, is a program that submits requests (SQL queries) to the server and gets results back from the server.

A client program is separate from the server program. In many cases, the client is also running on a separate computer. Using shared memory access or linked library access, you can link the client's code directly to the server's code so that both run as a single process. For more information, see *IBM solidDB Shared Memory Access and Linked Library Access User Guide*.

Since the client is a separate program, it cannot directly call functions in the server. Instead, it must use a communications protocol (such as TCP/IP or named pipes) to communicate with the server. Different platforms support different protocols. On some platforms, you may need to link a specific library file (which supports a specific protocol) to your application so that your application can communicate with the server.

1.5.2 How is the query passed to the server?

Queries are written using the SQL programming language.

One way that the server and client can exchange data is simply to pass literal strings back and forth. The client could send the server the string:

```
SELECT name FROM employees WHERE id = 12;
```

and the server could send back the string:

```
"Smith, Jane".
```

In practice, however, communication is usually done via a "driver", such as an ODBC driver or a JDBC driver. "ODBC" stands for "Open DataBase Connectivity" and is an API (Application Programming Interface) designed by Microsoft to make database access more consistent across vendors. If your client program follows the ODBC conventions, then your client program will be able to talk with any database server that follows those same conventions. Most major database vendors support ODBC to at least some extent. The ODBC standard is generally used by programs written in the C programming language.

"JDBC" stands for "Java DataBase Connectivity". It is based heavily on the ODBC standard and is essentially "ODBC for Java programs".

There are two major ways to pass specific data values (for example, "Smith, Jane" to the server. The first way is to simply embed the values as literals in the query. This can be seen in SQL statements like:

```
INSERT INTO employees (id, name) VALUES (12, 'Smith, Jane');
```

This works well if you have a single statement that you want to execute. There are times, however, that you may want to execute the same basic statement with different values. For example, if you want to insert data for 500 employees, you may not want to compose 500 separate statements such as

```
INSERT INTO employees (id, name) VALUES (12, 'Smith, Jane');  
INSERT INTO employees (id, name) VALUES (13, 'Jones, Sally');  
...
```

Instead, you might prefer to compose a single "generic" statement and then pass specific values for that statement. For example, you might want to compose the following statement:

```
INSERT INTO employees (id, name) VALUES (?, ?);
```

and have the question marks replaced with specific data values. This way you can easily execute all 500 INSERT statements inside a loop without composing a unique INSERT statement for each employee. By using parameters, you can specify different values each time a statement executes. A parameter allows you to specify a variable that will be used by the client program and the ODBC driver to store values that the client and server exchange. In essence, you pass a parameter for each place in the statement where you have a question mark.

Another situation where you might want to use parameters to exchange data values is when working with data that is difficult to represent as string literals. For example, if you want to insert a digitized copy of the song "American Pie" into your database, and you do not want to compose an SQL statement with a literal that contains a series of hexadecimal numbers to represent that digitized data, then you can store the digitized data in an array and notify the ODBC driver of the location of that array.

To use parameters with SQL statements, you go through a multistep process. The following shows the process of inserting data. The process is somewhat similar when you want to retrieve data.

1. Prepare the SQL statement. During the *prepare* phase, the server analyzes the statement and (among other things) looks to see how many parameters there will be. The number and meaning of the parameters is shown by the question marks that are included in the SQL statement.

2. Tell the ODBC driver which variables will be used as parameters. Telling the ODBC driver which variable is associated with which column or value is called "binding" the parameters.
3. Put values into the parameters (that is, set the values of the variables).
4. Execute the prepared statement.

During the *execution* phase, the ODBC driver will read the values you have stored in the parameters and will pass those values to the server to use with the statement that it has already prepared.

1.5.3 How are the results passed back to the client?

The result of a query is a set of 0 or more rows. If you are using an ODBC driver or JDBC driver, you retrieve each row by using the appropriate ODBC or JDBC functions.

As a general rule, you go through the following steps

1. Prepare the SQL statement. During the *prepare* phase, the server analyzes the statement and (among other things) looks to see how many parameters there will be. The number and meaning of the parameters is shown by the question marks that are included in the SQL statement.
2. Tell the ODBC driver which variables will be used as parameters. Telling the ODBC driver which variable is associated with which column or value is called "binding" the parameters.
3. Execute the prepared statement. This tells the server to execute the query and collect the result set. However, the result set is not passed to the client immediately.
4. Fetch the next row of the result set. When you do a *fetch*, you tell the server and the ODBC driver to retrieve one row of results from the result set and then store the values of that row into the parameters that you previously defined for the ODBC driver to share with your application.

Normally you will perform a loop, fetching one row at a time and reading the data from the parameters after each fetch.

1.5.4 Statement cache

Processing of queries is additionally optimized by a built-in *statement cache*.

Statement cache is an internal memory for storing a small number of previously prepared SQL statements. The statement cache operates in such a way that the prepare phase is omitted if the prepared statement is in the cache. If a connection is closed, the statement cache is purged.

In ODBC, the number of cached statements for a session can be set by using a client-side `solid.ini` configuration parameter **Client.StatementCache**.

In JDBC, the statement cache size can be dynamically set by using a non-standard `StatementCache` connection property.

Related concepts

3.5.3, “Non-standard JDBC connection properties,” on page 64

The following connection properties can be used to attain connection-specific behavior.

Related information

H.2, “Client section,” on page 294

2 Using solidDB ODBC API

This section contains solidDB-specific information and usage samples for developing applications that use the ODBC API.

In general, solidDB conforms to the Microsoft ODBC 3.51 standard. solidDB ODBC APIs are defined based on the function prototypes provided by Microsoft. This guide details those areas where solidDB-specific usage applies and where support for options, data types, and functions differ.

Note: This *IBM solidDB Programmer Guide* does not contain a full ODBC API reference. For details on developing applications with ODBC API, refer to the Microsoft ODBC Programmer's Reference.

solidDB provides two versions of the ODBC driver, one for Unicode and one for ASCII. The Unicode version is a superset of the ASCII version; you may use it with either Unicode or ASCII character sets.

2.1 Installing solidDB ODBC Driver

The solidDB installation program installs two ODBC Drivers: one for Unicode and one for ASCII. The Unicode version is a superset of the ASCII version; you can use it with either Unicode or ASCII character sets. You can also use the solidDB installation program to install only the ODBC driver.

Windows

In Windows environments, the solidDB installation program installs the ODBC drivers and the following system Data Source Names (DSN) automatically. You can also add your own user DSNs.

- Windows 32-bit operating systems:
 - IBM solidDB 7.0 32-bit – ANSI
 - IBM solidDB 7.0 32-bit – Unicode
- Windows 64-bit operating systems:
 - IBM solidDB 7.0 64-bit – ANSI
 - IBM solidDB 7.0 64-bit – Unicode

Linux and UNIX

In Linux and UNIX environments, the ODBC driver library files are installed to the following directories:

- <solidDB installation directory>/bin/: dynamic library files
 - sac<platform><version>.sa or sac<platform><version>.so – ANSI
 - soc<platform><version>.sa or soc<platform><version>.so – Unicode
- <solidDB installation directory>/lib/: static library files
 - solidodbca.sa or solidodbca.so – ANSI
 - solidodbcu.sa or solidodbcu.so – Unicode

The file extension .sa or .so depends on the operating system.

Installing ODBC drivers without solidDB installation

To install the ODBC drivers without installing solidDB:

1. Start the solidDB installation program.
2. Select **Custom** installation.
3. Select **ODBC** (unselect **Server** and **Samples**).
4. Follow the displayed instructions to complete the installation.

2.2 Using the ODBC driver library

The ODBC driver libraries must be linked with your client application program.

Static vs. Dynamic Libraries

You will then be able to call the functions that are defined in these libraries. For details about library names, see the SDK Notes in the solidDB package.

Some library files are static — i.e. they are linked to your client application's executable program at the time that you do a compile-and-link operation. Other library files are dynamic - these are stored separately from your executable and are loaded into memory at the time your program executes.

The advantage of a static library is that your application is largely self-contained; if you distribute the application to your customers, those customers do not have to install a separate shared library in addition to installing your application.

The advantage of a dynamic library is that on many systems it requires less disk space (and, on some platforms, less memory space) if more than one client uses that library. For example, if you have two client applications that each link to a 5 MB static library, you will need not only 5 MB of disk space to store the static library, but also 10 MB of additional disk space to store both copies of the library that are linked into the application. However, if you link two client applications to a dynamic library, no additional copies of that library will be required; each application does not keep its own copy.

For many libraries, solidDB provides both a static and a dynamic version on some or all platforms.

In addition, on Windows environments, solidDB provides an import library in some cases. Each import library is associated with a corresponding dynamic link library. Your application will link to the import library. When the application is actually loaded and executed, the operating system will load the corresponding dynamic link library.

2.3 solidDB ODBC Driver 3.51 Features Support

This section provides details about the ODBC Driver 3.51 features support for users who have migrated from a previous version (1.0, 2.0, and 3.0) of the solidDB ODBC Driver to solidDB ODBC Driver 3.51.

The following features are supported in this driver:

- Complete support of descriptors
- All catalog API support
- Unicode support

- Multithread support
- ADO/DAO/RDO/OLE DB support
- Data access through MS Access and MS Query
- Block cursor support

2.4 Overview of usage on Windows operating systems

On Windows operating systems, the solidDB ODBC Libraries are provided as .DLL files.

The files are named socw32VV.dll and sacw32VV.dll (where "VV" indicates the version number) for the Unicode and ASCII versions, respectively. For example, the Unicode ODBC driver in version 4.1 is named socw3241.dll. To call the functions in one of these .DLL files, you must link to a solidDB import library file. For the solidDB on Windows, this import library file is named solidimpodbcu.lib (Unicode) or solidimpodbca.lib (ASCII). This import library file contains the entry points to the corresponding solidDB ODBC DLL (for example, socw3241.dll).

Note: The library files have been produced with C++. Other development toolkit manufacturers' linkers may expect different library file formats. In such cases, the Import Library utility of the development toolkit should be used to build a library file that is compatible with your linker.

Instructions for usage of solidDB client DLLs (solidDB ODBC Driver files)

There are two alternatives to building application programs that use the solidDB ODBC driver:

1. Using Microsoft ODBC Driver Manager.

Microsoft ODBC software needs to be installed on all client workstations and a Data Source must be defined using solidDB ODBC Driver. If you use the Driver Manager, then any application that can use the solidDB ODBC driver will also work with any other ODBC compliant engine.

2. Using solidDB ODBC driver directly.

Connections are opened directly to a server process without using Microsoft ODBC Driver Manager. This usually makes embedded deployment of solidDB easier. However, the application can only use the functions provided by the solidDB library (that is, solidDB ODBC Driver); the application cannot use the ODBC functions that are implemented by the Microsoft ODBC Driver Manager or the Microsoft Cursor library.

solidDB provides some sample programs that can be used either with or without the Microsoft ODBC Driver Manager. These samples are in subdirectories of the samples directory in your solidDB installation directory. Below are brief instructions on how to build and run the provided samples in both of the alternative ways:

- Building the samples to use ODBC Driver Manager.

1. Create a new application project.
2. Add the C-source file (for example, sqled.c or embed.c) to the project.
3. Make the header files visible to the compiler.
4. Define SS_WINDOWS for the compiler.
5. Compile and link.

6. Make sure that you have installed the solidDB ODBC driver. Also, make sure that the connection string you intend to use is defined as the ODBC data source name.
 7. Run to connect to a listening solidDB server.
- Building the samples to use solidDB ODBC library directly.
The necessary changes to the ODBC Driver Manager configuration are listed below.
 1. Add solidDB ODBC driver library file (`solidimpodbcu.lib`) to the project.
 2. Remove ODBC Driver manager libraries `ODBC*.LIB` from the default library list.
 3. Compile and link.
 4. Now it is possible to connect to data sources bypassing ODBC Driver Manager. Make sure that the SQL API DLL `socw32<VV>.dll` (where "VV" indicates the version number) and the solidDB communication DLLs are available. Data Sources may be defined in `solid.ini` or in the ODBC Administration Window.
 5. Run the client to connect to a listening solidDB server.

2.5 Calling functions

This section provides information about how programs call functions in the ODBC driver.

Header files and function prototypes

If your program calls functions in the ODBC driver, your program must include the ODBC header files. These files define the ODBC functions, and the data types and constants that are used with ODBC functions. The header files are not solidDB-specific; they are standard header files provided by Microsoft. The solidDB ODBC driver (like any ODBC driver) implements the functions that are specified in these header files.

ASCII and Unicode

ODBC drivers come in two "flavors": ASCII and Unicode. The ASCII driver supports only ASCII character sets. The Unicode driver supports both the Unicode and the ASCII character sets.

For details on driver, API, and SQL conformance levels, refer to section Introduction to ODBC in the *Microsoft ODBC Programmer's Reference*.

Using the ODBC Driver Manager

An application may link directly to the solidDB ODBC driver, or the application may link to an ODBC Driver Manager.

On Windows systems, the Driver Manager is required if applications that connect to solidDB use OLE DB or ADO APIs, or you use database tools that require the Driver Manager, such as Microsoft Access, FoxPro, or Crystal Reports. In most other cases, you may link directly to the ODBC driver instead of linking to the Driver Manager.

On Windows systems, Microsoft supplies the Driver Manager, and you link to the Driver Manager import library (`ODBC32.LIB`) to gain access to the Driver Manager.

On other platforms, you can link to another vendor's Driver Manager. For example, on Linux systems, you can use unixODBC.

For basic application steps that occur whenever an application calls an ODBC function and details on calling ODBC functions, refer to section Introduction to ODBC in the *Microsoft ODBC Programmer's Reference*.

Data types

Appendix E, "Data types," on page 223 provides information about SQL data types that are supported by solidDB. The header files from Microsoft provide information about C-language data types used by your client program. To transfer data between the application program and the database server, you must use appropriate types. For example, on most 32-bit platforms, the C-language "int" data type corresponds to the SQL data type "INT". The C-language "float" data type corresponds to the SQL "REAL" data type.

Scalar functions

Scalar functions return a value for each row. For example, the "absolute value" scalar function takes a numeric column as an argument and returns the absolute value of each value in the column. Scalar functions are invoked with the following ODBC escape sequence:

```
{fn scalar-function}
```

Note: The starting and ending characters are the curly bracket characters, not parentheses.

For a list of scalar functions and a more complete example of their usage, refer to Appendix F, "Scalar functions," on page 269.

solidDB native scalar functions

solidDB provides the following native scalar functions, which cannot be invoked using the ODBC escape sequence.

- CURRENT_CATALOG() - returns a WVARCHAR string that contains the current active catalog name. This name is the same as ODBC scalar function {fn DATABASE()}.
- LOGIN_CATALOG() - returns a WVARCHAR string that contains the login catalog for the connected user (currently the login catalog is the same as the system catalog).
- CURRENT_SCHEMA() - returns a WVARCHAR string that contains the current active schema name.

Function return codes

When an application calls a function, the driver executes the function and returns a predefined code. These return codes indicate success, warning, or failure status. The return codes are:

- SQL_SUCCESS
- SQL_SUCCESS_WITH_INFO
- SQL_NO_DATA_FOUND
- SQL_ERROR
- SQL_INVALID_HANDLE

- SQL_STILL_EXECUTING
- SQL_NEED_DATA

If the function returns SQL_SUCCESS_WITH_INFO or SQL_ERROR, the application can call SQLError to retrieve additional information about the error.

2.6 Connecting to a data source

A data source can be a database server, a flat file, or another source of data.

To access the data source, you need to define the solidDB server's *network name* which the application will use in a SQLConnect() call (ServerName). The network name may be given in one of the three following ways:

- Include the network name in the function call directly
- Include the network name in the function call using a logical data source name mapping
- Give the network name in the function call as an empty string

There are three connectivity types for defining the network name:

- Basic connectivity
- Transparent connectivity
- SMA connectivity

The following sections describe how to connect to a data source using *basic connectivity*.

For information about *Transparent Connectivity (TC Info)*, see the sections *Using the Transparent Connectivity* and *Syntax of the Transparent Connectivity Info* in the *IBM solidDB High Availability User Guide*.

For information about SMA connectivity, see section *Establishing local connections for SMA* in the *IBM solidDB Shared Memory Access and Linked Library Access User Guide*.

2.6.1 Network name and connect string syntax

The solidDB network name syntax depends on the connectivity type: basic connectivity or transparent connectivity (TC Info). Both connectivity types can be used with network-based connections or local shared memory access (SMA) connections.

- **Basic connectivity**

Basic connectivity is the most commonly used connectivity type where the connect string defines the connection between the application and the solidDB server.

- **Transparent connectivity (TC Info)**

Transparent connectivity (TC Info) is used in High Availability (HA) configurations for specifying a (single) connection between the application and the solidDB HotStandby servers.

Both basic connectivity and transparent connectivity can be used in shared memory access (SMA) setups.

Network name syntax

The syntax of the network name is the following:

```
<network_name>::=<basic_connectivity>|<transparent_connectivity>
```

where

- <basic_connectivity>::=[<encryption attribute>]
 <network_connect_string> | <sma_connect_string>

```
<encryption attribute>::=USE_ENCRYPTION=YES|NO
```

```
<network_connect_string>::=protocol_name  
                  [options] [server_name]  
                  [port_number]
```

```
<sma_connect_string>::=sma protocol_name port_number  
                  | pipe_name
```

For example

```
USE_ENCRYPTION=YES tcp localhost 1315
```

```
sma tcp 2315
```

For more information about SMA connections, see section *Establishing local connections for SMA* in the *IBM solidDB Shared Memory Access and Linked Library Access User Guide*.

- <solidDB_TC_Info>::= {[<failure_transparency_level_attribute>]
 [<preferred_access_attribute>] [<encryption_attribute>]
 <connect_target_list>} | <cluster_info>

For details on the attributes, see *Syntax of the Transparent Connectivity Info - ODBC* in the *IBM solidDB High Availability User Guide*

For example

```
TF=SESSION USE_ENCRYPTION=YES SERVERS=tcp 2315 tcp 1315
```

solidDB connect string (<network_connect_string>)

The most commonly used network-based solidDB connect string consists of a *communication protocol*, a possible set of *special options*, an optional *host computer name* and a *server name*.

By this combination, the client specifies the server it will establish a connection to. The communication protocol and the server name must match the ones that the server is using in its network listening name. In addition, most protocols need a specified host computer name if the client and server are running on different machines. All components of the client's network name are case insensitive.

The format of a connect string is the following:

```
protocol_name [options] [host_computer_name] server_name
```

where

- *options* can be any combination of the following:

Table 4. Connect string options

Option	Description	Protocol
-4	Specifies that client connects using IPv4 protocol only.	TCP/IP
-6	Specifies that client connects using IPv6 protocol only. In Windows environments, this option is mandatory if IPv6 protocol is used.	TCP/IP

Table 4. Connect string options (continued)

Option	Description	Protocol
-i <i>source_address</i>	Specifies an explicit connecting socket source address for cases where the system default source IP address binding does not meet application needs. <i>source_address</i> can be an IP address or a host name.	TCP/IP
-z	Enables data compression for this connection	All
-c <i>milliseconds</i>	Specifies the login timeout (the default is operating-system-specific). A login request fails after the specified time has elapsed.	TCP/IP
-r <i>milliseconds</i>	Specifies the connection (or read) timeout. A network request fails when no response is received during the time specified. The value 0 (default) sets the timeout to infinite (operating system default timeout applies).	TCP/IP
-o <i>filename</i>	Turns on the Network trace facility and defines the name of the trace output file See <i>Network trace facility</i> in the <i>IBM solidDB Administrator Guide</i> for details.	All
-p <i>level</i>	Pings the server at the given level (0-5). Clients can always use the solidDB Ping facility at level 1 (0 is no operation/default). Levels 2, 3, 4 or 5 may only be used if the server is set to use the Ping facility at least at the same level. See <i>Ping facility</i> in the <i>IBM solidDB Administrator Guide</i> for details.	All
-t	Turns on the Network trace facility See <i>Network trace facility</i> in the <i>IBM solidDB Administrator Guide</i> for details.	All

- *host_computer_name* is needed with TCP/IP and Named Pipes protocols, if the client and server are running on different machines.
- *server_name* depends on the communication protocol:
 - In TCP/IP protocol, *server_name* is a service port number, such as '2315'.
 - In other protocols, *server_name* is a name, such as 'soliddb' or 'chicago_office'.
 For details on the syntax in different communication protocols, see *Communication protocols* in the *IBM solidDB Administrator Guide*.

Note:

- The *protocol_name* and the *server_name* must match the ones that the server is using in its network listening name.
- If given at the connection time, the connect string must be enclosed in quotes.
- All components of the connect string are case insensitive.

The same format of the connect string applies to the **Com.Connect** parameter as well as to the connect string used by solidDB tools or ODBC applications.

Examples

```
[Com]
Connect=tcp -z -c1000 1315

[Com]
Connect=nmpipe host22 SOLID
so1sql "tcp localhost 1315"
so1sql "tcp 192.168.255.1 1315"
rc = SQLConnect(hdbc, "upipe SOLID", (SWORD)SQL_NTS, "dba", 3, "dba", 3);
rc = SQLDriverConnect(hdbc,
                      (SQLHWND)NULL,
                      (SQLCHAR*)"DSN=tcp localhost 1964;UID=dba;PWD=dba",
                      38,
```

```
out_string,  
255,  
&out_length,  
SQL_DRIVER_NOPROMPT);
```

2.6.2 Using logical data source names

If the data source name is not a valid solidDB connect string, the driver assumes it is a logical data source name.

The logical data source name can be mapped to a data source as a 'logical name' and 'connect string' (network name) pair in the following ways:

- **Using the [Data Sources] section in the client-side solid.ini file**

The syntax of the parameters is the following:

```
[Data Sources]  
logical_name = connect_string; Description
```

where Description can be used for comments on the purpose of the logical name

Example:

To map a logical name **My_application** to a database that you want to connect using TCP/IP, include the following lines in the solid.ini file:

```
[Data Sources]  
My_application = tcpip irix 1313; Sample data source
```

When an application calls the data source 'My_application', the solidDB client maps this to a call to 'tcpip irix 1313'.

- **In Windows environments, using the registry settings (ODBC Driver Manager)**

You can use the **Control Panel > Administrative Tools > Data Sources (ODBC)** dialog or the **Registry Editor (regedit)** to add mappings.

For details, see *Configuring the solidDB ODBC Data Source for Windows* in the *IBM solidDB Programmer Guide*.

Tip: The solidDB data management tools use the solidDB ODBC API. If you have defined an ODBC Data Source, you can use the logical name source name also when connecting to solidDB server with the solidDB tools.

For example, if you have created a data source named 'solid_1' with ServerName 'tcp 2525', you can connect to solidDB with solidDB SQL Editor (**solsql**) with the following command:

```
solsql solid_1
```

When connecting to the solidDB server, if the network name is not a valid connect string, the solidDB tools and clients assume it is a logical data source name. To find a mapping between the logical data source name and a valid connect string, the solidDB tools and clients check the client-side solid.ini file.

In Windows environments, if the solid.ini file is not found or the logical data source name is not defined in the [Data Sources] section, the data source settings made with the Windows registry settings are checked in the following order.

1. Look for the Data Source Name from the following registry path:

```
HKEY_CURRENT_USER\software\odbc\odbc.ini\DSN
```

2. Look for the Data Source Name from the following registry path

```
HKEY_LOCAL_MACHINE\software\odbc\odbc.ini\DSN
```

The check for the logical data source mappings might impact performance:

- If the file system is particularly slow, for example, because the working directory is mapped to a network drive, checking the existence of the `solid.ini` file can have a measurable performance impact.
 - In Windows environments, all logical data source mappings in the ODBC registry are checked. The time consumed for this operation is proportional to the amount of defined data sources.
 - With only few (1 to 5) data sources, the connection time will be approximately 5 ms.
 - With 1000 data sources, the connection time will be approximately 200 ms.
- However, if the `solid.ini` file contains the logical data source name mapping, the tools and clients do not try to access the ODBC registry for the mapping.

2.6.3 Empty data source name

When an application uses the ODBC API directly and calls `SQLConnect()` without specifying a `solidDB` server network name (by giving an empty string), it is read from the parameter `Connect` in the `[Com]` section of the client application's `solid.ini` file.

The `solid.ini` file must reside in the current working directory of the application or in a path specified by the `SOLIDDIR` environment variable.

The following connect line in the `solid.ini` of the application workstation will connect an application (client) using the TCP/IP protocol to a `solidDB` server running on a host computer named 'spiff' and listening with the name (port number in this case) '1313'.

```
[Com]
Connect = tcpip spiff 1313
```

If the `Connect` parameter is not found in the `solid.ini` configuration file, then the client uses the environment-dependent default instead. The defaults for the `Listen` and `Connect` parameters are selected so that the application (client) will always connect to a local `solidDB` server listening with a default network name. So local communication (inside one machine) does not necessarily need a configuration file for establishing a connection.

2.6.4 Configuring the `solidDB` ODBC Data Source for Windows

To configure an ODBC data source for Windows platforms, you need to perform the steps described in this section.

Before you begin

To be able to configure `solidDB` ODBC data sources, the `solidDB` ODBC Driver must be installed.

Procedure

1. Invoke **Data Sources (ODBC)** from **Control Panel > Administrative Tools**
2. Open the **User DSN** tab.
3. Click the **Add...** button.
4. Select the `solidDB` ODBC Driver (ANSI or UNICODE according to your database requirements).

5. Enter the Data Source configuration in the solidDB ODBC Driver Setup box as shown in the following example.

Note: The **NetworkName** entry should be compliant with the database server listen addresses defined in `solid.ini`. The network name follows the connection string format presented in 1.1.3, “Format of the solidDB connect string,” on page 3.




Figure 1. ODBC driver setup

2.6.5 Using solidDB ODBC Driver with unixODBC

unixODBC is an ODBC driver manager for UNIX type environments. Instead of linking an application directly with the solidDB ODBC driver, a unixODBC DriverManager can be used.

For detailed information on unixODBC and the unixODBC DriverManager, see <http://www.unixodbc.org/>.

Configuration files

The unixODBC DriverManager loads the correct data source driver according to the specifications in the following two configuration files:

- `odbc.ini` or `.odbc.ini`: specifies the logical name of the data source and the actual ODBC driver
The `odbc.ini` file defines the system-level settings that are available to all users. The `.odbc.ini` file defines user-level settings.
- `odbcinst.ini`: connects the logical driver name with its physical location in the file system.
The `odbcinst.ini` is a system-level file.

In addition to the files above, the solidDB ODBC Driver needs a client-side `solid.ini` configuration file where the logical data source name is connected with the a valid solidDB connect string.

Syntax of the `odbc.ini` configuration files

The `odbc.ini` or `.odbc.ini` file must include at least the following two items for each data source:

- Logical name of the data source inside brackets, for example `[my_solid]`
- Logical name of the actual ODBC driver to be used by using the syntax `Driver=<driver name>`, for example `Driver = solid_odbc`.

An additional description can be added by using the syntax `Description=My first Solid`

All additional information is ignored.

Syntax of the `odbcinst.ini` configuration file

The logical name and the physical location of the ODBC driver must be specified in the `odbcinst.ini` file as follows:

- `[<the logical name of the driver>]`, for example, `[solid_odbc]`
- `Driver = <absolute path to the driver>`, for example, `Driver = /home/jsmith/sac12x64.so`

Syntax of the client-side `solid.ini` configuration file

In the client-side `solid.ini` file, the logical data source name must be connected to a valid `solidDB` connect string (network name) as follows:

- `[Data Sources]`
- `<the logical data source name> = <connect_string>` , for example, `my_solid=tcp my_machine 1964`

Location of the configuration files

The system-level configuration files, `odbc.ini` and `odbcinst.ini` are located in a system level configuration directory, such as `/etc/`. For example:

```
/usr/local/etc/odbc.ini
/usr/local/etc/odbcinst.ini
```

User-level data sources are specified in `~/.odbc.ini`.

The client-side `solid.ini` file can be located either in the directory set by the `SOLIDDIR` environment variable or in the current working directory.

Linking the driver

To link to the `unixODBC` driver instead of the `solidDB` ODBC Driver:

1. Copy the `unixODBC` driver to the location of your choice.
2. Replace the `solidDB` ODBC Driver library file with the `unixODBC` library file.

For example:

Direct linking: `LD_FLAGS = $(SOLID_LIB)/linux/sac12x65.so`

`unixODBC` driver manager: `LD_FLAGS = $(SOLID_LIB)/linux/libodbc.so`

Examples of the configuration files

```
~/odbc.ini      :  
  
    [foo]  
    Description      = Testing Solid  
    Driver           = solid_driver_65  
  
~/odbcinst.ini  :  
  
    [solid_driver_65]  
    Description      = The newest ODBC driver  
    Driver           = /home/jsmith/Solid6.50.0009/bin/sac12x64.so  
  
$SOLIDDIR/solid.ini  :  
  
[Data Sources]  
solid1 = tcp 1964
```

2.6.6 Retrieving user login information

This section describes how the Driver Manager retrieves login information.

If the application calls `SQLDriverConnect()` and requests that the user be prompted for information, the Driver Manager displays a dialog box similar to the following example:




Figure 2. ODBC data source administrator

On request from the application, the driver retrieves login information by displaying a dialog box.

2.7 ODBC handle validation

You can control ODBC handle validation with the client-side **Com.ODBCHandleValidation** parameter or dynamically with the non-standard ODBC attribute `SQL_ATTR_HANDLE_VALIDATION`. For performance reasons, ODBC handle validation in solidDB is switched off by default.

For example, in Windows environments with ODBC driver manager, the driver manager performs the handle validation and the solidDB ODBC driver does not need to repeat the same validation procedures. Also, a carefully written ODBC application does not normally cause invalid handles to be used; in that case, the handle validation in the ODBC driver is not needed. In both cases, the applications may benefit from performance improvements when skipping the handle validation in the driver.

To switch ODBC handle validation on or off:

- Set the client-side **Com.ODBCHandleValidation** to Yes or No. Default is No.
[Com]
`ODBCHandleValidation=Yes`
or
- Set the non-standard environment attribute `SQL_ATTR_HANDLE_VALIDATION` to 1 (on) or 0 (off). Default is 0.
 - To switch handle validation on:
`SQLSetEnvAttr(henv, SQL_ATTR_HANDLE_VALIDATION, (SQLPOINTER)1, 0);`
 - To switch handle validation off:
`SQLSetEnvAttr(henv, SQL_ATTR_HANDLE_VALIDATION, (SQLPOINTER)0, 0);`

Important: The `SQL_ATTR_HANDLE_VALIDATION` attribute must be set after creating the environment handle but before any other handle is created. The `SQL_ATTR_HANDLE_VALIDATION` attribute is global; when set, it affects all the solidDB ODBC handles initiated by the application. This ensures consistency by preventing the application from allocating both validated and non-validated handles.

When the handle validation is switched on, any ODBC function may fail with the standard return value `SQL_INVALID_HANDLE`.

If handle validation is turned off and invalid handle is used by the application, the ODBC driver behavior is unpredictable and most likely causes the application to crash.

Related concepts

2.11, “solidDB Extensions for ODBC API,” on page 29

The following functions and connection attributes are solidDB-specific extensions to ODBC API.

Appendix H, “Client-side configuration parameters,” on page 293

The client-side configuration parameters are stored in the client-side `solid.ini` configuration file and are read when the client starts.

2.8 Executing transactions

This section provides information about how transactions are committed.

In *auto-commit* mode, each SQL statement is a complete transaction, which is automatically committed when the statement finishes executing. Refer to the important notes in the *Committing Read-Only Transactions* section on committing read-only SELECTs.

In *manual-commit* mode, a transaction consists of one or more statements. In manual-commit mode, when an application submits an SQL statement and no transaction is open, the driver implicitly begins a transaction. The transaction remains open until the application commits or rolls back the transaction with `SQLEndTran`.

Committing Read-Only Transactions

Important:

- When the isolation level is other than `READ COMMITTED`, even read-only statements (for example, `SELECT`) must be committed. Furthermore, the user must commit `SELECT` statements even if the server is in autocommit mode. Failure to commit statements can reduce performance or cause the server to run out of memory. This is explained in more detail below.
- If the isolation level is `READ COMMITTED`, read-only statements need not be committed. In that case, the explanation below does not apply.

Even a read-only statement must be committed. The reason for this is that `solidDB` saves the 'read-level' of each transaction and until that transaction commits, all subsequent transactions from other connections are also maintained in memory. (This behavior is part of the row versioning performed by the Bonsai Tree technology. See *solidDB Administration Guide* for more details about the Bonsai Tree.) If a transaction is not committed, the server will need more and more memory as other transactions accumulate; this will reduce performance, and eventually the server may run out of available memory. For more details, read the Performance Tuning chapter in *solidDB Administration Guide*.

SELECT and autocommit

Using autocommit mode does not ensure that `SELECT` statements are committed. The server cannot automatically commit `SELECT`s because `SELECT`s do not execute as a single statement. Each `SELECT` involves opening a cursor, fetching rows, and then closing the cursor.

There are two possible ways that the server could automatically commit when fetching multiple rows: the server could commit after the final fetch, or the server could commit after each individual fetch. Unfortunately, neither of these is practical, and therefore the server cannot commit the `SELECT` statement even in autocommit mode.

The server cannot automatically commit after the final fetch because the server does not know which fetch is the final fetch — the server does not know how many rows the user will fetch. (Until the user closes the cursor, the server does not know that the user is done fetching.)

It is not practical to commit after each individual fetch because each transaction should see the data as it was at the time that the transaction started, and therefore if each fetch is in a different transaction then the data can be from a different "snapshot" of the database. Putting each fetch in a different transaction would also

make REPEATABLE READ and SERIALIZABLE transaction isolation levels confusing or meaningless for the cursor, even though the cursor is for a single SELECT statement.

To commit the SELECT statement, the user may:

- Execute an explicit COMMIT WORK statement.
- Execute a statement to which autocommit does apply (i.e. a statement other than SELECT).
- If the cursor is the only open cursor, then the user may commit by explicitly closing the cursor (the server automatically commits when a cursor is closed and there are no other open cursors (and the server is in autocommit mode). This is part of why we recommend that you explicitly close every cursor as soon as you are done with it.

Note: To ensure that the data in the cursor is consistent and recent, the server actually does an automatic commit immediately prior to opening the cursor (if autocommit is on). The server then immediately starts a new transaction to contain the subsequent FETCH statement(s). This new transaction, like any other transaction, must be committed (or rolled back).

Summary

All statements must be committed, even if they are read-only statements, if an isolation level other than READ COMMITTED is used.

In most cases when you are doing SELECT statements in autocommit mode, you should explicitly close each cursor as soon as you are done with it and then explicitly COMMIT, even though you are in autocommit mode.

2.9 Retrieving information about the data source's catalog

This section describes functions (known as *catalog functions*) that return information about a data source's catalog.

- SQLTables returns the names of tables stored in a data source.
- SQLTablePrivileges returns the privileges associated with one or more tables.
- SQLColumns returns the names of columns in one or more tables.
- SQLColumnPrivileges returns the privileges associated with each column in a single table.
- SQLPrimaryKeys returns the names of columns that comprise the primary key of a single table.
- SQLForeignKeys returns the names of columns in a single table that are foreign keys. It also returns the names of columns in other tables that refer to the primary key of the specified table.
- SQLSpecialColumns returns information about the optimal set of columns that uniquely identify a row in a single table or the columns in that table that are automatically updated when any value in the row is updated by a transaction.
- SQLStatistics returns statistics about a single table and the indexes associated with that table.
- SQLProcedures returns the names of procedures stored in a data source.
- SQLProcedureColumns returns a list of the input and output parameters, as well as the names of columns in the resultset, for one or more procedures.

Each function returns the information as a resultset. An application retrieves these results by calling `SQLBindCol()` and `SQLFetch()`.

Executing Functions Asynchronously

Note: ODBC drivers in all solidDB products do not support asynchronous execution.

2.10 Using ODBC extensions to SQL

ODBC defines extensions to SQL, which are common to most database management systems.

For details on SQL extensions, refer to Escape Sequences in ODBC in the *Microsoft ODBC Programmer's Reference*.

Included in the ODBC extensions to SQL are:

- Procedures
- Hints

Details on solidDB usage for these extensions are described in the following sections.

2.10.1 Procedures

Stored procedures are procedural program code containing one or more SQL statements and program logic.

Stored procedures are stored in the database and executed with one call from the application or another stored procedure.

An application can call a procedure in place of an SQL statement. The escape clause ODBC uses for calling a procedure is:

```
{call procedure-name [(parameter)[,parameter]...]}
```

where *procedure-name* specifies the name of a procedure stored on the data source and *parameter* specifies a procedure parameter.

Note: solidDB does not support the optional "?>" included in the ODBC standard:

```
{[?]= call procedure-name [(parameter)[,parameter]...]}
```

A procedure can have zero or more parameters.

- For input and input/output parameters, *parameter* can be a literal or a parameter marker. Because some data sources do not accept literal parameter values, be sure that interoperable applications use parameter markers.
- For output parameters, *parameter* must be a parameter marker.

If a procedure call includes parameter markers, the application must bind each marker by calling `SQLBindParameter()` prior to calling the procedure.

Procedure calls do not require input and input/output parameters; however, the following rules apply:

- A procedure called with parentheses but with parameters omitted, such as `{call procedure_name()}` may cause the procedure to fail.

- A procedure called without parentheses, such as {call *procedure_name*}, returns no parameter values.
- Input parameters may be omitted. Omitted input or input/output parameters cause the driver to instruct the data source to use the default value of the parameter. As an option, a parameter's default value can be set using the value of the length/indicator buffer bound to the parameter to SQL_DEFAULT_PARAM.
- When a parameter is omitted, the comma delimiting it from other parameters must be present.
- Omitted input/output parameters or literal parameter values cause the driver to discard the output value.
- Omitted parameter markers for a procedure's return value cause the driver to discard the return value.
- If an application specifies a return value parameter for a procedure that does not return a value, the driver sets the value of the length/indicator buffer bound to the parameter to SQL_NULL_DATA.

To determine if a data source supports procedures, an application calls SQLGetInfo() with the SQL_PROCEDURES information type.

For more information about procedures, see *Stored procedures* in the *IBM solidDB SQL Guide*.

2.10.2 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.

Hints syntax

Hints are detected through a pseudo comment syntax as specified in SQL-92.

```
--(* 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]
```

For more information on the hints syntax, see *HINT* in the *IBM solidDB SQL Guide*.

Enabling and disabling hints

Hints are enabled and disabled using the `SQL.EnableHints` parameter. By default, hints are enabled (`SQL.EnableHints=yes`).

2.10.3 Additional ODBC extension functions

ODBC provides the following functions related to SQL statements.

Refer to the API reference in the *Microsoft ODBC Programmer's Reference* for more information about these functions.

Table 5. Additional ODBC Extension Functions

Function	Description
SQLDescribeParam	Retrieves information about prepared parameters.
SQLNumParams	Retrieves the number of parameters in an SQL statement.
SQLSetStmtAttr SQLSetConnectAttr SQLGetStmtAttr	These functions set or retrieve statement options, such as asynchronous processing, orientation for binding rowsets, maximum amount of variable length data to return, maximum number of resultset rows to return, and query timeout value. Note that <code>SQLSetConnectAttr</code> sets options for all statements in a connection.

2.11 solidDB Extensions for ODBC API

The following functions and connection attributes are solidDB-specific extensions to ODBC API.

Non-standard ODBC functions

Table 6. *solidDB-specific ODBC functions to ODBC API*

Function	Description
SQLFetchPrev	This function is the same as the ODBC function SQLFetch, but for fetching previous record.
SQLSetParamValue	This function sets the value of a parameter marker in the SQL statement specified in SQLPrepare. Parameter markers are numbered sequentially from left-to-right, starting with one, and may be set in any order.
SQLGetCol	This function is the same as the ODBC function SQLGetData.
SQLGetAnyData	This function is the same as the ODBC function SQLGetData.

Non-standard ODBC attributes

The following connection attributes are specific to solidDB.

Note: If the attribute is marked as OUT, it is a read-only attribute and cannot be set through the ODBC interface.

- SQL_ATTR_TF_LEVEL
OUT: integer (TF level: 0=NONE, 1=CONNECTION, 3=SESSION)
The failure transparency level.
- SQL_ATTR_TF_RECONNECT_TIMEOUT
IN/OUT: integer
The time in milliseconds the driver should wait until it tries to reconnect to the primary in case of switchover or failover.
- SQL_ATTR_TF_WAIT_TIMEOUT
IN/OUT: integer
The time in milliseconds the driver should wait for the server to switch state.
- SQL_ATTR_TC_PRIMARY
OUT: string, Primary server connection string
There is always a value indicating the current Primary server.
- SQL_ATTR_TC_SECONDARY
OUT: string, Secondary server connection string
The value indicates the assigned workload server if:
 1. PA=READ_MOSTLY, and
 2. the Secondary is the designated workload server.Otherwise, the returned string is empty.
- SQL_ATTR_TF_WAITING
OUT: string, Secondary server connection string.

The value indicates the assigned watchdog (waiting) connection. Waiting connection is used by ODBC driver internally to detect possible losses (crashes, unavailability) of the primary server faster. The string is empty if the connection is not a TC connection.

- `SQL_ATTR_PA_LEVEL`

OUT: integer (Preferred Access level: 0=WRITE_MOSTLY, 1=READ_MOSTLY)

The attribute indicates whether the load balancing is used or not.

- `SQL_ATTR_TC_WORKLOAD_CONNECTION`

OUT: string, server name of the workload connection

The current workload connection server; if queried before the Commit, the value indicates the server the transaction will be committed on. It may be queried as the statement attribute as well. In that case, it indicates the server the next statement will be executed on.

- `SQL_ATTR_LOGIN_TIMEOUT_MS`

IN/OUT: integer, login timeout in milliseconds

Note: There is also a standard attribute `SQL_ATTR_LOGIN_TIMEOUT` that can be used to set the timeout in seconds.

- `SQL_ATTR_CONNECTION_TIMEOUT_MS`

IN/OUT: integer, connection timeout in milliseconds

Note: There is also a standard attribute `SQL_ATTR_CONNECTION_TIMEOUT` that can be used to set the timeout in seconds.

- `SQL_ATTR_QUERY_TIMEOUT_MS`

IN/OUT: integer, query timeout in milliseconds

Note: There is also a standard attribute `SQL_ATTR_QUERY_TIMEOUT` that can be used to set the timeout in seconds.

- `SQL_ATTR_IDLE_TIMEOUT`

IN/OUT: integer, connection idle timeout in minutes

Indicates the connection specific idle timeout to be used by the server. If there is no activity on the connection for specified time period, the server automatically shuts down the connection, effectively throwing out the user.

Special semantics:

- -1 (default) - the connection timeout is equal to the server default

- 0 - no idle timeout, connection is never closed

This property value can be set only before executing `SQLConnect()`.

- `SQL_ATTR_HANDLE_VALIDATION` (environment handle attribute)

IN/OUT: integer, turns ODBC standard handle validation on (1) or off (0).

Default is 0.

This attribute is global; when set, it affects all the solidDB ODBC connections initiated by the application. This ensures consistency by preventing the application from allocating both validated and non-validated handles.

In certain systems, for example in Windows with ODBC driver manager involved, the driver manager performs the handle validation and the solidDB ODBC driver does not have to repeat the same validation procedures by itself. Also, a carefully written ODBC application normally does not cause invalid handles to be used; in that case, the handle validation in the ODBC driver is not needed. In both cases, the applications may benefit from performance improvements when skipping the handle validation in the driver. In the case the

handle validation is turned off, and invalid handle is used by the application, the ODBC driver behavior is unpredictable and, most likely, it causes the application to crash.

- `SQL_ATTR_SET_CONNECTION_DEAD`
IN/OUT: integer, should be set to 1 when needed
When this attribute is set on a connection, it causes the driver to abort the connection forcibly, without a disconnecting handshake with the server. After the attribute is set to 1, the connection becomes unusable.
- `SQL_ATTR_PASSTHROUGH_READ`
IN: string, SQL passthrough mode for read-type statements ("NONE", "CONDITIONAL", "FORCE")
- `SQL_ATTR_PASSTHROUGH_WRITE`
IN: string, SQL passthrough mode for write-type statements ("NONE", "CONDITIONAL", "FORCE")

2.12 Using cursors

The ODBC Driver uses a *cursor* concept to keep track of its position in the resultset, that is, in the data rows retrieved from the database. A cursor is used for tracking and indicating the current position, as the cursor on a computer screen indicates current position.

Each time an application calls `SQLFetch`, the driver moves the cursor to the next row and returns that row. An application can also call `SQLFetchScroll` or `SQLExtendedFetch` (ODBC 2.x), which fetches more than one row with a single fetch or call into the application buffer. This is known as "block cursor" support. Note that the actual number of rows fetched depends upon the rowset size specified by the application.

An application can call `SQLSetPos` to position a cursor within a fetched block of data using the `SQL_POSITION` option. This allows an application to refresh data in the rowset. `SQLSetPos` is also called to update data with the `SQL_UPDATE` option or delete data in the resultset with the `SQL_DELETE` option.

The cursor supported by the core ODBC functions only scrolls forward, one row at a time. (To re-retrieve a row of data that it has already retrieved from the resultset, the application must close the cursor by calling `SQLFreeStmt` with the `SQL_CLOSE` option, re-execute the `SELECT` statement, and fetch rows with `SQLFetch`, `SQLFetchScroll`, or `SQLExtendedFetch` (ODBC 2.x) until the target row is retrieved.) If you need the ability to scroll backward as well as forward, use block cursors.

2.12.1 Assigning storage for rowsets (binding)

In addition to binding individual rows of data, an application can call `SQLBindCol` to assign storage for a *rowset* (one or more rows of data). By default, rowsets are bound in column-wise fashion. They can also be bound in row-wise fashion.

To specify how many rows of data are in a rowset, an application calls `SQLSetStmtAttr` with the `SQL_ROWSET_SIZE` option.

Column-wise binding

To assign storage for column-wise bound results, an application performs the following steps for each column to be bound:

1. Allocates an array of data storage buffers. The array has as many elements as there are rows in the rowset.
2. Allocates an array of storage buffers to hold the number of bytes available to return for each data value. The array has as many elements as there are rows in the rowset.
3. Calls `SQLBindCol` and specifies the address of the data array, the size of one element of the data array, the address of the number-of-bytes array, and the type to which the data will be converted. When data is retrieved, the driver will use the array element size to determine where to store successive rows of data in the array.

Row-wise binding

To assign storage for row-wise bound results, an application performs the following steps:

1. Declares a structure that can hold a single row of retrieved data and the associated data lengths. (For each column to be bound, the structure contains one field to contain data and one field to contain the number of bytes of data available to return.)
2. Allocates an array of these structures. This array has as many elements as there are rows in the rowset.
3. Calls `SQLBindCol` for each column to be bound. In each call, the application specifies the address of the column's data field in the first array element, the size of the data field, the address of the column's number-of-bytes field in the first array element, and the type to which the data will be converted.
4. Calls `SQLSetStmtAttr` with the `SQL_BIND_TYPE` option and specifies the size of the structure. When the data is retrieved, the driver will use the structure size to determine where to store successive rows of data in the array.

2.12.2 Cursor support

Applications require different means to sense changes in the tables underlying a resultset. Various cursor models are designed to meet these needs, each of which requires different sensitivities to changes in the tables underlying the resultset.

For example, when balancing financial data, an accountant needs data that appears static; it is impossible to balance books when the data is continually changing. When selling concert tickets, a clerk needs up-to-the minute, or dynamic, data on which tickets are still available.

solidDB cursors which are set with `SQLSetStmtAttr` as "dynamic" closely resemble static cursors, with some dynamic behavior. solidDB dynamic cursor behavior is static in the sense that changes made to the resultset by other users are not visible to the user, as opposed to ODBC dynamic cursors in which changes are visible to the user.

In solidDB, as long as the cursor scrolls forward from block to block and never scrolls backward or the cursors move back and forth within the same block after an update is done, then the user gets the dynamic cursor behavior. This means that all changes are visible. Note, however that this behavior is affected by the solidDB `AUTOCOMMIT` mode setting. For details, read "Cursors and autocommit" on page 34. For an example of cursor behavior when using `SQLSetPos`, read "Cursors and positioned operations" on page 36.

Another characteristic of solidDB's cursor behavior is that transactions are able to view their own data changes (with some limitations), but cannot view the changes made by other transactions that overlap in time. (For more details about the limitations on users seeing their own data changes, refer to "Cursors and positioned operations" on page 36). For example, once Transaction_A starts, it will not see any changes made by any other transaction that did not commit work before Transaction_A started. The conditions in solidDB that cause a user's own changes to be invisible to that user are:

- In a SELECT statement when an ORDER BY clause or a GROUP BY clause is used, solidDB caches the resultset, which causes the user's own change to be invisible to the user.
- In applications written using ADO or OLE DB, solidDB cursors are more like dynamic ODBC cursors to enable functions such as a rowset update.

Specifying the cursor type

To specify the cursor type, an application calls `SQLSetStmtAttr` with the `SQL_CURSOR_TYPE` option. The application can specify a cursor that only scrolls forward, a static cursor, or a dynamic cursor.

Unless the cursor is a forward-only cursor, an application calls `SQLExtendedFetch` (ODBC 2.x) or `SQLFetchScroll` (ODBC 3.x) to scroll the cursor backwards or forwards.

Cursor support

This section describes the cursor type supported by solidDB.

Three types of cursors are defined in ODBC 3.51:

- Driver Manager supported cursors
- Server supported cursors
- Driver supported cursors

solidDB cursors are server supported cursors.

Cursors and autocommit

This section provides information about cursors and autocommit.

For solidDB-specific information about cursors and autocommit, read "Committing Read-Only Transactions" on page 25.

There are also some limitations in using the solidDB Autocommit mode if your application uses block cursors and positioned updates and deletes. For a brief description of these cursor features, read 2.12, "Using cursors," on page 32.

When using block cursors and positioned updates and deletes, you must:

- In the application, set commit mode to `SQL_AUTOCOMMIT_OFF`.
- Commit changes in the application only when all the fetch and positioned operations are done.
- In between positioned operations, be sure not to commit the changes.

Attention:

If the application uses commit mode as `SQL_AUTOCOMMIT_ON` or commits the changes before it is done with all the positioned operation, then the application may experience unpredictable behavior while browsing through the resultset. Read the section below for details.

Positioned Cursor Operations and `SQL_AUTOCOMMIT_ON`

The solidDB ODBC Driver keeps a row number/counter for every row in the rowset, which is the data rows retrieved from the database. When an application has the commit mode set to `SQL_AUTOCOMMIT_ON` and then executes a positioned update or a delete on a row in the rowset, the row is immediately updated in the database. Depending on the new value of the row, the row may be moved from its original position in the resultset. Since the updated row has now moved and its new position is unpredictable (since it is totally dependent on the new value), the driver loses the counter for this row.

In addition, the counter for all other rows in the rowset may also become invalid because of a change in position of the updated row. Hence the application may see incorrect behavior when it does the next fetch or `SQLSetPos` operation.

Following is an example that explains this limitation.

Assume an application performs the following steps:

1. Sets the commit mode to `SQL_AUTOCOMMIT_ON`.
2. Sets the rowset size to 5.
3. Executes a query to generate a resultset containing n rows.
4. Fetches the first rowset of 5 rows with `SQLFetchScroll`.

A sample resultset is shown below. In the sample, the resultset has only 1 column (defined as `varchar(32)`). In the table below, the first column shows the row number maintained by the driver internally. The second column shows the actual row values.

Table 7. A Sample Resultset

Row Counter Stored Internally by the Driver	Row Value
1	Antony
2	Ben
3	Charlie
4	David
5	Edgar

Assume now that the application calls `SQLSetPos` to update the third row with a new value of Gerard. To perform the update, the new row value is moved and positioned as shown below:

Table 8. A sample resultset

Row Counter Stored Internally by the Driver	Row Value
1	Antony
2	Ben
Empty row	
4	David
5	Edgar
New row	Gerard

Now the row counter for "David" becomes 3 and not 4, while the counter for "Edgar" becomes 4 and not 5. Since some row counters are now invalid, they will give wrong results when used by the driver to do relative or absolute positioning of the cursor.

If the commit mode had been set to `SQL_AUTOCOMMIT_OFF`, the database is not updated until the `SQLEndTran` function is called to commit the changes.

For solidDB-specific information about cursors and autocommit, read "Committing Read-Only Transactions" on page 25.

Cursors and positioned operations

When an application is performing positioned operations (such as updates and deletes when calling `SQLSetPos`), there are limitations in resultset visibility.

Case 1 illustrates cursor behavior when using `SQLSetPos`. In Case 1, the cursor scrolls back and forth within the same block after the update is applied.

Although Case 1 is intended to illustrate the visibility of updates in the resultset when using cursors, the exact circumstances under which visibility occurs depends on several factors. These include the size of the resultset relative to the size of the memory buffer, the transaction isolation level, and the frequency with which you commit data, and so on.

Case 2 shows how cursor behavior is limited using `SQLSetPos` when the cursor scrolls backward within a rowset or the cursors move back and forth within a different rowset after an update is applied.

Case 1

Following is an example that shows cursor behavior using positioned operations and shows how positioned updates can be visible to users.

Assume an application performs the following steps:

1. Sets the commit mode to `SQL_AUTOCOMMIT_OFF`.

This is a requirement described in "Cursors and autocommit" on page 34.

2. Sets the rowset size to 5.

3. Executes a query to generate a resultset of n rows.
4. Fetches the first rowset of 5 rows with SQLFetchScroll.

A sample resultset is shown below. In the sample, the resultset has only 1 column (defined as varchar(32)). In the table below, the first column shows the row number maintained by the driver internally. The second column shows the actual row values.

Table 9. A sample resultset

Row Counter Stored Internally by the Driver	Row Value
1	Antony
2	Ben
3	Charlie
4	David
5	Edgar

Assume now that the application calls SQLSetPos to update the third and fourth rows of the resultset with the names Caroline and Debbie. After the updates, the actual row values now contain Caroline and Debbie, as shown below:

Table 10. A Sample Resultset

Row Counter Stored Internally by the Driver	Row Value
1	Antony
2	Ben
3	Caroline
4	Debbie
5	Edgar

Note: In some cases, the resultset for a SELECT statement may be too large to fit in memory. As the user scrolls back and forth within the resultset, the ODBC Driver may discard some rows from memory and read in others. This can cause unexpected results: in some situations, updates to data in the cursor may seem to "disappear" and then "reappear" if the cursor re-reads (for example, from disk) the original values for a row that it previously modified.

Case 2

Case 2 shows the limitations when using positioned operations. The following example shows cursor behavior using positioned operations and shows when position updates are not visible to users.

Assume an application performs the following steps:

1. Sets the commit mode to `SQL_AUTOCOMMIT_OFF`.
This is a requirement explained in "Cursors and autocommit" on page 34.
2. Sets the rowset size to 5.
3. Executes a query to generate a resultset of n rows.
4. Fetches the first rowset of 5 rows with `SQLFetchScroll`.

A sample resultset is shown below. In the sample, the first two rowsets are shown. The resultset has only 1 column (defined as `varchar(32)`). In the table below, the first column shows the row number maintained by the driver internally. The second column shows the actual row values.

Table 11. A Sample Resultset

Row Counter Stored Internally by the Driver	Row Value
1	Antony
2	Ben
3	Charlie
4	David
5	Edgar
6	Fred
7	Gough
8	Harry
9	Ivor
10	John

Assume that after the first 4 steps above, the application calls `SQLSetPos` to perform the following tasks:

5. Updates the third row of the resultset.
6. Scrolls to the next rowset by calling `SQLFetchScroll`. This will get rows 6 to 10 and the cursor will be pointing to row 6.
7. Scrolls backward one rowset to get to the first rowset. This is done by calling `SQLScrollFetch` with the `FETCH_PRIOR` option.

After these tasks are performed, the value of the third row that was updated in step 5 still has the old value rather than the updated value as in "Case 1". The updated value is only visible in the Case 2 situation when the change is committed. But due to the unpredictable behavior when setting `SQL_AUTOCOMMIT_ON` as described in section "Positioned Cursor Operations and `SQL_AUTOCOMMIT_ON`" on page 35, commits cannot be done until all work related to block cursors and positioned operations is completed.

2.13 Using bookmarks

A bookmark is a 32-bit value that an application uses to return to a row. solidDB provides no support for bookmarks.

2.14 Error text format

Error messages returned by `SQLException` come from two sources: data sources and components in an ODBC connection. The error text must use a specific format depending of where the error is issued.

Typically, data sources do not directly support ODBC. Consequently, if a component in an ODBC connection receives an error message from a data source, it must identify the data source as the source of the error. It must also identify itself as the component that received the error.

If the source of an error is the component itself, the error message must explain this. Therefore, the error text returned by `SQLException` has two different formats: one for errors that occur in a data source and one for errors that occur in other components in an ODBC connection.

For errors that do not occur in a data source, the error text must use the format:

```
[vendor_identifier] [ODBC_component_identifier]  
component_supplied_text
```

For errors that occur in a data source, the error text must use the format:

```
[vendor_identifier] [ODBC_component_identifier]  
[data_source_identifier] data_source_supplied_text
```

The following table shows the meaning of each element.

Table 12. Errors in a Data Source

Element	Meaning
<i>vendor_identifier</i>	Identifies the vendor of the component in which the error occurred or that received the error directly from the data source.
<i>ODBC_component_identifier</i>	Identifies the component in which the error occurred or that received the error directly from the data source.
<i>data_source_identifier</i>	Identifies the data source. For single-tier drivers, this is typically a file format. For multiple-tier drivers, this is the DBMS product.
<i>component_supplied_text</i>	Generated by the ODBC component.
<i>data_source_supplied_text</i>	Generated by the data source.

Note: The brackets ([]) are included in the error text; they do not indicate optional items.

Sample error messages

The following examples show how various components in an ODBC connection might generate the text of error messages and how solidDB returns them to the application with `SQLError`.

Table 13. Sample Error Messages

SQLSTATE	Error Message
01000	General warning
01S00	Invalid connection string attribute
08001	Client unable to establish connection

SQLSTATE values are strings that contain five characters; the first two are a class value, followed by a three-character subclass value. For example 01000 has 01 as its class value and 000 as its subclass value. Note that a subclass value of 000 means there is no subclass for that SQLSTATE. Class and subclass values are defined in SQL-92.

Table 14. SQLSTATE values

Class value	Meaning
01	Indicates a warning and includes a return code of <code>SQL_SUCCESS_WITH_INFO</code> . Note: Error class 01 returns both warnings and errors.
01, 07, 08, 21, 22, 23, 24, 25, 28, 34, 3C, 3D, 3F, 40, 42, 44, HY	Indicates an error that includes a return value of <code>SQL_ERROR</code> . Note: Error class 01 returns both warnings and errors.
IM	Indicates warning and errors that are derived from ODBC.

Related reference

Appendix C, "SQLSTATE error codes," on page 193

This topic contains an error codes table that provides possible SQLSTATE values that a driver returns for the `SQLGetDiagRec` function.

2.14.1 Processing error messages

Applications provide users with all the error information available through `SQLError`: the ODBC SQLSTATE, the native error code, the error text, and the source of the error.

The application may parse the error text to separate the text from the information identifying the source of the error. It is the application's responsibility to take appropriate action based on the error or provide the user with a choice of actions.

The ODBC interface provides functions that terminate statements, transactions, and connections, and free statement, connection, and environment handles.

2.15 Terminating transactions and connections

The ODBC interface provides functions that terminate statements, transactions, and connections, and free statement (hstmt), connection (hdbc), and environment (henv) handles.

Terminating Statement Processing

To free resources associated with a statement handle, an application calls `SQLFreeStmt` with the following options:

- `SQL_CLOSE` - Closes the cursor, if one exists, and discards pending results. The application can use the statement handle again later. In ODBC 3.51, `SQLCloseCursor` can also be used.
- `SQL_UNBIND` - Frees all return buffers bound by `SQLBindCol` for the statement handle.
- `SQL_RESET_PARAMS` - Frees all parameter buffers requested by `SQLBindParameter` for the statement handle.

`SQLFreeHandle` is used to close the cursor if one exists, discard pending results, and free all resources associated with the statement handle.

Terminating transactions

An application calls `SQLEndTran` to commit or roll back the current transaction.

Terminating connections

To terminate a connection to a driver and data source, an application performs the following steps:

1. Calls `SQLDisconnect` to close the connection. The application can then use the handle to reconnect to the same data source or to a different data source.
2. Calls `SQLFreeHandle` to free the connection or environment handle and free all resources associated with the handle.

2.16 Constructing an application

This section provides two examples of C-language source code for applications: an example that uses static SQL functions to create a table, add data to it, and select the inserted data; and another example of interactive, ad-hoc query processing.

Microsoft provides two types of header files, one for ASCII data and the other for Unicode data. This example can use either of the Microsoft ODBC header files.

Static SQL example

The following example constructs SQL statements within the application.

```
/******  
Sample Name: Example1.c  
Author      : IBM  
  
Location   : CONSTRUCTING AN APPLICATION-  
              Programmer Guide  
Purpose    : Sample example that uses static SQL  
              functions to  
              create a table,  
              add data to it, and
```

```

        select the inserted data.

*****/
#if (defined(SS_UNIX) || defined(SS_LINUX))
#include <solidodbc3.h>
#else
#include <windows.h>
#endif

#include <stdio.h>
#include <test_assert.h>

#define MAX_NAME_LEN 50
#define MAX_STMT_LEN 100

/*****
Function Name: PrintError
Purpose.....: To Display the error associated with
              the handle
*****/
SQLINTEGER PrintError(SQLSMALLINT handleType, SQLHANDLE handle)
{
    SQLRETURN rc = SQL_ERROR;
    SQLWCHAR  sqlState[6];
    SQLWCHAR  eMsg[SQL_MAX_MESSAGE_LENGTH];
    SQLINTEGER nError;

    rc = SQLGetDiagRecW(handleType, handle, 1,
        (SQLWCHAR *)&sqlState, (SQLINTEGER *)&nError,
        (SQLWCHAR *)&eMsg, 255, NULL);
    if (rc == SQL_SUCCESS || rc == SQL_SUCCESS_WITH_INFO) {
        printf("\n\t Error:%1s\n", eMsg);
    }
    return(SQL_ERROR);
}

/*****
Function Name: DrawLine
Purpose      : To Draw a specified charcter (chr) for
              specified number of times (len)
*****/
void DrawLine(SQLINTEGER len, SQLCHAR chr)
{
    printf("\n");
    while(len > 0) {
        printf("%c", chr);
        len--;
    }
    printf("\n");
}

/*****
Function Name: example1
Purpose      : Connect to the specified data source and
              execute the set of SQL Statements
*****/
SQLINTEGER example1(SQLCHAR *server, SQLCHAR *uid, SQLCHAR *pwd)
{
    SQLHENV    henv;
    SQLHDBC    hdbc;
    SQLHSTMT   hstmt;
    SQLRETURN  rc;

```

```

SQLINTEGER id;
SQLWCHAR drop[MAX_STMT_LEN];
SQLCHAR name[MAX_NAME_LEN+1];
SQLWCHAR create[MAX_STMT_LEN];
SQLWCHAR insert[MAX_STMT_LEN];
SQLWCHAR select[MAX_STMT_LEN];
SQLINTEGER namelen;

/* Allocate environment and connection handles. */
/* Connect to the data source. */
/* Allocate a statement handle. */

rc = SQLAllocHandle(SQL_HANDLE_ENV, SQL_NULL_HANDLE,
    &henv);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

rc = SQLSetEnvAttr(henv, SQL_ATTR_ODBC_VERSION,
    (SQLPOINTER)SQL_OV_ODBC3, SQL_NTS);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

rc = SQLAllocHandle(SQL_HANDLE_DBC, henv, &hdbc);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

rc = SQLConnect(hdbc, server, SQL_NTS, uid, SQL_NTS,
    pwd, SQL_NTS);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

rc = SQLAllocHandle(SQL_HANDLE_STMT, hdbc, &hstmt);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* drop table 'nameid' if exists, else continue*/
wcscpy(drop, L"DROP TABLE NAMEID");
printf("\n%ls", drop);
DrawLine(wcslen(drop), '-');

rc = SQLExecDirectW(hstmt, drop, SQL_NTS);
if (rc == SQL_ERROR) {
    PrintError(SQL_HANDLE_STMT, hstmt);
}

/* commit work*/
rc = SQLEndTran(SQL_HANDLE_DBC, hdbc, SQL_COMMIT);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* create the table nameid(id integer,name varchar(50))*/
wcscpy(create,
    L"CREATE TABLE NAMEID(ID INT,NAME VARCHAR(50))");
printf("\n%ls", create);
DrawLine(wcslen(create), '-');

rc = SQLExecDirectW(hstmt, create, SQL_NTS);
if (rc == SQL_ERROR)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

/* commit work*/
rc = SQLEndTran(SQL_HANDLE_DBC, hdbc, SQL_COMMIT);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* insert data through parameters*/

```

```

wscpy(insert, L"INSERT INTO NAMEID VALUES(?,?)");
printf("\n%ls", insert);
DrawLine(wcslen(insert), '-');

rc = SQLPrepareW(hstmt, insert, SQL_NTS);
if (rc == SQL_ERROR)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

/* integer(id) data binding*/
rc = SQLBindParameter(hstmt, 1, SQL_PARAM_INPUT,
    SQL_C_LONG, SQL_INTEGER, 0, 0, &id, 0, NULL);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* char(name) data binding*/
rc = SQLBindParameter(hstmt, 2, SQL_PARAM_INPUT,
    SQL_C_CHAR, SQL_VARCHAR, 0, 0, &name,
    sizeof(name), NULL);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

id = 100;
strcpy(name, "SOLID");

rc = SQLExecute(hstmt);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* commit work*/
rc = SQLEndTran(SQL_HANDLE_DBC, hdbc, SQL_COMMIT);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* free the statement buffers*/
rc = SQLFreeStmt(hstmt, SQL_RESET_PARAMS);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

rc = SQLFreeStmt(hstmt, SQL_CLOSE);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

/* select data from the table nameid*/
wscpy(select, L"SELECT * FROM NAMEID");
printf("\n%ls", select);
DrawLine(wcslen(select), '-');

rc = SQLExecDirectW(hstmt, select, SQL_NTS);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* bind buffers for output data*/
id = 0;
strcpy(name, "");

rc = SQLBindCol(hstmt, 1, SQL_C_LONG, &id, 0, NULL);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

rc = SQLBindCol(hstmt, 2, SQL_C_CHAR, &name,
    sizeof(name), &namelen);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

rc = SQLFetch(hstmt);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)

```



```

        return(PrintError(SQL_HANDLE_DBC, hdbc));

printf("\n Data ID      :%d", id);
printf("\n Data Name    :%s(%d)\n", name, namelen);

rc = SQLFetch(hstmt);
assert(rc == SQL_NO_DATA);

/* free the statement buffers*/
rc = SQLFreeStmt(hstmt, SQL_UNBIND);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

rc = SQLFreeStmt(hstmt, SQL_CLOSE);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

/* Free the statement handle. */
rc = SQLFreeHandle(SQL_HANDLE_STMT, hstmt);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

/* Disconnect from the data source. */
rc = SQLDisconnect(hdbc);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* Free the connection handle. */
rc = SQLFreeHandle(SQL_HANDLE_DBC, hdbc);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* Free the environment handle. */
rc = SQLFreeHandle(SQL_HANDLE_ENV, henv);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

return(0);
}

/*****
Function Name: main
Purpose      : To Control all operations
*****/
void main(SQLINTEGER argc, SQLCHAR *argv[])
{
    puts("\n\t SOLID ODBC Driver 3.51:");
    puts("\n\t -Usage of static SQL functions");
    puts("\n\t =====");

    if (argc != 4){
        puts("USAGE: Example1 <DSN name> <username> <passwd>");
        exit(0);
    }
    else {
        example1(argv[1], argv[2], argv[3]);
    }
}

```

Interactive ad hoc query example

The following example illustrates how an application can determine the nature of the resultset prior to retrieving results.

```

/*****
Sample Name      : Example2.c(ad-hoc query processing)
Author           : IBM

Location        : CONSTRUCTING AN APPLICATION-
                  Programmer Guide
Purpose         : To illustrate how an application determines
                  the nature of the result set prior to
                  retrieving results.
*****/

#ifdef SS_UNIX || defined(SS_LINUX)
#include <solidodbc3.h>
#else
#include <windows.h>
#endif

#include <stdio.h>

#ifdef TRUE
#define TRUE 1
#endif

#define MAXCOLS 100
#define MAX_DATA_LEN 255

SQLHENV  henv;
SQLHDBC  hdbc;
SQLHSTMT hstmt;

/*****
Function Name: PrintError
Purpose      : To Display the error associated with
              the handle
*****/
SQLINTEGER PrintError(SQLSMALLINT handleType, SQLHANDLE handle)
{
    SQLRETURN rc = SQL_ERROR;
    SQLCHAR  sqlState[6];
    SQLCHAR  eMsg[SQL_MAX_MESSAGE_LENGTH];
    SQLINTEGER nError;

    rc = SQLGetDiagRec(handleType, handle, 1,
                      (SQLCHAR *)&sqlState, (SQLINTEGER *)&nError,
                      (SQLCHAR *)&eMsg, 255, NULL);
    if (rc == SQL_SUCCESS || rc == SQL_SUCCESS_WITH_INFO) {
        printf("\n\t Error:%s\n",eMsg);
    }
    return(SQL_ERROR);
}

/*****
Function Name: DrawLine
Purpose      : To Draw a specified character (line) for
              specified number of times (len)
*****/
void DrawLine(SQLINTEGER len, SQLCHAR line)
{
    printf("\n");
    while(len > 0) {
        printf("%c",line);
        len--;
    }
    printf("\n");
}

```

```

}

/*****
Function Name: example2
Purpose      : Connect to the specified data source and
              execute the given SQL statement.
*****/
SQLINTEGER example2(SQLCHAR *sqlstr)
{
    SQLINTEGER i;

    SQLCHAR colname[32];
    SQLSMALLINT coltype;
    SQLSMALLINT colnamelen;
    SQLSMALLINT nullable;
    SQLINTEGER collen[MAXCOLS];
    SQLSMALLINT scale;
    SQLINTEGER outlen[MAXCOLS];
    SQLCHAR data[MAXCOLS][MAX_DATA_LEN];
    SQLSMALLINT nresultcols;
    SQLINTEGER rowcount, nRowCount=0, lineLength=0;
    SQLRETURN rc;

    printf("\n%s",sqlstr);
    DrawLine(strlen(sqlstr),'=');

    /* Execute the SQL statement. */
    rc = SQLExecDirect(hstmt, sqlstr, SQL_NTS);
    if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
        return(PrintError(SQL_HANDLE_STMT, hstmt));

    /* See what kind of statement it was. If there are */
    /* no result columns, the statement is not a SELECT */
    /* statement. If the number of affected rows is */
    /* greater than 0, the statement was probably an */
    /* UPDATE, INSERT, or DELETE statement, so print */
    /* the number of affected rows. If the number of */
    /* affected rows is 0, the statement is probably a */
    /* DDL statement, so print that the operation was */
    /* successful and commit it. */

    rc = SQLNumResultCols(hstmt, &nresultcols);
    if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
        return(PrintError(SQL_HANDLE_STMT, hstmt));

    if (nresultcols == 0) {
        rc = SQLRowCount(hstmt, &rowcount);
        if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO) {
            return(PrintError(SQL_HANDLE_STMT, hstmt));
        }
        if (rowcount > 0) {
            printf("%ld rows affected.\n", rowcount);
        }
        else {
            printf("Operation successful.\n");
        }

        rc = SQLEndTran(SQL_HANDLE_DBC, hdbc, SQL_COMMIT);
        if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
            return(PrintError(SQL_HANDLE_DBC, hdbc));

    }

    /* Otherwise, display the column names of the result */
    /* set and use the display_size() function to */

```

```

/* compute the length needed by each data type. */
/* Next, bind the columns and specify all data will */
/* be converted to char. Finally, fetch and print */
/* each row, printing truncation messages as */
/* necessary. */
else {
    for (i = 0; i < nresultcols; i++) {
        rc = SQLDescribeCol(hstmt, i + 1, colname,
            (SQLSMALLINT)sizeof(colname),
            &colnamelen, &coltype, &collen[i],
            &scale, &nullable);
        if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO){
            return(PrintError(SQL_HANDLE_STMT, hstmt));
        }
        /* print column names */
        printf("%s\t", colname);
        rc = SQLBindCol(hstmt, i + 1, SQL_C_CHAR,
            data[i], sizeof(data[i]), &outlen[i]);
        if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO){
            return(PrintError(SQL_HANDLE_STMT, hstmt));
        }
        lineLength += 6 + strlen(colname);
    }

    DrawLine(lineLength-6, '-');

    while (TRUE) {
        rc = SQLFetch(hstmt);
        if (rc == SQL_SUCCESS || rc == SQL_SUCCESS_WITH_INFO){
            nRowCount++;
            for (i = 0; i < nresultcols; i++) {
                if (outlen[i] == SQL_NULL_DATA) {
                    strcpy((char *)data[i], "NULL");
                }
                printf("%s\t", data[i]);
            }
            printf("\n");
        }
        else {
            if (rc == SQL_ERROR)
                PrintError(SQL_HANDLE_STMT, hstmt);
            break;
        }
    }
    printf("\n\tTotal Rows:%d\n", nRowCount);
}

SQLFreeStmt(hstmt, SQL_UNBIND);
SQLFreeStmt(hstmt, SQL_CLOSE);
return(0);
}

/*****
Function Name: main
Purpose      : To Control all operations
*****/
int __cdecl main(SQLINTEGER argc, SQLCHAR *argv[])
{
    SQLRETURN rc;

    printf("\n\t SOLID ODBC Driver 3.51-Interactive");
    printf("\n\t ad-hoc Query Processing");
    printf("\n\t =====\n");
}

```

```

if (argc != 4) {
    puts("USAGE: Example2 <DSN name> <username> <passwd>");
    exit(0);
}

/* Allocate environment and connection handles. */
/* Connect to the data source. */
/* Allocate a statement handle. */
rc = SQLAllocHandle(SQL_HANDLE_ENV, SQL_NULL_HANDLE, &henv);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

rc = SQLSetEnvAttr(henv, SQL_ATTR_ODBC_VERSION,
    (SQLPOINTER)SQL_OV_ODBC3, SQL_NTS);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

rc = SQLAllocHandle(SQL_HANDLE_DBC, henv, &hdbc);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

printf("\n Connecting to %s\n ", argv[1]);
rc = SQLConnect(hdbc, argv[1], SQL_NTS, argv[2], SQL_NTS,
    argv[3], SQL_NTS);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

rc = SQLAllocHandle(SQL_HANDLE_STMT, hdbc, &hstmt);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* execute the following SQL statements */
example2("SELECT * FROM SYS_TABLES");
example2("DROP TABLE TEST_TAB");
example2("CREATE TABLE TEST_TAB(F1 INT, F2 VARCHAR)");
example2("INSERT INTO TEST_TAB VALUES(10, 'SOLID')");
example2("INSERT INTO TEST_TAB VALUES(20, 'MVP')");
example2("UPDATE TEST_TAB SET F2='UPDATED' WHERE F1 = 20");
example2("SELECT * FROM TEST_TAB");

/* Free the statement handle. */
rc = SQLFreeHandle(SQL_HANDLE_STMT, hstmt);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_STMT, hstmt));

/* Disconnect from the data source. */
rc = SQLDisconnect(hdbc);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* Free the connection handle. */
rc = SQLFreeHandle(SQL_HANDLE_DBC, hdbc);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_DBC, hdbc));

/* Free the environment handle. */
rc = SQLFreeHandle(SQL_HANDLE_ENV, henv);
if (rc != SQL_SUCCESS && rc != SQL_SUCCESS_WITH_INFO)
    return(PrintError(SQL_HANDLE_ENV, henv));

return(0);
}

```

2.17 Testing and debugging an application

The Microsoft ODBC SDK provides tools for application development.

The tools included are as follows:

- ODBC Test, an interactive utility that enables you to perform ad hoc and automated testing on drivers. A sample test DLL (the Quick Test) is included, which covers basic areas of ODBC driver conformance.
- ODBC Spy, a debugging tool with which you can capture data source information, emulate drivers, and emulate applications.
- Sample applications, including source code and makefiles:
 - A `#define`, `ODBCVER`, to specify which version of ODBC you want to compile your application with. To use the ODBC 3.51 constants and prototypes, add the following line to your application code before providing the include files.

```
#define ODBCVER 0X0352
```
 - For solidDB specific ODBC extensions, include the `solidodbc3.h` header file.

For additional information about the ODBC SDK tools, see the *Microsoft ODBC SDK Guide*.

3 Using solidDB JDBC Driver

The solidDB JDBC Driver 2.0 is a JDBC type 4 driver. *Type 4* means that it is a 100% Pure Java implementation of the Java Database Connectivity (JDBC) 2.0 standard.

The JDBC API defines Java classes to represent database connections, SQL statements, result sets, database metadata, and so on. It allows a Java programmer to issue SQL statements and process the results. JDBC is the primary API for database access in Java. More information on the JDBC technology can be found at the JDBC Technology Homepage (<http://java.sun.com/products/jdbc/>).

The solidDB JDBC Driver is written entirely in Java and it communicates directly with the solidDB server using the TCP/IP network protocol. The solidDB JDBC Driver does not require any additional database access libraries. The driver requires that a Java Runtime Environment (JRE) or Java Development Kit (JDK) is available.

3.1 What is solidDB JDBC Driver

This topic provides information about the solidDB JDBC driver.

The JDBC API defines Java classes to represent database connections, SQL statements, result sets, database metadata, and so on. It allows a Java programmer to issue SQL statements and process the results. JDBC is the primary API for database access in Java. More information about the JDBC technology can be found at the JDBC Technology Homepage (<http://java.sun.com/products/jdbc/>).

solidDB's JDBC driver is written entirely in Java and communicates directly with the solidDB server using the TCP/IP network protocol. The solidDB driver does not require any additional database access libraries, such as ODBC. The driver requires that a JRE (Java Runtime Environment) or JDK (Java Development Kit) is available.

The solidDB JDBC Driver is a solidDB implementation of the JDBC 2.0 standard. It is usable in all Java environments supporting JDK 1.4.2 and above.

3.2 Getting started with solidDB JDBC Driver

The solidDB JDBC Driver (`SolidDriver2.0.jar`) is installed during solidDB installation. You can verify the installation using a sample Java program that is provided in the solidDB package. Depending on your environment, you may need to set various configuration settings before using the solidDB JDBC Driver.

Default installation directory

The solidDB JDBC Driver is installed during solidDB installation into the `jdbc` directory in the solidDB installation directory.

The `jdbc` directory contains also the solidDB Data Store Helper Class (`SolidDataStoreHelper.jar`) for use with WebSphere®.

The `samples/jdbc` directory in the solidDB installation directory contains Java code samples that use the solidDB JDBC Driver. Instructions for running the sample are available in the `readme.txt` file, which is located in the same directory.

Requirements for Java environment

- Ensure that you have a working Java runtime or development environment that supports JDBC API specification release 2.0.
- Check from your Java environment documentation whether it can use compressed bytecode. The `SolidDriver2.0.jar` contains the solidDB JDBC Driver classes in compressed bytecode format usable by most Java Virtual Machines. However, some environments (such as Microsoft J++) require uncompressed bytecode. If your environment requires uncompressed bytecode, you must extract the `SolidDriver2.0.jar` file using a tool that supports long filenames.

Setting the CLASSPATH environmental variable

The CLASSPATH environment variable for your environment needs to include the solidDB JDBC Driver `.jar` file installation path.

• Windows

The installation adds the default solidDB JDBC Driver installation path to the System CLASSPATH environment variable automatically.

You can check and set the System CLASSPATH environment variable through the Control Panel:

Control Panel > System > Advanced > Environment Variables

• Linux and UNIX

Set your CLASSPATH environment variable to include the solidDB JDBC Driver (`SolidDriver2.0.jar`) installation path.

For example, in Bourne shell, use the following command:

```
export CLASSPATH=<solidDB installation directory>/jdbc/SolidDriver2.0.jar:$CLASSPATH
```

If you are using another shell than the Bourne shell, modify this command to make it appropriate for your shell.

Verifying the installation with `sample1.java`

The `sample1.java` sample application (available in the `samples/jdbc` directory) can be used for validating the solidDB JDBC Driver installation.

The `sample1.java` sample application performs the following actions:

1. Registers the solidDB JDBC Driver using JDBC Driver Manager services.
2. Asks for the connect string for a running solidDB process.
3. Connects to solidDB using the driver.
4. Creates the following statement for one query for retrieving data from one of the solidDB system tables:

```
SELECT TABLE_CATALOG, TABLE_SCHEMA, TABLE_NAME,  
TABLE_TYPE FROM TABLES
```
5. Executes the query.
6. Fetches all the rows of a result set.

An empty solidDB database dictionary contains approximately 86 rows.

Running the sample

Before running the sample, check the following:

- Your PATH environment variable contains the directories that hold the Java compiler and JRE.
 - The samples/jdbc directory contains a valid solidDB (evaluation) license.
1. If you do not already have a solidDB process running, start it now and create an empty database.
 2. Change your working directory to /samples/jdbc, which is the directory that contains the sample Java program.
 3. Compile the Java sample program:

```
javac sample1.java
```
 4. Start the sample application with the command:

```
java sample1
```
 5. The application prompts for a valid connect string. The connect string format is:

```
jdbc:solid://<hostname>:<port>/<username>/<password>
```

For example, the following string attempts to connect to a solidDB server at host 'mymachine' which listens to TCP/IP protocol at port 2315.

```
jdbc:solid://localhost:2315/dba/dba
```

After entering the connect string, the sample application outputs the query results.

Troubleshooting the sample1.java sample application

Possible problems in running the sample1.java sample application and solutions for them are listed below:

1. The driver cannot be successfully registered.
 - The Java environment does not support java.sql classes.
 - SolidDriver2.0.jar is not in the CLASSPATH definition.
2. Unable to connect to solidDB process
 - The version of the solidDB server should be 7.0 or later.
Older solidDB versions may refuse connections from the driver that is provided with the current release.
 - The connect string may be wrong or solidDB may not be listening to the specified TCP/IP.
Check that solidDB is running and verify the listening information. You can use, for example, solidDB SQL Editor (solsql) to ensure that a connection can be established through the network.

3.2.1 Registering solidDB JDBC Driver

The JDBC driver manager handles loading and unloading drivers and interfacing connection requests with the appropriate driver.

The driver can be registered as shown below. After execution of this code, the driver registers itself in the DriverManager.

```
// registration using Class.forName service  
Class.forName("solid.jdbc.SolidDriver");
```

3.2.2 Connecting to the database

Once the driver is successfully registered with the driver manager, a connection is established by creating an instance of java.sql.Connection.

The following code example creates an instance of `java.sql.Connection`.

```
Connection conn = null;
// sCon is the JDBC connection string
(jdbc:solid://hostname:port/login/password)
String sCon = "jdbc:solid://fb9:1314/dba/dba";
try {
    conn = DriverManager.getConnection(sCon);
} catch (SQLException e) {
    System.out.println("Connect failed : " + e.getMessage());
}
```

The parameter required by the `DriverManager.getConnection` function is the JDBC connection string. The JDBC connection string identifies which computer the database server is running on; the string also contains other information required to connect to the server.

The syntax of the JDBC URL (connection string) for solidDB is:

```
jdbc:solid://<hostname>:<port>/<username>/<password>[?<property-name>=<value>]...
```

For example, the following connect string attempts to connect to a solidDB server in machine fb9 listening to the tcp/ip protocol at port 1314:

```
"jdbc:solid://fb9:1314/dba/dba"
```

The application can establish multiple connections to the database by using multiple `Connection` objects. You should manage connection lifecycle in a very accurate way, otherwise there may be conflicts between concurrent users and applications trying to access the database. For details and instructions, see 3.7, "Code examples," on page 78.

Note: The solidDB JDBC Driver only supports a connection for administration options, with no queries allowed. For this type of connection, set the `java.util.Properties` name `ADMIN_USER` to `true`. After it is set to `true` and a connection is established, only `ADMIN` commands are allowed.

Transactions and autocommit mode

As the JDBC specification defines, a connection to the solidDB database can be in either autocommit or non-autocommit mode.

- When not in autocommit mode, each transaction needs to be explicitly committed before the modifications it made can be seen by other database connections.
- The autocommit state can be monitored by `Connection.getAutoCommit()` method.
- The state can be set by `Connection.setAutoCommit()`. An solidDB server's default setting for autocommit state is `true`.
- If autocommit mode is off, then the transactions can be committed in two ways.
 - calling the `Connection.commit()` method, or
 - executing a statement for SQL `'COMMIT WORK'`

Handling database errors

Database errors in JDBC are handled and managed by the exception mechanism. Most of the methods specified in JDBC interfaces may throw an instance of `SQLException`. As these errors may appear in the normal application workflow (representing concurrency conflicts, for instance) your code should be tolerant to

such errors. Basically, you must not leave your connections in any other state than "closed" regardless of the result of your code's execution. This approach allows avoiding situations where all available connections remain open due to unhandled exceptions.

You can get an exception's error code by calling `e.getErrorCode()`. For listings of solidDB error codes, see the Appendix *Error codes in IBM solidDB Administrator Guide*.

The following code example shows a correct way of handling errors coming from the database:

```
Public void listTablesExample() {
    try {
        Class.forName("solid.jdbc.SolidDriver");
    } catch (ClassNotFoundException e) {
        System.err.println("Solid JDBC driver is not registered
in the classpath");
        return; //exit from the method
    }
    Connection conn = null;
    Statement stmt = null;
    ResultSet rs = null;
    try {
        conn = DriverManager.getConnection("jdbc:solid://
localhost:1313", "dba", "dba");
        stmt = conn.createStatement();
        rs = stmt.executeQuery("SELECT * FROM tables");
        while (rs.next()) {
            System.out.println(rs.getObject(0)); //printing
            out results
        }
    } catch (SQLException e) {
        e.printStackTrace();
    } finally {
        /* It's a good idea to release
resources in a finally{} block
in reverse-order of their creation
if they are no-longer needed
*/
        if (rs != null) {
            try {
                rs.close();
            } catch (SQLException sqlEx) { // ignore
                rs = null;
            }
        }
        if (stmt != null) {
            try {
                stmt.close();
            } catch (SQLException sqlEx) { // ignore
                stmt = null;
            }
        }
    }
    if (conn != null)
        try {
            conn.close();
        } catch (SQLException e) {
            e.printStackTrace();
        } finally {
            conn = null;
        }
}
```

3.3 Special notes about solidDB and JDBC

JDBC does not specify what SQL dialect you can use; it simply passes the SQL on to the driver and lets the driver either pass it on directly to the database, or parse the SQL itself. Because of this, the solidDB JDBC Driver behavior is particular to solidDB.

In some functions, the JDBC specification leaves some details open. For the details particular to solidDB's implementation of the methods, check 3.4, "JDBC driver interfaces and methods".

The solidDB JDBC Driver provides support for catalogs and schemas in solidDB.

Executing stored procedures

In solidDB databases, stored procedures can be called by executing statements `CALL proc_name [(parameter ...)]` as in any other SQL statement. Procedures can also be used in JDBC in the same way, through a standard `CallableStatement` interface.

Note: solidDB stored procedures can return result sets. Calling procedures through the JDBC `CallableStatement` interface is not necessary. For an example of calling solidDB procedures using JDBC, see the source code for the Sample 3 application in 3.7, "Code examples," on page 78.

3.4 JDBC driver interfaces and methods

solidDB JDBC Driver 2.0 is compatible with the JDBC 2.0 standard, with support to selected features of JDBC 2.0 Optional Package (known before as Standard Extension).

This topic describes solidDB-specific differences from the standard API. You can browse standard packages and interfaces in the `java.sql` and `javax.sql` packages, and see details of a particular implementation by checking the list of "All Known Implementing Classes".

For a description of how different data types are supported by solidDB JDBC Driver, see 3.8, "solidDB JDBC Driver type conversion matrix," on page 89.

Array

The `java.sql.Array` interface is not supported. This interface is used to map SQL type Array in the Java programming language. It reflects the SQL-99 standard that is currently unavailable in solidDB.

Blob

The `java.sql.Blob` interface is not supported. This interface is used to map SQL type Blob in the Java programming language. It reflects the SQL-99 standard that is currently unavailable in solidDB.

CallableStatement

A `java.sql.CallableStatement` interface is intended to support calling database stored procedures. Thus, solidDB stored procedures are used in JDBC in the same way as any statement; the use of class `CallableStatement` is not necessary when

you are writing applications on a solidDB server only. However, for portability reasons, using `CallableStatement` is a wise choice.

Note: The JDBC Driver allows for the creation of a `Statement` object that generates `ResultSet` objects with the given type and concurrency. This differs from the `createStatement` method in JDBC 1.0 because it allows the default result set type and result set concurrency type to be overridden.

Differences with the standard API

Following are the differences from the standard `CallableStatement` interface defined in the JDBC API.

Table 15. Differences to the Standard CallableStatement Interface

Method name	Notes
<code>getArray(int i)</code>	Not supported by solidDB.
<code>getBlob(int i)</code>	Not supported by solidDB.
<code>getClob(int i)</code>	Not supported by solidDB.
<code>getDate(int parameterIndex, Calendar cal)</code>	Works as specified in Java API. Note: Uses a given <code>Calendar</code> object to specify time zone and locale, different from default ones. The same rule corresponds to other similar methods operating with <code>Calendar</code> instances.
<code>getObject (int i, Map map)</code>	Not supported by solidDB.
<code>getRef(int i)</code>	Not supported by solidDB.
<code>registerOutParameter(int parameterIndex, int sqlType, String typeName)</code>	Not supported by solidDB. This method throws an exception with the following message: "This method is not supported"

Clob

The `java.sql.Clob` interface is not supported. This interface is used to map SQL type `Clob` in the Java programming language. It reflects the SQL-99 standard that is currently unavailable in solidDB.

Connection

The `java.sql.Connection` interface is a public interface. It is used to establish a connection (session) with a specified database. SQL statements are executed and results are returned within the context of a connection.

Differences with the standard API

Following are the differences from the standard `Connection` interface defined in the JDBC API.

Table 16. Differences to the Standard Connection Interface

Method name	Notes
getTypeMap()	solidDB provides this method, but it always returns null.
isReadOnly()	solidDB only supports read-only connections and read-only transactions if the database is declared as read-only. This method always returns false.
nativeSQL(String sql)	Works as specified in Java API. solidDB JDBC Driver does not change the SQL passed to the solidDB server. The SQL query the user passes is returned.
prepareCall(String sql)	Works as specified in Java API. Note: Note that the escape call syntax is not supported.
setReadOnly(boolean readOnly)	solidDB only supports read-only database and read-only transactions if the database is declared as read-only. This method exists but does not affect the connection behavior.
setTransactionIsolation(int level)	Works as specified in Java API.
setTypeMap(Map map)	Not supported by solidDB.

DatabaseMetaData

The `java.sql.DatabaseMetaData` interface is a public abstract interface. It provides general, comprehensive information about the database.

All methods for this interface are supported by solidDB.

For a description of how different data types are supported by solidDB JDBC Driver, see 3.8, “solidDB JDBC Driver type conversion matrix,” on page 89.

Driver

The `java.sql.Driver` interface is a public abstract interface. Every driver class implements this interface and all methods for this interface are supported by solidDB.

PreparedStatement

The `java.sql.PreparedStatement` interface is a public abstract interface. It extends the statement interface. It provides an object that represents a precompiled SQL statement.

Note: The JDBC Driver allows for the creation of a `PreparedStatement` object that generates `ResultSet` objects with the given type and concurrency. This differs from

the `prepareStatement` method in JDBC 1.0 because it allows the default result set type and result set concurrency type to be overridden.

Subinterfaces: `CallableStatement`

Differences with the standard API

Following are the differences from the standard `PreparedStatement` interface defined in the JDBC API.

Table 17. Differences to the Standard PreparedStatement Interface

Method name	Notes
<code>setArray(int i, Array x)</code>	Not supported by solidDB.
<code>setBlob(int I, Blob x)</code>	Not supported by solidDB.
<code>setClob(int I, Clob x)</code>	Not supported by solidDB.
<code>setObject(int parameterIndex, Object x)</code>	Works as specified in Java API. Note: The following objects are not supported by solidDB: BLOB, CLOB, ARRAY, REF, and object (USING <code>java.util.Map</code>).
<code>setObject(int parameterIndex, Object x, int targetSqlType)</code>	Not supported by solidDB. This method throws an exception with the following message: "This method is not supported"
<code>setObject(int parameterIndex, Object x, int targetSQLType, int scale)</code>	Not supported by solidDB. This method throws an exception with the following message: "This method is not supported"
<code>setRef(int I, Ref x)</code>	Not supported by solidDB.

Ref

The `java.sql.Ref` interface is a public abstract interface.

This interface is a reference to an SQL structured type value in the database. This interface is not supported by solidDB.

ResultSet

The `java.sql.ResultSet` interface is a table of data that represents a database result set from a query statement. This object includes a cursor that points to its current row of data. The cursor's initial position is before the first row. It is moved to the next row by the next method. When there are no more rows left in the result set, the method returns `false`; this allows the use of a `WHILE` loop to iterate through the result set.

Differences with the standard API

Following are the differences from the standard ResultSet interface defined in the JDBC API.

Table 18. Differences to the Standard ResultSet Interface

Method name	Notes
getArray(int i)	Not supported by solidDB.
getArray(String ColName)	Not supported by solidDB.
getBigDecimal(String columnName)	Works as specified in Java API.
getCharacterStream(int columnIndex)	Works as specified in Java API. NOTE: The JDBC Driver sets the designated parameter to the given Reader object at the given character length. When a large UNICODE value is input to a LONG VARCHAR/LONG WVARCHAR parameter, for convenience, you can send it via a java.io.Reader. The JDBC Driver reads the data from the stream as needed, until it reaches end-of-file. The driver does all the necessary conversion from UNICODE to the database CHAR format.
getCharacterStream(String columnName)	Works as specified in Java API. The note above also applies to this method.
getFetchSize()	Not supported by solidDB.
getObject(int columnIndex)	Works as specified in Java API. NOTE: following objects are not supported by solidDB: BLOB, CLOB, ARRAY, REF, and object (USING java.util.Map).
getObject(int i, Map map)	Not supported by solidDB.
getObject(String colName, Map map)	Not supported by solidDB. This method throws an exception with the following message: "This method is not supported"
getRef(int i)	Not supported by solidDB.
getRef(String colName)	Not supported by solidDB.
refreshRow()	Not supported by solidDB.
setFetchSize(int rows)	No operation in solidDB. Sets the value for the number of rows to be fetched from the database each time. The value a user sets with this method is ignored.

ResultSetMetaData

The `java.sql.ResultSetMetaData` interface is a public abstract interface. This interface is used to find out about the types and properties of the columns in a `ResultSet`.

SQLData

The `java.sql.SQLData` interface is not supported. This interface is used to custom map SQL user-defined types. It reflects the SQL-99 standard that is currently unavailable in solidDB.

SQLInput

The `java.sql.SQLInput` interface is not supported. This interface is an input stream that represents an instance of an SQL structured or distinct type. It reflects the SQL-99 standard that is currently unavailable in solidDB.

SQLOutput

The `java.sql.SQLOutput` interface is not supported. This interface is an output stream used to write the attributes of a user-defined type back to the database. It reflects the SQL-99 standard that is currently unavailable in solidDB.

Statement

The `java.sql.Statement` interface is a public abstract interface. It is the object used to execute a static SQL statement and obtain the results of the execution.

Note: The JDBC Driver allows for the creation of a `Statement` object that generates `ResultSet` objects with the given type and concurrency. This differs from the `CreateStatement` method in JDBC 1.0 because it allows the default result set type and result set concurrency type to be overridden.

Subinterfaces:

- `CallableStatement`
- `PreparedStatement`

Differences with the standard API

Following are the differences from the standard `Statement` interface defined in the JDBC API.

Table 19. Differences to the Standard Statement Interface

Method name	Notes
<code>getFetchSize()</code>	No operation in solidDB.
<code>getMaxFieldSize()</code>	Maxfield size does not affect the solidDB server's behavior.
<code>getMoreResults()</code>	solidDB does not support multiple result sets.

Table 19. Differences to the Standard Statement Interface (continued)

Method name	Notes
getResultSetType()	Not supported by solidDB.
setFetchSize(int rows)	No operation in solidDB. Sets the value for the number of rows to be fetched from the database each time. The value a user sets with this method is ignored.
setMaxFieldSize(int max)	Maxfield size does not affect the solidDB server's behavior.

Struct

The java.sql.Struct interface is not supported. This interface represents the standard mapping in the Java programming language for an SQL structured type. It reflects the SQL-99 standard that is currently unavailable in solidDB.

ResultSet (updateable)

The java.sql.ResultSet interface contains methods for producing ResultSet objects that are updateable. A result set is updateable if its concurrency type is CONCUR_UPDATABLE. Rows in the result set may be updated, deleted, or new rows inserted using methods update xxx, where xxx refers to the datatype and methods updateRow and deleteRow.

Differences with the standard API

Following are the differences from the standard ResultSet interface defined in the JDBC API.

Table 20. Differences to the Standard ResultSet Interface

Method name	Notes
getRef(int i)	This method is not supported.
getRef(String colName)	This method is not supported.
refreshRow()	This method is not supported.
rowDeleted()	This method is not supported.
rowInserted()	This method is not supported.
setFetchSize(int rows)	This method is not supported.

3.5 solidDB JDBC Driver extensions

This section describes the non-standard extensions included in the solidDB JDBC Driver.

3.5.1 WebSphere compatibility

The solidDB JDBC Driver includes features that improve WebSphere compatibility.

Java Transaction API (JTA) support

solidDB server can participate in distributed transactions using the Java Transaction API (JTA) interface. The following interfaces are supported, as described in the *Java Transaction API Specification 1.1*:

- XAResource Interface (javax.transaction.xa.XAResource)
- Xid Interface (javax.transaction.xa.Xid)

When using JTA with SQL passthrough in solidDB Universal Cache, only read statements (SELECT) are supported.

solidDB Data Store Helper Class in WebSphere

WebSphere needs an adapter class for those JDBC data sources that are to be used within WebSphere. The base class for these adapters is the `com.ibm.websphere.rsadapter.GenericDataStoreHelper` class; solidDB implements its own version of this adapter inside a class called: `com.ibm.websphere.rsadapter.SolidDataStoreHelper`.

This class is provided within the solidDB product as a separate archive file called `SolidDataStoreHelper.jar`. You can find this file in the `jdbc` directory in the solidDB installation directory.

When you are configuring a new solidDB data source in WebSphere, you need to

- give the class `com.ibm.websphere.rsadapter.SolidDataStoreHelper` in the data store helper field of the configuration, and
- specify the full path to the `SolidDataStoreHelper.jar` file in the data source configuration of WebSphere.

See the WebSphere documentation for further details how to define new data sources in WebSphere.

For an example of how to install a solidDB's WebSphere sample application in the Websphere Studio Application Developer's workspace, see the `samples/websphere` directory in your solidDB installation directory.

solidDB Data Source Properties and WebSphere

You need to define the following properties when configuring a new data source in the WebSphere:

URL

- type: `java.lang.String`
- value should use syntax similar to the following syntax: `'jdbc:solid://<hostname>:<port>'`

user

- type: `java.lang.String`
- value should be a valid user name

password

- type: java.lang.String
- value should be a valid password

3.5.2 Connection timeout in JDBC

This topic describes features available in solidDB to set a connection timeout.

Connection timeout means response timeout of any JDBC call invoking data transmission over a connection socket. If the response message is not received within the time specified, an I/O exception is thrown. The JDBC standard (2.0/3.0) does not support setting of the connection timeout. solidDB has introduced two ways for doing that: one using a non-standard driver manager extension method and the other one using the property mechanisms. The time unit in either case is one millisecond.

Driver Manager Method get/setConnectionTimeout()

The following example illustrates the solution. The effect of the setting is immediate. This allows to set the timeout to zero if you want to force-disconnect.

```
//Import Solid JDBC:
import solid.jdbc.*;

//Define the connection:
solid.jdbc.SolidConnection conn = null;

//Cast to SolidConnection in order to use Solid-specific methods:
conn = (SolidConnection)java.sql.DriverManager.getConnection(sCon);

//Set connection timeout in milliseconds:
conn.setConnectionTimeout(3000);
```

3.5.3 Non-standard JDBC connection properties

The following connection properties can be used to attain connection-specific behavior.

- “Statement cache property”
- “Timeout properties” on page 65
- “Appinfo property” on page 66
- “Transparent connectivity (TC) properties” on page 66
- “Shared memory access (SMA) connection property” on page 67
- “SQL passthrough properties” on page 67
- “Catalog and schema name properties” on page 67
- “Setting connection properties with the URL string” on page 67

Statement cache property

solidDB JDBC driver introduces a property for setting the value of the connection's statement cache.

The name of the property is "StatementCache" and the default size of the cache is 8. The valid value range is 1 to 512 (inclusive). If the value exceeds the range, the driver silently forces the value either to 1 or 512.

Below is an example on how to use the property.

```

// create a Solid JDBC driver instance
Class.forName("solid.jdbc.SolidDriver");

// create a new Properties instance and insert a value for
// StatementCache property
java.util.Properties props = new java.util.Properties();
props.put("StatementCache","32");

// define the connection string to be used
String sCon="jdbc:solid://localhost:1315/uname1/pwd1";

// get the Connection object with a statement cache of 32
java.sql.Connection conn = java.sql.DriverManager.getConnection(sCon, props);

```

Timeout properties

The timeouts listed in the sections below may be set as connection properties.

Connection timeout property (solid_connection_timeout_ms)

Using the property "solid_connection_timeout_ms", you can set the connection timeout value (in milliseconds). The property must be set before getting a new connection. Once a connection object is created, changing the property value has no effect.

Login timeout property (solid_login_timeout_ms)

Using the property "solid_login_timeout_ms", you can set the timeout (in milliseconds) for opening of a connection .

Note: You can use also the method `DriverManger.setLoginTimeout(seconds)` to set the login timeout. This is a standard-compliant method.

Idle Timeout Property (solid_idle_timeout_min)

Using the property "solid_idle_timeout_min", you can set the timeout (in minutes) that fires when a connection has been idle for a longer time than specified (in minutes).

If the property value is not set, the server-side setting of the configuration parameter **ConnectTimeOut** applies. The factory value is 480 minutes. 0 means 'infinite timeout' (never expires).

Example

The following example shows how to set a connect timeout using the "solid_connection_timeout_ms" property.

```

// Set connection timeout with "solid_connection_timeout_ms" property //
public class Test {

    public static void main( String args[] ){

        // create property object
        Properties props = new Properties();

        // put username and password in the properties
        props.put("user", "MYUSERNAME");
        props.put("password", "MYPASSWORD");

        //

```

```

// Put connection timeout in the property object
//
props.put("solid_connection_timeout_ms", "10000");

try {

    // create driver
    Driver d = (Driver)(
        Class.forName("solid.jdbc.SolidDriver").newInstance());

    // get connection with url and property info
    Connection c = DriverManager.getConnection(
        "jdbc:solid://localhost:1313", props );

    // close connection
    c.close();

} catch ( Exception e ) {
    ; // save the day
}
}
}

```

Appinfo property

The connection attribute called Appinfo can be used to uniquely identify applications running in the same computer and under the same username, for the purposes of tracing and management. It can be retrieved, on the server side, with the command ADMIN COMMAND 'userlist'. By default, the value (a string) is not set. It is possible to set the value with the connection property "solid_appinfo".

Note: In ODBC applications, the value of Appinfo is passed by way of the environmental variable SOLAPPINFO.

Transparent connectivity (TC) properties

Transparent connectivity (TC) is a connection mode that can be used with solidDB HA solution (HotStandby). In JDBC, the TC mode is enacted with the following connection properties:

- "solid_tf_level"

Sets the transparent failover level to "CONNECTION or "SESSION", or "NONE". The default is "NONE".
- "solid_preferred_access"

Sets the preferred access mode to "WRITE_MOSTLY" or "READ_MOSTLY". The value "READ_MOSTLY" enacts automatic load balancing of read-only transactions between the Primary and Secondary servers. The default is "WRITE_MOSTLY" that corresponds to a normal HotStandby operation whereby all the load is transferred to the Primary server.
- "solid_tf1_reconnect_timeout"

Sets the timeout for how long the driver should wait until it tries to reconnect to the primary in case of TF switchover or failover. The unit is millisecond. The default is 10000 (10 seconds).
- "solid_tf_wait_timeout"

Sets the timeout for how long the driver should wait for the server to switch state. The unit is millisecond. The default is 10000 (10 seconds).

For more information about using the TC connection properties, see section *Using the Transparent Connectivity* in the *IBM solidDB High Availability User Guide*.

Shared memory access (SMA) connection property

A local (non RPC-based) JDBC connection to a SMA server can be made by setting the following connection property to "yes". In addition, you also need to define that you are using a local server at a given port.

- "solid_shared_memory"
The only valid value is "yes".

For example:

```
Properties props = new Properties();
// enable the direct access property
props.put("solid_shared_memory", "yes");
// get connection
Connection c = DriverManager.getConnection
("jdbc:solid://localhost:1315", props);
```

For more details, see *Making JDBC connections for SMA* in the *IBM solidDB Shared Memory Access and Linked Library Access User Guide*.

SQL passthrough properties

The SQL passthrough mode can be set with the following connection properties:

- "solid_passthrough_read"
Sets the SQL passthrough mode for read-type statements to "NONE", "CONDITIONAL", or "FORCE".
- "solid_passthrough_write"
Sets the SQL passthrough mode for write-type statements to "NONE", "CONDITIONAL", or "FORCE".

For more information about the SQL passthrough and the passthrough modes, see section *Setting SQL passthrough mode* in the *IBM solidDB Universal Cache User Guide*.

Catalog and schema name properties

The catalog name and schema name can be set with the following connection properties:

- "solid_catalog"
Sets the solidDB catalog name.
- "solid_schema"
Sets the solidDB schema name.

Setting connection properties with the URL string

Any connection property can be also set, at connect time, within JDBC URL passed to the JDBC method `DriverManager.getConnection()`. In this case, the syntax of solidDB JDBC URL is the following:

```
"jdbc:solid://<hostname>:<port>/>username/<password>[?<property-name>=<value>]..."
```

Examples

```
"jdbc:solid://localhost:1964/dba/dba"
"jdbc:solid://server1.acme.com:1964/dba/dba?solid_login_timeout_ms=100"
"jdbc:solid://server1.acme.com:1964/dba/dba?solid_login_timeout_ms=100?solid_idle_timeout_min=5"
```

3.6 JDBC 2.0 optional package API support

The solidDB JDBC Driver 2.0 supports selected features of the JDBC 2.0 specification Optional Package (known before as Standard Extension).

3.6.1 JDBC connection pooling

The JDBC 2.0 Standard Extension API specifies that users can implement a pooling technique by using specific caching or pooling algorithms that best suit their needs. solidDB provides classes that implement the standard `ConnectionPoolDataSource` and `PooledConnection` interfaces.

solidDB implements these classes as follows:

- `ConnectionPoolDataSource`

A `javax.sql.ConnectionPoolDataSource` interface serves as a resource manager connection factory for pooled `java.sql.Connection` objects. solidDB provides the implementation for that interface in class `SolidConnectionPoolDataSource`. For API functions, see “`ConnectionPoolDataSource` API Functions”.

- `PooledConnection`

A `javax.sql.PooledConnection` interface encapsulates the physical connection to a database. solidDB provides the implementation for that interface in class `SolidPooledConnection`. For API functions, see “`PooledConnection` API Functions” on page 74.

Note: solidDB does not provide an implementation for the actual connection pool; the data structure and the logic to actually pool the `PooledConnection` instances are not available. You must implement your own connection pooling logic, that is, a class that actually pools the connections.

ConnectionPoolDataSource API Functions

The public class `SolidConnectionPoolDataSource` implements `javax.sql.ConnectionPoolDataSource`. The API functions for `javax.sql.ConnectionPoolDataSource` interface are described as follows:

Table 21. Constructor

Description type	Description
Function Name	Constructor
Function Type	solidDB proprietary API
Description	Initializes class variables
Parameters	None
Return value	None
Syntax and exceptions	<code>public SolidConnectionPoolDataSource()</code>

Table 22. Constructor

Description type	Description
Function Name	Constructor
Function Type	solidDB proprietary API
Description	Initializes class variables
Parameters	url As String which identifies the DB server
Return value	None
Syntax and exceptions	public SolidConnectionPoolDataSource(String urlStr)

Table 23. setDescription

Description type	Description
Function Name	setDescription
Function Type	solidDB proprietary API
Description	This function sets the description string.
Parameters	description string (descString)
Return value	None
Syntax and exceptions	public void setDescription(String descString)

Table 24. getDescription

Description type	Description
Function Name	getDescription
Function Type	solidDB proprietary API
Description	This function returns the description string.
Parameters	None
Return value	returns a String (description)
Syntax and exceptions	public String getDescription()

Table 25. *setURL*

Description type	Description
Function Name	setURL
Function Type	solidDB proprietary API
Description	This function sets the url string which points to an solidDB server.
Parameters	url string (urlStr)
Return value	None
Syntax and exceptions	public void setURL(String urlStr)

Table 26. *getURL*

Description type	Description
Function Name	getURL
Function Type	solidDB proprietary API
Description	This function returns the DB url string.
Parameters	None
Return value	returns a String (url)
Syntax and exceptions	public String getURL()

Table 27. *setUser*

Description type	Description
Function Name	setUser
Function Type	solidDB proprietary API
Description	This function sets the username string. (WebSphere compatibility)
Parameters	username string
Return value	None
Syntax and exceptions	public void setUser(String newUser)

Table 28. *getUser*

Description type	Description
Function Name	getUser
Function Type	solidDB proprietary API
Description	This function returns the username string. (WebSphere compatibility)
Parameters	None
Return value	returns a String (username)
Syntax and exceptions	public String getUser()

Table 29. *setPassword*

Description type	Description
Function Name	setPassword
Function Type	solidDB proprietary API
Description	This function sets the password string. (WebSphere compatibility)
Parameters	password string
Return value	None
Syntax and exceptions	public void setPassword(String newPassword)

Table 30. *getPassword*

Description type	Description
Function Name	getPassword
Function Type	solidDB proprietary API
Description	This function returns the password string. (WebSphere compatibility)
Parameters	None
Return value	returns a String (password)
Syntax and exceptions	public String getPassword()

Table 31. *setConnectionURL*

Description type	Description
Function Name	setConnectionURL
Function Type	solidDB proprietary API
Description	This function sets the url string which points to an solidDB server.
Parameters	url string
Return value	None
Syntax and exceptions	public void setConnectionURL(String newUrl)

Table 32. *getConnectionURL*

Description type	Description
Function Name	getConnectionURL
Function Type	solidDB proprietary API
Description	This function returns the url string.
Parameters	None
Return value	returns a String (url)
Syntax and exceptions	public String getConnectionURL()

Table 33. *getLoginTimeout*

Description type	Description
Function Name	getLoginTimeout
Function Type	javax.sql.ConnectionPoolDataSource API
Description	This function returns the login timeout value.
Parameters	None
Return value	returns a timeout value as an integer (seconds)
Syntax and exceptions	public int getLoginTimeout() throws java.sql.SQLException

Table 34. *getLogWriter*

Description type	Description
Function Name	getLogWriter
Function Type	javax.sql.ConnectionPoolDataSource API
Description	This function returns the handle to a writer used for printing debugging messages.
Parameters	None
Return value	returns a handle to java.io.PrintWriter
Syntax and exceptions	public java.io.PrintWriter getLogWriter() throws java.sql.SQLException

Table 35. *getPooledConnection*

Description type	Description
Function Name	getPooledConnection
Function Type	javax.sql.ConnectionPoolDataSource API
Description	This function returns a PooledConnection object from the connection pool. This object has a valid connection to the database server.
Parameters	None
Return value	returns a PooledConnection object.
Syntax and exceptions	public javax.sql.PooledConnection getPooledConnection() throws java.sql.SQLException

Table 36. *getPooledConnection*

Description type	Description
Function Name	getPooledConnection
Function Type	javax.sql.ConnectionPoolDataSource API
Description	This function returns a PooledConnection object from the connection pool. This object has a valid connection to the database server.
Parameters	user (username as String), password (password as String)
Return value	returns a PooledConnection object.

Table 36. *getPooledConnection* (continued)

Description type	Description
Syntax and exceptions	public javax.sql.PooledConnection getPooledConnection(String user, String password) throws java.sql.SQLException

Table 37. *setLoginTimeout*

Description type	Description
Function Name	setLoginTimeout
Function Type	javax.sql.ConnectionPoolDataSource API
Description	This function sets the login timeout value in seconds
Parameters	seconds (as integer)
Return value	None
Syntax and exceptions	public void setLoginTimeout(int seconds)

Table 38. *setLogWriter*

Description type	Description
Function Name	setLogWriter
Function Type	javax.sql.ConnectionPoolDataSource API
Description	This function sets the handle to a writer object that will be used to print/log debug messages.
Parameters	handle to java.io.PrintWriter
Return value	None
Syntax and exceptions	public void setLogWriter(java.io.PrintWriter out) throws java.sql.SQLException

PooledConnection API Functions

The public class `SolidPooledConnection` implements `javax.sql.PooledConnection`. The API functions for `javax.sql.PooledConnection` interface are:

Table 39. *addConnectionEventListener*

Description type	Description
Function Name	addConnectionEventListener
Function Type	javax.sql.PooledConnection API
Description	Adds an event listener to whom this object should notify when it wants to release the connection. This listener is generally the connection pooling module.
Parameters	listener (handle to javax.sql.ConnectionEventListener)
Return value	None
Syntax and exceptions	public void addConnectionEventListener(javax.sql.ConnectionEventListener listener)

Table 40. *close*

Description type	Description
Function Name	close
Function Type	javax.sql.PooledConnection API
Description	This function closes the physical connection.
Parameters	None
Return value	None
Syntax and exceptions	public void close() throws java.sql.SQLException

Table 41. *getConnection*

Description type	Description
Function Name	getConnection
Function Type	javax.sql.PooledConnection API
Description	returns a handle to java.sql.Connection
Parameters	None
Return value	java.sql.Connection

Table 41. *getConnection* (continued)

Description type	Description
Syntax and exceptions	public java.sql.Connection getConnection() throws java.sql.SQLException

Table 42. *removeConnectionEventListener*

Description type	Description
Function Name	removeConnectionEventListener
Function Type	javax.sql.PooledConnection API
Description	This function removes the reference to the listener
Parameters	listener
Return value	None
Syntax and exceptions	public void removeConnectionEventListener(javax.sql.ConnectionEventListener listener)

3.6.2 solidDB Connected RowSet Class: SolidJDBCRowSet

The RowSet described in this topic, extends solid.jdbc.SolidBaseRowSet (which implements javax.sql.RowSet) constructors.

```
/**
 * Create a SolidJDBCRowSet with an existing Connection handle */
public SolidJDBCRowSet(java.sql.Connection conn)

/**
 * Create a SolidJDBCRowSet with an existing ResultSet handle */
public SolidJDBCRowSet(java.sql.ResultSet rset)

/**
 * Create a new SolidJDBCRowSet with given url, username and
 * password.
 */
public SolidJDBCRowSet(String url, String uname, String pwd)

/**
 * Create a new SolidJDBCRowSet with given url, username,
 * password and JNDI naming context.
 */
public SolidJDBCRowSet(String dsname,
                       String username,
                       String password,
                       Context namingcontext)
```

For examples, see the method interface description in the *Java 2 Platform, Standard Edition, v 1.4.2 API Specification*: <http://java.sun.com/j2se/1.4.2/docs/api/javax/sql/RowSet.html>

Considerations about the usage of SolidJDBCRowSet

There are certain methods that you can call (usually for setting parameters for commands to be executed or setting the properties of the RowSet instance) before a connection to the database has been made. However, most of the RowSet interface methods can be called only after a connection to the database has been made. This means that method a command has been set with method `setCommand(String)` and method `execute()` has been called. If the SolidJDBCRowSet instance has no previous `java.sql.Connection` handle, the connection will be established during `execute()` call. After the `execute()` call, the row set instance contains a `java.sql.Connection` object, a `java.sql.PreparedStatement` object, and if the command `execute` was a query statement, it contains also a `java.sql.ResultSet` handle. It also contains all parameter setting methods: `setString`, `setObject`, and so on.

The following example describes the proper use of SolidJDBCRowSet class.

```
/**
 * A simple example on how to use SolidJDBCRowSet
 * First: create an instance of a connected RowSet class.
 * You can give the url, username and password
 * right away in the constructor below, but null parameters
 * for the corresponding values have been given in the example
 * just to show how to use setUrl, setUsername (and so on) methods of the
 * RowSet class.
 */
SolidJDBCRowSet rs = new SolidJDBCRowSet(null, null, null);

// Set the url for the connection
rs.setUrl("jdbc:solid://localhost:1313");

// set the username
rs.setUsername("user1");

// set the passwd
rs.setPassword("pwd1");

/**
 * Note! You can set command parameters and other properties
 * in any order you like, for example, you can set the parameters
 * before you have defined the command to be executed. You can
 * also define the command parameters in any order, since the
 * command statement as well as the given parameters will not be
 * parsed until a connection to the database has been made in
 * the execute() method call.
 */

// set parameter #2 for the command
rs.setString(2, "SYS_SYNC%");

// set the command string
rs.setCommand("select table_name from tables where table_name like ?
and table_name not like ?;");

// set the parameter #1
rs.setString(1, "SYS_%");

// execute the command. The connection to the database is not
// established before this call.
rs.execute();

// now you can browse the ResultSet
while( rowset.next() ){
    // do stuff
}
}
```

```
// close the result set. This method call closes the connection
// to the database as well.
rs.close()
```

3.6.3 Java Naming and Directory Interface (JNDI)

The solidDB JDBC Driver supports the Java Naming and Directory Interface (JNDI).

JNDI allows applications to access naming and directory services through a common interface. JNDI is not a service, but a set of interfaces. These interfaces allow applications to access many different directory services including: file systems, directory services such as Lightweight Directory Access Protocol (LDAP), Network Information System (NIS), and distributed object systems such as the Common Object Request Broker Architecture (CORBA), Java Remote Method Invocation (RMI), and Enterprise JavaBeans (EJB).

3.7 Code examples

This topic contains four Java code samples that use the solidDB JDBC driver.

Java Code Example 1: sample1.java

```
/**
 *      sample1 JDBC sample application
 *
 *
 *      This simple JDBC application does the following using
 *      Solid JDBC driver.
 *
 * 1. Registers the driver using JDBC driver manager services
 * 2. Prompts the user for a valid JDBC connect string
 * 3. Connects to Solid using the driver
 * 4. Creates a statement for one query,
 *    'SELECT TABLE_SCHEMA, TABLE_NAME, TABLE_TYPE FROM TABLES'
 *    for reading data from one of the Solid system tables.
 * 5. Executes the query
 * 6. Fetches and dumps all the rows of a result set.
 * 7. Closes connection
 *
 * To build and run the application
 *
 * 1. Make sure you have a working Java Development environment
 * 2. Install and start Solid to connect. Ensure that the
 *    server is up and running.
 * 3. Append SolidDriver2.0.jar into the CLASSPATH definition used
 *    by your development/running environment.
 * 4. Create a java project based on the file sample1.java.
 * 5. Build and run the application.
 *
 * For more information read the readme.txt file contained in the
 * solidDB package.
 */

import java.io.*;

public class sample1 {

    public static void main (String args[]) throws Exception
    {
        java.sql.Connection conn;
        java.sql.ResultSetMetaData meta;
```

```

java.sql.Statement stmt;
java.sql.ResultSet result;
int i;

System.out.println("JDBC sample application starts...");
System.out.println("Application tries to register the driver.");

// this is the recommended way for registering Drivers
java.sql.Driver d =
    (java.sql.Driver)Class.forName("solid.jdbc.SolidDriver").newInstance();

System.out.println("Driver succesfully registered.");

// the user is asked for a connect string
System.out.println(
    "Now sample application needs a connectstring in format:\n"
);
System.out.println(
    "jdbc:solid://<host>:<port>/<user name>/<password>\n"
);
System.out.print("\nEnter the connect string >");
BufferedReader reader =
    new BufferedReader(new InputStreamReader(System.in));
String sCon = reader.readLine();

// next, the connection is attempted
System.out.println("Attempting to connect :" + sCon);
conn = java.sql.DriverManager.getConnection(sCon);

System.out.println("SolidDriver succesfully connected.");

String sQuery = "SELECT TABLE_SCHEMA, TABLE_NAME, TABLE_TYPE FROM TABLES";

stmt= conn.createStatement();

result = stmt.executeQuery(sQuery);
System.out.println("Query executed and result set obtained.");

// we get a metadataobject containing information about the
// obtained result set
System.out.println("Obtaining metadata information.");
meta = result.getMetaData();
int cols = meta.getColumnCount();

System.out.println("Metadata information for columns is as follows:");
// we dump the column information about the result set
for (i=1; i <= cols; i++)
{
    System.out.println("Column i:"+i+" "+meta.getColumnName(i)+ "," +
        meta.getColumnType(i) + "," + meta.getColumnTypeName(i));
}

// and finally, we dump the result set
System.out.println("Starting to dump result set.");
int cnt = 1;
while(result.next())
{
    System.out.print("\nRow "+cnt+" : ");
    for (i=1; i <= cols; i++) {
        System.out.print(result.getString(i)+"\t");
    }
    cnt++;
}

stmt.close();

conn.close();

```

```

    // and not it is all over
    System.out.println("\nResult set dumped. Sample application finishes.");
}
}

```

Java Code Example 1 Output

```

Solid\DatabaseEngine4.1\jdbc\samples>java sample1.java
JDBC sample application starts...
Application tries to register the driver.
Driver succesfully registered.
Now sample application needs a connectstring in format:

jdbc:solid://<host>:<port>/<user name>/<password>

Enter the connect string >jdbc:solid://localhost:1313/dba/dba
Attempting to connect :jdbc:solid://localhost:1313/dba/dba
SolidDriver succesfully connected.
Query executed and result set obtained.
Obtaining metadata information.
Metadata information for columns is as follows:
Column i:1  TABLE_SCHEMA,12,VARCHAR
Column i:2  TABLE_NAME,12,VARCHAR
Column i:3  TABLE_TYPE,12,VARCHAR
Starting to dump result set.

Row 1 : _SYSTEM SYS_TABLES      BASE TABLE
Row 2 : _SYSTEM SYS_COLUMNS     BASE TABLE
Row 3 : _SYSTEM SYS_USERS       BASE TABLE
Row 4 : _SYSTEM SYS_URole       BASE TABLE
Row 5 : _SYSTEM SYS_RELAuth     BASE TABLE
Row 6 : _SYSTEM SYS_ATTAuth     BASE TABLE
Row 7 : _SYSTEM SYS_VIEWS       BASE TABLE
Row 8 : _SYSTEM SYS_KEYPARTS    BASE TABLE
Row 9 : _SYSTEM SYS_KEYS        BASE TABLE
Row 10 : _SYSTEM      SYS_CARDINAL  BASE TABLE
Row 11 : _SYSTEM      SYS_INFO      BASE TABLE
Row 12 : _SYSTEM      SYS_SYNONYM   BASE TABLE
Row 13 : _SYSTEM      TABLES VIEW
Row 14 : _SYSTEM      COLUMNS VIEW
Row 15 : _SYSTEM      SQL_LANGUAGES  BASE TABLE
Row 16 : _SYSTEM      SERVER_INFO   VIEW
Row 17 : _SYSTEM      SYS_TYPES     BASE TABLE
Row 18 : _SYSTEM      SYS_FORKEYS   BASE TABLE
Row 19 : _SYSTEM      SYS_FORKEYPARTS  BASE TABLE
Row 20 : _SYSTEM      SYS_PROCEDURES  BASE TABLE
Row 21 : _SYSTEM      SYS_TABLEMODES  BASE TABLE
Row 22 : _SYSTEM      SYS_EVENTS     BASE TABLE
Row 23 : _SYSTEM      SYS_SEQUENCES  BASE TABLE
Row 24 : _SYSTEM      SYS_TMP_HOTSTANDBY  BASE TABLE
Result set dumped. Sample application finishes.

```

Java Code Example 2: sample2.java

```

/**
 *      sample2 JDBC sample applet
 *
 *
 *      This simple JDBC applet does the following using
 *      Solid native JDBC driver.
 *
 *      1. Registers the driver using JDBC driver manager services
 *      2. Connects to Solid using the driver.
 *         Used url is read from sample2.html
 *      3. Executes given SQL statements
 *
 *      To build and run the application
 *
 */

```

```

* 1. Make sure you have a working Java Development environment
* 2. Install and start Solid to connect. Ensure that
* the server is up and running.
* 3. Append SolidDriver2.0.jar into the CLASSPATH definition used
* by your development/running environment.
* 4. Create a java project based on the file sample2.java.
* 5. Build and run the application. Check that sample2.html
* defines valid url to your environment.
*
* For more information read the readme.txt file contained
* in the IBM SolidDB Development Kit package.
*
*/

import java.util.*;
import java.awt.*;
import java.applet.Applet;
import java.net.URL;
import java.sql.*;

public class sample2 extends Applet {
    TextField textField;
    static TextArea textArea;

    String url = null;
    Connection con = null;

    public void init() {
        // a valid value for url could be
        // url = "jdbc:solid://localhost:1313/dba/dba";

        url = getParameter("url");

        textField = new TextField(40);
        textArea = new TextArea(10, 40);
        textArea.setEditable(false);

        Font font = textArea.getFont();
        Font newfont = new Font("Monospaced", font.PLAIN, 12);
        textArea.setFont(newfont);

        // Add Components to the Applet.
        GridBagLayout gridBag = new GridBagLayout();
        setLayout(gridBag);
        GridBagConstraints c = new GridBagConstraints();
        c.gridwidth = GridBagConstraints.REMAINDER;

        c.fill = GridBagConstraints.HORIZONTAL;
        gridBag.setConstraints(textField, c);
        add(textField);

        c.fill = GridBagConstraints.BOTH;
        c.weightx = 1.0;
        c.weighty = 1.0;
        gridBag.setConstraints(textArea, c);
        add(textArea);

        validate();

        try {
            // Load the Solid JDBC Driver
            Driver d =
                (Driver)Class.forName ("solid.jdbc.SolidDriver").newInstance();

            // Attempt to connect to a driver.
            con = DriverManager.getConnection (url);

```

```

// If we were unable to connect, an exception
// would have been thrown. So, if we get here,
// we are successfully connected to the url

// Check for, and display and warnings generated
// by the connect.
checkForWarning (con.getWarnings ());

// Get the DatabaseMetaData object and display
// some information about the connection
DatabaseMetaData dma = con.getMetaData ();

textArea.appendText("Connected to " + dma.getURL() + "\n");
textArea.appendText("Driver      " + dma.getDriverName() + "\n");
textArea.appendText("Version    " + dma.getDriverVersion() + "\n");
}
catch (SQLException ex) {
    printSQLException(ex);
}
catch (Exception e) {
    textArea.appendText("Exception: " + e + "\n");
}
}

public void destroy() {
    if (con != null) {
        try {
            con.close();
        }
        catch (SQLException ex) {
            printSQLException(ex);
        }
        catch (Exception e) {
            textArea.appendText("Exception: " + e + "\n");
        }
    }
}

public boolean action(Event evt, Object arg) {
    if (con != null) {
        String sqlstmt = textField.getText();
        textArea.setText("");
        try {
            // Create a Statement object so we can submit
            // SQL statements to the driver
            Statement stmt = con.createStatement ();
            // set row limit
            stmt.setMaxRows(50);
            // Submit a query, creating a ResultSet object
            ResultSet rs = stmt.executeQuery (sqlstmt);

            // Display all columns and rows from the result set
            textArea.setVisible(false);
            dispResultSet (stmt,rs);
            textArea.setVisible(true);

            // Close the result set
            rs.close();

            // Close the statement
            stmt.close();
        }
        catch (SQLException ex) {
            printSQLException(ex);
        }
        catch (Exception e) {
            textArea.appendText("Exception: " + e + "\n");
        }
    }
}

```

```

        }
        textField.selectAll();
    }
    return true;
}

//-----
// checkForWarning
// Checks for and displays warnings. Returns true if a warning
// existed
//-----

private static boolean checkForWarning (SQLWarning warn)
    throws SQLException
{
    boolean rc = false;

    // If a SQLWarning object was given, display the
    // warning messages. Note that there could be
    // multiple warnings chained together

    if (warn != null) {
        textArea.appendText("\n*** Warning ***\n");
        rc = true;
        while (warn != null) {
            textArea.appendText("SQLState: " +
                warn.getSQLState () + "\n");
            textArea.appendText("Message: " +
                warn.getMessage () + "\n");
            textArea.appendText("Vendor: " +
                warn.getErrorCode () + "\n");
            textArea.appendText("\n");
            warn = warn.getNextWarning ();
        }
    }
    return rc;
}

//-----
// dispResultSet
// Displays all columns and rows in the given result set
//-----

private static void dispResultSet (Statement sta, ResultSet rs)
    throws SQLException
{
    int i;

    // Get the ResultSetMetaData. This will be used for
    // the column headings
    ResultSetMetaData rsmd = rs.getMetaData ();

    // Get the number of columns in the result set
    int numCols = rsmd.getColumnCount ();
    if (numCols == 0) {
        textArea.appendText("Updatecount is "+sta.getUpdateCount());
        return;
    }

    // Display column headings
    for (i=1; i<=numCols; i++) {
        if (i > 1) {
            textArea.appendText("\t");
        }
        try {
            textArea.appendText(rsmd.getColumnLabel(i));
        }
    }
}

```

```

        catch(NullPointerException ex) {
            textArea.appendText("null");
        }
    }
    textArea.appendText("\n");

    // Display data, fetching until end of the result set
    boolean more = rs.next ();
    while (more) {

        // Loop through each column, get the
        // column data and display it
        for (i=1; i<=numCols; i++) {
            if (i > 1) {
                textArea.appendText("\t");
            }
            try {
                textArea.appendText(rs.getString(i));
            }
            catch(NullPointerException ex) {
                textArea.appendText("null");
            }
        }
        textArea.appendText("\n");

        // Fetch the next result set row
        more = rs.next ();
    }
}

private static void printSQLException(SQLException ex)
{
    // A SQLException was generated. Catch it and
    // display the error information. Note that there
    // could be multiple error objects chained
    // together

    textArea.appendText("\n*** SQLException caught ***\n");

    while (ex != null) {
        textArea.appendText("SQLState: " +
            ex.getSQLState () + "\n");
        textArea.appendText("Message: " +
            ex.getMessage () + "\n");
        textArea.appendText("Vendor: " +
            ex.getErrorCode () + "\n");
        textArea.appendText("\n");
        ex = ex.getNextException ();
    }
}
}

```

Java Code Example 3: sample3.java

```

/**
 * sample3 JDBC sample application
 *
 *
 * This simple JDBC application does the following using
 * Solid JDBC driver.
 *
 * 1. Registers the driver using JDBC driver manager services
 * 2. Prompts the user for a valid JDBC connect string
 * 3. Connects to Solid using the driver
 * 4. Drops and creates a procedure sample3. If the procedure
 * does not exist dumps the related exception.
 * 5. Calls that procedure using java.sql.Statement

```



```

* 6. Fetches and dumps all the rows of a result set.
* 7. Closes connection
*
* To build and run the application
*
* 1. Make sure you have a working Java Development environment
* 2. Install and start Solid to connect. Ensure that the
*    server is up and running.
* 3. Append SolidDriver2.0.jar into the CLASSPATH definition used
*    by your development/running environment.
* 4. Create a java project based on the file sample3.java.
* 5. Build and run the application.
*
* For more information read the readme.txt
* file contained in the solidDB Development Kit package.
*/

import java.io.*;
import java.sql.*;

public class sample3 {

    static Connection conn;
    public static void main (String args[]) throws Exception
    {
        System.out.println("JDBC sample application starts...");
        System.out.println("Application tries to register the driver.");

        // this is the recommended way for registering Drivers
        Driver d = (Driver)Class.forName("solid.jdbc.SolidDriver").newInstance();

        System.out.println("Driver succesfully registered.");

        // the user is asked for a connect string
        System.out.println(
            "Now sample application needs a connectstring in format:\n"
        );
        System.out.println(
            "jdbc:solid://<host>:<port>/<user name>/<password>\n"
        );
        System.out.print("\nEnter the connect string >");
        BufferedReader reader =
            new BufferedReader(new InputStreamReader(System.in));
        String sCon = reader.readLine();

        // next, the connection is attempted
        System.out.println("Attempting to connect :" + sCon);
        conn = DriverManager.getConnection(sCon);

        System.out.println("SolidDriver succesfully connected.");

        DoIt();

        conn.close();
        // and now it is all over
        System.out.println(
            "\nResult set dumped. Sample application finishes."
        );
    }

    static void DoIt() {
        try {
            createprocs();
            PreparedStatement pstmt = conn.prepareStatement("call sample3(?)");
            // set parameter value
            pstmt.setInt(1,10);
        }
    }
}

```

```

        ResultSet rs = pstmt.executeQuery();
        if (rs != null) {
            ResultSetMetaData md = rs.getMetaData();
            int cols = md.getColumnCount();
            int row = 0;
            while (rs.next()) {
                row++;
                String ret = "row "+row+": ";
                for (int i=1;i<=cols;i++) {
                    ret = ret + rs.getString(i) + " ";
                }
                System.out.println(ret);
            }
        }
        conn.commit();
    }
    catch (SQLException ex) {
        printexp(ex);
    }
    catch (java.lang.Exception ex) {
        ex.printStackTrace ();
    }
}

static void createprocs() {
    Statement stmt = null;
    String proc = "create procedure sample3 (limit integer)" +
        "returns (c1 integer, c2 integer) " +
        "begin " +
        "  c1 := 0;" +
        "  while c1 < limit loop " +
        "    c2 := 5 * c1;" +
        "    return row;" +
        "    c1 := c1 + 1;" +
        "  end loop;" +
        "end";

    try {
        stmt = conn.createStatement();
        stmt.execute("drop procedure sample3");
    } catch (SQLException ex) {
        printexp(ex);
    }

    try {
        stmt.execute(proc);
    } catch (SQLException ex) {
        printexp(ex);
        System.exit(-1);
    }
}

public static void printexp(SQLException ex) {
    System.out.println("\n*** SQLException caught ***");
    while (ex != null) {
        System.out.println("SQLState: " + ex.getSQLState());
        System.out.println("Message: " + ex.getMessage());
        System.out.println("Vendor: " + ex.getErrorCode());
        ex = ex.getNextException ();
    }
}
}
}

```

Java Code Example 4: sample4.java

```
/**
 *      sample4 JDBC sample application
 *
 *
 *      This simple JDBC application does the following using
 *      Solid JDBC driver.
 *
 * 1. Registers the driver using JDBC driver manager services
 * 2. Prompts the user for a valid JDBC connect string
 * 3. Connects to Solid using the driver
 * 4. Drops and creates a table sample4. If the table
 *    does not exist dumps the related exception.
 * 5. Inserts file given as an argument to database (method Store)
 * 6. Reads this 'blob' back to file out.tmp (method Restore)
 * 7. Closes connection
 *
 * To build and run the application
 *
 * 1. Make sure you have a working Java Development environment
 * 2. Install and start Solid to connect. Ensure that
 *    the server is up and running.
 * 3. Append SolidDriver2.0.jar into the CLASSPATH definition used
 *    by your development/running environment.
 * 4. Create a java project based on the file sample4.java.
 * 5. Build and run the application.
 *
 * For more information read the readme.txt file
 * contained in the solidDB Development Kit package.
 */

import java.io.*;
import java.sql.*;

public class sample4 {

    static Connection conn;
    public static void main (String args[]) throws Exception
    {
        String filename = null;
        String tmpfilename = null;

        if (args.length < 1) {
            System.out.println("usage: java sample4 <infile>");
            System.exit(0);
        }
        filename = args[0];
        tmpfilename = "out.tmp";
        System.out.println("JDBC sample application starts...");
        System.out.println("Application tries to register the driver.");

        // this is the recommended way for registering Drivers
        Driver d = (Driver)Class.forName("solid.jdbc.SolidDriver").newInstance();

        System.out.println("Driver succesfully registered.");

        // the user is asked for a connect string
        System.out.println(
            "Now sample application needs a connectstring in format:\n"
        );
        System.out.println(
            "jdbc:solid://<host>:<port>/<user name>/<password>\n"
        );
        System.out.print("\nEnter the connect string >");
        BufferedReader reader =
            new BufferedReader(new InputStreamReader(System.in));
```

```

String sCon = reader.readLine();

// next, the connection is attempted
System.out.println("Attempting to connect :" + sCon);
conn = DriverManager.getConnection(sCon);

System.out.println("SolidDriver succesfully connected.");

// drop and create table sample4
createsample4();
// insert data into it
Store(filename);
// and restore it
Restore(tmpfilename);

conn.close();
// and it is all over
System.out.println("\nSample application finishes.");
}

static void Store(String filename) {
String sql = "insert into sample4 values(?,?)";
FileInputStream inFileStream ;
try {
File f1 = new File(filename);
int blobsize = (int)f1.length();
System.out.println("Inputfile size is "+blobsize);
inFileStream = new FileInputStream(f1);

PreparedStatement stmt = conn.prepareStatement(sql);
stmt.setLong(1, System.currentTimeMillis());
stmt.setBinaryStream(2, inFileStream, blobsize);
int rows = stmt.executeUpdate();
stmt.close();
System.out.println(""+rows+" inserted.");
conn.commit();
}
catch (SQLException ex) {
printexp(ex);
}
catch (java.lang.Exception ex) {
ex.printStackTrace ();
}
}

static void Restore(String filename) {
String sql = "select id,blob from sample4";
FileOutputStream outFileStream ;
try {
File f1 = new File(filename);
outFileStream = new FileOutputStream(f1);

PreparedStatement stmt = conn.prepareStatement(sql);
ResultSet rs = stmt.executeQuery();
int readsize = 0;
while (rs.next()) {
InputStream in = rs.getBinaryStream(2);
byte bytes[] = new byte[8*1024];
int nRead = in.read(bytes);
while (nRead != -1) {
readsize = readsize + nRead;
outFileStream.write(bytes,0,nRead);
nRead = in.read(bytes);
}
}
}
}

```

```

        }
        stmt.close();
        System.out.println("Read "+readsize+" bytes from database");
    }
    catch (SQLException ex) {
        printexp(ex);
    }
    catch (java.lang.Exception ex) {
        ex.printStackTrace ();
    }
}

static void createsample4() {
    Statement stmt = null;
    String proc = "create table sample4 (" +
        "id numeric not null primary key, "+
        "blob long varbinary)";

    try {
        stmt = conn.createStatement();
        stmt.execute("drop table sample4");
    } catch (SQLException ex) {
        printexp(ex);
    }

    try {
        stmt.execute(proc);
    } catch (SQLException ex) {
        printexp(ex);
        System.exit(-1);
    }
}

static void printexp(SQLException ex) {
    System.out.println("\n*** SQLException caught ***");
    while (ex != null) {
        System.out.println("SQLState: " + ex.getSQLState());
        System.out.println("Message: " + ex.getMessage());
        System.out.println("Vendor: " + ex.getErrorCode());
        ex = ex.getNextException ();
    }
}
}

```

3.8 solidDB JDBC Driver type conversion matrix

The conversion matrix included in this topic shows how the Java data type to SQL data type conversion is supported by solidDB JDBC Driver.

This matrix applies to both `ResultSet.getXXX` and `ResultSet.setXXX` methods for getting and setting data. An X indicates that the method is supported by solidDB JDBC Driver.

Table 43. Java data type to SQL data type conversion

Java Data Type applies to getting and setting data	TINYINT	SMALLINT	INTEGER	REAL	FLOAT	DOUBLE	DECIMAL	NUMERIC	CHAR	VARCHAR	LONGVARCHAR	WCHAR	WVARCHAR	LONGVARCHAR	BINARY	VARBINARY	LONGVARBINARY	DATE	TIME	TIMESTAMP
getArray/setArray																				
getBlob/setBlob																				
getBytes/setByte	X	X	X	X	X	X	X	X	X	X	X	X	X	X						
getCharacterStream/ setCharacterStream									X	X	X	X	X	X	X	X	X	X	X	X
getClob/setClob																				
getShort/setShort	X	X	X	X	X	X	X	X	X	X	X									
getInt/setInt	X	X	X	X	X	X	X	X	X	X	X									
getlong/setLong	X	X	X	X	X	X	X	X	X	X	X									
getfloat/setfloat	X	X	X	X	X	X	X	X	X	X	X									
getDouble/setDouble	X	X	X	X	X	X	X	X	X	X	X									
getBigDecimal/setBigDecimal	X	X	X	X	X	X	X	X	X	X	X									
getRef/setRef																				
getBoolean/setBoolean	X	X	X	X	X	X	X	X	X	X	X									
getString/setString	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
getBytes/setBytes									X	X	X	X	X	X	X	X				
getDate/setDate									X	X	X	X	X	X				X		X
getTime/setTime									X	X	X	X	X	X					X	X
getTimeStamp/setTimeStamp									X	X	X	X	X	X				X		X
getAsciiStream/setAsciiStream									X	X	X	X	X	X	X	X				
getUnicodeStream/setUnicodeStream									X	X	X	X	X	X	X	X				
getBinaryStream/setBinaryStream									X	X	X	X	X	X	X	X				
getObject/setObject	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

4 Using solidDB SA

This section describes how to use the solidDB Application Programming Interface (API) also known as *solidDB SA*

solidDB SA is a low level C-language client library to access solidDB database management products. solidDB SA is a layer that resides internally in solidDB products. Normally, the use of an industry standards based interface, such as ODBC or JDBC, is recommended. However, in environments with heavy write load (BATCH INSERTS AND UPDATES), solidDB SA can provide a significant performance advantage.

4.1 What is solidDB SA?

solidDB SA is a C-language client library to connect solidDB database products. This library is used internally in solidDB products and provides access to data in solidDB database tables. The library contains 90 functions providing low-level mechanisms for connecting the database and running cursor-based operations.

Compared to industry standard interfaces, such as ODBC or JDBC, solidDB SA offers better flexibility in constructing network messages that are sent to the database server. In applications whose read or write performance (for example, batch inserts or primary key lookups) needs to be optimized, using solidDB SA can provide a significant performance advantage. For example, if your site experiences performance bottlenecks during batch inserts, solidDB SA can reduce the bottleneck because solidDB SA lets you pass several rows for insertion inside a single network message or remote procedure call.

Note: If the performance bottleneck is in read operations, using solidDB SA will provide only minor improvement in performance.

The solidDB SA interface bypasses the SQL parser, interpreter, and optimizer. With solidDB SA you can access result sets, as long as you are not using SQL through solidDB SA. If retrieval of result sets is necessary through SQL, then you must use industry standard APIs such as ODBC or the solidDB Light Client based on ODBC.

To use solidDB SA requires that you convert your existing interface. This is why we recommend that you use solidDB SA only after you have already attempted (with little success) to use other means for improving performance, which include:

- Writing or indexing columns by primary key for the most appropriate row order. solidDB, otherwise, stores rows on disk in the order they are inserted into the database.
- Eliminating unnecessary indexes. For example, a query that selects more than 15% of a table's rows may be performed faster by a full table scan.
- Optimizing the transaction size by committing transactions after every 100-200 rows inserted.
- Using stored procedures.

4.2 Getting started with solidDB SA

The solidDB SA can be used with shared memory access (SMA) and linked library access (LLA). This topic describes the steps that you should take before using solidDB SA.

Before getting started with solidDB SA, be sure you have:

1. If you are building a local application, you need the SMA or LLA library file. The SMA and LLA library files are installed during solidDB installation. The libraries include the solidDB SA functions as well as full server functionality.
2. If you are building a remote user application, you need the solidDB SA library so that you can link it into your application, for example `solidimpsa.lib` for Windows operating systems.
3. Started solidDB. If necessary, create a new database before using solidDB SA.

Setting up the development environment and building a sample program

Building an application program using the solidDB SA library in the linked library access or the SA client library is identical to building any normal C/C++ program:

1. Insert the linked library access library file or SA client library into your project. Refer to section *Creating and running LLA applications* in the *IBM solidDB Shared Memory Access and Linked Library Access User Guide* for the correct filenames.
2. Include the following solidDB SA header file, which is required in applications using solidDB SA library in the linked library access or the solidDB SA client library:

```
#include "sa.h"
```

Insert the directory containing all the other necessary solidDB SA headers into your development environment's include directories setting.

3. Compile the source code.
4. Link the program.

Verifying the development environment setup

You can verify the development setup with the solidDB SA sample program. This enables you to verify your development environment without writing any code.

Verify the following in your development environment:

- In the Windows environment, the TCP/IP services are provided by standard DLL `wsock32.dll`. To link these services into your project, add `wsock32.lib` into linker's lib file list.

Connecting to a database by using the sample application

In solidDB SA, a connection to a database is represented by the `SaConnectT` structure. This structure is established by calling the function `SaConnect`. The following sample code establishes a connection to a database listening TCP/IP protocol at local machine port 1313. User account DBA with password DBA has been defined in the database.

```
SaConnectT* scon;
```

```
scon = SaConnect("tcp localhost 1313", "dba", "dba");  
if (scon == NULL)  
{
```



```

    /* Connect failed, display connect error text. */
    char* errstr;
    SaErrorInfo(NULL, &errstr, NULL);
    printf("%s\n", errstr);
    return(1);
}

```

4.3 Writing data by using solidDB SA without SQL

With solidDB SA, data is written using cursors.

For delete and update operations, after the cursor is created, a search is performed so that the cursor points to the row that is to be updated and deleted. For insert operations, after the cursor is created, the insertion row(s) are immediately written to the cursor. solidDB SA also enables passing several rows for insertion inside a single network message.

Performing insert operations

solidDB SA functions required for insert operations are listed in the table below.

After solidDB creates a cursor to a certain table, variables are bound to columns, rows are written to the cursor, and then the cursor is closed.

Note: If you use SaArrayInsert to insert more than one row in a single message, then you must perform an explicit flush to send the rows to the database.

Table 44. Insert operation steps

Steps	SA Function(s)	Comment
1. Create a cursor	SaCursorCreate	
2. Binding variables to cursor	SaCursorColData, SaCursorColDate, SaCursorColDateFormat, SaCursorColDFloat, SaCursorColDouble, SaCursorColDynData, SaCursorColDynstr, SaCursorColFloat, SaCursorColInt, SaCursorColLong, SaCursorColStr, SaCursorColTime, SaCursorColTimestamp	
3. Open the cursor	SaCursorOpen	
4. Write a row(s) to the cursor	SaArrayInsert for more than one row or SaCursorInsert for a single row	Perform this in a loop if necessary
5. Free the cursor	SaCursorFree	
6. Flush the network message to the server	SaArrayFlush	Necessary only if using SaArrayInsert.

The following code sample excerpt demonstrates how to write four rows of data in a single network message using the SaArrayInsert function. In the code, a call to SaArrayFlush flushes all rows to the server so they are passed in the same network message.

```

scur = SaCursorCreate(scon, "SAEXAMPLE");

/* Bind variables to columns. */
SaCursorColInt(scur, "INTC", &intc);
SaCursorColStr(scur, "CHARC", &str);

/* Open the cursor. */
SaCursorOpen(scur);

/* Insert values to the table. The column values are taken
 * from the user variables that are bound to columns.
 */
for (intc = 2; intc <= 5; intc++) {
    switch (intc) {
        case 2:
            str = "B";
            break;
        case 3:
            str = "C";
            break;
        case 4:
            str = "D";
            break;
        case 5:
            str = "E";
            break;
    }
    SaArrayInsert(scur);
}

/* Close the cursor. */
SaCursorFree(scur);

/* Flush the inserts to the server. */
SaArrayFlush(scon, NULL);

```

Performing update and delete operations

solidDB SA functions required for basic update and delete operations are listed in the table below.

After solidDB creates a cursor to a specific table, variables are bound to columns of the table, and the cursor is opened. Before the actual search begins, the constraints for finding the row for deletion are set. If there are more rows to be updated, each of the rows requires a separate fetch before they are updated or deleted. After the operation, the cursor is freed.

Table 45. Update and delete operation steps

Steps	SA Function(s)	Comment
1. Create a cursor	SaCursorCreate	
2. Binding variables to cursor	SaCursorColData, SaCursorColDate, SaCursorColDateFormat, SaCursorColDFloat, SaCursorColDouble, SaCursorColDynData, SaCursorColDynstr, SaCursorColFloat, SaCursorColInt, SaCursorColLong, SaCursorColStr, SaCursorColTime, SaCursorColTimestamp	
3. Open the cursor	SaCursorOpen	

Table 45. Update and delete operation steps (continued)

Steps	SA Function(s)	Comment
4. Set the search constraint for the row to be updated or deleted	SaCursorEqual, SaCursorAtleast, SaCursorAtmost	
5. Start a search for the row to be updated or deleted	SaCursorSearch	
6. Fetch the row to be updated or deleted	SaCursorNext	
7. Perform actual update or delete	SaCursorUpdate or SaCursorDelete	For updates, the new values need to be in variables bound in step 2.
8. Free the cursor	SaCursorFree	

The following code sample excerpt demonstrates how to update a row in a table using SaCursorUpdate. Note that in the code, the new values for the update are in variables which are bound to the columns of the table using SaCursorColInt and SaCursorColStr after the cursor is created.

```

scur = SaCursorCreate(scon, "SAEXAMPLE");

/* Bind variables to columns INTC and CHARC of test table. */
SaCursorColInt(scur, "INTC", &intc);
SaCursorColStr(scur, "CHARC", &str);

/* Open the cursor. */
SaCursorOpen(scur);

/* Set search constraint. */
str = "D";
SaCursorEqual(scur, "CHARC");

/* Start a search. */
SaCursorSearch(scur);

/* Fetch the column. */
SaCursorNext(scur);

/* Update the current row in the cursor. */
intc = 1000;
str = "D Updated";
SaCursorUpdate(scur);

/* Close the cursor. */
SaCursorFree(scur);

```

4.4 Reading data by using solidDB SA without SQL

solidDB SA functions required for query operations are listed in this topic.

With solidDB SA, data is queried using cursors. The query data is found in a way similar to update and delete operations. A cursor is created to a specific table, variables are bound to columns of the table, and the cursor is then opened. The constraints for finding the rows for the query are set before starting the actual search. If more than one row is found, each row must be fetched separately. After all the rows are fetched, the cursor needs to be freed.

Basically, all solidDB SA queries use the solidDB optimizer in a way similar to SQL-based queries. The index selection strategy is the same as in SQL. The only exception is that the solidDB SA search uses ORDER BY for selecting an index. This means that an index that best fits ORDER BY is the one selected. If two indices are equally good, then the one with a smaller cost is selected. The query is optimized each time SaCursorSearch is called.

Note: There is no way to use optimizer hints functionality when solidDB SA is used.

Table 46. Query Operation Steps

Steps	SA Function(s)	Comment
1. Create a cursor	SaCursorCreate	
2. Bind variables to cursor	SaCursorColInt, SaCursorColStr and others for other data types	
3. Open the cursor	SaCursorOpen	
4. Set the search constraint for the row to be queried	SaCursorEqual, SaCursorAtleast, SaCursorAtmost	
5. Start a search for the row to be queried	SaCursorSearch	
6. Fetch the row(s) that match the given criteria	SaCursorNext	Perform this in a loop if necessary
7. Free the cursor	SaCursorFree	

Example

```

/* Create cursor to a database table. */
scur = SaCursorCreate(scon, "SAEXAMPLE");

/* Bind variables to columns of test table. */
rc = SaCursorColInt(scur, "INTC", &intc);
rc = SaCursorColStr(scur, "CHARC", &str);

/* Open the cursor. */
rc = SaCursorOpen(scur);
assert(rc == SA_RC_SUCC);

/* Set search constraints. */
str = "A";
rc = SaCursorAtleast(scur, "CHARC");
str = "C";
rc = SaCursorAtmost(scur, "CHARC");

```

```

/* Set ordering criteria. */
rc = SaCursorAscending(scur, "CHARC");

/* Start a search. */
rc = SaCursorSearch(scur);

/* Fetch the rows. */
for (i = 1; i <= 3; i++) {
    rc = SaCursorNext(scur);
    switch (intc) {
        case 1:
            assert(strcmp(str, "A") == 0);
            break;
        case 2:
            assert(strcmp(str, "B") == 0);
            break;
        case 3:
            assert(strcmp(str, "C") == 0);
            break;
    }
}

/* Close the cursor. */
SaCursorFree(scur);
}

```

4.5 Running SQL Statements by Using solidDB SA

In addition to bypassing SQL Parser and SQL interpretation, solidDB SA also allows limited execution of SQL statements directly using function SaSQLExecDirect.

This function is designed to execute simple SQL statements, such as CREATE TABLE. If you need to retrieve SQL result sets, you must use another programming interface such as ODBC or solidDB Light Client.

Example

```

/* Create test table and index. */
SaSQLExecDirect(scon,
    "CREATE TABLE SAEXAMPLE(INTC INTEGER, CHARC VARCHAR)");
SaSQLExecDirect(scon,
    "CREATE INDEX SAEXAMPLE_I1 ON SAEXAMPLE (CHARC)");

```

4.6 Transactions and autocommit mode

By default, solidDB SA runs in autocommit mode.

Autocommit mode is switched off by calling the function SaTransBegin, which explicitly begins a transaction. In this mode, the transaction is committed using the SaTransCommit function or rolled back using SaTransRollback.

Note: After the transaction is committed, solidDB SA returns to its autocommit mode setting.

In autocommit mode, the transaction is committed immediately after an insert (SaCursorInsert), update (SaCursorUpdate), or delete (SaCursorDelete). Note that even when using SaArrayInsert, each individual record is inserted in a separate transaction if autocommit is used (see 4.9.2, "SaArrayInsert," on page 104 for more

details). To improve performance when inserting multiple rows with the SaArrayInsert function, put multiple inserts into a single transaction by using SaTransBegin and SaTransCommit.

4.7 Handling database errors

This section contains information about database error handling.

solidDB SA does not provide ODBC-like error processing capability. Generally, solidDB SA functions return SA_RC_SUCC or a pointer to the requested object if successful. If not successful, then the return value is either one of the solidDB SA error codes (see the table in the following section) or NULL. If the error is a database error, the error text is returned by the SaErrorInfo function.

```
if (scon == NULL) {
    /* Connect failed, display connect error text. */
    char* errstr;
    SaErrorInfo(NULL, &errstr, NULL);
    printf("%s\n", errstr);
    return(1);
}
```

The function SaCursorErrorInfo returns error text if the last cursor operation failed. Note that SaErrorInfo has a connection parameter and thus returns the last error applicable to that connection, while SaCursorErrorInfo has a cursor parameter and thus returns the last error of that cursor.

Error Code and Messages for solidDB SA Functions

Following are the possible return codes for solidDB SA functions. All of these error codes are defined in the sa.h file.

Table 47. solidDB SA Function Return Codes

Error Code	Meaning
SA_RC_SUCC	Operation was successful
SA_RC_END	Operation has completed
SA_ERR_FAILED	Operation failed
SA_ERR_CURNOTOPENED	Cursor is not open
SA_ERR_CUROPENED	Cursor is open
SA_ERR_CURNOSEARCH	No active search in cursor
SA_ERR_CURSEARCH	There is active search in cursor
SA_ERR_ORDERBYILL	Illegal "order by" specification
SA_ERR_COLNAMEILL	Illegal column name
SA_ERR_CONSTRILL	Illegal constraint
SA_ERR_TYPECONVILL	Illegal type conversion

Table 47. solidDB SA Function Return Codes (continued)

Error Code	Meaning
SA_ERR_UNIQUE	Unique constraint violation
SA_ERR_LOSTUPDATE	Concurrency conflict, two transactions updated or deleted the same row
SA_ERR_SORTFAILED	Failed to sort the search result set
SA_ERR_CHSETUNSUPP	Unsupported character set
SA_ERR_CURNOROW	No current row in cursor
SA_ERR_COLISNOTNULL	NULL value given for a NOT NULL column
SA_ERR_LOCALSORT	Result set is sorted locally, cannot update or delete the row
SA_ERR_COMERROR	Communication error, connection is lost
SA_ERR_NOSTRCONSTR	String for constraint is missing.
SA_ERR_ILLENUMVAL	Illegal numeric value
SA_ERR_COLNOTBOUND	Column is not bound
SA_ERR_CALLNOSUP	Operation is not supported*
SA_ERR_RPCPARAM	RPC parameter error
SA_ERR_TABLENOTFOUND	Table not found
SA_ERR_READONLY	Connection is read only
SA_ERR_ILLPARAMCOUNT	Wrong number of parameters
SA_ERR_INVARG	Invalid argument
SA_ERR_INVCALLSEQ	Invalid call sequence

Note: SaArray* functions are not supported in linked library access; they work only with the network client library. They return SA_ERR_CALLNOSUP with linked library access.

4.8 Special notes about solidDB SA

This topic contains important information and restrictions about solidDB SA.

solidDB SA and Binary Large Objects (BLOBs)

Currently, solidDB SA does not support BLOB streams and the maximum size of an attribute value is limited to 32K.

SaCursorCol* Functions and solidDB SQL Supported Datatypes

The SaCursorColXXX() functions bind a variable of type XXX to a specified column. For example, the SaCursorColInt function binds a variable of type int to a specified column. When you bind a variable to a column, the variable and column usually have corresponding types; for example, you usually bind an int C variable to an INT SQL column. However, it is not absolutely required that the data type of the column and the data type of the bound variable be equivalent. For example, you could bind a C int variable to an SQL FLOAT, but you would risk losing precision (or even overflowing or underflowing) as data was transferred back and forth.

The SaCursorCol* functions support the SQL datatypes listed in the following table.

Table 48. Supported SQL Datatype

SaCursorCol* Function	T I N Y I N T	S M A L L I N T	I N T E R E A L	F L O A T	D O U B L E	D E C I M A L	N U M E R I C	C H A R	V A R C H A R	L O N G V A R C H A R	W V A R C H A R	L O N G V A R C H A R	B I N A R Y	V A R B I N A R Y	L O N G V A R B I N A R Y	D A T E	T I M E	T I M E S T A M P
SaCursorColInt	X	X	X	X	X	X	X	X	X	X	X	X						
SaCursor ColLong	X	X	X	X	X	X	X	X	X	X	X	X						
SaCursor ColFloat	X	X	X	X	X	X	X	X	X	X	X	X						
SaCursor ColDouble	X	X	X	X	X	X	X	X	X	X	X	X						
SaCursorColStr								X	X	X								
SaCursorCol Date								X	X	X	X	X				X		X
SaCursor ColTime								X	X	X	X	X					X	X
SaCursor ColTimestamp								X	X	X	X	X					X	X
SaCursor ColData													X	X	X			
SaCursor ColDynData		X	X	X	X	X	X	X	X	X			X	X	X			
SaCursor ColFixStr		X	X	X	X	X	X	X	X	X			X	X	X	X	X	X

Table 48. Supported SQL Datatype (continued)

	TINYINT	SMALLINT	INTEGER	REAL	FLOAT	DOUBLE	DECIMAL	NUMERIC	CHAR	VARCHAR	LONGVARCHAR	WVARCHAR	LONGVARCHAR	BINARY	VARBINARY	LONGVARBINARY	DATE	TIME	TIMESTAMP
SaCursorCol*																			
Function																			
SaCursor ColDynStr		X	X	X	X	X	X	X	X	X	X	X							

Note: Keep in mind that, as in other APIs, the success of some conversions in solidDB SA depend on declared values. For example, SaCursorCollInt is only able to handle the SQL datatype CHAR (as in 'foo') if the actual value of the field is an integer (as in '123').

4.9 solidDB SA Function Reference

This topic contains the list of the solidDB SA functions in alphabetic order.

Each description includes the purpose, synopsis, parameters, return value, and comments.

Function Synopsis

The declaration synopsis for the function is:

```
SA_EXPORT_H function(modifier parameter [...]);
```

where modifier can be:

```
SaConnectT*
SaColSearchT*
SaCursorT*
SaDataTypeT*
SaDateT*
SaDfloatT*
SaDynDataT*
SaDynStrT*
SaChSetT
char*
char**
double*
long*
float*
int
int*
unsigned*
void
```

Parameters are in italics and are described below.

Parameter Description

In each function description, parameters are described in a table format. Included in the table is the general usage type of the parameter (described in the next section), as well as the use of the parameter variable in the specific function.

Parameter Usage Type

The table below shows the possible usage type for solidDB SA parameters. Note that if a parameter is used as a pointer, it contains a second category of usage to specify the ownership of the parameter variable after the call.

Table 49. solidDB SA Parameter Usage Types

Usage Type	Meaning
in	Indicates the parameter is input.
output	Indicates the parameter is output.
in out	Indicates the parameter is input/output.
take	Applies only to a pointer parameter. It means that the parameter value is taken by the function. The caller cannot reference to the parameter after the function call. The function or an object created in the function is responsible for releasing the parameter when it is no longer needed.
hold	Applies only to a pointer parameter. It means that the function holds the parameter value even after the function call. The caller can reference to the parameter value after the function call and is responsible for releasing the parameter. Typically, this kind of parameter is passed to the constructor of some object which holds the pointer value inside a local data structure. The caller cannot release the parameter until the object that holds the parameter is deleted.
use	Applies only to a pointer parameter. It means that the parameter is just used during the function call. The caller can do whatever it wants with the parameter after the function call. This is the most common type of parameter passing.
ref	Applies only to out parameters. See "Return Value" below for details.
give	Applies only to out parameters. See "Return Value" below for details.

Return value

Each function description indicates if the function returns a value and the type of value that is returned. Return Values can be one of the following values:

- Boolean (TRUE, FALSE),
- int (such as 1, 0),

- SaRetT (error return codes) such as SA_RC_SUCC. For a list of valid error codes, refer to 4.7, “Handling database errors,” on page 98.
- Pointer (out parameter)

The possible return usage types for pointers are:

Table 50. Return Usage Types for Pointers

Usage Type	Meaning
ref	Indicates the caller can only reference the returned value, but cannot release it. Ensure that the returned value is not used after it is released by the object that returned it.
give	Indicates the function gives the returned value to the caller. The caller is responsible for releasing the returned value.

4.9.1 SaArrayFlush

SaArrayFlush flushes the array operation buffer (that is, it sends the data to the server) after a series of calls to SaArrayInsert fills that buffer.

By default, all SA operations, even SaArrayFlush operations, are done in autocommit mode. In autocommit mode, the SaArrayFlush function does not automatically insert all of the array's records in a single transaction; instead, when SaArrayFlush is called, each record's insertion is treated as a separate transaction. To maximize performance, you may want to do an explicit SaTransBegin before you call SaArrayFlush, and do an explicit SaTransCommit after you call SaArrayFlush.

SaArray* functions are not supported in linked library access; they work only with the network client library. They return SA_ERR_CALLNOSUP with linked library access.

Synopsis

```
SaRetT SA_EXPORT_H SaArrayFlush(SaConnectT* scon, SaRetT* rctab)
```

The SaArrayFlush function accepts the following parameters:

Table 51. SaArrayFlush Parameters

Parameters	Usage Type	Description
<i>scon</i>	use	Pointer to a connection object
<i>rctab</i>	use	Array of return codes for each array operation If this parameter is non-NULL, the return code of each array operation is returned in rctab[i], where <i>i</i> is the order number of the array operation since the last SaArrayFlush.

Return Value

SA_RC_SUCC or error code of first failed array operation.

See also

4.9.2, “SaArrayInsert.”

4.9.2 SaArrayInsert

SaArrayInsert inserts an array of values on one network message. This function places the inserted value in the array insert buffer. You can flush the buffer (that is, send the data to the server) using function SaArrayFlush.

SaArrayInsert may also perform an implicit flush if the internal cache becomes full. However, to ensure that all rows are sent to the server, you should call SaArrayFlush after you insert the last record using SaArrayInsert.

Note:

1. By default, all SA operations, even SaArrayInsert and SaArrayFlush operations, are done in autocommit mode. See 4.9.1, “SaArrayFlush,” on page 103 for an important note about performance.
2. SaArray* functions are not supported in linked library access; they work only with the network client library. They return SA_ERR_CALLNOSUP with linked library access.

Synopsis

```
SaRetT SA_EXPORT_H SaArrayInsert(SaCursorT* scur)
```

The SaArrayInsert function accepts the following parameters:

Table 52. SaArrayInsert Parameters

Parameters	Usage Type	Description
<i>scur</i>	use	Pointer to a cursor object

Return Value

SA_RC_SUCC or error code

See Also

4.9.1, “SaArrayFlush,” on page 103.

4.9.3 SaColSearchCreate

SaColSearchCreate starts a column information search for a specified table.

Synopsis

```
SaColSearchT* SA_EXPORT_H SaColSearchCreate(  
    SaConnectT* scon,  
    char* tablename)
```

The SaColSearchCreate function accepts the following parameters:

Table 53. SaColSearchCreate Parameters

Parameters	Usage Type	Description
<i>scon</i>	In	Pointer to a connection object
<i>tablename</i>	In	Table name

Return Value

Pointer to the column search object, or NULL if table does not exist.

4.9.4 SaColSearchFree

SaColSearchFree releases the column search object.

Synopsis

```
void SA_EXPORT_H SaColSearchFree(SaColSearchT* colsearch)
```

The SaColSearchCreate function accepts the following parameters:

Table 54. SaColSearchCreate Parameters

Parameters	Usage Type	Description
<i>colsearch</i>	In, take	Column search pointer

Return Value

None

4.9.5 SaColSearchNext

SaColSearchNext returns information about the next column in the table.

Synopsis

```
int SA_EXPORT_H SaColSearchNext(
    SaColSearchT* colsearch,
    char** p_colname,
    SaDataTypeT* p_coltype)
```

The SaColSearchNext function accepts the following parameters:

Table 55. SaColSearchNext Parameters

Parameters	Usage Type	Description
<i>colsearch</i>	in, use	Column search pointer
<i>p_colname</i>	out, ref	Pointer to the local copy of the column name is stored into * p_colname
<i>p_coltype</i>	out	Type of column is stored into * p_coltype. See the sa.h file for a description of the SaDataTypeT data type and the valid values that it can hold.

Return Value

Table 56. SaColSearchNext Return Value

Value	Description
1	Next column found, parameters updated.
0	0 No more columns, parameters not updated. The function also will return 0 if the input parameters are invalid.

4.9.6 SaConnect

SaConnect creates a connection to the solidDB server. Several connections can be active at the same time, but operations in different connections are executed in separate transactions.

Synopsis

```
SaConnectT* SA_EXPORT_H SaConnect(  
    char* servername,  
    char* username,  
    char* password)
```

The SaConnect function accepts the following parameters:

Table 57. SaConnect Parameters

Parameters	Usage Type	Description
<i>servername</i>	in, use	Server name. An empty servername connects to the linked server.
<i>username</i>	in, use	User name
<i>password</i>	in, use	Password

Return Value

Table 58. SaConnect Return Value

Return Usage Type	Description
give	Connect pointer, or if connection failed, NULL.

4.9.7 SaCursorAscending

SaCursorAscending specifies ascending order criteria for a column.

To sort by more than one column, you must call this function once for each column. If there is no key (primary key or index) on the column, then the rows are sorted locally (on the client) rather than on the server side.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorAscending(  
    SaCursorT* scur,  
    char* colname)
```

The SaCursorAscending function accepts the following parameters:

Table 59. SaCursorAscending parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name

Return Value

SA_RC_SUCC or error code

4.9.8 SaCursorAtleast

SaCursorAtleast specifies the Atleast criterion for a column. Atleast criterion means that the column value must be greater than or equal to the Atleast value. The Atleast value is taken from the user variable currently bound to the column.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorAtleast(  
    SaCursorT* scur,  
    char* colname)
```

The SaCursorAtleast function accepts the following parameters:

Table 60. SaCursorAtleast Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name

Return Value

SA_RC_SUCC or error code

4.9.9 SaCursorAtmost

SaCursorAtmost specifies the Atmost criterion for a column. Atmost criterion means that the column value must be less than or equal to the Atmost value. The Atmost value is taken from the user variable currently bound to the column.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorAtmost(  
    SaCursorT* scur,  
    char* colname)
```

The SaCursorAtmost function accepts the following parameters:

Table 61. SaCursorAtmost Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name

Return Value

SA_RC_SUCC or error code

4.9.10 SaCursorBegin

SaCursorBegin positions the cursor to the beginning of the set. The subsequent call to the SaCursorNext function returns the first row.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorBegin(
    SaCursorT* scur)
```

The SaCursorBegin function accepts the following parameters:

Table 62. SaCursorBegin Parameters

Parameters	Usage Type	Description
<i>scur</i>	use	Pointer to a cursor object

Return Value

SA_RC_SUCC or error code

4.9.11 SaCursorClearConstr

SaCursorClearConstr clears all search constraints from a cursor.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorClearConstr(
    SaCursorT* scur)
```

The SaCursorClearConstr function accepts the following parameters:

Table 63. SaCursorClearConstr Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return Value

SA_RC_SUCC or error code

4.9.12 SaCursorColData

SaCursorColData binds a user variable to a database column.

The bound variable may be used either as an "input" parameter or an "output" parameter. An "input" parameter passes data values from the client to the server for operations such as inserts and updates, and for search constraints. An "output" parameter holds values read by the server during searches. For example, to INSERT data, the user first binds the variables, then stores values into the bound variables before the actual INSERT; those values are then copied into the database when the INSERT is performed. Similarly, during a fetch operation, when the next row is retrieved, the values in the columns of that row are copied into bound variables so that the client program can see them.

A variable may be used multiple times after a single binding. For example, if you wanted to insert multiple rows, you might create a loop in which you store appropriate values in the bound variable and then invoke the INSERT operation. The "bind" operation would only need to be done once before the loop; it would not need to be executed inside the loop for each INSERT operation. Similarly, after binding the variables once, you could retrieve many rows (one at a time) using the SaCursorNext function. Each time that you retrieved a row, its values would be copied into the bound variables. Note that the address of the data buffer does not change; only the value stored there changes each time what you call SaCursorNext.

If the column has been set as a search constraint (rather like using a WHERE clause in a SELECT statement), then the value for this constraint is set to the value pointed to by the user data variable whose address is passed as dataptr. For example, if the function SaCursorEquals has been called for the column, then the server retrieves only the rows whose value exactly matches the current value of the bound variable. Note that the search constraints are set up for the search operations (SaCursorSearch, followed by calls to SaCursorNext) but may actually be used to set the cursor to the correct position for other operations (such as SaCursorUpdate or SaCursorDelete). Typically, updates are combined with searches to update only some of the rows. This means that the values for columns which have search constraints are used to define the affected rows (in effect the "WHERE" clause in SQL) and other bound variables are used to define the new values for the rest of the columns. Note that the same bound variable can be used in both the search constraint and in the update/insert operation (just as the same column may be used in both the WHERE clause and the "UPDATE ... SET col = value" clause of an SQL UPDATE statement). If the same bound variable is used in both the search constraint and to convey data back and forth between the client and the server, the search constraint does not change each time that the data in the bound variable is updated; the server uses the value that was in the bound variable at the time that the search constraint was created (for example, when functions like SaCursortAtmost() were called).

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations the new value for the column is taken from the user variable.

When the bound variable is used as an "in" parameter (for example, in INSERT or UPDATE operations), the user is responsible for the allocation and freeing of the buffer. When a bound variable is used as an "out" parameter, the SA layer allocates and frees the buffers. When the variable is used as an "out" parameter, the value stored to the user variable is a pointer to a buffer that contains a local copy of the

column data. After each row is retrieved, that row's value will be copied to this buffer. The pointer to this buffer is valid until the next SaCursorOpen or SaCursorFree call, after which the pointer should not be referenced.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColData(
    SaCursorT* scur,
    char* colname,
    char** dataptr,
    unsigned* lenptr)
```

The SaCursorColData function accepts the following parameters:

Table 64. SaCursorColData Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>dataptr</i>	in, hold	Pointer to the user variable
<i>lenptr</i>	in, hold	Pointer to variable used to hold length of data

Return Value

SA_RC_SUCC or error code.

4.9.13 SaCursorColDate

SaCursorColDate binds a user variable of type SaDateT to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations the new value for the column is taken from the user variable.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColDate(
    SaCursorT* scur,
    char* colname,
    SaDateT* dateptr)
```

The SaCursorColDate function accepts the following parameters:

Table 65. SaCursorColDate Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>dateptr</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.12, “SaCursorColData,” on page 109 for a more detailed discussion of binding variables.

4.9.14 SaCursorColDateFormat

SaCursorColDateFormat binds date format string to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. Depending on the column data type, the format string should be date, time, or timestamp format.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColDateFormat(  
    SaCursorT* scur,  
    char* colname,  
    char* dtformat)
```

The SaCursorColDateFormat function accepts the following parameters:

Table 66. SaCursorColDateFormat parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>dtformat</i>	in, hold	Date/time/timestamp format string

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables. For explanation of possible date/time/timestamp formats, see 4.9.49, “SaDateSetAscii,” on page 133.

4.9.15 SaCursorColDfloat

SaCursorColDfloat binds a user variable of type SaDfloatT to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations, the new value for the column is taken from the user variable.

Note: SaDfloatT corresponds to the SQL data type DECIMAL (not FLOAT).

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColDfloat(  
    SaCursorT* scur,  
    char* colname,  
    SaDfloatT* dfloatptr)
```

The SaCursorColDfloat function accepts the following parameters:

Table 67. SaCursorColDfloat Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>dfloatptr</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

4.9.16, "SaCursorColDouble."

4.9.19, "SaCursorColFloat," on page 115.

See 4.9.12, "SaCursorColData," on page 109 for a more detailed discussion of binding variables.

4.9.16 SaCursorColDouble

SaCursorColDouble binds a user variable of type double to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations the new value for the column is taken from the user variable.

Note: The C-language data type "double" is equivalent to SQL data type "FLOAT".

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColDouble(  
    SaCursorT* scur,  
    char* colname,  
    double* doubleptr)
```

The SaCursorColDouble function accepts the following parameters:

Table 68. SaCursorColDouble Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Table 68. SaCursorColDouble Parameters (continued)

Parameters	Usage Type	Description
<i>colname</i>	in, use	Column name
<i>doubleptr</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

4.9.19, “SaCursorColFloat,” on page 115.

4.9.15, “SaCursorColDfloat,” on page 111.

See 4.9.17, “SaCursorColDynData” for a more detailed discussion of binding variables.

4.9.17 SaCursorColDynData

SaCursorColDynData binds a user variable of type SaDynDataT to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations, the new value for the column is taken from the user variable.

In search operations, the column data is stored to the SaDynDataT variable using function SaDynDataMove, which overwrites the old data. The user is responsible for releasing the SaDynDataT variable after the search ends using function SaDynDataFree.

Dynamic data objects (SaDynDataT) are an abstraction that simplifies the handling of variable length data. Although dynamic data can be used with all types of data, it is best fit for variable length data (VARBINARY, LONG VARBINARY, VARCHAR, LONG VARCHAR, and so on).

The memory management of the data object is hidden inside the object. Dynamic data objects have two externally-visible attributes: the data and the length. Typically, the functions SaDynDataMove and SaDynDataAppend are used to set and modify the data value inside the dynamic data object. More memory will be automatically allocated when necessary and all the associated memory will be automatically deallocated when the dynamic data object is disposed of using SaDynDataFree. The user can access the data or the length using the respective functions SaDynDataGetData and SaDynDataGetLen.

The use of SaDynDataMove and SaDynDataAppend may not be feasible when the data already exists completely in a memory buffer. In addition to increasing the memory usage by keeping two copies of the same data, the overhead of the memory copy may be significant if the buffers are large. Therefore, it may be wise to directly assign the data pointer by using SaDynDataMoveRef (rather than

copying by using SaDynDataMove). In this case, the user may modify or deallocate the memory buffer only after the dynamic data object itself has been freed.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColDynData(
    SaCursorT* scur,
    char* colname,
    SaDynDataT* dd)
```

The SaCursorColDynData function accepts the following parameters:

Table 69. SaCursorColDynData Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>dd</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

4.9.18 SaCursorColDynStr

SaCursorColDynStr binds a user variable of type SaDynStrT to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations, the new value for the column is taken from the user variable.

In search operations, the column data is stored to the SaDynStrT variable using function SaDynStrMove, which overwrites the old data. The user is responsible for releasing the SaDynStrT variable after the search ends using function SaDynStrFree.

The user may bind an SaDynStrT variable to any type of column (not just character columns) and the data will be converted back and forth between the column type and the Dynamic String type.

Dynamic String objects (SaDynStrT) are an abstraction that simplifies the handling of variable length strings. Typically, the functions SaDynStrMove and SaDynStrAppend are used to set and modify the data value inside the dynamic string object. More memory will be automatically allocated when necessary and all the associated memory will be automatically deallocated when the dynamic data object is disposed of using SaDynStrFree.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColDynStr(
    SaCursorT* scur,
    char* colname,
    SaDynStrT* ds)
```

The SaCursorColDynStr function accepts the following parameters:

Table 70. SaCursorColDynStr Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>ds</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of binding variables.

4.9.19 SaCursorColFloat

SaCursorColFloat binds a user variable of type float to a database column.

After the variable has been bound, it can be used to hold a value that will be written to or read from a column, or that will be used to constrain a search operation (for example, as part of the equivalent of a WHERE clause in SQL). In search operations, the user variable is updated to contain the value read from the current row that has been retrieved. Also, if search criteria are involved, this function can be used to pass the values for them. In update and insert operations, the new value is taken from the bound user variable and then written to the column in the database.

Note: The C-language "float" data type corresponds to the SQL "SMALLFLOAT" data type, not the SQL "FLOAT" data type.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColFloat(
    SaCursorT* scur,
    char* colname,
    float* floatptr)
```

The SaCursorColFloat function accepts the following parameters:

Table 71. SaCursorColFloat Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>floatptr</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

4.9.16, "SaCursorColDouble," on page 112.

4.9.15, "SaCursorColDfloat," on page 111.

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of binding variables.

4.9.20 SaCursorCollnt

SaCursorCollnt binds a user variable of type int to a database column.

After the variable has been bound, it can be used to hold a value that will be written to or read from a column, or that will be used to constrain a search operation (for example, as part of the equivalent of a WHERE clause in SQL). In search operations, the user variable is updated to contain the value read from the current row that has been retrieved. Also, if search criteria are involved, this function can be used to pass the values for them. In update and insert operations, the new value is taken from the bound user variable and then written to the column in the database.

Note: The C-language "int" data type is platform-dependent, while the SQL data types (TINYINT, SMALLINT, INT, and BIGINT) are platform-independent. You must be careful to map the appropriate C-language data type and value to the corresponding SQL data type.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorCollnt(  
SaCursorT* scur,  
char* colname,  
int* intptr)
```

The SaCursorCollnt function accepts the following parameters:

Table 72. SaCursorCollnt Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>intptr</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.21 SaCursorColLong

SaCursorColLong binds a user variable to a database column.

After the variable has been bound, it can be used to hold a value that will be written to or read from a column, or that will be used to constrain a search operation (for example, as part of the equivalent of a WHERE clause in SQL). In search operations, the user variable is updated to contain the value read from the current row that has been retrieved. Also, if search criteria are involved, this function can be used to pass the values for them. In update and insert operations, the new value is taken from the bound user variable and then written to the column in the database.

Note: The C-language “long” data type is platform-dependent, while the SQL data types (TINYINT, SMALLINT, INT, and BIGINT) are platform-independent. You must be careful to map the appropriate C-language data type and value to the corresponding SQL data type.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColLong(  
    SaCursorT* scur,  
    char* colname,  
    long* longptr)
```

The SaCursorColLong function accepts the following parameters:

Table 73. SaCursorColLong Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>longptr</i>	in, hold	Pointer to the user variable

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.22 SaCursorColNullFlag

SaCursorColNullFlag binds a NULL value flag to a column.

If the column value is NULL, then * p_isnullflag has a value 1, otherwise the value is 0. The * p_isnullflag value is updated automatically during fetch operations. In

search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. During insert and update, a NULL value is inserted to the database if *p_isnullflag is not zero.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColNullFlag(
SaCursorT* scur,
char* colname,
int* p_isnullflag)
```

The SaCursorColNullFlag function accepts the following parameters:

Table 74. SaCursorColNullFlag Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>p_isnullflag</i>	in, hold	Pointer to an integer variable into where the NULL status is stored during fetch operations, and from where the NULL status is taken during insert and update operations.

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.23 SaCursorColStr

SaCursorColStr binds a user variable to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations the new value for the column is taken from the user variable.

In search operations, the value stored to the user variable is a pointer to a local copy of the column data. The data pointer is valid until the next SaCursorOpen or SaCursorFree call, after which the pointer should not be referenced.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColStr(
SaCursorT* scur,
char* colname,
char** strptr)
```

The SaCursorColStr function accepts the following parameters:

Table 75. SaCursorColStr Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>strptr</i>	in, hold	Pointer to the user variable.

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.24 SaCursorColTime

SaCursorColTime binds a user variable of type SaDateT to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations, the new value for the column is taken from the user variable.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColTime(
    SaCursorT* scur,
    char* colname,
    SaDateT* timeptr)
```

Note: The data type of timeptr is indeed SaDateT; there is no separate SaTimeT for time data.

The SaCursorColTime function accepts the following parameters:

Table 76. SaCursorColTime parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>timeptr</i>	in, hold	Pointer to the user variable.

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.25 SaCursorColTimestamp

SaCursorColTimestamp binds a user variable of type SaDateT to a database column.

In search operations, the user variable is updated to contain the value of the current row. Also, if search criteria are involved, this function is used to pass the values for them. In insert and update operations the new value for the column is taken from the user variable.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorColTimestamp(  
    SaCursorT* scur,  
    char* colname,  
    SaDateT* timestampptr)
```

Note: The data type of timeptr is indeed SaDateT; there is no separate SaTimestampT for timestamp data.

The SaCursorColTimestamp function accepts the following parameters:

Table 77. SaCursorColTimestamp parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>timestampptr</i>	in, hold	Pointer to the user variable.

Return Value

SA_RC_SUCC or error code.

See Also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.26 SaCursorCreate

SaCursorCreate creates a cursor to a table specified by table name. The operation fails if the table does not exist.

Synopsis

```
SaCursorT* SA_EXPORT_H SaCursorCreate(  
    SaConnectT* scon,  
    char* tablename)
```

The SaCursorCreate function accepts the following parameters:

Table 78. SaCursorCreate Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, hold	Pointer to a connection object
<i>tablename</i>	in, use	Table name

Return Value

The parameter *scon* has the Usage Type "hold" because the created cursor object keeps referencing the *scon* object even after the function call has returned.

Table 79. Return Value

Return Usage Type	Description
give	Pointer to the cursor object, or NULL if table does not exist.

4.9.27 SaCursorDelete

SaCursorDelete deletes the current row in a cursor from the database. The cursor must be positioned to a row.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorDelete(SaCursorT* scur)
```

The SaCursorDelete function accepts the following parameters:

Table 80. SaCursorDelete parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC or error code.

4.9.28 SaCursorDescending

SaCursorDescending specifies descending sorting criterion for a column.

To sort by more than one column, you must call this function once for each column.

If there is no key (primary key, or index) on the column, then the rows are sorted locally (on the client) rather than on the server side.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorDescending(
    SaCursorT* scur,
    char* colname)
```

The SaCursorDescending function accepts the following parameters:

Table 81. SaCursorDescending parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name

Return value

SA_RC_SUCC or error code.

4.9.29 SaCursorEnd

SaCursorEnd positions the cursor to the end of the set. A subsequent call to SaCursorPrev will position the cursor to the last row in the set.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorEnd(
    SaCursorT* scur)
```

The SaCursorEnd function accepts the following parameters:

Table 82. SaCursorEnd parameters

Parameters	Usage Type	Description
<i>scur</i>	use	Pointer to a cursor object

Return value

SA_RC_SUCC or error code.

4.9.30 SaCursorEqual

SaCursorEqual specifies an equal search criterion for a column.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorEqual(
    SaCursorT* scur,
    char* colname)
```

The SaCursorEqual function accepts the following parameters:

Table 83. SaCursorEqual parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name

Return value

SA_RC_SUCC or error code.

4.9.31 SaCursorErrorInfo

SaCursorErrorInfo returns error information from the last operation in the cursor.

Synopsis

```
bool SA_EXPORT_H SaCursorErrorInfo(  
SaCursorT* scur,  
char** errstr,  
int* errcode)
```

The SaCursorErrorInfo function accepts the following parameters:

Table 84. SaCursorErrorInfo parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>errstr</i>	out, ref	If non-NULL, pointer to a local copy of an error string is stored into * errstr.
<i>errcode</i>	out	If non-NULL, error code is stored into * errcode.

Return value

TRUE If there are errors, errstr and errcode are updated.

FALSE If there are no errors, errstr and errcode are not updated.

4.9.32 SaCursorFree

SaCursorFree releases a cursor. After this call the cursor pointer is invalid.

Synopsis

```
void SA_EXPORT_H SaCursorFree(SaCursorT* scur)
```

The SaCursorFree function accepts the following parameters:

Table 85. SaCursorFree parameters

Parameters	Usage Type	Description
<i>scur</i>	in, take	Pointer to a cursor object

Return value

None.

4.9.33 SaCursorInsert

SaCursorInsert inserts a new row into the database. Column values for the new row are taken from the user variables bound to columns. The cursor must be opened before new rows can be inserted.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorInsert(SaCursorT* scur)
```

The SaCursorInsert function accepts the following parameters:

Table 86. SaCursorInsert parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC or error code.

4.9.34 SaCursorLike

SaCursorLike specifies a like criterion for a column.

The value cannot contain any wild card characters like '_' or '%' in SQL. If such characters exist in the column value, they are quoted with escape characters by the system. Thus, the like value is effectively the same as the SQL like with no wild card characters ending with a '%' character. For example, if you specify that the engine should search for "MARK" in the column, then the engine will find all values that start with "MARK", such as "MARK", "MARK SMITH", and "MARKETING".

The like value is taken from the user variable bound to the column.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorLike(  
    SaCursorT* scur,  
    char* colname,  
    int likelen)
```

The SaCursorLike function accepts the following parameters:

Table 87. SaCursorLike parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name
<i>likelen</i>	in	Length of like part (excluding the string terminator)

Return value

SA_RC_SUCC or error code.

4.9.35 SaCursorNext

SaCursorNext fetches the next row from the database. All user variables bound to columns are updated.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorNext(SaCursorT* scur)
```

The SaCursorNext function accepts the following parameters:

Table 88. SaCursorNext parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC Next row found

SA_RC_END End of search

4.9.36 SaCursorOpen

SaCursorOpen opens a cursor.

All SaCursorColXXX operations must be done before the cursor is opened. When the cursor is opened, possible existing search is terminated. Also, all search criteria specified for the cursor are cleared.

After the cursor is opened, user can insert new rows to the cursor or specify search criteria. Cursor must be opened before a search can be started.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorOpen(SaCursorT* scur)
```

The SaCursorOpen function accepts the following parameters:

Table 89. SaCursorOpen parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC or error code.

4.9.37 SaCursorOrderbyVector

SaCursorOrderbyVector is used to specify the order of columns used in a search.

The initial values are used as a vector of values to specify the starting position for the search in the key. The initial value is used only for the starting point selection in the key; after that, the initial values are not checked against the column values. If several criteria are given, they are solved in the given order. A proper key must exist for the ordering.

The initial value is taken from the user variable bound to the column.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorOrderByVector(
    SaCursorT* scur,
    char* colname)
```

The SaCursorOrderByVector function accepts the following parameters:

Table 90. SaCursorOrderByVector parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>colname</i>	in, use	Column name

Return Value

SA_RC_SUCC or error code.

SaCursorOrderByVector example

```
/* These variables will be bound to the columns named "I" and "J" */
int i, j;
/* Bind variables to columns in this cursor. */
SaCursorColStr(scur, "I", 'i');
SaCursorColStr(scur, "J", 'j');
/* Set the values that we want to use in the search. */
i = 2;
j = 1;
/* Specify the order of the columns. */
SaCursorOrderByVector(scur, "I");
SaCursorOrderByVector(scur, "J");
/* Search the cursor for matching values. */
SaCursorSearch(scur);
```

The preceding would be the equivalent of the following SQL WHERE clause:

```
...WHERE (i,j) >= (2,1)
```

4.9.38 SaCursorPrev

SaCursorPrev fetches the previous row from the database. All user variables currently bound to columns are updated.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorPrev(SaCursorT* scur)
```

The SaCursorPrev function accepts the following parameters:

Table 91. SaCursorPrev parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC Previous row found

SA_RC_END Beginning of search (we are already at the first row, so there is no previous row).

Note: SA_RC_END can apply to either end (start or finish) of the cursor.

4.9.39 SaCursorReSearch

SaCursorReSearch starts a new search using old search criteria.

Synopsis

SaRetT SA_EXPORT_H SaCursorReSearch(SaCursorT* *scur*)

The SaCursorReSearch function accepts the following parameters:

Table 92. SaCursorReSearch Parameters

Parameters	Usage Type	Description
<i>scur</i>	use	Pointer to a cursor object

Return value

SA_RC_SUCC, SA_RC_END, or error code. See 4.7, "Handling database errors," on page 98 for a list of error codes.

4.9.40 SaCursorSearch

SaCursorSearch starts a search in a cursor. After the search is started, the user can fetch rows from the database. Every search is executed as a separate transaction and it does not see any changes made by the current user or any other user after the search is started.

Synopsis

SaRetT SA_EXPORT_H SaCursorSearch(SaCursorT* *scur*)

The SaCursorSearch function accepts the following parameters:

Table 93. SaCursorSearch parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC, SA_RC_END or error code

4.9.41 SaCursorSearchByRowid

SaCursorSearchByRowid starts a new search where the row specified by rowid belongs to the search set.

SaCursorSearchByRowid searches only according to rowid, so it returns one row or zero rows. Previous search constraints are not removed, and they become effective in the next SaCursorReSearch call.

To get the rowid for a particular record, read the value of the rowid column. Every table has a rowid column; you do not need to explicitly create a rowid column within a CREATE TABLE or ALTER TABLE statement.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorSearchByRowid(  
    SaCursorT* scur,  
    void* rowid,  
    int rowidlen)
```

The SaCursorSearchByRowid function accepts the following parameters:

Table 94. SaCursorSearchByRowid parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>rowid</i>	in, use	Pointer to a data area containing rowid. The rowid should be in the form of a string (char *) despite the fact that it is declared as "void *".
<i>rowidlen</i>	in	Length of data (string) pointed to by the rowid parameter

Return Value

SA_RC_SUCC, SA_RC_END, or error code.

4.9.42 SaCursorSearchReset

SaCursorSearchReset resets a search cursor.

The old search constraints are used, but their values are read again from the user buffers (i.e. the parameters). This allows you to increase performance in situations where you want to repeat a search using a query that is identical except for the specific values used.

For example, suppose that a particular user or connection always searches a particular table based on the ID column in that table, but uses a different ID value in each search. Instead of creating a "new" query to search for the next ID, you can reset the existing query and use a new value.

As an example, suppose that your existing code looks similar to the following:

```

...
/* Bind variable(s) to column(s). */
SaCursorColInt(scur, "MY_COL_NAME", &search_parameter1);

/* Repeat a query using different values each time. */
while (there_are_more_values_to_look_for) {
    /* Set the parameter to the value that you want to search for. */
    search_parameter1 = some_value;
    /* Specify the search criterion. */
    rc = SaCursorEqual(scur, "MY_COL_NAME");
    /* Create new query that uses that search criterion and param value*/
    rc = SaCursorSearch(scur);
    /* Get the row (or rows) that match the search criteria. */
    rc = SaCursorNext(scur);
    /* Process the retrieved data... */
    foo();
    ...
    /* Get rid of the old query before the next loop iteration. */
    rc = SaCursorClearConstr(scur);
}
...

```

You can improve performance in most cases by changing your code to look like the following example:

```

...
/* Bind variable(s) to column(s). */
SaCursorColInt(scur, "MY_COL_NAME", &search_parameter1);

/* Create a new query. */
rc = SaCursorEqual(scur, "MY_COL_NAME");
rc = SaCursorSearch(scur);
/* Set the parameter to the value that you want to search for. */
search_parameter1 = some_value;

/* Repeat a query using different values each time. */
while (there_are_more_values_to_look_for) {
    /* Get the row (or rows) that match the search criteria. */
    rc = SaCursorNext(scur);
    /* Process the retrieved data... */
    foo();
    ...
    /* Set the param to the next value that you want to search for. */
    search_parameter1 = some_value;
    /* Reset the existing query to use the latest value in the param. */
    rc = SaCursorSearchReset(scur);
}
...

```

When you use `SaCursorSearchReset()`, you no longer have to re-specify the constraint condition ("Equal", in the example above) and call `SaCursorSearch()` each time.

`SaCursorSearchReset` resets the cursor to the beginning of the new result set. For example, if you reset a search with no constraints at all, it will reposition the cursor to the beginning of the table.

Note: Ensure that you update the values of the search parameters in the buffers before you call this function; the new values are read during this function call.

Limitations

1. `SaCursorSearchReset()` can not be used in the following scenarios:

- The search has a local sort, i.e. not all sorting criteria could be solved by the index used for the search
- The search is done by rowid with SaCursorSearchByRowid

In these cases, SaCursorSearchReset returns SA_ERR_NORESETSEARCH.

- Each "like" value that you use in constraints must be the same length. The reason for this is that SaCursorLike() takes the length of the "like" constraint as an argument, but it is not possible to change this length when SaCursorSearchReset() is called. For example, the function will work correctly if you use the following sequence of "like" values, because they are all the same length:

```
"SMITH"
"JONES"
```

However, the function will not work correctly if you use the following sequence of "like" values:

```
"SMITH"
"JOHNSON"
```

- Using SaCursorSearchReset is usually impractical if you set multiple constraints using the same column binding. For example, suppose that you want to search for values of "col" in the range between 1 and 10 (inclusive). Your code would look like the following example::

```
SaCursorColInt(scur, "col", &i);
i = 1;
SaCursorAtleast(scur, "col");
i = 10;
SaCursorAtmost(scur, "col");
```

If you reset a search like this, the new value for the column is read from the variable i only once. Therefore, the server reads one value and uses it as both the upper and lower bound. For example, suppose that you use the following code:

```
i = 5;
SaCursorSearchReset(scur);
```

This code makes the search $5 \leq i \leq 5$, which is not the desired result.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorSearchReset(
    SaCursorT* scur
```

The SaCursorSearchReset function accepts the following parameters:

Table 95. SaCursorSearchReset Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return Value

SA_RC_SUCC or error code.

4.9.43 SaCursorSetLockMode

SaCursorSetLockMode sets the cursor search mode.

This setting affects the possible locking modes in the server. If a search is already active, the setting will affect only the next search done in the same cursor. By default the search mode is SA_LOCK_SHARE.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorSetLockMode(
    SaCursorT* scur,
    sa_lockmode_t lockmode)
```

The SaCursorSetLockMode function accepts the following parameters:

Table 96. SaCursorSetLockMode Parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>lockmode</i>	in	Search mode that can be one of the following mode: <ul style="list-style-type: none"> • SA_LOCK_SHARE • SA_LOCK_FORUPDATE • SA_LOCK_EXCLUSIVE

The meanings of the various modes are as follows:

- SA_LOCK_SHARE: default optimistic concurrency control.
- SA_LOCK_FORUPDATE: locks the row for update; others can only read, not write.
- SA_LOCK_EXCLUSIVE: locks the row exclusively; others cannot read or write this record.

Note: This function applies to any table; the table does not need to have a particular lock mode for this function to apply.

Return Value

SA_RC_SUCC

SA_ERR_ILLENUMVAL

4.9.44 SaCursorSetPosition

SaCursorSetPosition positions the cursor to a row specified by a key value. The key value is taken from user bound column variables which have a constraint specification.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorSetPosition(
    SaCursorT* scur)
```

The SaCursorSetPosition function accepts the following parameters:

Table 97. SaCursorSetPosition parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC or error code.

4.9.45 SaCursorSetRowsPerMessage

SaCursorSetRowsPerMessage sets the number of rows to be sent in one network message from the server to the client.

The setting has no effect after the search has been started by function SaCursorSearch.

Synopsis

```
SaRetT SA_EXPORT_H
SaCursorSetRowsPerMessage(
    SaCursorT* scur,
    int rows_per_message)
```

The SaCursorSetRowsPerMessage function accepts the following parameters:

Table 98. SaCursorSetRowsPerMessage parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object
<i>rows_per_message</i>	in	Number of rows to send in one network message

Return value

SA_RC_SUCC Success

SA_ERR_FAILED Error, rows_per_message < 1

4.9.46 SaCursorUpdate

SaCursorUpdate updates the current row in a cursor in the database.

The cursor must be positioned to a row. Column values for the new row are taken from the user variables bound to columns.

Synopsis

```
SaRetT SA_EXPORT_H SaCursorUpdate(SaCursorT* scur)
```

The SaCursorUpdate function accepts the following parameters:

Table 99. SaCursorUpdate parameters

Parameters	Usage Type	Description
<i>scur</i>	in, use	Pointer to a cursor object

Return value

SA_RC_SUCC or error code.

4.9.47 SaDateCreate

SaDateCreate creates a new date object.

The date stored in the date object is undefined.

Synopsis

```
SaDateT* SA_EXPORT_H SaDateCreate(void)
```

The SaDateCreate function accepts no parameters.

Return value

Table 100. SaDateCreate Return Values

Return Usage Type	Description
give	A new date object

4.9.48 SaDateFree

SaDateFree releases a date object.

After this call, the date object is invalid and cannot be used.

Synopsis

```
void SA_EXPORT_H SaDateFree(SaDateT* date)
```

The SaDateFree function accepts the following parameters:

Table 101. SaDateFree parameters

Parameters	Usage Type	Description
<i>date</i>	in, take	Date object

Return value

None.

4.9.49 SaDateSetAsciiz

SaDateSetAsciiz sets ASCII zero string date to a date object.

The following special characters are recognized in the format string:

YYYY	year including century
YY	year with default century 1900
MM	month
M	month
DD	day of month
D	day of month
HH	hours
H	hours

NN minutes
 N minutes
 SS seconds
 S seconds
 FFF fractions of a second, 1/1000 seconds

All fields are optional. The fields are scanned from the format string, and when a match is found, the field is replaced with the proper value. All other characters in the format are treated literally.

Double letters (for example, "MM", "DD", and so on) indicate that the values should be expressed with two digits (values 1-9 will be preceded with the 0 character, for example, 01). Single letters indicate that the value should be expressed with one digit if possible. For example, if you define the date format as "YY-M-D" then the date January 2, 1999 will look like "99-1-2". If you define the date format as "YY-MM-DD", then the date will look like "99-01-02"

The following examples show the usage of date formats:

```

SaDateSetAsciiz(date, "YY-MM-DD", "94-09-13");
SaDateSetAsciiz(date, "MM/DD/YY HH.NN", "09/13/94 19.20");
  
```

The default date format is YYYY-MM-DD HH:NN:SS, where time fields are optional.

Synopsis

```

SaRetT SA_EXPORT_H SaDateSetAsciiz(
    SaDateT* date,
    char* format,
    char* asciiz)
  
```

The SaDateSetAsciiz function accepts the following parameters:

Table 102. SaDateSetAsciiz Parameters

Parameters	Usage Type	Description
<i>date</i>	in, out	Date object
<i>format</i>	in, use	Format of date in asciiz (zero-terminated ASCII) buffer, or NULL if default format is used
<i>asciiz</i>	in, use	Buffer containing the data in asciiz (zero-terminated ASCII) string format

Return Value

SA_RC_SUCC
 SA_ERR_FAILED

4.9.50 SaDateSetTimet

SaDateSetTimet copies the input value from the variable named "timet" to the variable named "date". The value is automatically converted from the format time_t (the format returned by C-library function time()) to the format SaDateT.

Synopsis

```
SaRetT SA_EXPORT_H SaDateSetTimet(  
SaDateT* date,  
long timet)
```

The SaDateSetTimet function accepts the following parameters:

Table 103. SaDateSetTimet parameters

Parameters	Usage Type	Description
<i>date</i>	use	Date object
<i>timet</i>	in	New date value in time_t format

Return value

SA_RC_SUCC

SA_ERR_FAILED

4.9.51 SaDateToAsciiz

SaDateToAsciiz stores the date in an ASCII zero-terminated string format.

For an explanation of different date formats, see 4.9.49, “SaDateSetAsciiz,” on page 133.

Synopsis

```
SaRetT SA_EXPORT_H SaDateToAsciiz(  
SaDateT* date,  
char* format,  
char* asciiz)
```

The SaDateToAsciiz function accepts the following parameters:

Table 104. SaDateToAsciiz parameters

Parameters	Usage Type	Description
<i>date</i>	in, use	Date object
<i>format</i>	in, use	Format of date in asciiz (zero-terminated ASCII) buffer, or NULL if default format is used
<i>asciiz</i>	out	Buffer where date is stored. Note: Note that the caller must allocate a sufficiently large buffer before calling this function, and is also responsible for deallocating the buffer when done with it.

Return Value

SA_RC_SUCC

SA_ERR_FAILED

4.9.52 SaDateToTimet

SaDateToTimet stores the date in a time_t format. The time_t date is the same value as returned by C-library function time().

Synopsis

```
SaRetT SA_EXPORT_H SaDateToTimet(  
SaDateT* date,  
long* p_timet)
```

The SaDateToTimet function accepts the following parameters:

Table 105. SaDateToTimet parameters

Parameters	Usage Type	Description
<i>date</i>	in, use	Date object
<i>timet</i>	out	Pointer to a long variable into where the date is stored in time_t format.

Return value

SA_RC_SUCC

SA_ERR_FAILED

4.9.53 SaDefineChSet

SaDefineChSet defines the client character set.

Synopsis

```
SaRetT SA_EXPORT_H SaDefineChSet(  
SaConnectT* scon,  
SaChSetT chset)
```

The SaDefineChSet function accepts the following parameters:

Table 106. SaDefineChSet parameters

Parameters	Usage Type	Description
<i>scon</i>	in out	Pointer to a connection object
<i>chset</i>	in	Enumerated charset specification. The valid character sets are listed in the sa.h file and include SA_CHARSET_DEFAULT, SA_CHARSET_ANSI, and so on.

Note: The usage type of scon includes "out" because the scon parameter is modified by this function call.

Return value

SA_RC_SUCC when OK or SA_ERR_CHARSETUNSUPP when specified character set is not supported.

4.9.54 SaDfloatCmp

SaDfloatCmp compares two dfloat values.

Synopsis

```
int SA_EXPORT_H SaDfloatCmp(  
SaDfloatT* p_dfl1,  
SaDfloatT* p_dfl2)
```

The SaDfloatCmp function accepts the following parameters:

Table 107. SaDfloatCmp parameters

Parameters	Usage Type	Description
<i>p_dfl1</i>	in, use	Pointer to dfloat variable
<i>p_dfl2</i>	in, use	Pointer to dfloat variable

Return value

```
< -1 if p_dfl1 < p_dfl2  
=  0 if p_dfl1 = p_dfl2  
>  1 if p_dfl1 > p_dfl2
```

This parallels the strcmp() function in C, which returns a negative number if the first parameter is less than the second, zero if the two are equal, and a positive number (greater than zero) if the first parameter is greater than the second.

4.9.55 SaDfloatDiff

SaDfloatDiff calculates the difference of two dfloat values (that is, *p_dfl1* - *p_dfl2*). The result is stored into **p_result_dfl*.

Synopsis

```
SaRetT SA_EXPORT_H SaDfloatDiff(  
SaDfloatT* p_result_dfl,  
SaDfloatT* p_dfl1,  
SaDfloatT* p_dfl2)
```

The SaDfloatDiff function accepts the following parameters:

Table 108. SaDfloatDiff parameters

Parameters	Usage Type	Description
<i>p_result_dfl</i>	out	Pointer to dfloat variable where the result is stored.
<i>p_dfl1</i>	in, use	Pointer to dfloat variable
<i>p_dfl2</i>	in, use	Pointer to dfloat variable

Return value

SA_RC_SUCC

SA_ERR_FAILED

4.9.56 SaDfloatOverflow

SaDfloatOverflow checks if the dfloat contains an overflow value.

Synopsis

```
int SA_EXPORT_H SaDfloatOverflow(  
    SaDfloatT* p_dfl)
```

The SaDfloatOverflow function accepts the following parameters:

Table 109. SaDfloatOverflow parameters

Parameters	Usage Type	Description
<i>p_dfl</i>	in, use	Pointer to dfloat variable

Return value

1: dfloat value is an overflow value

0: dfloat value is not an overflow value

4.9.57 SaDfloatProd

SaDfloatProd calculates the product of two dfloat values. The result is stored into *
p_result_dfl.

Synopsis

```
SaRetT SA_EXPORT_H SaDfloatProd(  
    SaDfloatT* p_result_dfl,  
    SaDfloatT* p_dfl1,  
    SaDfloatT* p_dfl2)
```

The SaDfloatProd function accepts the following parameters:

Table 110. SaDfloatProd Parameters

Parameters	Usage Type	Description
<i>p_result_dfl</i>	out	Pointer to dfloat variable where the result is stored.
<i>p_dfl1</i>	in	Pointer to dfloat variable.
<i>p_dfl2</i>	in	Pointer to dfloat variable.

Return Value

SA_RC_SUCC

SA_ERR_FAILED

4.9.58 SaDfloatQuot

SaDfloatQuot calculates the quotient of two dfloat values (that is, p_df1 / p_df2). The result is stored into $* p_result_df1$.

Synopsis

```
SaRetT SA_EXPORT_H SaDfloatQuot(  
    SaDfloatT* p_result_dfl,  
    SaDfloatT* p_df1,  
    SaDfloatT* p_df2)
```

The SaDfloatQuot function accepts the following parameters:

Table 111. SaDfloatQuot parameters

Parameters	Usage Type	Description
<i>p_result_dfl</i>	out	Pointer to dfloat variable where the result is stored
<i>p_df1</i>	in	Pointer to dfloat variable.
<i>p_df2</i>	in	Pointer to dfloat variable.

Return value

SA_RC_SUCC

SA_ERR_FAILED

4.9.59 SaDfloatSetAsciiiz

SaDfloatSetAsciiiz sets the value of the dfloat from a zero-terminated ASCII string.

Synopsis

```
SaRetT SA_EXPORT_H SaDfloatSetAsciiiz(  
    SaDfloatT* p_dfl,  
    char* asciiiz)
```

The SaDfloatSetAsciiiz function accepts the following parameters:

Table 112. SaDfloatSetAsciiiz parameters

Parameters	Usage Type	Description
<i>p_dfl</i>	out	Pointer to dfloat variable where the result is stored.
<i>asciiiz</i>	in	Buffer where the dfloat value is read as a zero-terminated ASCII string.

Return value

SA_RC_SUCC

SA_ERR_FAILED

4.9.60 SaDfloatSum

SaDfloatSum calculates the sum of two dfloat values. The result is stored into * p_result_dfl.

Synopsis

```
SaRetT SA_EXPORT_H SaDfloatSum(  
    SaDfloatT* p_result_dfl,  
    SaDfloatT* p_dfl1,  
    SaDfloatT* p_dfl2)
```

The SaDfloatSum function accepts the following parameters:

Table 113. SaDfloatSum parameters

Parameters	Usage Type	Description
<i>p_result_dfl</i>	out	Pointer to dfloat variable where the result is stored
<i>p_dfl1</i>	in	Pointer to dfloat variable.
<i>p_dfl2</i>	in	Pointer to dfloat variable.

Return value

SA_RC_SUCC or error code.

4.9.61 SaDfloatToAsciiz

SaDfloatToAsciiz stores the dfloat value as an asciiz (zero-terminated ASCII) string.

Synopsis

```
SaRetT SA_EXPORT_H SaDfloatToAsciiz(  
    SaDfloatT* p_dfl,  
    char* asciiz)
```

The SaDfloatToAsciiz function accepts the following parameters:

Table 114. SaDfloatToAsciiz parameters

Parameters	Usage Type	Description
<i>p_dfl</i>	in	Pointer to dfloat variable.
<i>asciiz</i>	out	Buffer where the dfloat is stored in asciiz (zero-terminated ASCII) string format. The memory for this must already be allocated by the caller.

Return value

SA_RC_SUCC

SA_ERR_FAILED

4.9.62 SaDfloatUnderflow

SaDfloatUnderflow checks if the dfloat contains an underflow value.

Synopsis

```
int SA_EXPORT_H SaDfloatUnderflow(  
    SaDfloatT* p_dfl)
```

The SaDfloatUnderflow function accepts the following parameters:

Table 115. SaDfloatUnderflow parameters

Parameters	Usage Type	Description
<i>p_dfl</i>	in, use	Pointer to dfloat variable.

Return value

1: dfloat value is an underflow value

0: dfloat value is not an underflow value

4.9.63 SaDisconnect

SaDisconnect disconnects the user from the solidDB server.

Synopsis

```
void SA_EXPORT_H SaDisconnect(SaConnectT* scon)
```

The SaDisconnect function accepts the following parameters:

Table 116. SaDisconnect Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, take	Pointer to a connection object.

Return Value

None.

4.9.64 SaDynDataAppend

SaDynDataAppend appends data to the dynamic data object.

Synopsis

```
void SA_EXPORT_H SaDynDataAppend(  
    SaDynDataT* dd,  
    char* data,  
    unsigned len)
```

The SaDynDataAppend function accepts the following parameters:

Table 117. SaDynDataAppend parameters

Parameters	Usage Type	Description
<i>dd</i>	use	Dynamic data object.
<i>data</i>	in out, use	Data that is appended to the dd.
<i>len</i>	in	Length of the data to be appended.

Return Value

None.

See also

See 4.9.17, “SaCursorColDynData,” on page 113 for a more detailed discussion of binding variables.

4.9.65 SaDynDataChLen

SaDynDataChLen changes the data area length of dynamic data object. It allocates and deallocates memory as necessary.

If the new length is smaller than the current length, the data area is truncated. If the new length is greater than the current length, the new data area content is initialized with space characters.

Synopsis

```
void SA_EXPORT_H SaDynDataChLen(  
    SaDynDataT* dd,  
    unsigned len)
```

The SaDynDataChLen function accepts the following parameters:

Table 118. SaDynDataChLen Parameters

Parameters	Usage Type	Description
<i>dd</i>	in out, use	Dynamic data object.
<i>len</i>	in	New data area length of dynamic data object.

Return Value

None.

See Also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.66 SaDynDataClear

SaDynDataClear deallocates the memory allocations from an SaDynDataT object.

SaDynDataClear deallocates the data; it does not deallocate the SaDynDataT object itself. The result of SaDynDataClear leaves an "empty" dynamic data object returned by SaDynDataCreate. The SaDynDataT object itself must be deallocated separately using the SaDynDataFree function.

Synopsis

```
void SA_EXPORT_H SaDynDataClear(  
    SaDynDataT* dd)
```

The SaDynDataClear function accepts the following parameters:

Table 119. SaDynDataClear Parameters

Parameters	Usage Type	Description
<i>dd</i>	in out, use	Dynamic data object.

Return Value

None.

See Also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.67 SaDynDataCreate

SaDynDataCreate creates a new dynamic data object. A dynamic data object is an object that can hold variable amounts of any type of data.

Dynamic data objects can be manipulated using other SaDynDataXXX functions.

Synopsis

```
SaDynDataT* SA_EXPORT_H SaDynDataCreate(void)
```

SaDynDataCreate accepts no parameters.

Return Value

Table 120. SaDynDataCreate Return Value

Return Usage Type	Description
give	A new empty dynamic data object. Returns NULL in case of an error.

See Also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.68 SaDynDataFree

SaDynDataFree releases a dynamic data object. After this call, the dynamic data object pointer is invalid and cannot be used.

Synopsis

```
void SA_EXPORT_H SaDynDataFree(  
    SaDynDataT* dd)
```

The SaDynDataFree function accepts the following parameters:

Table 121. SaDynDataFree parameters

Parameters	Usage Type	Description
<i>dd</i>	in, take	Dynamic data object.

Return Value

None

See also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.69 SaDynDataGetData

SaDynDataGetData returns the pointer to the data area of the dynamic data object.

Synopsis

```
char* SA_EXPORT_H SaDynDataGetData(  
    SaDynDataT* dd)
```

The SaDynDataGetData function accepts the following parameters:

Table 122. SaDynDataGetData parameters

Parameters	Usage Type	Description
<i>dd</i>	in, use	Dynamic data object.

Return Value

A reference to the local data area of dynamic data object.

See also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.70 SaDynDataGetLen

SaDynDataGetLen returns the length of the data area of the dynamic data object.

Synopsis

```
unsigned SA_EXPORT_H SaDynDataGetLen(  
    SaDynDataT* dd)
```

The SaDynDataGetData function accepts the following parameters:

Table 123. SaDynDataGetData Parameters

Parameters	Usage Type	Description
<i>dd</i>	in, use	Dynamic data object.

Return Value

Data area length. The function returns 0 if there is an error, or if the actual length of the data area is 0.

See Also

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.71 SaDynDataMove

SaDynDataMove copies data from the parameter named "data" to a dynamic data object (named dd). This function overwrites possible existing data.

The parameter dd must point to a Dynamic Data Object previously created with the SaDynDataCreate function.

Note: SaDynDataMove copies the data. To copy just the reference rather than the data, see 4.9.72, "SaDynDataMoveRef," on page 146.

Typically, the functions SaDynDataMove and SaDynDataAppend are used to set and modify the data value inside the dynamic data object. More memory will be automatically allocated when necessary and all the associated memory will be automatically deallocated when the dynamic data object is disposed of using SaDynDataFree. The user can access the data or the length using the respective functions SaDynDataGetData and SaDynDataGetLen.

The use of SaDynDataMove and SaDynDataAppend may not be feasible when the data already exists completely in a memory buffer. In addition to increasing the memory usage by keeping two copies of the same data, the overhead of the

memory copy may be significant if the buffers are large. Therefore, it may be wise to directly assign the data pointer by using SaDynDataMoveRef (rather than copying by using SaDynDataMove). In this case, the user may modify or deallocate the memory buffer only after the dynamic data object itself has been freed.

Synopsis

```
void SA_EXPORT_H SaDynDataMove(
SaDynDataT* dd,
char* data,
unsigned len)
```

The SaDynDataMove function accepts the following parameters:

Table 124. SaDynDataMove Parameters

Parameters	Usage Type	Description
<i>dd</i>	in out, use	Dynamic data object.
<i>data</i>	in, use	New data
<i>len</i>	in	Length of data (if the data is a string, this length should include the string terminator)

Return Value

None.

See Also

4.9.72, "SaDynDataMoveRef."

See 4.9.17, "SaCursorColDynData," on page 113 for more information about "Dynamic Data" (SaDynDataT).

4.9.72 SaDynDataMoveRef

SaDynDataMoveRef moves a data reference to a dynamic dataobject.

SaDynDataMoveRef copies the pointer (address) from the parameter named "data" to the appropriate field of the parameter named "dd". The caller must guarantee that the input data is alive as long as the dynamic data object refers to that data.

Note: This function copies only the reference, not the data. To copy the data rather than just the reference, see 4.9.71, "SaDynDataMove," on page 145.

Typically, the functions SaDynDataMove and SaDynDataAppend are used to set and modify the data value inside the dynamic data object. More memory will be automatically allocated when necessary and all the associated memory will be automatically deallocated when the dynamic data object is disposed of using SaDynDataFree. The user can access the data or the length using the respective functions SaDynDataGetData and SaDynDataGetLen.

The use of SaDynDataMove and SaDynDataAppend may not be feasible when the data already exists completely in a memory buffer. In addition to increasing the

memory usage by keeping two copies of the same data, the overhead of the memory copy may be significant if the buffers are large. Therefore, it may be better to directly assign the data pointer by using SaDynDataMoveRef (rather than copying by using SaDynDataMove). In this case, the user may modify or deallocate the memory buffer only after the dynamic data object itself has been freed.

Synopsis

```
void SA_EXPORT_H SaDynDataMoveRef(
    SaDynDataT* dd,
    char* data,
    unsigned len)
```

The SaDynDataMoveRef function accepts the following parameters:

Table 125. SaDynDataMoveRef Parameters

Parameters	Usage Type	Description
<i>dd</i>	in out, use	Dynamic data object
<i>data</i>	in, hold	Data
<i>len</i>	in	Length of data (if the data is a string, this length should include the string terminator)

Return Value

None

See Also

4.9.71, "SaDynDataMove," on page 145.

See 4.9.17, "SaCursorColDynData," on page 113 for a more detailed discussion of "Dynamic Data" (SaDynDataT).

4.9.73 SaDynStrAppend

SaDynStrAppend appends another string at the end of a dynamic string.

Synopsis

```
void SA_EXPORT_H SaDynStrAppend(
    SaDynStrT* p_ds,
    char* str)
```

The SaDynStrAppend function accepts the following parameters:

Table 126. SaDynStrAppend parameters

Parameters	Usage Type	Description
<i>p_ds</i>	out	Dynamic string
<i>str</i>	in, use	String that is appended at p_ds

Return Value

None.

4.9.74 SaDynStrCreate

SaDynStrCreate creates (initializes) a new dynamic string object.

Synopsis

```
SaDynStrT SA_EXPORT_H SaDynStrCreate(void)
```

SaDynDataGetData accepts no parameters.

Return Value

Table 127. SaDynStrCreate Return Value

Return Usage Type	Description
give	Dynamic string object initialized with empty data. Returns NULL if running out of memory.

4.9.75 SaDynStrFree

SaDynStrFree frees the SaDynStrT variable.

In search operations, the column data is stored to the SaDynStrT variable using function SaDynStrMove, which overwrites the old data. The user is responsible for releasing the SaDynStrT variable after the search ends using function SaDynStrFree.

Synopsis

```
void SA_EXPORT_H SaDynStrFree(  
    SaDynStrT* p_ds)
```

The SaDynStrFree function accepts the following parameters:

Table 128. SaDynStrFree parameters

Parameters	Usage Type	Description
<i>p_ds</i>	in, take	Dynamic string.

Note: Because the function deallocates the memory, the pointer *p_ds* is no longer valid after the function call and thus the usage type is "take".

Return Value

None.

4.9.76 SaDynStrMove

SaDynStrMove copies the value of the string (the second parameter) to the SaDynStrT (the first parameter).

SaDynStrMove copies the string, not the pointer.

The SaDynStrT must be initialized with SaDynStrCreate before SaDynStrT is set with SaDynStrMove.

CAUTION:

Do not copy a SaDynStrT to another SaDynStrT (for example with memcpy). This would result in two SaDynStrT pointers pointing at the same allocated area.

Synopsis

```
void SA_EXPORT_H SaDynStrMove(  
    SaDynStrT* p_ds,  
    char* str)
```

The SaDynStrMove function accepts the following parameters:

Table 129. SaDynStrMove Parameters

Parameters	Usage Type	Description
<i>p_ds</i>	out	Pointer to a dynamic string variable.
<i>str</i>	in, use	New value of a dynamic string.

Return Value

None.

4.9.77 SaErrorInfo

SaErrorInfo returns error information from the last operation in a server connection.

Cursor errors cannot be checked with this function; instead function SaCursorErrorInfo must be used.

Synopsis

```
bool SA_EXPORT_H SaErrorInfo(  
    SaConnectT* scon,  
    char** errstr,  
    int* errcode)
```

The SaErrorInfo function accepts the following parameters:

Table 130. SaErrorInfo Parameters

Parameters	Usage Type	Description
<i>scon</i>	use	Pointer to a connection object.
<i>errstr</i>	out, ref	If there was an error, and if this parameter is non-NULL, then a pointer to a local copy of an error string is stored into *errstr.

Table 130. SaErrorInfo Parameters (continued)

Parameters	Usage Type	Description
<i>errcode</i>	out	If there was an error, and if this parameter is non-NULL, then an error code is stored into *errcode.

Return Value

TRUE There was an error, so errstr and errcode were updated.

FALSE There were no errors, so errstr and errcode were not updated.

4.9.78 SaGlobalInit

SaGlobalInit performs some global initialization in the SA system.

This function must be called before any other SA function except SaConnect. If the SaConnect function is called before any other SA function, then you do not need to call SaGlobalInit because SaConnect will call it for you.

Synopsis

```
void SA_EXPORT_H SaGlobalInit(void)
```

SaGlobalInit accepts no parameters.

Return Value

None.

4.9.79 SaSetDateFormat

SaSetDateFormat defines default date format.

For explanation of the possible date formats, see 4.9.51, "SaDateToAscii," on page 135.

Synopsis

```
SaRetT SA_EXPORT_H SaSetDateFormat(  
SaConnectT* scon,  
char* dateformat)
```

The SaSetDateFormat function accepts the following parameters:

Table 131. SaSetDateFormat Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, out, use	Pointer to a connection object.
<i>dateformat</i>	in, use	Default date format for connection.

Note: The usage type includes "out" because the `scon` parameter is modified by this function call.

Return Value

`SA_RC_SUCC` if success.

`SA_ERR_COMERROR` if the connection to the server is broken.

See Also

For explanation of the possible date, time, and timestamp formats, see 4.9.51, "SaDateToAscii," on page 135.

4.9.80 SaSetSortBufSize

`SaSetSortBufSize` sets the amount of memory that a connection uses for local sorts (sorts that are done on the client side by the SA library).

Synopsis

```
SaRetT SA_EXPORT_H SaSetSortBufSize(  
    SaConnectT* scon,  
    unsigned long size)
```

The `SaSetSortBufSize` function accepts the following parameters:

Table 132. *SaSetSortBufSize* Parameters

Parameters	Usage Type	Description
<code>scon</code>	in, out, use	Pointer to a connection object.
<code>size</code>	in	Memory buffer size in bytes.

Note: The usage type includes "out" because the `scon` parameter is modified by this function call.

Return Value

`SA_RC_SUCC` when OK or `SA_ERR_FAILED` when specified memory size was too small (< 10KB)

4.9.81 SaSetSortMaxFiles

`SaSetSortMaxFiles` sets the maximum number of files that the connection uses for local sorts (sorts that are done on the client side by the SA library).

Synopsis

```
SaRetT SA_EXPORT_H SaSetSortMaxFiles(  
    SaConnectT* scon,  
    unsigned int nfiles)
```

The `SaSetSortMaxFiles` function accepts the following parameters:

Table 133. SaSetSortMaxFiles parameters

Parameters	Usage Type	Description
<i>scon</i>	in, out, use	Pointer to a connection object.
<i>nfiles</i>	in	Maximum number of files

Note: The usage type includes "out" because the *scon* parameter is modified by this function call.

Return Value

SA_RC_SUCC when OK or SA_ERR_FAILED when specified number of files is too small (< 3).

4.9.82 SaSetTimeFormat

SaSetTimeFormat defines the default time format.

For an explanation of the possible formats, see the time portion documentation of SaDateSetAscii in 4.9.49, "SaDateSetAscii," on page 133.

Synopsis

```
SaRetT SA_EXPORT_H SaSetTimeFormat(
    SaConnectT* scon,
    char* timeformat)
```

The SaSetTimeFormat function accepts the following parameters:

Table 134. SaSetTimeFormat Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, out, use	Pointer to a connection object.
<i>timeformat</i>	in	Default time format for connection.

Note: The usage type of *scon* includes "out" because the *scon* parameter is modified by this function call.

Return Value

SA_RC_SUCC

SA_ERR_COMERROR if the connection to the server is broken.

See Also

For an explanation of the possible date, time, and timestamp formats, see 4.9.49, "SaDateSetAscii," on page 133.

4.9.83 SaSetTimestampFormat

SaSetTimestampFormat defines the default timestamp format.

For an explanation of the possible date, time, and timestamp formats, see 4.9.49, “SaDateSetAscii,” on page 133.

Synopsis

```
SaRetT SA_EXPORT_H SaSetTimestampFormat(  
    SaConnectT* scon,  
    char* timestampformat)
```

The SaSetTimestampFormat function accepts the following parameters:

Table 135. SaSetTimestampFormat parameters

Parameters	Usage Type	Description
<i>scon</i>	in, out, use	Pointer to a connection object.
<i>timestampformat</i>	in	Default timestamp format for connection.

Return Value

SA_RC_SUCC

See Also

For an explanation of the possible date, time, and timestamp formats, see 4.9.49, “SaDateSetAscii,” on page 133.

4.9.84 SaSQLExecDirect

SaSQLExecDirect allows you to execute simple SQL statements such as CREATE TABLE, DROP TABLE, INSERT, and DELETE.

You cannot do SELECT operations because it is not possible to fetch the data.

Synopsis

```
SaRetT SA_EXPORT_H SaSQLExecDirect(SaConnectT* scon,  
    char *sqlstr)
```

The SaSQLExecDirect function accepts the following parameters:

Table 136. SaSQLExecDirect Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, use	Pointer to a connection object.
<i>sqlstr</i>	in, use	Pointer to a string containing the SQL statement to execute.

Return Value

SA_RC_SUCC

The possible error codes are as follows:

- 15001: SAP_ERR_SYNTAXERROR_SD. Syntax error: <error>, <line>.
- 15002: SAP_ERR_ILLCOLNAME_S. Illegal column name <name>.
- 15003: SAP_ERR_TOOMANYPARAMS. Too many parameters for string constraints.
- 15004: SAP_ERR_TOOFEWPARAMS. Too few parameters for string constraints.

4.9.85 SaTransBegin

SaTransBegin starts a new transaction. After this call, all select, insert, update and delete operations are executed in the same transaction, and the changes are not visible in the database until SaTransCommit is called.

Without the SaTransBegin call, the server is in autocommit mode by default and therefore each select, insert, update, and delete operation is executed in a separate transaction. No explicit commit (SaTransCommit) is required when in autocommit mode.

The transaction is run in a mode where write operations are validated for lost updates and unique errors.

Synopsis

```
void SA_EXPORT_H SaTransBegin(SaConnectT* scon)
```

The SaTransBegin function accepts the following parameters:

Table 137. SaTransBegin parameters

Parameters	Usage Type	Description
<i>scon</i>	in, use	Pointer to a connection object.

Return Value

None.

4.9.86 SaTransCommit

SaTransCommit commits the current transaction started by SaTransBegin.

After calling this function, all changes are made persistent in the database. After the current transaction is completed, the database server returns to autocommit mode until the next call to SaTransBegin.

Synopsis

```
SaRetT SA_EXPORT_H SaTransCommit(SaConnectT* scon)
```

The SaTransCommit function accepts the following parameters:

Table 138. SaTransCommit Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, use	Pointer to a connection object

Return Value

SA_RC_SUCC or error code.

4.9.87 SaTransRollback

SaTransRollback rolls back the current transaction started by SaTransBegin. No changes are made to the database.

After the current transaction is completed, the database server returns to autocommit mode until the next call to SaTransBegin.

Synopsis

```
SaRetT SA_EXPORT_H SaTransRollback(SaConnectT* scon)
```

The SaTransRollback function accepts the following parameters:

Table 139. SaTransRollback Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, use	Pointer to a connection object

Return Value

SA_RC_SUCC or error code.

4.9.88 SaUserId

SaUserId returns current user id of a connection.

Synopsis

```
int SA_EXPORT_H SaUserId(SaConnectT* scon)
```

The SaUserId function accepts the following parameters:

Table 140. SaUserId Parameters

Parameters	Usage Type	Description
<i>scon</i>	in, use	Pointer to a connection object

Return Value

User id in the server.

5 Using Unicode

solidDB supports the Unicode standard, providing the capability to encode characters used in the major languages of the world. To use Unicode encoded data, you do not need to use any non-standard or solidDB-specific implementations for application development; standard ODBC API or JDBC API can be used, as well as solidDB tools. solidDB also supports heterogeneous multi-client environments where each application can be set to use different encoding.

Unicode database modes

Starting from version 6.5, the solidDB databases can be created in two modes: *Unicode* mode or *partial Unicode* mode. This database mode is based on the encoding of character data types (CHAR, VARCHAR and so on) in the solidDB server. Wide character data types (WCHAR, WVARCHAR and so on) are Unicode encoded in both modes.

- Unicode mode

In the Unicode mode, the internal representation for character data types is UTF-8.

The internal representation for wide character data types is UTF-16.

- partial Unicode mode

In the partial Unicode mode, the internal representation for character data types uses no particular encoding; instead, the data is stored in byte strings with the assumption that user applications are aware of this and handle the conversion as necessary.

The internal representation for wide character data types is UTF-16.

The databases created with solidDB version 6.3 or earlier are of the partial Unicode type.

Important: The default database mode in 6.5 is partial Unicode.

Note: Unicode applications can be built on both Unicode and partial Unicode databases. However, the instructions in this section assume that the Unicode support is based on the Unicode database mode.

Key features of solidDB Unicode databases

- **Storing and retrieving of Unicode data**

The internal representation of Unicode data is based on UTF-8 and UTF-16 encoding. Data in wide character column types is represented internally in UTF-16 and data in character column types is represented in UTF-8.

This means that both single and multi-byte data can be stored in character column types; if mainly multi-byte data is expected, you can optimize space-efficiency by choosing to store the multi-byte data into wide character column types.

- **No restrictions on the encoding used in the applications**

solidDB ODBC/JDBC drivers handle the conversion of data between the application encoding and the UTF-8/UTF-16 format in the solidDB server.

- **Standard ODBC API and JDBC API available for application development**

There are no non-standard or solidDB-specific requirements for application development; standard ODBC API or JDBC API can be used.

5.1 What is Unicode?

The Unicode Standard is the universal character representation standard for text in computer processing. Unicode provides a consistent way of encoding multilingual plain text making it easier to exchange text files internationally.

The Unicode Standard defines code points (unique numbers) for characters used in the major languages written today. This includes punctuation marks, diacritics, mathematical symbols, technical symbols, arrows, dingbats, and so on. In all, the Unicode Standard provides codes for over 100,000 characters from the world's alphabets, ideograph sets, and symbol collections, including classical and historical texts of many written languages. The characters can be represented in different encoding forms, such as UTF-8 and UTF-16.

The Unicode Standard is fully compatible with the International Standard ISO/IEC 10646; it contains all the same characters and code points as ISO/IEC 10646. This code-for-code identity is true for all encoded characters in the two standards, including the East Asian (Han) ideographic characters. The Unicode Standard also provides additional information about the characters and their use. Any implementation that conforms to Unicode also conforms to ISO/IEC 10646.

Encoding forms

Unicode characters are represented in one of three encoding forms: a 32-bit form (UTF-32), a 16-bit form (UTF-16), and an 8-bit form (UTF-8). These character encoding standards define not only the identity of each character and its numeric value (code position), but also how this value is represented in bits.

Starting from version 6.5, solidDB can be configured to use the UTF-8 encoding for representing character data and UTF-16 encoding for wide character data. The database mode is controlled with the parameter **General.InternalCharEncoding**.

If a database is created in Unicode mode (**General.InternalCharEncoding=UTF8**), the following applies:

- Data in wide character column types is represented internally in UTF-16.
- Data in character column types is represented in UTF-8.

If a database is created in the partial Unicode mode (**General.InternalCharEncoding=Raw**), the following applies

- Data in wide character column types is represented internally in UTF-16.
- Data in character column types is not encoded in any particular encoding; instead, the data is stored in byte strings, with the assumption that user applications are aware of this and handle the conversion as necessary.

The UTF-8 and UTF-16 encodings are essentially ways of turning the encoding into the actual bits that are used in implementation; UTF-8 and UTF-16 encodings share the same character set, but the data size of each character differs.

- UTF-16

UTF-16 assumes 16-bit characters and allows for a certain range of characters to be used as an extension mechanism in order to access an additional million characters using 16-bit character pairs.

- UTF-8

UTF-8 is a way of transforming all Unicode characters into a variable length encoding of bytes. It has the advantages that the Unicode characters corresponding to the familiar ASCII set end up having the same byte values as ASCII, and that Unicode characters transformed into UTF-8 can be used with much existing software without extensive software rewrites.

The Unicode Consortium also endorses the use of UTF-8 as a way of implementing the Unicode standard. Any Unicode character expressed in the 16-bit UTF-16 form can be converted to the UTF-8 form and back without loss of information.

5.2 Designing Unicode databases

This section contains information on how to setup solidDB databases for use with Unicode.

Note: Unicode applications can be built on both Unicode and partial Unicode databases. However, the instructions in this section assume that the Unicode support is based on the Unicode database mode.

Creating Unicode databases

The solidDB database mode is controlled with the parameter **General.InternalCharEncoding**.

- Unicode mode: **General.InternalCharEncoding=UTF8**

When the **InternalCharEncoding** is set to UTF8, the internal representation for character data types is UTF-8. Both character data types and wide character data types are converted between the solidDB server and the application.

- partial Unicode mode: **General.InternalCharEncoding=Raw**

When the **InternalCharEncoding** is set to Raw, the internal representation for character data types uses no particular encoding; instead, the data is stored in byte strings with the assumption that user applications are aware of this and handled the conversion as necessary. Wide character data types are converted between the solidDB server and the application.

The databases created with solidDB version 6.3 or earlier are of the partial Unicode type.

Important: The database mode must be defined when the database is created and it cannot be changed later.

If the database already exists in either mode and the database mode contradicts the value of the parameter, the server startup fails with the following error message in the `solerr.out`:

```
Parameter General.InternalCharEncoding contradicts the existing database mode
```

Determining which data types to use in Unicode databases

Both character and wide character data types can be used to store Unicode data in Unicode databases. If mainly multi-byte data is expected, you can optimize space-efficiency by choosing to store the multi-byte data into wide character column types. This is because even though UTF-8 and UTF-16 encodings share the same character set, the data size of each character differs.

- Data in wide character column types (WCHAR/WVARCHAR/LONG VARCHAR) is represented internally in UTF-16: each character is represented in two or four bytes.
 - Characters in the Basic Multilingual Plane (BMP): two bytes
 - Characters outside the BMP (surrogate characters): four bytes
- Data in character column types (CHAR/VARCHAR/LONG VARCHAR) is represented in UTF-8: each character is represented in one to four bytes.

The size depends on the code point:

Simplified example:

- ASCII characters: one byte
- Cyrillic, Arabic, Hebrew, Latin-1 supplement and so on characters: two bytes
- Asian characters/rest of BMP characters: three bytes
- Characters outside the BMP (surrogate characters): four bytes

For example, Asian languages are stored more efficiently on wide character data types (UTF-16) since most characters are part of BMP, which requires two bytes. European languages are stored more efficiently on character data types (UTF-8) since most common characters are represented in one byte.

Wide character data requires also less processing; using wide character data types may improve performance.

The Unicode data types are interoperable; because UTF-16 and UTF-8 share the same character set, there is no risk of data loss when using either data type. All string operations are possible between character and wide character data types with implicit type conversions.

Creating columns for storing Unicode data

In order to start storing Unicode data in a Unicode database, tables with Unicode data columns need to be created first as follows:

```
CREATE TABLE customer1 (c_id INTEGER, c_name VARCHAR,...)
CREATE TABLE customer2 (c_id INTEGER, c_name WVARCHAR,...)
```

Ordering data columns (collation)

The character data columns are ordered based on the binary values of the UTF-8 and wide character data columns on the UTF-16 format (using most significant byte order). If the binary order is different than what the national language users expect, you need to provide a separate column to store the correct ordering information.

Using Unicode in database entity names

It is possible to name database entities such as tables, columns, and procedures with Unicode strings simply by enclosing the Unicode names with double quotes in all the SQL statements.

solidDB tools can handle Unicode strings according to the default locale of the environment or according to a specified locale.

For more details, see 5.3, “Using solidDB tools with Unicode,” on page 161.

Using Unicode in user names and passwords

User names and passwords can also be Unicode strings. However, to avoid access problems from different tools, the original database administrator account information must be given as pure ASCII strings.

Using Unicode in file names

Unicode strings cannot be used in any file names.

5.3 Using solidDB tools with Unicode

This section contains information about how to use the solidDB tools with Unicode and partial Unicode databases.

The following solidDB tools can be used to output and import data in the system default locale or a specified locale in both Unicode and partial Unicode databases.

- solidDB SQL Editor (**solsql**)
- solidDB Data Dictionary (**soldd**)
- solidDB Export (**solexp**)
- solidDB Speed Loader (**solloado**)

solidDB Remote Control (**solcon**) does not support conversions of data to UTF-8. For example, if an error message that is output to **solcon** contains Unicode encoded data, it is not displayed correctly in the console.

The locale to be used in conversions is defined with the command line options when starting the tool.

Important:

- The solidDB tools use the solidDB ODBC API 3.5.1; this means that if the binding method for character data types is defined with the server-side **Srv.ODBCDefaultCharBinding** or client-side **Client.ODBCCharBinding** parameters, this setting also impacts the behavior of the solidDB tools.
- The Unicode and partial Unicode databases behave differently in reference to conversions of CHAR and WCHAR data types:
 - **Unicode databases**
Both CHAR and WCHAR data types are converted between the UTF-8/UTF-16 format in solidDB and the locale/codepage defined with the chosen binding method.
 - **partial Unicode databases**
CHAR data types are not converted; instead, they are handled in the raw (binary) format that is used to store CHAR data in partial Unicode databases. WCHAR data types are converted between the UTF-16 format in solidDB and the locale/codepage defined with the chosen binding method.

Table 141. Command line options for solidDB tools for partial Unicode and Unicode databases

Option	Description
No option/Factory setting	The console locale setting is used, unless overridden with the server-side or client-side parameters in the <code>solid.ini</code> file. Note: If the server-side <code>Srv.ODBCDefaultCharBinding</code> or client-side <code>Client.ODBCCharBinding</code> parameter is set to UTF8, the locale of the console must support UTF-8.
<code>-m</code>	The console locale setting is used, despite the server-side or client-side parameters in the <code>solid.ini</code> file.
<code>-M<locale_name></code>	The locale console setting is overridden with the locale defined with <code><locale_name></code> . The <code><locale_name></code> depends on the operating system. For example, in Linux environments, the locale name for the code page GB18030 in Chinese/China is <code>zh_CN.gb18030</code> . In Windows environments, the locale name for Latin1 code page in Finnish/Finland is <code>fin_fin.1252</code> .
<code>-u</code>	Input/output is forced to UTF-8.

Note: If the server-side or client-side parameters in the `solid.ini` file are set to use 'Raw' binding, you should always use the `-m`, `-M` or `-u` option to override the `solid.ini` settings.

5.4 Compatibility between Unicode and partial Unicode databases

If a database has been created in Unicode mode, it cannot be changed to partial Unicode mode, and vice versa. If you need to convert a partial Unicode database to a Unicode (or vice versa), you can use the solidDB tools to export and reload your database.

If the database already exists in either mode and the database mode contradicts the value of the parameter, the server startup fails with the following error message in `solerr.out`.

Parameter General.InternalCharEncoding contradicts the existing database mode

5.4.1 Converting partial Unicode databases to Unicode

To convert a partial Unicode database to a Unicode database, use the solidDB tools to export and reload your database.

Before you begin

- Create a backup of your database.
- Verify the locale/codepage that is used in the application side for encoding data in CHAR data type columns in the partial Unicode database.

During the export phase, the data in CHAR data type columns is not converted by solidDB tools; instead, it is output as such. This means that the underlying locale/codepage for CHAR data types becomes the locale/codepage format for the output file. To enable the output files to contain data in a single locale/codepage, solidDB must be able to convert the data in the WCHAR data type columns from UTF-16 into the exact same locale/codepage format as the CHAR data.

In the import phase, the solidDB tools convert the data from the locale/codepage format of the output file into the UTF-8 (CHAR) and UTF-16 (WCHAR) encoding used in Unicode databases.

About this task

In this procedure, the following setup is used as an example:

- The server name is solidDB and the protocol used for connections is TCP/IP, using port 1964 (network name is "tcpip 1964").
- The partial Unicode database has been created with the username "dbadmin" and password "password".
- CHAR data types in the partial Unicode database are encoded in the application side with the locale zh_CN.gb18030 (Chinese/China and code page GB18030).

Tip: If you have your database creation scripts available, you can use them to create the new database table definitions, instead of using soldd and solsql for exporting and importing them.

Procedure

1. Extract data definitions with solidDB Data Dictionary (soldd).

Use the following command to extract an SQL script containing definitions for all tables, views, triggers, indexes, procedures, sequences, and events.

```
soldd -Mzh_CN.gb18030 "tcpip 1964" dbadmin password
```

The default file name soldd.sql is used.

Note: User and role definitions are not listed for security reasons. If the database contains users or roles, add CREATE statements for them to the extracted SQL file manually.

Important: To preserve referential integrity, you may need to reorganize the table definition statements to ensure that the referenced tables are created before the referencing tables.

2. Extract data from you database with solidDB Export (solexp).

Use the following command to extract the control and data files for all tables.

```
solexp -Mzh_CN.gb18030 "tcpip 1964" dbadmin password *
```

This export creates control files (<table_name>.ctr) and data files (<table_name>.dat) for each table. The default file name is the same as the exported table name.

3. Create a new Unicode database.

- a. Set the **General.InternalCharEncoding** parameter to UTF8.

```
[General]  
InternalCharEncoding=UTF8
```

- b. Create a new database by starting solidDB in the working directory for the new Unicode database.

4. Import data definitions into the new database using the solidDB SQL Editor (solsql).

Use the following command to execute the SQL script created by solidDB Data Dictionary (soldd).

```
solsql -fsoldd.sql -Mzh_CN.gb18030 "tcpip 1964" dbadmin password
```

5. Load the data into the new database using the solidDB Speed Loader (solloado).

For each table, use the following command to load data into the new database:

```
solloado -Mzh_CN.gb18030 "tcpip 1964" dbadmin password <table_name>.ctr
```

Related topics

- *Using solidDB management tools in the IBM solidDB Administrator Guide*

5.5 Developing applications for Unicode

This section contains information on how to design your applications for use with solidDB databases in Unicode mode.

Supported interfaces

- **ODBC**

The solidDB ODBC Driver is Unicode compliant; it conforms to the Microsoft ODBC 3.51 standard.

Note: solidDB provides two versions of the ODBC driver, one for Unicode and one for ASCII. The Unicode version is a superset of the ASCII version; you can use it with either Unicode or ASCII character sets.

- **JDBC**

Unicode is supported in the solidDB JDBC Driver which is a solidDB implementation of the JDBC 2.0 standard.

As Java uses natively Unicode strings, supporting Unicode means primarily that when accessing character data in solidDB, no data type conversions are necessary. Additionally, JDBC ResultSet Class methods `getUnicodeStream` and `setUnicodeStream` are supported for handling large Unicode texts stored in solidDB.

Multi-client environments with different locale settings

In Unicode databases, the solidDB ODBC and JDBC drivers handle the conversion of data between the application encoding and the UTF-8/UTF-16 format in the solidDB server.

In ODBC environments, the conversions can be set to expect the application default locale or a user-defined locale for the encoding in the application buffer. This is controlled with server-side **Srv.ODBCDefaultCharBinding** and client-side **Client.ODBCCharBinding** configuration parameters.

SQL string functions

SQL string functions work as expected. Conversions are provided implicitly, when necessary. If either of the operands is of wide character type, the result is always of wide character type.

The functions `UPPER()` and `LOWER()` perform the uppercase or lowercase conversions on Unicode strings only when the characters are part of Latin 1 code page. If the Unicode character cannot be converted to uppercase or lowercase, the input string is returned as it is.

Character padding

The `SQL.CharPadding=yes` parameter setting is not effective in Unicode databases; blank characters in CHAR values are always discarded.

5.5.1 ODBC applications and Unicode databases

In ODBC environments, the solidDB ODBC driver handles the conversion of data between the encoding used in the application (client) and the UTF-8/UTF-16 format in the solidDB Unicode database. The binding of character data can be set for all clients using the server-side parameter `Srv.ODBCDefaultCharBinding` or per client using the client-side parameter `Client.ODBCCharBinding`. In both cases, the standard C type identifier `SQL_C_CHAR` is used.

For binding of character data, you can set the ODBC driver to use one of the following methods:

- the current client locale encoding
- specific encoding as defined with a locale name
- no encoding
- UTF-8 encoding

Two use cases are supported for all methods:

- The same binding method is set for all clients using the server-side parameter `ODBCDefaultCharBinding`.

```
[Srv]
ODBCDefaultCharBinding=raw|locale|locale:|locale:<locale name>|UTF8
```

- The binding method is set per client by using the client-side parameter `ODBCCharBinding`.

```
[Client]
ODBCCharBinding=raw|locale|locale:<locale name>
```

The `ODBCCharBinding` parameter overrides the server-side settings set by `ODBCDefaultCharBinding`.

The factory value for both is `locale:`.

- `raw` — no data conversion takes place between solidDB server and the client
The value `raw` can be used when you want your database to use the binding used in the 6.3 or earlier versions of solidDB.
- `locale` — the current client locale setting is used, also if set by the client system
- `locale:` — the current client setting are overridden with a default locale set of the client system

The driver calls `setlocale()` with an empty string which effectively searches for the locale setting set in the system.

For example, in Linux environments, the environmental variable `LC_CTYPE` is checked first and if that is not defined, the environmental variable `LANG` is searched.

- `locale:<locale name>` — the current client systems setting are overridden and the given locale is used

The convention for `<locale name>` depends on the operating system.

For example, in Linux environments, the locale name for the code page GB18030 in Chinese/China is `zh_CN.gb18030`. In Windows environments, the locale name for Latin1 code page in Finnish/Finland is `fin_fin.1252`.

- UTF8 — UTF-8 binding is enforced regardless of the locale set in the client-side system

Note:

- If the value in `Srv.ODBCDefaultCharBinding` is other than `locale`, it overrides any current system locale setting for all clients.
- If the value in `Client.ODBCCharBinding` is other than `locale`, it overrides both the server side value (if set) and the current system locale setting.

Using the current client locale encoding (`locale`)

To use the current client locale encoding:

1. **Configure the parameter setting:**

- **All clients use the same binding method (server-side parameter)**

In the server-side `solid.ini`, section `[Server]`, set the `ODBCDefaultCharBinding` parameter.

```
[Srv]
ODBCDefaultCharBinding=locale
```

- **Some or all clients require different binding methods (client-side parameter)**

In the client-side `solid.ini`, section `[Client]`, set the `ODBCCharBinding` parameter.

```
[Client]
ODBCCharBinding=locale
```

The client-side parameter overrides the server-side settings.

2. **Set application to call `setlocale()`.**

Using a specific locale encoding (`locale:<locale_name>`)

To use a specific locale encoding:

Define the locale in `solid.ini`:

- **All clients use the same binding method (server-side parameter)**

In the server-side `solid.ini`, section `[Server]`, set the `ODBCDefaultCharBinding` parameter.

```
[Srv]
ODBCDefaultCharBinding=locale:<locale name>
```

For example in Linux environments:

```
[Srv]
ODBCDefaultCharBinding=locale:zh_CN.gb18030
```

- **Some or all clients require different binding methods (client-side parameter)**

In the client-side `solid.ini`, section `[Client]`, set the `ODBCCharBinding` parameter.

```
[Client]
ODBCCharBinding=locale:<locale name>
```

The client-side parameter overrides the server-side settings.

For example in Linux environments:

```
[Client]
ODBCCharBinding=locale:zh_CN.gb18030
```

Note: Setting a specific locale overrides the application settings defined with `setlocale()`.

Example 1

All clients use the current locale of the client. Different clients can use different code pages.

The server-side `solid.ini` is used:

```
[Srv]
ODBCDefaultCharBinding=locale
```

Example 2

Some clients use the current locale of the client, some clients use Latin1 code page:

The server-side `solid.ini` is used:

```
[Srv]
ODBCDefaultCharBinding=locale
```

In those clients that require the Latin1 code page, the client-side `solid.ini` is used:

```
[Client]
ODBCCharBinding=locale:fin_fin.1252
```

5.5.2 JDBC applications and Unicode databases

In JDBC environments, the `solidDB` JDBC driver handles the conversion of data between the encoding used in the application (client) and the UTF-8/UTF-16 format in the Unicode database. You do not need to make any `solidDB` specific settings to use Unicode with JDBC.

6 Using Transaction Log Reader

The solidDB Transaction Log Reader is a solution that makes it possible to read log records from the solidDB transaction log transaction by transaction. Using the Log Reader interface, you can, for example, write an application that listens to and displays log traffic in the solidDB server.

The Log Reader is based on a read-only system table called *SYS_LOG* where each row corresponds to a single log entry. 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. Each log read can be started from different log record.

For each entry in the transaction log, the *SYS_LOG* table contains data for identifying the log record, the type of transaction and statement executed, as well as the row with the changed data itself.

The *SYS_LOG* table can be accessed with ODBC and JDBC drivers using SQL statements. For example, an application could be written to read the solidDB transaction log and extract records pertaining to SQL DML statements. The application could then reconstruct the statements into plain text SQL strings and print them to a desired type of output.

Applications can read the *SYS_LOG* table both locally and remotely. Several applications can read the *SYS_LOG* table concurrently without interference.

A sample application that demonstrates the use of the Log Reader interface is included in the solidDB package, available in `samples/logreader` directory.

For a detailed description of the *SYS_LOG* table, see the section *SYS_LOG* in the Appendix *Database virtual tables* in the *IBM solidDB SQL Guide*.

6.1 Considerations for developing applications with Log Reader

Supported table types

- Both in-memory and disk-based tables are supported.
- Transient and temporary tables are not supported.
Transient and temporary tables are not logged and thus data in them is not returned through the Log Reader.

Supported database operations

- Only committed transactions are returned by the Log Reader.
All events for one transaction are returned at once. Each transaction is returned in full at the time of commit. Log Reader returns full committed transactions in the order they are committed to the log. Overlapping transactions are returned transaction by transaction.
- Triggers are supported.
Operations in the trigger action part are logged as normal user operations. Write operations in triggers are logged when they are executed: before-triggers are logged before user data operation, and after-triggers are logged after user data operation.

- Cascading actions are supported.
Operations resulted from cascading referential actions are logged as normal user operations. Cascading operations are logged after the actual user data operation.
- DDL operations are supported.
For DDL operations, the Log Reader will return a special DBE_LOGREADER_LOG_REC_DDL record that contains the original SQL statement.

Catchup and live data modes

When a log read is started from SYS_LOG, the read goes first into *catchup* mode. The catchup mode means that the log read start position is searched from the log and the read is then started from that position. When log read reaches the end of the current log, it starts to read *live* data. In *live* mode transactions are returned as they are executed.

When several Log Readers are used, each Log Reader has its own data mode.

In live data mode the cursor returns every second even if there is no data available. In that case, the FLAGS field in the SYS_LOG table is zero.

Primary keys

Primary key is not mandatory for the tables. System generated internal and hidden primary key value is not returned through the Log Reader.

When designing your database, you must decide how rows are identified if primary key is not defined.

High Availability

High Availability (HotStandby) is supported so that the log file contents and log address are compatible between Primary and Secondary servers. If the log is read from Primary server and there is a failover, a new read from SYS_LOG can be started using the last LOGADDR received from the old Primary server.

The log can be read also from the Secondary server. This can be useful, for example, for load balancing reasons.

Throttling

If the server can generate log records faster than the client can read, throttling can occur. This means that user transactions writing to the server are slowed down to make sure the Log Readers are not too much behind the live data. The **LogReader.MaxSpace** parameter can be used to control the buffering after which throttling can occur.

Applications that start a read but then stop reading can also cause the server to stop.

Log maximum size

If the application using the Log Reader is stopped or terminated for longer time, the log maximum size may be reached. In such a case, no error message is

produced. Also, the position stored by the application for catchup is not available any more and the effort to catchup fails.

Access rights

Administrator rights are needed to access the SYS_LOG table, to add tables to a partition or to remove tables from a partition.

Stopping the Log Reader

The log reading can be stopped at any time. No data is lost even if some undelivered data are left in the log.

The log reading may be resumed without any loss of information if the last read position is known. By using the current log position, the application will be able to continue reading the log, without any loss of data, upon the next SYS_LOG query. If the SYS_LOG table is accessed without specifying the log position, the reading starts from the live data.

The Log Reader can be stopped in the solidDB server with the ADMIN COMMAND 'LOGREADER STOP' command.

6.2 Configuring the Log Reader

The Log Reader is configured with the server-side configuration parameters in the LogReader section of the solid.ini configuration file.

About this task

Important: The **LogReaderEnabled**, **MaxSpace**, and **MaxLogSize** parameters are also used with solidDB Universal Cache and InfoSphere™ CDC Replication.

Procedure

- Enable the Log Reader by setting **LogReaderEnabled** to 'yes'.
This enables the Log Reader, allowing reads from SYS_LOG. The transaction logging mode is also more verbose.
- As necessary for your environment, set the following parameters:
 - Set **MaxSpace** value to define the maximum number of log records buffered into memory before throttling occurs.
 - Set **MaxLogSize** value to define the maximum size of the log available for a catchup.
When the log reaches the defined size, old log data is deleted and catchup is not possible from the older LOGADDR log positions.
 - Set **MaxMemLogSize** value to define the maximum size of the Log Reader logfile in memory, when logging is not enabled (**Logging.LogEnabled=No**).

Example

```
[LogReader]
LogReaderEnabled=yes ;Default: no
;
;MaxLogSize=100000 ;default: 10240 (MB)
; The amount of the log files (in MB) that will be always maintained
; for the sake of a possible catchup. The size should be adjusted to
; the biggest size of a catchup that is reasonable. The space declared
; is always fully occupied.
```

```

;
MaxSpace=500000          ;default:100000
; The size of the in-memory log reader buffer used in throttling,
; in records. When the buffer fills up, the throttling (slowing down)
; is enacted. If the buffer is used, the size adds up to the
; footprint of the solidDB server process.

```

6.3 Reading log data with the Log Reader

The solidDB transaction log can be read with a Log Reader specific SELECT statement.

About this task

The transaction log can be read by multiple concurrently active SELECT statements at the same time. Each log read can be started from different log position, independently from the others.

Tip:

The solidDB package contains a sample application that demonstrates the use of the Log Reader interface. The sample application is available in the `samples/logreader` directory in the solidDB installation directory.

Procedure

1. Read log data from the SYS_LOG table with a SELECT statement.

The basic syntax to read log is:

```
SELECT RECID, RELID, FLAGS, LOGADDR, DATA FROM SYS_LOG WHERE LOGADDR > ?;
```

The WHERE condition is allowed only for LOGADDR field. Only constraint that is allowed is *greater than* (>). Constraints to other fields will result in an error.

Alternatively, you can start the log read from a specified log position, defined with the LOGADDR field.

- a. Retrieve the current LOGADDR value.

```
SELECT LOGADDR FROM SYS_LOG LIMIT 1;
```

For example:

```
SELECT LOGADDR FROM SYS_LOG LIMIT 1;
LOGADDR
```

```
-----
```

```
0000000000000001FFFFFFFF0000029500000295
```

```
1 rows fetched.
```

- b. Define the LOGADDR in the SELECT statement.

For example:

```
SELECT RECID,RELID,FLAGS,LOGADDR,DATA FROM SYS_LOG WHERE LOGADDR > '0000000000000001FFFFFFFF0000029500000295';
```

ResultWhen a read is started, the fetch calls will start returning rows. The user data for the log record is included in the DATA column in binary format.

When no data is available, the server will return empty data in one second intervals. In such a case, you can issue a new fetch call to see if new data is available.

For more details on the DATA column and other columns in the SYS_LOG table, see *SYS_LOG table definition* in the *IBM solidDB SQL Guide*.

2. Reconstruct the user data in the rows returned by the Log Reader.

Note: The exact steps for handling the log records depends on the application design. The following steps are the basic steps that are needed to output the user data in a format that can be used to process the log records further.

- a. Get metadata for the column by querying the solidDB system tables.

For example, in Java environments, the column metadata can be read from the solidDB system tables using the `ResultSet.getMetaData()` call.

- b. Using the metadata, parse the data in the DATA column to produce output in the format of your choice.

For example, in the sample application (`samples/logreader`), the log records are converted into plain text SQL statements.

6.4 Partitioning and filtering log records

By default, all the log records generated by any application are returned. If you only want to access the log records for a subset of the database, you can specify log reader partitions. A Log Reader partition is a named collection of tables. The partitions may overlap.

information about the partitions is stored in `SYS_FEDT_DB_PARTITION` and `SYS_FEDT_TABLE_PARTITION` system tables.

6.4.1 Creating and deleting partitions

The Log Reader partitions are be created, modified, and deleted with SQL statements.

About this task

The partition settings are transactional and persistent.

Procedure

- **Creating partitions**

Create partitions with the following command:

```
CREATE LOGREADER PARTITION <partition-name>
```

- **Deleting partitions**

Delete (drop) partitions with the following command:

```
DROP LOGREADER PARTITION <partition-name>
```

- **Modifying partitions**

Add or remove tables from a partition with the following command:

```
ALTER LOGREADER PARTITION <partition-name> {ADD | DROP} TABLE <table-name>
```

Note: When a table is part of a partition, it cannot be dropped or altered. The only ALTER statement available is `ALTER LOGREADER PARTITION ... DROP TABLE`.

6.4.2 Using partition filters

A session-specific partition filter can be set to read log records only from a specific partition.

About this task

The setting for the partition filter applies only to the current session and is valid for all reads from `SYS_LOG` that are started after the setting. The partition data is

stored in the SYS_FEDT_DB_PARTITION system table.

Procedure

- Set partition filters with the following statement:

```
SET LOGREADER PARTITION { <partition-name> | NONE }
```

When set to NONE (default), all log records are read.
- View the existing partitions in the SYS_FEDT_DB_PARTITION system table with a SELECT statement.
For example:

```
SELECT * FROM sys_fedt_db_partition
```

6.5 Setting transaction batches

With transaction batching, multiple transactions from the log can be returned as a single transaction. Transaction batch sizes are set for a session. If the logged transactions are being fed into another database, transaction batching can improve read performance over the network. For example, a number of inserts could be batched into one transaction so that when the transactions need to be executed in another database, only one transaction needs to be executed.

About this task

The transaction batch setting applies only to the current session and is valid for all reads from SYS_LOG that are started after the setting.

Procedure

Set transaction batch size with the following command:

```
SET LOGREADER BATCH <size>
```

The default batch size is 1: no batching of transactions is done.

Results

Setting the batch transaction size does not alter the catchup position; the catchup position can be used to read all transaction in the batch again.

Appendix A. solidDB supported ODBC functions

This topic describes the ODBC functions supported by solidDB.

Table 142. solidDB supported ODBC functions

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
<i>Connecting to a Data Source</i>			
SQLAllocEnv (1.0)	N/A	Deprecated (replaced by SQLAllocHandle)	N/A
SQLAllocConnect (1.0)	N/A	Deprecated (replaced by SQLAllocHandle)	N/A
SQLAllocHandle (3.0)	Returns the list of supported data source attributes. Returns the list of installed drivers and their attributes.	Supported Supported	ISO 92 ODBC
SQLConnect (1.0)	Establishes connections to a driver and a data source. The connection handle references storage of all information about the connection to the data source, including status, transaction state, and error information.	Supported	ISO 92
SQLDriverConnect (1.0)	This function is an alternative to SQLConnect. It supports data sources that require more connection information than the three arguments in SQLConnect, including dialog boxes to prompt the user for all connection information, and data sources that are not defined in the system information.	Supported (including Unicode version of this function).	ODBC
SQLBrowseConnect (1.0)	Returns successive levels of attributes and attribute values. When all levels have been enumerated, a connection to the data source is completed and a complete connection string is returned. A return of SQL_SUCCESS_WITH_INFO indicates that all connection information has been specified and the application is now connected to the data source.	Not supported	ISO 92
SQLGetInfo (1.0)	Returns general information about the driver and data source associated with a connection.	Supported	ISO 92

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
SQLGetFunctions (1.0)	Returns information about whether a driver supports a specific ODBC function.	Supported; this function is implemented in the ODBC Driver Manager. It can also be implemented in drivers. If a driver implements SQLGetFunctions, the Driver manager calls the function in the driver. Otherwise, it executes the function itself. In the case of solidDB, the function is implemented in the driver so that the application linked to the driver can also call this function from the application.	ISO 92
SQLGetTypeInfo (1.0)	Returns information about data types supported by the data source. The driver returns the information in the form of an SQL result set. The data types are intended for use in Data Definition Language (DDL) statements.	Supported	ISO 92
<i>Obtaining Information about a Driver and Data Source</i>			
SQLDataSources (1.0)	Returns information about a data source.	Supported; this function is implemented in the ODBC Driver Manager. For platforms other than Windows which do not have the Microsoft ODBC Driver manager, this function is not supported.	ISO 92
SQLDrivers (2.0)	Lists driver descriptions and driver attribute keywords.	Supported; this function is implemented in the ODBC Driver Manager. For Windows, the Driver Manager is required if applications that connect to solidDB use OLE DB or ADO APIs or if database tools that require the Driver Manager, such as Microsoft Access, FoxPro, or Crystal Reports are to be used. For platforms other than Windows, the Driver Managers are provided by vendors such as iODBC, Merant, or UnixODBC.	ODBC
SQLGetConnectAttr (3.0)	Returns the value of a connection attribute.	Supported	ISO 92
SQLSetConnectAttr (3.0)	Sets a connection attribute.	Supported	ISO 92

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
SQLGetEnvAttr (3.0)	Returns the value of an environment attribute.	Supported	ISO 92
SQLSetEnvAttr (3.0)	Sets an environment attribute.	Supported	ISO 92
SQLGetStmtAttr (3.0)	Returns the value of a statement attribute.	Supported	ISO 92
SQLSetStmtAttr (3.0)	Sets a statement attribute.	Supported	ISO 92
SQLSetConnectOption (1.0)	N/A	Deprecated (replaced by SQLSetConnectAttr)	N/A
SQLGetConnectOption (1.0)	N/A	Deprecated (replaced by SQLGetConnectAttr)	N/A
SQLGetStmtOption (1.0)	N/A	Deprecated (replaced by SQLGetStmtAttr)	N/A
SQLSetStmtOption (1.0)	N/A	Deprecated (replaced by SQLSetStmtAttr)	N/A
<i>Setting and Retrieving Descriptor Fields</i>			
SQLGetDescField (3.0)	Returns the current setting or value of a single descriptor field.	Supported	ISO 92
SQLSetDescField (3.0)	Sets the value of a single field of a descriptor record.	Supported	ISO 92
SQLGetDescRec (3.0)	Returns the current settings or values of multiple fields of a descriptor record. The fields returned describe the name, data type, and storage of column or parameter data.	Supported	ISO 92
SQLSetDescRec (3.0)	Sets multiple descriptor fields that affect the data type and buffer bound to a column or parameter data.	Supported	ISO 92
SQLCopyDesc (3.0)	Copies descriptor information from one descriptor handle to another.	Supported	ISO 92
<i>Preparing SQL Requests</i>			
SQLAllocStmt (1.0)	N/A	Deprecated (replaced by SQLAllocHandle)	N/A
SQLPrepare (1.0)	Prepares an SQL statement for later execution.	Supported	ISO 92
SQLBindParameter (2.0)	Assigns storage for a parameter in an SQL statement.	Supported Note: This function replaces SQLBindParam which did not exist in ODBC 2.x, although it is in the X/Open and ISO standards.	ODBC

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
SQLGetCursorName (1.0)	Returns the cursor name associated with a statement handle.	Supported	ISO 92
SQLSetCursorName (1.0)	Specifies a cursor name with an active statement. If an application does not call SQLSetCursorName, the driver generates cursor names as needed for SQL statement processing.	Supported	ISO 92
SQLParamOptions (1.0)	N/A	Deprecated (replaced by SQLSetStmtAttr)	N/A
SQLSetParam (1.0)	N/A	Deprecated (replaced by SQLBindParameter)	N/A
SQLSetScrollOptions (1.0)	Sets options that control cursor behavior.	Deprecated (replaced by SQLGetInfo and SQLSetStmtAttr)	ODBC
<i>Submitting Requests</i>			
SQLExecute (1.0)	Executes a prepared statement using the current values of the parameter marker variables if any parameter markers exist in the statement.	Supported	ISO 92
SQLExecDirect (1.0)	Executes a preparable statement using the current values of the parameter marker variables if any parameters exist in the statement. SQLExecDirect is the fastest way to submit an SQL statement for one-time execution.	Supported	ISO 92
SQLNativeSQL (1.0)	Returns the SQL string as modified by the driver. SQLNativeSQL does not execute the SQL statement.	Not implemented; solidDB does not support this functionality.	N/A
SQLDescribeParam (1.0)	Returns the text of an SQL statement as translated by the driver. This information is also available in the fields of the IPD.	Supported	ODBC
SQLNumParams (1.0)	Returns the number of parameters in an SQL statement.	Supported	ISO 92
SQLParamData (1.0)	Used in conjunction with SQLPutData to supply parameter data at execution time. (Useful for long data values.)	Supported	ISO 92
SQLPutData (1.0)	Allows an application to send data for a parameter or column to the driver at statement execution time. This function can be used to send character or binary data values in parts to a column with a character, binary, or data source-specific data type (for example, parameters of the SQL_LONGVARBINARY or SQL_LONGVARCHAR types).	Supported	ISO 92

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
<i>Retrieving Results and Information about Results</i>			
SQLRowCount (1.0)	Returns the number of rows affected by an UPDATE, INSERT, or DELETE statement.	Supported	ISO 92
SQLNumResultCols (1.0)	Returns the number of columns in a result set.	Supported	ISO 92
SQLDescribeCol (1.0)	<p>Returns the result descriptor (column name, type, column size, decimal digits, and nullability) for one column in the result set. This information is also available in the fields of the IRD.</p> <p>NOTE: The driver now returns the number of characters instead of the number of bytes for the following attributes: SQL_DESC_LABEL, SQL_DESC_NAME, SQL_DESC_SCHEMA_NAME, SQL_DESC_CATALOG_NAME, SQL_DESC_BASE_COLUMN_NAME, and SQLDESC_BASE_TABLE_NAME</p> <p>This conforms more closely to the ODBC standard and works correctly using ADO, VB, OLE-DB, and ODBC calls. Note, however, that this causes failure of the Microsoft Visual DataBase Project. After updating/inserting the record, the record is not saved and the following error is displayed: "the table does not exist."</p>	Supported.	ISO 92
SQLColAttributes (1.0)	N/A	Deprecated (replaced by SQLColAttribute)	N/A
SQLColAttribute (3.0)	<p>Describes attributes of a column in the result set.</p> <p>Note: The driver now returns the number of characters instead of the number of bytes for the following attributes: SQL_DESC_LABEL, SQL_DESC_NAME, SQL_DESC_SCHEMA_NAME, SQL_DESC_CATALOG_NAME, SQL_DESC_BASE_COLUMN_NAME, and SQLDESC_BASE_TABLE_NAME</p> <p>This conforms more closely to the ODBC standard and works correctly using ADO, VB, OLE-DB, and ODBC calls. Note, however, that this causes failure of the Microsoft Visual DataBase Project. After updating/inserting the record, the record is not saved and the following error is displayed: "the table does not exist."</p>	Supported.	ISO 92
SQLBindCol (1.0)	Assigns storage for a result column and specifies the data type.	Supported	ISO 92
SQLFetch (1.0)	Returns multiple result rows, fetching the next rowset of data from the result set and returning data for all bound columns.	Supported	ISO 92

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
SQLExtendedFetch (2.0)	N/A	Replaced by SQLFetchScroll	N/A
SQLFetchScroll (3.0)	Returns scrollable result rows, fetching the specified rowset of data from the result set and returning data for all bound columns. Block cursor support enables an application to fetch more than one row with a single fetch into the application buffer. When working with an ODBC 2.x driver, the Driver Manager maps this function to SQLExtendedFetch.	Supported Note: Since the solidDB ODBC Driver currently has no support for bookmarks, it is not possible to support the SQL_FETCH_BOOKMARK option in SQLFetchScroll.	ISO 92
SQLGetData (1.0)	Returns part or all of one column of one row of a result set. It can be called multiple times to retrieve variable length data in parts, making it useful for long data values.	Supported	ISO 92
SQLSetPos (1.0)	Positions a cursor within a fetched block of data and allows an application to refresh data in the rowset or to update or delete data in the result set.	Supported, along with all the options, that is, SQL_POSITION, SQL_DELETE, and SQL_UPDATE	ODBC
SQLBulkOperations (3.0)	Performs bulk insertions and bulk bookmark operations, including update, delete, and fetch by bookmark.	solidDB supports this, but only when using the SQL_ADD option.	ODBC
SQLMoreResults (1.0)	Determines whether there are more results available on a statement containing SELECT, UPDATE, INSERT, or DELETE statement and, if so, initializes processing for those results.	Not supported solidDB does not support multiple results.	ODBC
SQLGetDiagField (3.0)	Returns additional diagnostic information (a single field of the diagnostic data structure associated with a specified handle). This information includes error, warning, and status information.	Supported	ISO 92
SQLGetDiagRec (3.0)	Returns additional diagnostic information (multiple fields of the diagnostic data structure). Unlike SQLGetDiagField, which returns one diagnostic field per call, SQLGetDiagRec returns several commonly used fields of a diagnostic record, including the SQLSTATE, the native error code, and the diagnostic message text.	Supported	ISO 92
SQLError (1.0)	N/A	Deprecated (replaced by SQLGetDiagRec)	N/A
<i>Obtaining Information about the Data Source's System Tables</i>			

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
SQLColumnPrivileges (1.0)	Returns a list of columns and associated privileges for the specified table. The driver returns the information as a result set on the specified StatementHandle. This function is supported via an appropriate SQL execution.	Supported	ODBC
SQLColumns (1.0)	Returns a list of columns and associated privileges for the specified table. The driver returns the information as a result set on the specified StatementHandle. This function is supported via an appropriate SQL execution.	Supported	X/Open
SQLForeignKeys (1.0)	Returns two type of lists: <ul style="list-style-type: none"> • Foreign keys in the specified table (columns in the specified table that refer to primary keys in other tables). • Foreign keys in other tables that refer to the primary key in the specified table. The driver returns each list as a result set on the specified statement.	Supported	ODBC
SQLPrimaryKeys (1.0)	Returns the list of column names that make up the primary key for a table. The driver returns the information as a result set. This function does not support returning primary keys from multiple tables in a single call.	Supported	ODBC
SQLProcedureColumns (1.0)	Returns the list of input and output parameters, as well as the columns that make up the result set for the specified procedures. The driver returns the information as a result set on the specified statement.	Supported.	ODBC
SQLProcedures (1.0)	Returns the list of procedure names stored in a specific data source. Procedure is a generic term used to describe an executable object, or a named entity that can be invoked using input and output parameters.	Supported	ODBC
SQLSpecialColumns (1.0)	Returns the following information about columns within a specified table: <ul style="list-style-type: none"> • The optimal set of columns that uniquely identifies a row in the table. • Columns that are automatically updated when any value in the row is updated by a transaction. 	Supported	X/Open
SQLStatistics (1.0)	Returns statistics about a single table and the list of indexes associated with the table. The driver returns the information as a result set.	Supported	ISO 92

Table 142. solidDB supported ODBC functions (continued)

Function Names/Version Introduced ¹	Purpose	Availability when using ODBC	Conformance ²
SQLTablePrivileges (1.0)	Returns a list of tables and the privileges associated with each table. The driver returns the information as a result set on the specified statement.	Supported	ODBC
SQLTables (1.0)	Returns the list of table, catalog, or schema names, and table types, stored in a specific data source.	Supported	X/Open
<i>Terminating a statement</i>			
SQLFreeStmt (1.0)	Ends statement processing, discards pending results, and optionally, frees all resources associated with the statement handle.	Supported Note: The SQLFreeStmt with an option of SQL_DROP is replaced by SQLFreeHandle.	ISO 92
SQLCloseCursor (3.0)	Closes a cursor that has been opened on a statement, and discards pending results.	Supported	ISO 92
SQLCancel (1.0)	Cancels the processing on an SQL statement.	Supported	ISO 92
SQLEndTran (3.0)	Requests a transaction commit or rollback on all statements associated with a connection. SQLEndTran can also request that a commit or rollback operation be performed for all connections associated with an environment.	Supported	ISO 92
SQLTransact (1.0)	N/A	Deprecated (replaced by SQLEndTran)	N/A
<i>Terminating a Connection</i>			
SQLDisconnect (1.0)	Closes the connection associated with a specific connection handle.	Supported	ISO 92
SQLFreeConnect (1.0)	N/A	Deprecated (replaced by SQLFreeHandle)	N/A
SQLFreeEnv (1.0)	N/A	Deprecated (replaced by SQLFreeHandle)	N/A
SQLFreeHandle (3.0)	Frees resources associated with a specific environment, connection, statement, or descriptor handle	Supported	ISO 92

¹ Version introduced is the version when the function was first added to the ODBC API.

² Conformance level can be:

- ISO 92 (also appears in X/Open version 1 because X/Open is a pure superset of ISO 92)
- X/Open (also appears in ODBC 3.x because ODBC 3.x is a pure superset of X/Open version 1)

- ODBC (appears in neither ISO 92 or X/Open)
- N/A (Deprecated in ODBC 3.x)

Appendix B. solidDB ODBC Driver 3.5.1 attributes support

This topic provides information about the solidDB ODBC Driver 3.5.1 attributes.

The attributes are grouped in the following categories:

- Environment-level attributes
- Connection-level attributes
- Statement-level attributes
- Column-level attributes

Table 143. 001 Environment Level

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_CONNECTION_POOLING	SQL_CP_OFF SQL_CP_ONE_PER_DRIVER SQL_CP_ONE_PER_HENV	All values are supported	All values are not applicable to the driver	All values are not applicable to ODBC drivers, handled by Driver Manager, so this attribute will be supported if the application links to ODBC DM and can't be simulated by the driver itself.
SQL_ATTR_CP_MATCH	SQL_CP_STRICT_MATCH SQL_CP_RELAXED_MATCH	All values supported	All values are not applicable to the driver	All values are not applicable to ODBC drivers, handled by Driver Manager, so this attribute will be supported if the application links to ODBC DM and can't be simulated by the driver itself.
SQL_ATTR_ODBC_VERSION	SQL_OV_ODBC3 SQL_OV_ODBC2	Supported Not Supported	Supported Not Supported	Allows user to set and get the version to 2, but the behavior is a per 3.0 and above.
SQL_ATTR_OUTPUT_NTS	SQL_TRUE SQL_FALSE	Supported Not Supported	Supported Not Supported	

Table 144. 002 Connection Level

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_ODBC_CURSORS	SQL_CUR_IF_NEEDED SQL_FETCH_PRIOR SQL_CUR_USE_ODBC SQL_CUR_USE_DRIVER	All values not supported	All values not supported	
SQL_ATTR_ACCESS_MODE	SQL_MODE_READ_ONLY SQL_MODE_READ_WRITE	All values not supported	All values not supported	
SQL_ATTR_ASYNC_ENABLE	SQL_ASYNC_ENABLE_OFF SQL_ASYNC_ENABLE_ON	All values not supported	All values not supported	
SQL_ATTR_AUTO_IPD	SQL_TRUE SQL_FALSE	All values not supported	All values not supported	
SQL_ATTR_AUTOCOMMIT	SQL_ATTR_AUTOCOMMIT_OFF SQL_ATTR_AUTOCOMMIT_ON	All values are supported	All values are supported	
SQL_ATTR_CONNECTION_TIMEOUT	Timeout value in sec	Supported	Supported	
SQL_ATTR_CURRENT_CATALOG	CatalogName	Supported	Supported	
SQL_ATTR_LOGIN_TIMEOUT	Timeout value in sec	Supported	Supported	
SQL_ATTR_METADATA_ID	SQL_TRUE SQL_FALSE	All values not supported	All values not supported	
SQL_ATTR_PACKET_SIZE	Packet size in bytes	the desired size	Not supported	
SQL_ATTR_QUIET_MODE	Set to NULL	can set and get	Not supported	
SQL_ATTR_TRACE	SQL_TRACE_OFF SQL_TRACE_ON	All values supported	All values not supported	All values handled by DM, not by driver

Table 144. 002 Connection Level (continued)

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_TRACEFILE	Pointer to trace file name	Supported	Not supported	Handled by DM, not by driver
SQL_ATTR_TRANSLATE_LIB	Pointer Name of lib	Supported	Not supported	Handled by DM, not by driver
SQL_ATTR_TXN_ISOLATION	SQL_TXN_SERIALIZABLE SQL_TXN_READ_UNCOMMITTED SQL_TXN_READ_COMMITTED SQL_TXN_REPEATABLE_READ	All values are supported, except SQL_TXN_READ_UNCOMMITTED	All values are supported, except SQL_TXN_READ_UNCOMMITTED	An solidDB server does not support the READ_UNCOMMITTED feature.

Table 145. 03 Statement Level

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_CONCURRENCY	SQL_CONCUR_READ_ONLY SQL_CONCUR_LOCK SQL_CONCUR_ROWVER SQL_CONCUR_VALUES	All values supported	All values supported	For the value SQL_CONCUR_READ_ONLY, set and get are supported. For all other values, set is supported and get returns READ_ONLY.
SQL_ATTR_CURSOR_TYPE	SQL_CURSOR_FORWARD_ONLY SQL_CURSOR_KEYSET_DRIVEN SQL_CURSOR_DYNAMIC SQL_CURSOR_STATIC	Supported Forced to dynamic Forced to dynamic Forced to dynamic	Supported Forced to dynamic Forced to dynamic Forced to dynamic	
SQL_ATTR_MAX_LENGTH	Length in bytes	Not supported	Not supported	Whatever the length, sets only to default (0).
SQL_ATTR_MAX_ROWS	Maximum number of rows	Not supported	Not supported	Whatever the length, sets only to default (0).
SQL_ATTR_RETRIEVE_DATA	SQL_RD_OFF SQL_RD_ON	Not supported Supported	Not supported Supported	Sets to SQL_RD_ON only

Table 145. 03 Statement Level (continued)

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_USE_BOOKMARKS	SQL_UB_OFF SQL_UB_ON	All values not supported	All values not supported	
SQL_ATTR_ROW_ARRAY_SIZE	An SQLUSMALLINT* value that points to an array of SQLUSMALLINT values containing row status values after a call to SQLFetch or SQLFetchScroll.	Supported	Supported	The array has as many elements as there are rows in the rowset.
SQL_ATTR_ROWS_FETCHED_PTR	An SQLUIINTEGER* value that points to a buffer in which to return the number of rows fetched after a call to SQLFetch or SQLFetchScroll.	Supported	Supported	
SQL_ATTR_ROW_STATUS_PTR	An SQLUIINTEGER value that specifies the number of rows returned by each call to SQLFetch or SQLFetchScroll.	Supported	Supported	
SQL_ROWSET_SIZE	Number of rows to return	Supported	Supported	Allows the ODBC application to set its value to greater than 1.
SQL_ASYNC_ENABLE	SQL_ASYNC_ENABLE_ON SQL_ASYNC_ENABLE_OFF	All values not supported	All values not supported	
SQL_BIND_TYPE	SQL_BIND_BY_COLUMN	Not supported	Not supported	
SQL_ATTR_KEYSET_SIZE	Size	Not supported	Not supported	Whatever the size, sets only to default (0).
SQL_ATTR_NOSCAN	SQL_NOSCAN_OFF SQL_NOSCAN_ON	Not supported Not supported	Not supported Not supported	Sets to SQL_NOSCAN_OFF only
SQL_ATTR_SIMULATE_CURSOR	SQL_SC_NON_UNIQUE SQL_SC_TRY_UNIQUE SQL_SC_UNIQUE	All values not supported	All values not supported	All values are not relevant to solidDB Driver

Table 145. 03 Statement Level (continued)

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_APP_PARAM_DESC	SQL_NULL_HDESC	Not supported	Not supported	
SQL_ATTR_APP_ROW_DESC	SQL_NULL_HDESC	Not supported	Not supported	
SQL_ATTR_CURSOR_SCROLLABLE	SQL_SCROLLABLE SQL_NONSCROLLABLE	Not supported Not supported	Not supported Not supported	Sets to SQL_NONSCROLLABLE only
SQL_ATTR_CURSOR_SENSITIVITY	SQL_UNSPECIFIED SQL_INSENSITIVE SQL_SENSITIVE	Not supported Not supported Not supported	Not supported Not supported Not supported	Sets to SQL_UNSPECIFIED only Sets to SQL_UNSPECIFIED only
SQL_ATTR_ROW_NUMBER	Number of current row	Supported	Supported	User can get the number of rows; cannot set because of read-only property
SQL_ATTR_ENABLE_AUTO_IPD	SQL_TRUE SQL_FALSE	Both values not supported	Both values not supported	
SQL_ATTR_METADATA_ID	SQL_TRUE SQL_FALSE	Both values not supported	Both values not supported	
SQL_ATTR_PARAM_BIND_OFFSET_PTR	SQL_DESC_DATA_PTR SQL_DESC_INDICATOR_PTR SQL_DESC_OCTET_LENGTH_PTR SQL_DESC_BIND_OFFSET_PTR	All values are supported	All values are supported	
SQL_ATTR_PARAM_OPERATION_PTR	Pointer to array containing list of parameters to be ignored	Supported	Supported	

Table 145. 03 Statement Level (continued)

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_ATTR_PARAMS_PROCESSED_PTR	Unsigned integer pointer to return the number of sets of parameters that have been processed by the SQL statement executed through SQLExecute or SQLExecDirect.	Supported	Supported	

Table 146. 04 Column Attributes

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_DESC_BASE_COLUMN_NAME		Supported	Supported	
SQL_DESC_BASE_TABLE_NAME		Supported	Supported	
SQL_DESC_DISPLAY_SIZE		Supported	Supported	
SQL_DESC_NAME				
SQL_DESC_NULLABLE		Supported	Supported	
SQL_DESC_OCTET_LENGTH		Supported	Supported	
SQL_DESC_PRECISION		Supported	Supported	
SQL_DESC_SCALE		Supported	Supported	
SQL_DESC_UPDATABLE		Supported	Supported	
SQL_DESC_FIXED_PREC_SCALE		Supported	Supported	
SQL_DESC_TABLE_NAME		Supported	Supported	
SQL_DESC_TYPE		Supported	Supported	
SQL_DESC_UNNAMED		Supported	Supported	
SQL_DESC_SCHEMA_NAME		Supported	Supported	
SQL_DESC_LOCAL_TYPE_NAME		Supported	Supported	
SQL_DESC_LABEL		Supported	Supported	
SQL_DESC_TYPE_NAME		Supported	Supported	
SQL_DESC_AUTO_UNIQUE_VALUE		Supported	Supported	

Table 146. 04 Column Attributes (continued)

Attribute	Value (Option)	Driver Manager	Driver Alone	Comments
SQL_DESC_CONCISE_TYPE		Supported	Supported	
SQL_DESC_LITERAL_PREFIX		Supported	Supported	
SQL_DESC_UNSIGNED		Supported	Supported	
SQL_DESC_LITERAL_PREFIX		Supported	Supported	
SQL_DESC_UNSIGNED		Supported	Supported	
SQL_DESC_LITERAL_SUFFIX		Supported	Supported	
SQL_DESC_CATALOG_NAME		Supported	Supported	
SQL_DESC_COUNT		Supported	Supported	
SQL_DESC_SEARCHABLE		Supported	Supported	
SQL_DESC_LENGTH		Supported	Supported	
SQL_DESC_CASE_SENSITIVE		Supported	Supported	
SQL_DESC_NUM_PREX_RADIX		Supported	Supported	

Appendix C. SQLSTATE error codes

This topic contains an error codes table that provides possible SQLSTATE values that a driver returns for the SQLGetDiagRec function.

Note: The SQLGetDiagRec and SQLGetDiagField return SQLSTATE values that conform to the X/Open (The Open Group) *Data Management: Structured Query Language (SQL), Version 2 (3/95)*.

Error codes table convention

Table 147. Error code class values

Class value	Meaning
01	Indicates a warning and includes a return code of SQL_SUCCESS_WITH_INFO. Note: Error class 01 returns both warnings and errors.
01, 07, 08, 21, 22, 23, 24, 25, 28, 34, 3C, 3D, 3F, 40, 42, 44, HY	Indicates an error that includes a return value of SQL_ERROR. Note: Error class 01 returns both warnings and errors.
IM	Indicates warning and errors that are derived from ODBC.

Note: Typically, when a function successfully executes, it returns a value of SQL_SUCCESS; in some cases, however, the function may also return the SQLSTATE 00000, which also indicates successful execution.

SQLSTATE codes

Table 148. SQLSTATE codes

SQLSTATE	Error	Can be returned from
01000	General warning	All ODBC functions except: SQLGetDiagField SQLGetDiagRec
01001	Cursor operation conflict	SQLExecDirect SQLExecute SQLParamData SQLSetPos
01002	Disconnect error	SQLDisconnect

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
01003	NULL value eliminated in set function	SQLExecDirect SQLExecute SQLParamData
01004	String data, right truncated	<ul style="list-style-type: none"> • SQLColAttribute • SQLDataSources • SQLDescribeCol • SQLDriverConnect • SQLDrivers • SQLExecDirect • SQLExecute • SQLExtendedFetch • SQLFetch • SQLFetchScroll • SQLGetConnectAttr • SQLGetCursorName • SQLGetData • SQLGetDescField • SQLGetDescRec • SQLGetEnvAttr • SQLGetInfo • SQLGetStmtAttr • SQLParamData • SQLPutData • SQLSetCursorName
01006	Privilege not revoked	SQLExecDirect SQLExecute SQLParamData
01007	Privilege not granted	SQLExecDirect SQLExecute SQLParamData
01S00	Invalid connection string attribute	SQLDriverConnect SQLSetPos
01S01	Error in row	SQLExtendedFetch
01S02	Option value changed	<ul style="list-style-type: none"> • SQLConnect • SQLDriverConnect • SQLExecDirect • SQLExecute • SQLParamData • SQLPrepare • SQLSetConnectAttr • SQLSetDescField • SQLSetEnvAttr • SQLSetStmtAttr

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
01S06	Attempt to fetch before the result set returned the first rowset	SQLExtendedFetch SQLFetchScroll
01S07	Fractional truncation	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData SQLSetPos
01S08	Error saving File DSN	SQLDriverConnect
01S09	Invalid keyword	SQLDriverConnect
07001	Wrong number of parameters	SQLExecDirect SQLExecute
07002	COUNT field incorrect	SQLExecDirect SQLExecute SQLParamData
07005	Prepared statement not a cursor_specification	SQLColAttribute SQLDescribeCol
07006	Restricted data type attribute violation	SQLBindCol SQLBindParameter SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData SQLPutData

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
07009	Invalid descriptor index	SQLBindCol SQLBindParameter SQLColAttribute SQLDescribeCol SQLDescribeParam SQLFetch SQLFetchScroll SQLGetData SQLGetDescField SQLParamData SQLSetDescField SQLSetDescRec SQLSetPos
07S01	Invalid use of default parameter	SQLExecDirect SQLExecute SQLParamData SQLPutData
08001	Client unable to establish connection	SQLConnect SQLDriverConnect
08002	Connection name in use	SQLConnect SQLDriverConnect SQLSetConnectAttr
08003	Connection does not exist	SQLAllocHandle SQLDisconnect SQLEndTran SQLGetConnectAttr SQLGetInfo SQLSetConnectAttr
08004	Server rejected the connection	SQLConnect SQLDriverConnect
08007	Connection failure during transaction	SQLEndTran

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
08S01	Communication link failure	<ul style="list-style-type: none"> • SQLColumnPrivileges • SQLColumns • SQLConnect • SQLConnect • SQLCopyDesc • SQLDescribeCol • SQLDescribeParam • SQLDriverConnect • SQLExecDirect • SQLExecute • SQLExtendedFetch • SQLFetch • SQLFetchScroll • SQLForeignKeys • SQLGetConnectAttr • SQLGetData • SQLGetDescField • SQLGetDescRec • SQLGetFunctions • SQLGetInfo • SQLGetTypeInfo • SQLMoreResults • SQLNumParams • SQLNumResultCols • SQLParamData • SQLPrepare • SQLPrimaryKeys • SQLProcedureColumns • SQLProcedures • SQLPutData • SQLSetConnectAttr • SQLSetDescField • SQLSetDescRec • SQLSetEnvAttr • SQLSetStmtAttr • SQLSpecialColumns • SQLStatistics • SQLTablePrivileges • SQLTables
21S01	Insert value list does not match column list	SQLExecDirect SQLPrepare
21S02	Degree of derived table does not match column list	SQLExecDirect SQLExecute SQLParamData SQLPrepare SQLSetPos

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
22001	String data, right truncated	SQLExecDirect SQLExecute SQLFetch SQLFetchScroll SQLParamData SQLPutData SQLSetDescField SQLSetPos
22002	Indicator variable required but not supplied	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData
22003	Numeric value out of range	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLGetInfo SQLParamData SQLPutData SQLSetPos
22007	Invalid datetime format	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData SQLPutData SQLSetPos

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
22008	Datetime field overflow	SQLExecDirect SQLExecute SQLParamData SQLPutData
22012	Division by zero	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData SQLPutData
22015	Interval field overflow	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData SQLPutData SQLSetPos
22018	Invalid character value for cast specification	SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLGetData SQLParamData SQLPutData SQLSetPos
22019	Invalid escape character	SQLExecDirect SQLExecute SQLPrepare

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
22025	Invalid escape sequence	SQLExecDirect SQLExecute SQLPrepare
22026	String data, length mismatch	SQLParamData
23000	Integrity constraint violation	SQLExecDirect SQLExecute SQLParamData SQLSetPos
24000	Invalid cursor state	SQLCloseCursor SQLColumnPrivileges SQLColumns SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLForeignKeys SQLGetData SQLGetStmtAttr SQLGetTypeInfo SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLConnectAttr SQLSetCursorName SQLSetPos SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
25000	Invalid transaction state	SQLDisconnect
25S01	Transaction state	SQLEndTran
25S02	Transaction is still active	SQLEndTran

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
25S03	Transaction is rolled back	SQLEndTran
28000	Invalid authorization specification	SQLConnect SQLDriverConnect
34000	Invalid cursor name	SQLExecDirect SQLPrepare SQLSetCursorName
3C000	Duplicate cursor name	SQLSetCursorName
3D000	Invalid catalog name	SQLExecDirect SQLPrepare SQLSetConnectAttr
3F000	Invalid schema name	SQLExecDirect SQLPrepare
40001	Serialization failure	SQLColumnPrivileges SQLColumns SQLEndTran SQLExecDirect SQLExecute SQLFetch SQLFetchScroll SQLForeignKeys SQLGetTypeInfo SQLMoreResults SQLParamData SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLSetPos SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
40002	Integrity constraint violation	SQLEndTran

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
40003	Statement completion unknown	SQLColumnPrivileges SQLColumns SQLExecDirect SQLExecute SQLFetch SQLFetchScroll SQLGetTypeInfo SQLForeignKeys SQLMoreResults SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLParamData SQLSetPos SQLSpecialColumns SQLStatistics SQLTables
42000	Syntax error or access violation	SQLExecDirect SQLExecute SQLParamData SQLPrepare SQLSetPos
42S01	Base table or view already exists	SQLExecDirect SQLPrepare
42S02	Base table or view not found	SQLExecDirect SQLPrepare
42S11	Index already exists	SQLExecDirect SQLPrepare
42S12	Index not found	SQLExecDirect SQLPrepare
42S21	Column already exists	SQLExecDirect SQLPrepare
42S22	Column not found	SQLExecDirect SQLPrepare

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
44000	WITH CHECK OPTION violation	SQLExecDirect SQLExecute SQLParamData
HY000	General Error	All ODBC functions except: SQLGetDiagField SQLGetDiagRec
HY001	Memory allocation error	All ODBC function except: SQLGetDiagField SQLGetDiagRec
HY003	Invalid application buffer type	SQLBindCol SQLBindParameter SQLGetData
HY004	Invalid SQL data type	SQLBindParameter SQLGetTypeInfo
HY007	Associated statement is not prepared	SQLCopyDesc SQLGetDescField SQLGetDescRec

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HY008	Operation canceled	All ODBC functions that can be processed asynchronously: SQLColAttribute SQLColumnPrivileges SQLColumns SQLDescribeCol SQLDescribeParam SQLExecDirect SQLExecute SQLExtendedFetch SQLFetch SQLFetchScroll SQLForeignKeys SQLGetData SQLGetTypeInfo SQLMoreResults SQLNumParams SQLNumResultCols SQLParamData SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLPutData SQLSetPos SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HY009	Invalid use of null pointer	SQLAllocHandle SQLBindParameter SQLColumnPrivileges SQLColumns SQLExecDirect SQLForeignKeys SQLGetCursorName SQLGetData SQLGetFunctions SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLPutData SQLSetConnectAttr SQLSetCursorName SQLSetEnvAttr SQLSetStmtAttr SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HY010	Function sequence error	<ul style="list-style-type: none"> • SQLAllocHandle • SQLBindCol • SQLBindParameter • SQLCloseCursor • SQLColAttribute • SQLColumnPrivileges • SQLColumns • SQLCopyDesc • SQLDescribeCol • SQLDescribeParam • SQLDisconnect • SQLEndTran • SQLExecDirect • SQLExecute • SQLExtendedFetch • SQLFetch • SQLFetchScroll • SQLForeignKeys • SQLFreeHandle • SQLFreeStmt • SQLGetConnectAttr • SQLGetCursorName • SQLGetData • SQLGetDescField • SQLGetDescRec • SQLGetFunctions • SQLGetStmtAttr • SQLGetTypeInfo • SQLMoreResults • SQLNumParams • SQLNumResultCols • SQLParamData • SQLPrepare • SQLPrimaryKeys • SQLProcedureColumns • SQLProcedures • SQLPutData • SQLRowCount • SQLSetConnectAttr • SQLSetCursorName • SQLSetDescField • SQLSetEnvAttr • SQLSetDescRec • SQLSetPos • SQLSetStmtAttr • SQLSpecialColumns • SQLStatistics • SQLTablePrivileges • SQLTables

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HY011	Attribute cannot be set now	SQLParamData SQLSetConnectAttr SQLSetPos SQLSetStmtAttr
HY012	Invalid transaction operation code	SQLEndTran
HY013	Memory Management err	All ODBC functions except: SQLGetDiagField SQLGetDiagRec
HY014	Limit on the number of handles exceeded	SQLAllocHandle
HY015	No cursor name available	SQLGetCursorName
HY016	Cannot modify an implementation row descriptor	SQLCopyDesc SQLSetDescField SQLSetDescRec
HY017	Invalid use of an automatically allocated descriptor handle	SQLFreeHandle SQLSetStmtAttr
HY018	Server declined cancel request	SQLCancel
HY019	Non-character and non-binary data sent in pieces	SQLPutData
HY020	Attempt to concatenate a null value	SQLPutData
HY021	Inconsistent descriptor information	SQLBindParameter SQLCopyDesc SQLGetDescField SQLSetDescField SQLSetDescRec
HY024	Invalid attribute value	SQLSetConnectAttr SQLSetEnvAttr SQLSetStmtAttr

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HY090	Invalid string or buffer length	<ul style="list-style-type: none"> • SQLBindCol • SQLBindParameter • SQLBrowseConnect • SQLColAttribute • SQLColumnPrivileges • SQLColumns • SQLConnect • SQLDataSources • SQLDescribeCol • SQLDriverConnect • SQLDrivers • SQLExecDirect • SQLExecute • SQLFetch • SQLFetchScroll • SQLForeignKeys • SQLGetConnectAttr • SQLGetCursorName • SQLGetData • SQLGetDescField • SQLGetInfo • SQLGetStmtAttr • SQLParamData • SQLPrepare • SQLPrimaryKeys • SQLProcedureColumns • SQLProcedures • SQLPutData • SQLSetConnectAttr • SQLSetCursorName • SQLSetDescField • SQLSetDescRec • SQLSetEnvAttr • SQLSetStmtAttr • SQLSetPos • SQLSpecialColumns • SQLTablePrivileges • SQLStatistics • SQLTables
HY091	Invalid descriptor field identifier	<ul style="list-style-type: none"> SQLColAttribute SQLGetDescField SQLSetDescField

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HY092	Invalid attribute/option identifier	SQLAllocHandle SQLCopyDesc SQLDriverConnect SQLEndTran SQLFreeStmt SQLGetConnectAttr SQLGetEnvAttr SQLGetStmtAttr SQLParamData SQLSetConnectAttr SQLSetDescField SQLSetEnvAttr SQLSetPos SQLSetStmtAttr
HY095	Function type out of range	SQLGetFunctions
HY096	Invalid information type	SQLGetInfo
HY097	Column type out of range	SQLSpecial Columns
HY098	Scope type out of range	SQLSpecial Columns
HY099	Nullable type out of range	SQLSpecial Columns
HY100	Uniqueness option type out of range	SQLStatistics
HY101	Accuracy option type out of range	SQLStatistics
HY103	Invalid retrieval code	SQLDataSources SQLDrivers
HY104	Invalid precision or scale value	SQLBindParameter
HY105	Invalid parameter type	SQLBindParameter SQLExecDirect SQLExecute SQLParamData SQLSetDescField
HY106	Fetch type out of range	SQLExtendedFetch SQLFetchScroll

Table 148. *SQLSTATE* codes (continued)

SQLSTATE	Error	Can be returned from
HY107	Row value out of range	SQLExtendedFetch SQLFetch SQLFetchScroll SQLSetPos
HY109	Invalid cursor position	SQLExecDirect SQLExecute SQLGetData SQLGetStmtAttr SQLParamData SQLSetPos
HY110	Invalid driver completion	SQLDriverConnect
HY111	Invalid bookmark value	SQLExtendedFetch SQLFetchScroll

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HYC00	Optional feature not implemented	<ul style="list-style-type: none"> • SQLBindCol • SQLBindParameter • SQLColAttribute • SQLColumnPrivileges • SQLColumns • SQLDriverConnect • SQLEndTran • SQLConnect • SQLExecDirect • SQLExecute • SQLExtendedFetch • SQLFetch • SQLFetchScroll • SQLForeignKeys • SQLGetConnectAttr • SQLGetData • SQLGetEnvAttr • SQLSetPos • SQLGetInfo • SQLGetStmtAttr • SQLGetTypeInfo • SQLParamData • SQLPrepare • SQLPrimaryKeys • SQLProcedureColumns • SQLProcedures • SQLSetConnectAttr • SQLSetEnvAttr • SQLSetStmtAttr • SQLSpecialColumns • SQLStatistics • SQLTablePrivileges • SQLTables

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
HYT00	Timeout expired	SQLBrowseConnect SQLColumnPrivileges SQLColumns SQLConnect SQLDriverConnect SQLExecDirect SQLExecute SQLExtendedFetch SQLForeignKeys SQLGetTypeInfo SQLParamData SQLPrepare SQLPrimaryKeys SQLProcedureColumns SQLProcedures SQLSetPos SQLSpecialColumns SQLStatistics SQLTablePrivileges SQLTables
HYT01	Connection timeout expired	All ODBC functions except: SQLDrivers SQLDataSources SQLGetEnvAttr SQLSetEnvAttr
IM001	Driver does not support this function	All ODBC functions except: SQLAllocHandle SQLDataSources SQLDrivers SQLFreeHandle SQLGetFunctions
IM002	Data source name not found and no default driver specified	SQLConnect SQLDriverConnect
IM003	Specified driver could not be loaded.	SQLConnect

Table 148. SQLSTATE codes (continued)

SQLSTATE	Error	Can be returned from
IM004	Driver's SQLAllocHandle on SQL_HANDLE_ENV failed	SQLDriverConnc SQLConnect SQLDriverConnect
IM005	Driver's SQLAllocHandle on SQL_HANDLE_DBC failed	SQLConnect SQLDriverConnect
IM006	Driver's SQLSetConnectAttr Failed	SQLConnect SQLDriverConnect
IM007	No data source or driver specified; dialog prohibited	SQLDriverConnect
IM008	Dialog failed	SQLDriverConnect
IM009	Unable to load translation DLL	SQLConnect SQLDriverConnect SQLSetConnectAttr
IM010	Data source name too long	SQLConnect SQLDriverConnect
IM011	Driver name too long	SQLDriverConnect
IM012	DRIVER keyword syntax error	SQLDriverConnect
IM013	Trace file error	All ODBC functions
IM014	Invalid name of File DSN	SQLDriverConnect
IM015	Corrupt file data source	SQLDriverConnect

Appendix D. Minimum SQL grammar requirements for ODBC

This section describes the minimum subset of SQL-92 Entry level syntax that an ODBC driver must support. An application that uses this syntax will be supported by any ODBC-compliant driver.

Applications can call SQLGetInfo with the SQL_SQL_CONFORMANCE to determine if additional features of SQL-92, not covered in this section, are supported.

Note: If the driver supports only read-only data sources, the SQL syntax that applies to changing data may not apply to the driver. Applications need to call SQLGetInfo with the SQL_DATA_SOURCE_READ_ONLY information type to determine if a data source is read-only.

D.1 SQL statements

This section describes the subset of SQL statements and elements.

```
create-table-statement ::=  
    CREATE TABLE base_table_name  
    (column_identifier data_type [, column_identifier data_type]...)
```

Important: As the *data_type* in a *create_table_statement*, applications require a data type from the TYPE_NAME column of the result set returned by SQLGetTypeInfo.

```
delete_statement_searched ::=  
    DELETE FROM table_name [WHERE search_condition]
```

```
drop_table_statement ::=  
    DROP TABLE base_table_name
```

```
select_statement ::=  
    SELECT [ALL | DISTINCT] select_list  
    FROM table_reference_list  
    [WHERE search_condition]  
    [order_by_clause]
```

```
statement ::= create_table_statement |  
    delete_statement_searched |  
    drop_table_statement |  
    insert_statement |  
    select_statement |  
    update_statement_searched
```

```
Update_statement_searched ::=  
    UPDATE table_name  
    SET column_identifier = {expression |  
        NULL}  
    [, column_identifier = {expression |  
        NULL}]...  
    [WHERE search_condition]
```

SQL statement elements

```
base_table_identifier ::= user_defined_name  
base_table_name ::= base_table_identifier  
boolean_factor ::= [NOT] boolean_primary  
boolean_primary ::= predicate | ( search_condition )  
boolean_term ::= boolean_factor [AND boolean_term]  
character_string_literal ::= "{character}..."  
(character is any character in the character set  
of the driver/data source. To include a single  
literal quote character (') in a character_string_literal,  
use two literal quote characters [""].)
```

```

column_identifier ::= user_defined_name
column_name ::= [table_name.]column_identifier
comparison_operator ::= < | > | <= | >= | = | <>
comparison_predicate ::= expression comparison_operator expression
data_type ::= character_string_type
(character_string_type is any data type for which the
"DATA_TYPE" column in the result set returned by SQLGetTypeInfo
is either SQL_CHAR or SQLVARCHAR.)
digit ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
dynamic_parameter ::= ?
expression ::= term | expression {+|-} term
factor ::= [+|-]primary
insert_value ::= dynamic_parameter | literal | NULL | USER
letter ::= lower_case_letter | upper_case_letter
literal ::= character_string_literal
lower_case_letter ::= a | b | c | d | e | f | g |
  h | i | j | k | l | m | n | o | p | q | r | s |
  t | u | v | w | x | y | z
order_by_clause ::= ORDER BY sort_specification [, sort_specification]...
primary ::= column_name | dynamic_parameter | literal | ( expression )
search_condition ::= boolean_term [OR search_condition]
select_list ::= * | select_sublist [, select_sublist]...
(select_list cannot contain parameters.)
select_sublist ::= expression
sort_specification ::= {unsigned_integer | column_name } [ASC | DESC]
table_identifier ::= user_defined_name
table_name ::= table_identifier
table_reference ::= table_name
table_reference ::= table_name [,table_reference]...
term ::= factor | term {*/} factor
unsigned_integer ::= {digit}
upper_case_letter ::= A | B | C | D | E | F | G |
  H | I | J | K | L | M | N | O | P | Q | R | S |
  T | U | V | W | X | Y | Z
user_defined_name ::= letter[ digit | letter| _ ]...

```

D.1.1 Control statements (logical condition)

This topic provides a summary of control statements that are available in solidDB database procedures.

For a more detailed description of these control statements, see the discussion on stored procedures in the *solidDB SQL Guide*.

Table 149. Control Statements

Control statement	Description
<i>set variable = expression</i>	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.
<i>variable := expression</i>	Alternate syntax for assigning values to variables.
<i>boolean_expr</i>	A boolean expression which evaluates to "true" or "false". The expression can include comparison operators, such as =, >, and <) and logical operators and, or, and not.

Table 149. Control Statements (continued)

Control statement	Description
<i>statement_list</i>	A valid procedure statement that executes as a result of a boolean expression.
while <i>boolean_expr</i> loop <i>statement_list</i> end loop	This loops while the expression is true. For examples of valid parentheses use in WHILE loops, see the discussion of stored procedures in <i>solidDB SQL Guide</i> .
leave	Leaves the innermost while loop and continues executing the procedure from the next statement after the keyword end loop.
if <i>boolean_expr</i> then <i>statement_list1</i> else <i>statement_list2</i> end if	Executes <i>statement_list1</i> if <i>boolean_expr</i> is true; otherwise, executes <i>statement_list2</i> . For examples of valid parentheses use in IF statements, see the discussion on stored procedures in <i>solidDB SQL Guide</i> .
if <i>boolean_expr1</i> then <i>statement_list1</i> elseif <i>boolean_expr2</i> then <i>statement_list2</i> end if	If <i>boolean_expr1</i> is true, executes <i>statement_list1</i> . If <i>boolean_expr2</i> is true, executes <i>statement_list2</i> . The statement can optionally contain multiple elseif statements and also an else statement. For examples of valid parentheses use in IF statement, see the discussion of stored procedures in <i>solidDB SQL Guide</i> .
return	Returns the current values of output parameters and exits the procedure. If a procedure has a <i>return row</i> statement, <i>return</i> behaves like <i>return norow</i> .
return sqlerror of <i>cursor_name</i>	Returns the sqlerror associated with the cursor and exits the procedure.
return row	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.
return norow	Returns the end of the set and exits the procedure.

D.2 Data type support

At minimum, ODBC drivers must support either SQL_CHAR or SQL_VARCHAR.

Other data types support is determined by the driver's or data source's SQL-92 conformance level. To determine the SQL-92 conformance level for a driver or data source, applications need to call SQLGetTypeInfo.

D.3 Parameter data types

This topic describes how data types are determined for parameters and the parameter markers support.

Even though each parameter specified with SQLBindParameter is defined using an SQL data type, the parameters in an SQL statement have no intrinsic data type. 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. An application 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 according to SQL-92 standards. For comprehensive information about inferring the parameter type, see the SQL-92 specification.

Table 150. Determining Data Ttype for Several Types of 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 second 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 or the first value if there is a parameter marker in the expression
A pattern value used with LIKE	VARCHAR
An update value used with UPDATE	Same as the update column

Parameter markers

According to the SQL-92 specification, an application cannot place parameter markers in the following locations:

- 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 a comprehensive list and more details, see the SQL-92 specification.

D.4 Literals in ODBC

This section contains information that will help driver writers who are converting a character string type to a numeric or interval type, or from a numeric or interval type to a character string type.

Interval literal syntax

The following syntax is used for interval literals in ODBC.

```

interval_literal ::= INTERVAL [+|_] interval_string interval_qualifier
interval_string ::= quote { year_month_literal
    | day_time_literal } quote
year_month_literal ::= years_value | [years_value] months_value
day_time_literal ::= day_time_interval | time_interval
day_time_interval ::= days_value [hours_value
    [:minutes_value[:seconds_value]]]
time_interval ::= hours_value [:minutes_value [:seconds_value ] ]
    | minutes_value [:seconds_value ]
    | seconds_value
years_value ::= datetime_value
months_value ::= datetime_value
days_value ::= datetime_value
hours_value ::= datetime_value
minutes_value ::= datetime_value
seconds_value ::= seconds_integer_value [.[seconds_fraction] ]
seconds_integer_value ::= unsigned_integer
seconds_fraction ::= unsigned_integer
datetime_value ::= unsigned_integer
interval_qualifier ::= start_field TO end_field
    | single_datetime_field
start_field ::= non_second_datetime_field
    [(interval_leading_field_precision )]
end_field ::= non_second_datetime_field
    | SECOND[(interval_fractional_seconds_precision)]
single_datetime_field ::= non_second_datetime_field
    [(interval_leading_field_precision)]
    | SECOND[(interval_leading_field_precision
    [, (interval_fractional_seconds_precision)]
datetime_field ::= non_second_datetime_field | SECOND
non_second_datetime_field ::= YEAR | MONTH | DAY | HOUR | MINUTE
interval_fractional_seconds_precision ::= unsigned_integer
interval_leading_field_precision ::= unsigned_integer
quote ::= '
unsigned_integer ::= digit...

```

Numeric Literal Syntax

The following syntax is used for numeric literals in ODBC:

```

numeric_literal ::= signed_numeric_literal | unsigned_numeric_literal
signed_numeric_literal ::= [sign] unsigned_numeric_literal
unsigned_numeric_literal ::= exact_numeric_literal
    | approximate_numeric_literal
exact_numeric_literal ::= unsigned_integer [period[unsigned_integer]]
    | period unsigned_integer
sign ::= plus_sign | minus_sign
approximate_numeric_literal ::= mantissa E exponent

```

```

mantissa ::= exact_numeric_literal
exponent ::= signed_integer
signed_integer ::= [sign] unsigned_integer
unsigned_integer ::= digit...
plus_sign ::= +
minus_sign ::= -
digit ::= 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0
period ::= .

```

D.5 List of reserved keywords

To ensure compatibility with drivers that support the core SQL grammar, there are keywords that applications should avoid using.

These words do not constrain the minimum SQL grammar. The #define value `SQL_ODBC_KEYWORDS` contains a comma-separated list of these keywords.

For a complete list of reserved keywords in several SQL standards and solidDB ODBC API, see “Reserved Words” in *IBM solidDB SQL Guide*.

Table 151. List of Reserved Keywords

Keyword	Keyword	Keyword	Keyword
ABSOLUTE	ACTION	ADA	ADD
ALL	ALLOCATE	ALTER	AND
ANY	ARE	AS	ASC
ASSERTION	AT	AUTHORIZATION	AVG
BEGIN	BETWEEN	BIT	BIT_LENGTH
BOTH	BY	CASCADE	CASCADED
CASE	CAST	CATALOG	CHAR
CHAR_LENGTH	CHARACTER	CHARACTER_LENGTH	CHECK
CLOSE	COALESCE	COLLATE	COLLATION
COLUMN	COMMIT	CONNECT	CONNECTION
CONSTRAINT	CONSTRAINTS	CONTINUE	CONVERT
CORRESPONDING	COUNT	CREATE	CROSS
CURRENT	CURRENT_DATE	CURRENT_TIME	CURRENT_TIMESTAMP
CURRENT_USER	CURSOR	DATE	DAY
DEALLOCATE	DEC	DECIMAL	DECLARE
DEFAULT	DEFERRABLE	DEFERRED	DELETE

Table 151. List of Reserved Keywords (continued)

Keyword	Keyword	Keyword	Keyword
DESC	DESCRIBE	DESCRIPTOR	DIAGNOSTICS
DISCONNECT	DISTINCT	DOMAIN	DOUBLE
DROP	ELSE	END	END-EXEC
ESCAPE	EXCEPT	EXCEPTION	EXEC
EXECUTE	EXISTS	EXTERNAL	EXTRACT
FALSE	FETCH	FIRST	FLOAT
FOR	FOREIGN	FORTRAN	FOUND
FROM	FULL	GET	GLOBAL
GO	GOTO	GRANT	GROUP
HAVING	HOUR	IDENTITY	IMMEDIATE
IN	INCLUDE	INDEX	INDICATOR
INITIALLY	INNER	INPUT	INSENSITIVE
INSERT	INT	INTEGER	INTERSECT
INTERVAL	INTO	IS	ISOLATION
JOIN	KEY	LANGUAGE	LAST
LEADING	LEFT	LEVEL	LIKE
LOCAL	LOWER	MATCH	MAX
MIN	MINUTE	MODULE	MONTH
NAMES	NATIONAL	NATURAL	NCHAR
NEXT	NO	NONE	NOT
NULL	NULLIF	NUMERIC	OCTET_LENGTH
OF	ON	ONLY	OPEN
OPTION	OR	ORDER	OUTER
OUTPUT	OVERLAPS	PASCAL	POSITION
PRECISION	PREPARE	PRESERVE	PRIMARY

Table 151. List of Reserved Keywords (continued)

Keyword	Keyword	Keyword	Keyword
PRIOR	PRIVILEGES	PROCEDURE	PUBLIC
READ	REAL	REFERENCES	RELATIVE
RESTRICT	REVOKE	RIGHT	ROLLBACK
ROWS	SCHEMA	SCROLL	SECOND
SECOND	SECTION	SELECT	SESSION
SESSION_USER	SET	SIZE	SMALLINT
SOME	SPACE	SQL	SQLCA
SQLCODE	SQLERROR	SQLSTATE	SQLWARNING
SUBSTRING	SUM	SYSTEM_USER	TABLE
TEMPORARY	THEN	TIME	TIMESTAMP
TIMEZONE_HOUR	TIMEZONE_MINUTE	TO	TRAILING
TRANSACTION	TRANSLATE	TRANSLATION	TRIM
TRUE	UNION	UNIQUE	UNKNOWN
UPDATE	UPPER	USAGE	USER
USING	VALUE	VALUES	VARCHAR
VARYING	VIEW	WHEN	WHENEVER
WHERE	WITH	WORK	WRITE
YEAR	ZONE		

Appendix E. Data types

This section describes the ODBC data types.

ODBC defines the following sets of data types:

- SQL data types, which indicate the data type of data stored at the data source (for example, the solidDB server).
- C data types, which indicate the data type of data stored in application buffers.

Each SQL data type corresponds to an ODBC C data type. Before returning data from the data source, the driver converts it to the specified C data type. Before sending data to the data source, the driver converts it from the specified C data type.

For information about driver-specific SQL data types, see the driver's documentation.

E.1 SQL data types

In accordance with the SQL-92 standard, each DBMS defines its own set of SQL data types. For each SQL data type in the SQL-92 standard, a #define value, known as a type identifier, is passed as an argument in ODBC functions or returned in the metadata of a result set.

Drivers map data source-specific SQL data types to ODBC SQL data type identifiers and driver-specific SQL data type identifiers. The SQL_DESC_CONCISE_TYPE field of an implementation descriptor is where the SQL data type is stored.

solidDB's ODBC driver does not support the following SQL_92 data types:

- BIT
- BIT_VARYING
- TIME_WITH_TIMEZONE
- TIMESTAMP_WITH_TIMEZONE
- NATIONAL_CHARACTER

E.2 C data types

ODBC defines the C data types and their corresponding ODBC type identifiers.

Applications call one of the following functions:

- SQLBindCol or SQLGetData to pass an applicable C type identifier in the TargetType argument. In this way, applications specify the C data type of the buffer that receives result set data.
- SQLBindParameter to pass the appropriate C type identifier in the ValueType argument. In this way, applications specify the C data type of the buffer containing a statement parameter.

The SQL_DESC_CONCISE_TYPE field of an application descriptor is where the C data type is stored.

Note: Driver-specific C data types do not exist.

E.3 Data type identifiers

Data type identifiers are stored in the `SQL_DESC_CONCISE_TYPE` field of a descriptor. Data type identifiers in applications describe their buffers to the driver.

They also retrieve metadata about the result set from the driver so applications know what type of C buffers to use for data storage. Applications use data type identifiers to perform these tasks by calling these functions:

- To describe the C data type of application buffers, applications call `SQLBindParameter`, `SQLBindCol`, and `SQLGetData`.
- To describe the SQL data type of dynamic parameters, applications call `SQLBindParameter`.
- To retrieve the SQL data types of result set columns, applications call `SQLColAttribute` and `SQLDescribeCol`.
- To retrieve the SQL data types of parameters, applications call `SQLDescribeParameter`.
- To retrieve the SQL data types of various schema information, applications call `SQLColumns`, `SQLProcedureColumns`, and `SQLSpecialColumns`.
- To retrieve a list of supported data types, applications call `SQLGetTypeInfo`.

In addition, the `SQLSetDescField` and `SQLSetDescRec` descriptor functions are also used to perform the above tasks. For details, see the `SQLSetDescField` and `SQLSetDescRec` functions.

E.4 SQL data types

A given driver and data source do not necessarily support all of the SQL data types defined in the ODBC grammar. Furthermore, they may support additional, driver-specific SQL data types.

A driver's support is determined by the level of SQL-92 conformance. To determine which data types a driver supports, an application calls `SQLGetTypeInfo`. See "SQLGetTypeInfo Result Set Example" on page 227. For information about driver-specific SQL data types, see the driver's documentation.

A driver also returns the SQL data types when it describes the data types of columns and parameters using the following functions:

- `SQLColAttribute`
- `SQLColumns`
- `SQLDescribeCol`
- `SQLDescribeParam`
- `SQLProcedureColumns`
- `SQLSpecialColumns`

Note:

For details on fields that store SQL data type values and characteristics, see E.8, "Data type identifiers and descriptors," on page 236.

The following table is not a comprehensive list of SQL data types, but offers commonly used names, ranges, and limits. A data source may only support some

of the data types that are listed in the table and depending on your driver, the characteristics of the data types can differ from this table's description. The table includes the description of the associated data type from SQL-92 (if applicable)

Table 152. Common SQL Data Type Names, Ranges, and Limits

SQL Type Identifier [1]	Typical SQL Data Type [2]	Typical Type Description
SQL_CHAR	CHAR(n)	Character string of fixed string length <i>n</i> .
SQL_VARCHAR	VARCHAR(n)	Variable-length character string with a maximum string length <i>n</i> .
SQL_LONGVARCHAR	LONG VARCHAR	Variable length character data. Maximum length is data source-dependent. [3]
SQL_WCHAR	WCHAR(n)	Unicode character string of fixed string length <i>n</i> .
SQL_WVARCHAR	VARWCHAR(n)	Unicode variable-length character string with a maximum string length <i>n</i> .
SQL_WLONGVARCHAR	LONGWVARCHAR	Unicode variable-length character data. Maximum length is data source-dependent.
SQL_DECIMAL	DECIMAL(<i>p</i> , <i>s</i>)	Signed, exact, numeric value with a precision <i>p</i> and scale <i>s</i> . (The maximum precision is driver-defined.) ($1 \leq p \leq 16$; $s \leq p$). [4]
SQL_NUMERIC	NUMERIC(<i>p</i> , <i>s</i>)	Signed, exact, numeric value with a precision <i>p</i> and scale <i>s</i> . ($1 \leq p \leq 16$; $s \leq p$). [4]
SQL_SMALLINT	SMALLINT	Exact numeric value with precision 5 and scale 0. (signed: $-32,768 \leq n \leq 32,767$, unsigned: $0 \leq n \leq 65,535$) solidDB supports only signed, not unsigned, SMALLINT. [5]
SQL_INTEGER	INTEGER	Exact numeric value with precision 10 and scale 0. (signed: $-2^{31} \leq n \leq 2^{31} - 1$, unsigned: $0 \leq n \leq 2^{32} - 1$) solidDB supports only signed, not unsigned, INTEGER. [5]
SQL_REAL	REAL	Signed, approximate, numeric value with a binary precision 24 (zero or absolute value 10^{-38} to 1038).
SQL_FLOAT	FLOAT(<i>p</i>)	Signed, approximate, numeric value with a binary precision of at least <i>p</i> . (The maximum precision is driver defined.) [6]
SQL_DOUBLE	DOUBLE PRECISION	Signed, approximate, numeric value with a binary precision 53 (zero or absolute value 10^{-308} to 10^{308}).
SQL_BIT	BIT	Single bit binary data. NOTE: solidDB does not support BIT/SQL_BIT. [7]

Table 152. Common SQL Data Type Names, Ranges, and Limits (continued)

SQL Type Identifier [1]	Typical SQL Data Type [2]	Typical Type Description
SQL_TINYINT	TINYINT	Exact numeric value with precision 3 and scale 0 (signed: $-128 \leq n \leq 127$ unsigned: $0 \leq n \leq 255$) solidDB supports only signed, not unsigned, TINYINT. [5].
SQL_BIGINT	BIGINT	Exact numeric value with precision 19 (if signed) or 20 (if unsigned) and scale 0 (signed: $-2^{63} \leq n \leq 2^{63} - 1$, unsigned: $0 \leq n \leq 2^{64} - 1$) solidDB supports only signed, not unsigned, BIGINT. [3], [5].
SQL_BINARY	BINARY(n)	Binary data of fixed length <i>n</i> . [3]
SQL_VARBINARY	VARBINARY(n)	Variable length binary data of maximum length <i>n</i> . The maximum is set by the user. [3]
SQL_LONGVARBINARY	LONG VARBINARY	Variable length binary data. Maximum length is data source-dependent. [3]
SQL_TYPE_DATE [8]	DATE	Year, month, and day fields, conforming to the rules of the Gregorian calendar. (See E.11, "Constraints of the gregorian calendar," on page 240.)
SQL_TYPE_TIME [8]	TIME(p)	Hour, minute, and second fields. Valid values for hours are 00 to 23. Valid values for minutes are 00 to 59. Valid values for seconds are 00 to 61 (60 and 61 are to handle "leap seconds" (see http://tycho.usno.navy.mil/leapsec.html). Precision <i>p</i> indicates the precision of the seconds field.
SQL_TYPE_TIMESTAMP [8]	TIMESTAMP(p)	Year, month, day, hour, minute, and second fields, with valid values as defined for the DATE and Time data types.

Note:

[1] This is the value returned in the DATA_TYPE column by a call to SQLGetTypeInfo.

[2] This is the value returned in the NAME and CREATE PARAMS column by a call to SQLGetTypeInfo. The NAME column returns the designation - for example, CHAR - while the CREATE PARAMS column returns a comma-separated list of creation parameters such as precision, scale, and length.

[3] This data type has no corresponding data type in SQL-92.

[4] SQL_DECIMAL and SQL_NUMERIC data types differ only in their precision. The precision of a DECIMAL(p,s) is an implementation-defined decimal precision that is no less than *p*, while the precision of a NUMERIC(p,s) is exactly equal to *p*.

[5] An application uses SQLGetTypeInfo or SQLColAttribute to determine if a particular data type or a particular column in a result set is unsigned.

[6] Depending on the implementation, the precision of SQL_FLOAT can be either 24 or 53: if it is 24, the SQL_FLOAT data type is the same as SQL_REAL; if it is 53, the SQL_FLOAT data type is the same as SQL_DOUBLE.

[7] The SQL_BIT data type has different characteristics than the BIT type in SQL-92.

[8] This data type has no corresponding data type in SQL-92.

SQLGetTypeInfo Result Set Example

Applications call SQLGetTypeInfo result set for a list of supported data types and their characteristics for a given data source.

The example below shows the data types that SQLGetTypeInfo returns for a data source; all data types under "DATA_TYPE" are supported in this data source.

The example below is divided into 3 sections so that it fits the width of a page. In fact, it is all one example.

Table 153. Data types SQLGetTypeInfo returns (1)

TYPE_NAME	DATA_TYPE	COLUMN_SIZE	LITERAL_PREFIX	LITERAL_SUFFIX	CREATE_PARAMS	NULLABLE
"char"	SQL_CHAR	255	""	""	"length"	SQL_TRUE
"text"	SQL_LONG VARCHAR	2147483647	""	""	<Null>	SQL_TRUE
"decimal"	SQL_DECIMAL	18 [a]	<Null>	<Null>	"precision, scale"	SQL_TRUE
"real"	SQL_REAL	7	<Null>	<Null>	<Null>	SQL_TRUE
"datetime"	SQL_TYPE_TIMESTAMP	29 [b]	""	""	<Null>	SQL_TRUE

Table 154. Data Types SQLGetTypeInfo Returns (2)

(continued)	CASE_SENSITIVE	SEARCHABLE	UNSIGNED_ATTRIBUTE	FIXED_PREC_SCALE	AUTO_UNIQUE_VALUE	LOCAL_TYPE_NAME
SQL_CHAR	SQL_FALSE	SQL_SEARCHABLE	<Null>	SQL_FALSE	<Null>	"char"
SQL_LONG VARCHAR	SQL_FALSE	SQL_PRED_CHAR	<Null>	SQL_FALSE	<Null>	"text"
SQL_DECIMAL	SQL_FALSE	SQL_PRED_BASIC	SQL_FALSE	SQL_FALSE	SQL_FALSE	"decimal"
SQL_REAL	SQL_FALSE	SQL_PRED_BASIC	SQL_FALSE	SQL_FALSE	SQL_FALSE	"real"
SQL_TYPE_TIMESTAMP	SQL_FALSE	SQL_SEARCHABLE	<Null>	SQL_FALSE	<Null>	"datetime"

Table 155. Data Types SQLGetTypeInfo Returns (3)

(continued)	MINIMUM_SCALE	MAXIMUM_SCALE	SQL_DATA_TYPE	SQL_DATETIME_SUB	NUM_PREC_RADIX	INTERVAL_PRECISION
SQL_CHAR	<Null>	<Null>	SQL_CHAR	<Null>	<Null>	<Null>

Table 155. Data Types SQLGetTypeInfo Returns (3) (continued)

(continued)	MINIMUM_SCALE	MAXIMUM_SCALE	SQL_DATA_TYPE	SQL_DATETIME_SUB	NUM_PREC_RADIX	INTERVAL_PRECISION
SQL_LONG_VARCHAR	<Null>	<Null>	SQL_LONG_VARCHAR	<Null>	<Null>	<Null>
SQL_DECIMAL	0	16	SQL_DECIMAL	<Null>	10	<Null>
SQL_REAL	<Null>	<Null>	SQL_REAL	<Null>	10	<Null>
SQL_TYPE_TIMESTAMP	3	3	SQL_DATETIME	SQL_CODE_TIMESTAMP	<Null>	12

Explanations of Footnote Numbering in the Table Above

[a] 16 digits, 1 decimal point, and an optional sign character for negative numbers

[b] 29 characters to display yyyy-mm-dd hh:MM:ss.nnnnnnnnn

E.5 C data types

The solidDB ODBC Driver supports all C data types in keeping with the need for character SQL type conversion to and from all C types.

The C data type is specified in the following functions:

- SQLBindCol and SQLGetData functions with the targetType argument.
- SQLBindParameter with the valueType argument.
- SQLSetDescField to set the SQL_DESC_CONCISE_TYPE field of an ARD1 or APD2
- SQLSetDescRec with the Type argument, SubType argument (if needed), and the DescriptorHandle argument set to the handle of an ARD or APD.
 - ARD: Application Row Descriptors contain information about application variables that are bound to columns returned by an SQL statement. The information includes the addresses, lengths, and C data types of the bound variables.
 - APD: Application Parameter Descriptors contain information about application variables that are bound to the parameter markers ("?) used in an SQL statement, for example:


```
SELECT * FROM table1 WHERE id = ?
```

 The information in the descriptors includes the addresses, lengths, and C data types of the bound variables.

The table below contains the following three columns:

C Type Identifiers

The first column shows the C Type Identifiers that are passed to functions like SQLBindCol to indicate the type of the variable that will be bound to the column. In the following example, SQL_C_DECIMAL is a C Type Identifier:


```
// Bind MySharedVariable to column 1 of the result set. (Column 1 is a
// DECIMAL column.) The C Type Identifier SQL_C_DECIMAL shows that the
// variable MySharedVariable is of a type equivalent to DECIMAL.
SQLBindCol(..., 1, SQL_C_DECIMAL, &MySharedVariable, ...);
```

ODBC C Data Type

The second column shows the ODBC C Data Type that is associated with each C Type Identifier. This ODBC C data type is a "typedef" that you use to define variables in your ODBC program. This helps insulate your program from platform-specific requirements. For example, if you have a column of type SQL FLOAT and you want to bind a variable to that column, you can declare your variable to be of type SQLFLOAT in the following way:

```
SQLFLOAT MySharedVariable; // Can be bound to a column of type SQL FLOAT.
```

C Type

The third column contains an example of a C type definition that corresponds to the ODBC C Data Type "typedef". The examples in this column show the most frequently used definitions on 32-bit platforms.

Note: The data types specified in this column are not platform-independent; they are examples.

```
// A portable way to declare a variable that will be bound to a column of
// type SQL FLOAT.
SQLFLOAT MySharedSQLFLOATVariable = 0.0;
// A non-portable way to declare a variable that will be bound to a column
// of type SQL INTEGER. This declaration works properly on most 32-bit
// platforms, but may fail on 64-bit platforms.
long int MySharedSQLINTEGERVariable = 0;
// Bind MySharedSQLFLOATVariable to column 1 of the result set.
SQLBindCol(..., 1, SQL_C_DOUBLE, &MySharedSQLFLOATVariable, ...);
// Bind MySharedSQLINTEGERVariable to column 2 of the result set.
SQLBindCol(..., 2, SQL_C_SLONG, &MySharedSQLINTEGERVariable, ...);
```

The C Type Identifier and the ODBC C Type do not always have similar names. The C Type Identifier has a name based on the C language data type (for example, "float"), while the ODBC C Typedef has a name that is based on the SQL data type. Since C-language "float" corresponds to SQL "REAL", the table lists "SQL_C_FLOAT" as the C Type Identifier that corresponds to the ODBC C Typedef "SQLREAL".

Table 156. C vs ODBC Naming Correspondence

C Type identifier	ODBC C Typedef	C Type
SQL_C_CHAR	SQLCHAR	unsigned char
SQL_C_TINYINT	SCHAR	char
SQL_C_UTINYINT [i]	UCHAR	unsigned char
SQL_C_SHORT [h]	SQLSMALLINT	short int
SQL_C_USHORT [h] [i]	SQLUSMALLINT	unsigned short int
SQL_C_SLONG [h]	SQLINTEGER	int
SQL_C_ULONG [h] [i]	SQLINTEGER	unsigned int

Table 156. C vs ODBC Naming Correspondence (continued)

C Type identifier	ODBC C Typedef	C Type
SQL_C_SBIGINT	SQLBIGINT	_int64 [g]
SQL_C_UBIGINT [i]	SQLUBIGINT	unsigned _int64 [g] solidDB does not support unsigned data types such as this.
SQL_C_FLOAT	SQLREAL	float
SQL_C_DOUBLE	SQLDOUBLE SQLFLOAT	double
SQL_C_NUMERIC	SQLNUMERIC	unsigned char [f]
SQL_C_DECIMAL	SQLDECIMAL	unsigned char [f]
SQL_C_BINARY	SQLCHAR *	unsigned char *
SQL_C_TYPE_DATE [c]	SQL_DATE_STRUCT	struct tagDATE_STRUCT { SQLSMALLINT year; SQLUSMALLINT month; SQLUSMALLINT day; } DATE_STRUCT; [a]
SQL_C_TYPE_TIME [c]	SQL_TIME_STRUCT	struct tagTIME_STRUCT { SQLUSMALLINT hour; SQLUSMALLINT minute; [d] SQLUSMALLINT second; [e] }
SQL_C_TYPE_TIMESTAMP [c]	SQL_TIMESTAMP_STRUCT	struct tagTIMESTAMP_STRUCT { SQLSMALLINT year; [a] SQLUSMALLINT month; [b] SQLUSMALLINT day; [c] SQLUSMALLINT hour; SQLUSMALLINT minute; [d] SQLUSMALLINT second; [e] SQLINTEGER fraction; }

Table 156. C vs ODBC Naming Correspondence (continued)

C Type identifier	ODBC C Typedef	C Type
<p>Note:</p>		
<p>[a] The values of the year, month, day, hour, minute, and second fields in the datetime C data types must conform to the constraints of the Gregorian calendar. (See E.11, "Constraints of the gregorian calendar," on page 240.)</p>		
<p>[b] The value of the fraction field is the number of nanoseconds (billionths of a second) and ranges from 0 through 999,999,999 (1 less than 1 billion). For example, the value of the fraction field for a half-second is 500,000,000, for a thousandth of a second (one millisecond) is 1,000,000, for a millionth of a second (one microsecond) is 1,000, and for a billionth of a second (one nanosecond) is 1.</p>		
<p>[c] In ODBC 2.x, the C date, time, and timestamp data types are SQL_C_DATE, SQL_C_TIME, and SQL_C_TIMESTAMP.</p>		
<p>[d] A number is stored in the val field of the SQL_NUMERIC_STRUCT structure as a scaled integer, in little endian mode (the leftmost byte being the least-significant byte). For example, the number 10.001 base 10, with a scale of 4, is scaled to an integer of 100010. Because this is 186AA in hexadecimal format, the value in SQL_NUMERIC_STRUCT would be "AA 86 01 00 00 ... 00", with the number of bytes defined by the SQL_MAX_NUMERIC_LEN #define.</p>		
<p>[e] The precision and scale fields of the SQL_C_NUMERIC data type are never used for input from an application, only for output from the driver to the application. When the driver writes a numeric value into the SQL_NUMERIC_STRUCT, it will use its own driver-specific default as the value for the precision field, and it will use the value in the SQL_DESC_SCALE field of the application descriptor (which defaults to 0) for the scale field. An application can provide its own values for precision and scale by setting the SQL_DESC_PRECISION and SQL_DESC_SCALE fields of the application descriptor.</p>		
<p>[f] The DECIMAL and NUMERIC data types take up more than one byte/character. The data types will actually be declared as arrays based on the precision required for the column. For example, a column of type SQL DECIMAL(10,4) might be declared as SQL_DECIMAL[13] to take into account the 10 digits, the sign character, the decimal point character, and the string terminator.</p>		
<p>[g] _int64 might not be supplied by some compilers.</p>		
<p>[h] _SQL_C_SHORT, SQL_C_LONG, and SQL_C_TINYINT have been replaced in ODBC by signed and unsigned types: SQL_C_SSHORT and SQL_C_USHORT, SQL_C_SLONG and SQL_C_ULONG, and SQL_C_STINYINT and SQL_C_UTINYINT. An ODBC 3.x driver that should work with ODBC 2.x applications should support SQL_C_SHORT, SQL_C_LONG, and SQL_C_TINYINT, because when they are called, the Driver Manager passes them through to the driver.</p>		
<p>[i] solidDB does not support unsigned SQL data types. You may bind an unsigned C data type to a signed SQL column, but you should not do this unless the values stored in the SQL column and the C variable are within the valid range for both data types. For example, since signed TINYINT columns hold values from -128 to +127, while unsigned SQL_C_UTINYINT variables hold values from 0 to 255, you may only store values between 0 and +127 in the column and bound variable if you want the values to be interpreted properly.</p>		

64-Bit Integer Structures

On Microsoft C compilers, the C data type identifiers `SQL_C_SBIGINT` and `SQL_C_UBIGINT` are defined as `_int64`. When a non-Microsoft C compiler is used, the C type may differ. If the compiler in use is supporting 64-bit integers natively, then define the driver or application `ODBCINT64` as the native 64-bit integer type. If the compiler in use does not support 64-bit integers natively, define the following structures to ensure access to these C types:

```
typedef struct{
SQLUIINTEGER dwLowWord;
SQLUIINTEGER dwHighWord;
} SQLUBIGINT
```

```
typedef struct {
SQLUIINTEGER dwLowWord;
SQLINTEGER sdwHighWord;
} SQLBIGINT
```

Because a 64-bit integer is aligned to the 8-byte boundary, be sure to align these structures to an 8-byte boundary.

Note:

solidDB supports signed `BIGINT`, but not unsigned `BIGINT`.

Default C data types

In applications that specify `SQL_C_DEFAULT` in `SQLBindCol`, `SQLGetData`, or `SQLBindParameter`, the driver assumes that the C data type of the output or input buffer corresponds to the SQL data type of the column or parameter to which the buffer is bound.

Important: To avoid compatibility problems when using different platforms, we strongly recommend that you avoid using `SQL_C_DEFAULT`. Instead, specify the C type of the buffer in use.

Drivers cannot always determine the correct default C type for these reasons:

- The DBMS may have promoted an SQL data type of a column or a parameter; in this case, the driver is unable to determine the original SQL data type and consequently, cannot determine the corresponding default C data type.
- The DBMS determined whether the data type of a column or parameter is signed or unsigned; in this case, the driver is unable to determine this for a particular SQL data type and consequently, cannot determine this for the corresponding default C data type.

See E.12, “Converting data from SQL to C data types,” on page 241.

SQL_C_TCHAR

The `SQL_C_TCHAR` type identifier is used for Unicode purposes. Use this identifier in applications that transfer character data and are compiled to use both ASCII and Unicode character sets. Note that the `SQL_C_TCHAR` is not a type identifier in the conventional sense; instead, it is a macro contained in the header file for Unicode conversion. `SQL_C_CHAR` or `SQL_C_WCHAR` replaces `SQL_C_TCHAR` depending on the setting of the `UNICODE` #define.

E.6 Numeric literals

To store numeric data values in character strings, you use numeric literals.

Numeric literal syntax specifies what is stored in the target during the following conversions:

- SQL data to an SQL_C_CHAR string
- C data to an SQL_CHAR or SQL_VARCHAR string

The syntax also validates what is stored in the source during the following conversions:

- numeric stored as an SQL_C_CHAR string to numeric SQL data
- numeric stored as an SQL_CHAR string to numeric C data

For more information, see “Numeric Literal Syntax” on page 219.

Conversion Rules

The following rules apply to conversions involving numeric literals. Following are terms used in this topic:

Table 157. Conversions Involving Numeric Literals

Term	Meaning
Store assignment	Refers to sending data into a table column in a database when calling SQLExecute and SQLExecDirect. During store assignment, "target" refers to a database column and "source" refers to data in application buffers.
Retrieval assignment	Refers to retrieving data from the database into application buffers when calling SQLFetch, SQLGetData, and SQLFetchScroll. During retrieval assignment, "target" refers to the application buffers and "source" refers to the database column.
CS	Value in the character source.
NT	Value in the numeric target.
NS	Value in the numeric source.
CT	Value in the character target.
Precision of an exact numeric literal	Number of digits that the literal contains.
Scale of an exact numeric literal	Number of digits to the right of the expressed or implied decimal point.
Precision of an approximate numeric literal	Precision of the literal's mantissa.

Rules for Character Source to Numeric Target

Following are the rules for converting from a character source (CS) to a numeric target (NT):

1. Replace CS with the value obtained by removing any leading or trailing spaces in CS. If CS is not a valid numeric-literal, SQLSTATE 22018 (Invalid character value for cast specification) is returned.
2. Replace CS with the value obtained by removing leading zeroes before the decimal point, trailing zeroes after the decimal point, or both.
3. Convert CS to NT. If the conversion results in a loss of significant digits, SQLSTATE 22003 (Numeric value out of range) is returned. If the conversion results in the loss of nonsignificant digits, SQLSTATE 01S07 (Fractional truncation) is returned.

Following are the rules for converting from a numeric source (NS) to a character target (CT):

1. Let LT be the length in characters of CT.
For retrieval assignment, LT is equal to the length of the buffer in characters minus the number of bytes in the null-termination character for this character set.
2. Take one of the following actions depending on the type of NS.
 - If NS is an exact numeric type, then let YP equal the shortest character string that conforms to the definition of exact-numeric-literal such that the scale of YP is the same as the scale of NS, and the interpreted value of YP is the absolute value of NS.
 - If NS is an approximate numeric type, then let YP be a character string as follows:
Case:
 - a. If NS is equal to 0, then YP is the string "0".
 - b. Let YSN be the shortest character string that conforms to the definition of exact-numeric-literal and whose interpreted value is the absolute value of NS. If the length of YSN is less than the (precision + 1) of the data type of NS, then let YP equal YSN.
 - c. Otherwise, YP is the shortest character string that conforms to the definition of approximate-numeric-literal whose interpreted value is the absolute value of NS and whose mantissa consists of a single digit that is not '0', followed by a period and an unsigned-integer.
3. If NS is less than 0, then let Y be the result of:
'-' || YP
where '||' is the string concatenation operator.
Otherwise, let Y equal YP.
4. Let LY be the length in characters of Y.

5. Take one of the following action depending on the value of LY.
 - If LY equals LT, then CT is set to Y.
 - If LY is less than LT, then CT is set to Y extended on the right by appropriate number of spaces.
 - Otherwise (LY > LT), copy the first LT characters of Y into CT.
 Case:
 - If this is a store assignment, return the error SQLSTATE 22001 (String data, right-truncated).
 - If this is retrieval assignment, return the warning SQLSTATE 01004 (String data, right-truncated). When the copy results in the loss of fractional digits (other than trailing zeros), depending on the driver definition, one of the following actions occurs:
 - a. The driver truncates the string in Y to an appropriate scale (which can be zero also) and writes the result into CT.
 - b. The driver rounds the string in Y to an appropriate scale (which can be zero also) and writes the result into CT.
 - c. The driver neither truncates nor rounds, but just copies the first LT characters of Y into CT.

E.7 Overriding default precision and scale for numeric data types

The following table provides the override default precision and scale values for numeric data type.

Table 158. Override Default Precision and Scale Values for Numeric Data Type

Function calls to	Setting	Override
SQLBindCol or SQLSetDescField	SQL_DESC_TYPE field in an ARD is set to SQL_C_NUMERIC	SQL_DESC_SCALE field in the ARD is set to 0 and the SQL_DESC_PRECISION field is set to a driver-defined default precision. [a]
SQLBindParameter or SQLSetDescField	SQL_DESC_SCALE field in an APD is set to SQL_C_NUMERIC	SQL_DESC_SCALE field in the ARD is set to 0 and the SQL_DESC_PRECISION field is set to a driver-defined default precision. This is true for input, input/output, or output parameters. [a]
SQLGetData	Data is returned into an SQL_C_NUMERIC structure	Default SQL_DESC_SCALE and SQL_DESC_PRECISION fields are used. [b]

Explanations of Footnote Numbering in the Table Above

[a] If the defaults are not acceptable for an application, the application can call the SQLSetDescField or SQLSetDescRec to set the SQL_DESC_SCALE or SQL_DESC_PRECISION field.

[b] If the defaults are not acceptable, the application must call SQLSetDescRec or SQLSetDescField to set the fields and then call SQLGetData with a targetType of SQL_ARD_TYPE to use the values in the descriptor fields.

E.8 Data type identifiers and descriptors

Unlike the "concise" SQL and C data types, where each identifier refers to a single data type, descriptors do not in all cases use a single value to identify data types. In some cases, descriptors use a verbose data type and a type subcode. For most data types, the verbose data type identifier matches the concise type identifier.

The exception, however, is the datetime and interval data types. For these data types:

- `SQL_DESC_TYPE` contains the verbose type (`SQL_DATETIME`)
- `SQL_DESC_CONCISE_TYPE` contains a concise type

For details on setting fields and a setting's effect on other fields, see the `SQLSetDescField` function description on the Microsoft ODBC web site.

When the `SQL_DESC_TYPE` or `SQL_DESC_CONCISE_TYPE` field is set for some data types, the following fields are set to default values appropriate for the data type:

- `SQL_DESC_DATETIME_INTERVAL_PRECISION`
- `SQL_DESC_LENGTH`
- `SQL_DESC_PRECISION`
- `SQL_DESC_SCALE`

For more information, see the `SQL_DESC_TYPE` field under `SQLSetDescField` function description on the Microsoft ODBC web site.

Note: If the default values set are not appropriate, you can explicitly set the descriptor field in the application by calling `SQLSetDescField`.

The following table lists for each SQL and C type identifier, the concise type identifier, verbose identifier, and type subcode for each datetime.

For datetime data types, the `SQL_DESC_TYPE` have the same manifest constants for both SQL data types (in implementation descriptors) and for C data types (in application descriptors):

Table 159. Concise Type Identifier, Verbose Identifier, and Type Subcode for Each Datetime

Concise SQL Type	Concise C Type	Verbose Type	DATETIME_INTERVAL_CODE (also called "type subcode")
<code>SQL_TYPE_DATE</code>	<code>SQL_C_TYPE_DATE</code>	<code>SQL_DATETIME</code>	<code>SQL_CODE_DATE</code>
<code>SQL_TYPE_TIME</code>	<code>SQL_C_TYPE_TIME</code>	<code>SQL_DATETIME</code>	<code>SQL_CODE_TIME</code>
<code>SQL_TYPE_TIMESTAMP</code>	<code>SQL_C_TYPE_TIMESTAMP</code>	<code>SQL_DATETIME</code>	<code>SQL_CODE_TIME STAMP</code>

Pseudo-type identifiers

ODBC defines a number of pseudo-type identifiers, which depending on the situation, resolve to existing data types. These identifiers do not correspond to actual data types, but are provided for your application programming convenience.

E.9 Decimal digits

Decimal digits apply to decimal and numeric data types. They refer to the maximum number of digits to the right of the decimal point, or the scale of the data.

Because the number of digits to the right of the decimal point is not fixed, the scale is undefined for approximate floating-point number columns or parameters. When datetime data contains a seconds component, the decimal digits are the number of digits to the right of the decimal point in the seconds component of the data.

Typically, the maximum scale matches the maximum precision for SQL_DECIMAL and SQL_NUMERIC data types. Some data sources, however, have their own maximum scale limit. An application can call SQLGetTypeInfo to determine the minimum and maximum scales allowed for a data type.

The following ODBC functions return parameter decimal attributes in an SQL statement data type or decimal attributes on a data source:

Table 160. ODBC Functions' Return Parameter

ODBC Function	Returns...
SQLDescribeCol	Decimal digits of the columns it describes.
SQLDescribeParam	Decimal digits of the parameters it describes.
SQLProcedureColumns	Decimal digits in a column of a procedure.
SQLColumns	Decimal digits in specified tables (such as the base table, view, or a system table).
SQLColAttribute	Decimal digits of columns at the data source.
SQLGetTypeInfo	Minimum and maximum decimal digits of an SQL data type on a data source.

Note: The SQLBindParameter sets the decimal digits for a parameter in an SQL statement.

The values returned by ODBC functions for decimal digits correspond to "scale" as defined in ODBC 2.x.

Descriptor fields describe the characteristics of a result set. They do not contain valid data values before statement execution. However, the decimal digits values returned by SQLColumns, SQLProcedureColumns, and SQLGetTypeInfo, do represent the characteristics of database objects, such as table columns and data types from the data source's catalog.

Each concise SQL data type has the following decimal digits definition as noted in the following table:

Table 161. SQL data type decimal digits

SQL Type Identifier	Decimal Digits
All character and binary types [a]	N/A
SQL_DECIMAL SQL_NUMERIC	The defined number of digits to the right of the decimal point. For example, the scale of a column defined as NUMERIC(10,3) is 3. (In some implementations, this can be a negative number to support storage of very large numbers without using exponential notation; for example, "12000" could be stored as "12" with a scale of -3. However, solidDB does not support negative scale.)
All exact numeric types other than SQL_DECIMAL and SQL_NUMERIC [a]	0
All approximate data types [a]	N/A
Note: [a] SQLBindParameter's DecimalDigits argument is ignored for this data type.	

For decimal digits, the values returned do not correspond to the values in any one descriptor field. The values returned (for example, in SQLColAttribute) for the decimal digits can come from either the SQL_DESC_SCALE or the SQL_DESC_PRECISION field, depending on the data type, as shown in the following table:

Table 162. Descriptor field corresponding to decimal digits

SQL Type Identifier	Descriptor field corresponding to decimal digits
All character and binary types	N/A
All exact numeric types	SCALE
All approximate numeric types	N/A
All datetime types	PRECISION

E.10 Transfer octet length

When data is transferred to its default C data type, an application receives a maximum number of bytes. This maximum is known as the transfer octet length of a column.

For character data, space for the null-termination character is not included in the transfer octet length. Note that the transfer octet length in bytes can differ from the number of bytes needed to store the data on the data source.

The following ODBC functions return parameter decimal attributes in an SQL statement data type or decimal attributes on a data source:

Table 163. ODBC Functions' Return parameter Decimal Attributes

ODBC Function	Returns
SQLColumns	Transfer octet length of a column in specified tables (such as the base table, view, or a system table).
SQLColAttribute	Transfer octet length of columns at the data source.
SQLProcedureColumns	Transfer octet length of a column in a procedure.

The values returned by ODBC functions for the transfer octet length may not correspond to the values returned in SQL_DESC_LENGTH. For all character and binary types, the values come from a descriptor field's SQL_DESC_OCTET_LENGTH. For other data types, there is no descriptor field that stores this information.

Descriptor fields describe the characteristics of a result set. They do not contain valid data values before statement execution. In its result set, SQLColAttribute returns the transfer octet length of columns at the data source; these values may not match the values in the SQL_DESC_OCTET_LENGTH descriptor fields. For more information about descriptor fields, see SQLSetDescField function description on the Microsoft ODBC Web site.

Each concise SQL data type has the following transfer octet length definition as noted in the table below.

Table 164. Transfer Octet Lengths

SQL Type Identifier	Transfer Octet Length
All character and binary types [a]	The defined or the maximum (for variable type) length of the column in bytes. This value matches the one in the SQL_DESC_OCTET_LENGTH descriptor field.
SQL_DECIMAL SQL_NUMERIC	The number of bytes required to hold the character representation of this data if the character set is ASCII, and twice this number if the character set is UNICODE. The character representation is the maximum number of digits plus two; the data is returned as a character string, where the characters are needed for digits, a sign, and a decimal point. For example, the transfer length of a column defined as NUMERIC(10,3) is 12 because there are 10 bytes for the digits, 1 byte for the sign, and 1 byte for the decimal point.
SQL_TINYINT	1
SQL_SMALLINT	2
SQL_INTEGER	4

Table 164. Transfer Octet Lengths (continued)

SQL Type Identifier	Transfer Octet Length
SQL_BIGINT	The number of bytes required to hold the character representation of this data if the character set is ASCII, and twice this number if the character set is UNICODE. This data type is returned as a character string by default. The character representation consists of 20 characters for 19 digits and a sign (if signed), or 20 digits (if unsigned). The length is 20. solidDB supports only signed, not unsigned, BIGINT.
SQL_REAL	4
SQL_FLOAT	8
SQL_DOUBLE	8
All binary types [a]	The number of bytes required to store the defined (for fixed types) or maximum (for variable types) number of characters.
SQL_TYPE_DATE SQL_TYPE_TIME	6 (size of the structures SQL_DATE_STRUCT or SQL_TIME_STRUCT).
SQL_TYPE_TIMESTAMP	16 (size of the structure SQL_TIMESTAMP_STRUCT).

Explanations of Footnote Numbering in the Table Above

[a] SQL_NO_TOTAL is returned when the driver cannot determine the column or parameter length for variable types.

E.11 Constraints of the gregorian calendar

The following table contains the Gregorian calendar constraints for date and datetime data types.

Table 165. Constraints of the Gregorian Calendar

Value	Requirement
month field	Must be between 1 and 12, inclusive.
day field	Range must be from 1 through the number of days in the month, which is determined from the values of the year and months fields and can be 28, 29, 30, or 31. A leap year can also affect the number of days in the month.
hour field	Must be between 0 and 23, inclusive.
minute field	Must be between 0 and 59, inclusive.

Table 165. Constraints of the Gregorian Calendar (continued)

Value	Requirement
trailing seconds field	Must be between 0 and 61.9(n), inclusive, where n specifies the number of digits at the place of "9" and the value of n is the fractional seconds precision. The range of seconds permits a maximum of two leap seconds to maintain synchronization of sidereal time.

E.12 Converting data from SQL to C data types

This section provides information about converting data from SQL to C data types.

When an application calls `SQLFetch`, `SQLFetchScroll`, or `SQLGetData`, the driver retrieves the data from the data source. If necessary, it converts the data from the data type in which the driver retrieved it to the data type specified by the `TargetType` argument in `SQLBindCol` or `SQLGetData`. Finally, it stores the data in the location pointed to by the `TargetValuePtr` argument in `SQLBindCol` or `SQLGetData` (and the `SQL_DESC_DATA_PTR` field of the ARD).

The following table shows the supported conversions from ODBC SQL data types to ODBC C data types. A solid circle indicates the default conversion for an SQL data type (the C data type to which the data will be converted when the value of `TargetType` is `SQL_C_DEFAULT`). A hollow circle indicates a supported conversion.

For an ODBC 3.x application working with an ODBC 2.x driver, conversion from driver-specific data types might not be supported.

The format of the converted data is not affected by the Windows country setting.

solidDB supports only signed, not unsigned, integer data types (`SQL_TINYINT`, `SQL_SMALLINT`, `SQL_INTEGER`, `SQL_BIGINT`). You may bind an unsigned C variable to a signed SQL column, but you must make sure that the values you store fit within the range supported by both data types.

solidDB does not support the `BIT/SQL_BIT` data type for SQL columns. However, you may bind a numeric SQL column to a `BIT` data type in your C application. For example, you may use a `TINYINT` column in your database and bind that column to a C variable of type `SQL_C_BIT`. The solidDB ODBC driver will try to convert numeric types in the database to `BIT` data types for the C variables. The numeric data values must be 1 or 0 or `NULL`; other values cause a data conversion error. The table below does not discuss `BIT/SQL_BIT` data types.

CAUTION:

Although the table shows a wide range of ODBC conversions, including conversions involving unsigned data types, solidDB supports only signed integer data types (for example, `TINYINT`, `SMALLINT`, `INTEGER`, and `BIGINT`).

Table 166. C Data Type — SQL_C_datatype where Datatype Is:

SQL Data Type	C	W	S	U	T	S	U	S	S	S	U	L	L	S	U	D	N	B	D	T	T	
	H	H	T	T	T	H	S	H	H	L	L	L	L	B	B	F	D	N	B	D	T	T
	A	A	I	I	I	O	O	O	O	O	O	O	O	I	I	L	O	U	I	A	I	M
	R	R	N	N	N	R	R	R	R	O	O	O	O	N	N	O	B	E	R	T	M	P
SQL_CHAR	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_VARCHAR	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_LONGVARCHAR	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_WCHAR	o	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_WVARCHAR	o	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_WLONGVARCHAR	o	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_TINYINT (signed)	o	o	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o			
SQL_TINYINT (unsigned)	o	o	o	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o			
SQL_SMALLINT (signed)	o	o	o	o	o	*	o	o	o	o	o	o	o	o	o	o	o	o	o			
SQL_SMALLINT (unsigned)	o	o	o	o	o	o	*	o	o	o	o	o	o	o	o	o	o	o	o			
SQL_INTEGER (signed)	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o	o	o	o	o			
SQL_INTEGER (unsigned)	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o	o	o	o			
SQL_BIGINT (signed)	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o			
SQL_BIGINT (unsigned)	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o			
SQL_REAL	o	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o			
SQL_FLOAT	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o			
SQL_DOUBLE	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o			
SQL_DECIMAL	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o			
SQL_NUMERIC	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o			
SQL_BINARY	o	o																	*			
SQL_VARBINARY	o	o																	*			
SQL_LONGVARBINARY	o	o																	*			

Table 166. C Data Type — SQL_C_datatype where Datatype Is: (continued)

SQL Data Type	C	W	S	U	T	S	S	S	S	U	L	L	L	S	U	D	N	B	D	T	T	
	H	H	I	I	I	H	H	H	O	O	O	O	O	I	I	O	E	I	A	A	I	M
	A	A	N	N	N	R	R	R	R	R	N	N	N	N	N	B	R	R	T	T	M	P
	R	R	T	T	T	T	T	T	T	T	G	G	G	T	T	E	C	Y	E	E	M	P
SQL_TYPE_DATE	o	o																o	*		o	
SQL_TYPE_TIME	o	o																o		*	o	
SQL_TYPE_TIMESTAMP	o	o																o	o	o	*	

* These datatypes have the word "TYPE" in the datatype name. For example, SQL_C_TYPE_DATE, SQL_C_TYPE_TIME, and SQL_C_TYPE_TIMESTAMP.

Legend:

* Default Conversion

o Supported Conversion

E.12.1 Data conversion tables from SQL to C

The tables in the following sections describe how the driver or data source converts data retrieved from the data source; drivers are required to support conversions to all ODBC C data types from the ODBC SQL data types that they support.

Conversion Table Description (SQL to C)

The following columns are included in the tables:

- For a given ODBC SQL data type, the first column of the table lists the legal input values of the TargetType argument in SQLBindCol and SQLGetData.
- The second column lists the outcomes of a test, often using the BufferLength argument specified in SQLBindCol or SQLGetData, which the driver performs to determine if it can convert the data.
- For each outcome, the third and fourth columns list the values placed in the buffers specified by the TargetValuePtr and StrLen_or_IndPtr arguments specified in SQLBindCol or SQLGetData after the driver has attempted to convert the data. (The StrLen_or_IndPtr argument corresponds to the SQL_DESC_OCTET_LENGTH_PTR field of the ARD.)
- The last column lists the SQLSTATE returned for each outcome by SQLFetch, SQLFetchScroll, or SQLGetData.

If the TargetType argument in SQLBindCol or SQLGetData contains a value for an ODBC C data type not shown in the table for a given ODBC SQL data type, SQLFetch, SQLFetchScroll, or SQLGetData returns SQLSTATE 07006 (Restricted data type attribute violation). If the TargetType argument contains a value that specifies a conversion from a driver-specific SQL data type to an ODBC C data

type and this conversion is not supported by the driver, SQLFetch, SQLFetchScroll, or SQLGetData returns SQLSTATE HYC00 (Optional feature not implemented).

Although it is not shown in the tables, the driver returns SQL_NULL_DATA in the buffer specified by the StrLen_or_IndPtr argument when the SQL data value is NULL. The length specified by StrLen_or_IndPtr does not include the null-termination byte. If TargetValuePtr is a null pointer, SQLGetData returns SQLSTATE HY009 (Invalid use of null pointer); in SQLBindCol, this unbinds the columns.

The following terms and conventions are used in the tables:

- Byte length of data is the number of bytes of C data available to return in *TargetValuePtr, whether or not the data will be truncated before it is returned to the application. For string data, this does not include the space for the null-termination character.
- Character byte length is the total number of bytes needed to display the data in character format.
- Words in italics represent function arguments or elements of the SQL grammar. See Appendix D, “Minimum SQL grammar requirements for ODBC,” on page 215 for the syntax of grammar elements.

SQL to C: Character

The character ODBC SQL data types are:

SQL_CHAR
 SQL_VARCHAR
 SQL_LONGVARCHAR
 SQL_WCHAR
 SQL_WVARCHAR
 SQL_WLONGVARCHAR

The following table shows the ODBC C data types to which character SQL data can be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (SQL to C)” on page 243.

Table 167. Character SQL Data to ODBC C Data Types

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_CHAR	Byte length of data < BufferLength	Data	Length of data in bytes	N/A
	Byte length of data >= BufferLength	Truncated data	Length of data in bytes	01004
SQL_C_WCHAR	Character length of data < BufferLength	Data	Length of data in characters	N/A
	(Character length of data) >= BufferLength	Truncated data	Length of data in characters	01004

Table 167. Character SQL Data to ODBC C Data Types (continued)

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
EXACT NUMERIC TYPES [h]	Data converted without truncation [b]	Data	Number of bytes of the C data type	N/A
SQL_C_STINYINT	Data converted with truncation of fractional digits [a]	Truncated data	Number of bytes of the C data type	01S07
SQL_C_UTINYINT		Undefined	Number of bytes of the C data type	22003
SQL_C_TINYINT	Conversion of data would result in loss of whole (as opposed to fractional) digits [b]	Undefined	Undefined	22018
SQL_C_SSHORT			Undefined	
SQL_C_USHORT				
SQL_C_SHORT	Data is not a numeric-literal [b]			
SQL_C_SLONG				
SQL_C_ULONG				
SQL_C_LONG				
SQL_C_SBIGINT				
SQL_C_UBIGINT				
SQL_C_NUMERIC				
APPROXIMATE NUMERIC TYPES [h]	Data is within the range of the data type to which the number is being converted [a]	Data	Size of the C data type	N/A
SQL_C_FLOAT		Undefined	Undefined	2003
SQL_C_DOUBLE	Data is outside the range of the data type to which the number is being converted [a]	Undefined	Undefined	22018
	Data is not a numeric-literal [b]			
SQL_C_BINARY	Byte length of data <= BufferLength	Data	Length of data	N/A
	Byte length of data > BufferLength	Truncated data	Length of data	01004

Table 167. Character SQL Data to ODBC C Data Types (continued)

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_TYPE_DATE	Data value is a valid date-value [a]	Data	6 [b]	N/A
		Data	6 [b]	N/A
	Data value is a valid timestamp-value; time portion is zero [a]	Truncated data	6 [b]	01S07
		Undefined	Undefined	22018
	Data value is a valid timestamp-value; time portion is nonzero [a], [c], Data value is not a valid date-value or timestamp_value [a]			
SQL_C_TYPE_TIME	Data value is a valid time-value and the fractional seconds value is 0 [a]	Data	6 [b]	N/A
		Data	6 [b]	N/A
	Data value is a valid timestamp-value or a valid time_value; fractional seconds portion is zero [a],[d]	Truncated data	6 [b]	01S07
		Undefined	Undefined	22018
	Data value is a valid timestamp-value ; fractional seconds portion is nonzero [a], [d], [e] Data value is not a valid timestamp-value or time_value [a]			
SQL_C_TYPE_TIMESTAMP	Data value is a valid timestamp-value or a valid time_value; fractional seconds portion not truncated [a], [d]	Data	16 [b]	N/A
		Truncated data	16 [b]	01S07
		Data [f]	16 [b]	N/A
	Data value is a valid timestamp-value or a valid time_value; fractional seconds portion truncated [a]	Data [g]	16 [b]	N/A
		Undefined	Undefined	22018
	Data value is a valid date-value [a]			
	Data value is a valid time_value [a] Data value is not a valid date_value, time_value, or timestamp_value [a]			

Table 167. Character SQL Data to ODBC C Data Types (continued)

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
Note:				
[a] The value of BufferLength is ignored for this conversion. The driver assumes that the size of *TargetValuePtr is the size of the C data type.				
[b] This is the size of the corresponding C data type.				
[c] The time portion of the timestamp-value is truncated.				
[d] The date portion of the timestamp-value is ignored.				
[e] The fractional seconds portion of the timestamp is truncated.				
[f] The time fields of the timestamp structure are set to zero.				
[g] The date fields of the timestamp structure are set to the current date.				
[h] The exact numeric types include NUMERIC/DECIMAL as well as integer. These data types store the exact value that you specify, as long as it is within the precision of the data type. The approximate data types include FLOAT/REAL, which store only approximately the value that you specify (in some cases, the least significant digit may be slightly different from what you specified).				

When character SQL data is converted to numeric, date, time, or timestamp C data, leading and trailing spaces are ignored.

SQL to C: Numeric

SQL_DECIMAL	SQL_BIGINT
SQL_NUMERIC	SQL_REAL
SQL_TINYINT	SQL_FLOAT
SQL_SMALLINT	SQL_DOUBLE
SQL_INTEGER	

The following table shows the ODBC C data types to which numeric SQL data may be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (SQL to C)” on page 243.

Table 168. SQL Data to ODBC C Data Types

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_CHAR	Character byte length < BufferLength	Data	Length of data in bytes	N/A
	Number of whole (as opposed to fractional) digits < BufferLength	Truncated data	Length of data in bytes	01004
	Number of whole (as opposed to fractional) digits ≥ BufferLength	Undefined	Undefined	22003

Table 168. SQL Data to ODBC C Data Types (continued)

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_WCHAR	Character length < BufferLength Number of whole (as opposed to fractional) digits < BufferLength Number of whole (as opposed to fractional) digits ≥ BufferLength	Data Truncated data Undefined	Length of data in bytes Length of data in bytes Undefined	N/A 01004 22003
EXACT NUMERIC TYPES [c] SQL_C_STINYINT SQL_C_UTINYINT SQL_C_TINYINT SQL_C_SBIGINT SQL_C_UBIGINT SQL_C_SSHORT SQL_C_USHORT SQL_C_SHORT a SQL_C_SLONG SQL_C_ULONG SQL_C_LONG SQL_C_NUMERIC	Data converted without truncation [a] Data converted with truncation of fractional digits [a] Conversion of data would result in loss of whole (as opposed to fractional) digits [a]	Data Truncated data Undefined	Size of the C data type Size of the C data type Undefined	N/A 01S07 22003
APPROXIMATE NUMERIC TYPES [c] SQL_C_FLOAT SQL_C_DOUBLE	Data is within the range of the data type to which the number is being converted [a] Data is outside the range of the data type to which the number is being converted [a]	Data Undefined	Size of the C data type Undefined	N/A 22003
SQL_C_BINARY	Length of data ≤ BufferLength Length of data > BufferLength	Data Undefined	Length of data Undefined	N/A 22003

Table 168. SQL Data to ODBC C Data Types (continued)

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
Note:				
[a] The value of BufferLength is ignored for this conversion. The driver assumes that the size of *TargetValuePtr is the size of the C data type.				
[b] This is the size of the corresponding C data type.				
[c] The exact numeric types include NUMERIC/DECIMAL as well as integer. These data types store the exact value that you specify, as long as it is within the precision of the data type. The approximate data types include FLOAT/REAL, which store only approximately the value that you specify (in some cases, the least significant digit may be slightly different from what you specified).				

SQL to C: Binary

The binary ODBC SQL data types are:

SQL_BINARY
 SQL_VARBINARY
 SQL_LONGVARBINARY

The following table shows the ODBC C data types to which binary SQL data may be converted. For an explanation of the columns and terms in the table, see "Conversion Table Description (SQL to C)" on page 243.

Table 169. Binary SQL Data to ODBC C Data Types

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_CHAR	(Byte length of data) * 2 < BufferLength	Data	Length of data in bytes	N/A
	(Byte length of data) * 2 >= BufferLength	Truncated data	Length of data in bytes	01004
SQL_C_WCHAR	(Character length of data) * 2 < BufferLength	Data	Length of data in bytes	N/A
	(Character length of data) * 2 >= BufferLength	Truncated data	Length of data in bytes	01004
SQL_C_BINARY	Byte length of data <= BufferLength	Data	Length of data in bytes	N/A
	Byte Length of data > BufferLength	Truncated data	Length of data in bytes	01004

When binary SQL data is converted to character C data, each byte (8 bits) of source data is represented as two ASCII characters. These characters are the ASCII character representation of the number in its hexadecimal form. For example, a binary 00000001 is converted to "01" and a binary 11111111 is converted to "FF".

The driver always converts individual bytes to pairs of hexadecimal digits and terminates the character string with a null byte. Because of this, if BufferLength is even and is less than the length of the converted data, the last byte of the

*TargetValuePtr buffer is not used. (The converted data requires an even number of bytes, the next-to-last byte is a null byte, and the last byte cannot be used.)

Application developers are discouraged from binding binary SQL data to a character C data type. This conversion is usually inefficient and slow.

SQL to C: Date

The date ODBC SQL data type is:

SQL_DATE

The following table shows the ODBC C data types to which date SQL data may be converted. For an explanation of the columns and terms in the table, see "Conversion Table Description (SQL to C)" on page 243.

Table 170. Date SQL Data to ODBC C Data Types

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_CHAR	BufferLength > Character byte length	Data	10	N/A
	11 <= BufferLength <= Character byte length	Truncated data	Length of data in bytes	01004
	BufferLength < 11	Undefined	Undefined	22003
SQL_C_WCHAR	BufferLength > Character length	Data	10	N/A
	11 <= BufferLength <= Character length	Truncated data	Length of data in bytes	01004
	BufferLength < 11	Undefined	Undefined	22003
SQL_C_BINARY	Byte length of data <= BufferLength	Data	Length of data in bytes	N/A
	Byte length of data > BufferLength	Undefined	Undefined	22003
SQL_C_DATE	None [a]	Data	6 [c]	N/A
SQL_C_TIMESTAMP	None [a]	Data [b]	16 [c]	N/A
Note:				
[a] The value of BufferLength is ignored for this conversion. The driver assumes that the size of *TargetValuePtr is the size of the C data type.				
[b] The time fields of the timestamp structure are set to zero.				
[c] This is the size of the corresponding C data type.				

When date SQL data is converted to character C data, the resulting string is in the "yyyy-mm-dd" format. This format is not affected by the Windows country setting.

SQL to C: Time

The time ODBC SQL data type is:

SQL_TIME

The following table shows the ODBC C data types to which time SQL data may be converted. For an explanation of the columns and terms in the table, see "Conversion Table Description (SQL to C)" on page 243.

Table 171. Time SQL Data to ODBC C Data Types

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_CHAR	BufferLength > Character byte length	Data	Length of data in bytes	N/A
	9 <= BufferLength <= Character byte length	Truncated data [a]	Length of data in bytes	01004
	BufferLength < 9	Undefined	Undefined	22003
SQL_C_WCHAR	BufferLength > Character byte length	Data	Length of data in characters	N/A
	9 <= BufferLength <= Character byte length	Truncated data [a]	Length of data in characters	01004
	BufferLength < 9	Undefined	Undefined	22003
SQL_C_BINARY	Byte length of data <= BufferLength	Data	Length of data in bytes	N/A
	Byte length of data > BufferLength	Undefined	Undefined	22003
SQL_C_DATE	None [a]	Data	6 [c]	N/A
SQL_C_TIMESTAMP	None [a]	Data [b]	16 [c]	N/A
<p>Note:</p> <p>[a]: The fractional seconds of the time are truncated.</p> <p>[b]: The value of BufferLength is ignored for this conversion. The driver assumes that the size of *TargetValuePtr is the size of the C data type.</p> <p>[c]: The date fields of the timestamp structure are set to the current date and the fractional seconds field of the timestamp structure is set to zero.</p> <p>[d]: This is the size of the corresponding C data type.</p>				

When time SQL data is converted to character C data, the resulting string is in the "hh:mm:ss" format.

SQL to C: Timestamp

The timestamp ODBC SQL data type is:

SQL_TIMESTAMP

The following table shows the ODBC C data types to which timestamp SQL data may be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (SQL to C)” on page 243.

Table 172. Timestamp SQL Data to ODBC C Data Types

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
SQL_C_CHAR	BufferLength > Character byte length	Data	Length of data in bytes	N/A
	20 <= BufferLength <= Character byte length	Truncated data [b]	Length of data in bytes	01004
	BufferLength < 20	Undefined	Undefined	22003
SQL_C_WCHAR	BufferLength > Character byte length	Data	Length of data in characters	N/A
	20 <= BufferLength <= Character byte length	Truncated data [b]	Length of data in characters	01004
	BufferLength < 20	Undefined	Undefined	22003
SQL_C_BINARY	Byte length of data <= BufferLength	Data	Length of data in bytes	N/A
	Byte length of data > BufferLength	Undefined	Undefined	22003
SQL_C_TYPE_DATE	Time portion of timestamp is zero [a]	Data	6 [f]	N/A
	Time portion of timestamp is non-zero [a]	Truncated data [c]	6 [f]	01S07
SQL_C_TYPE_TIME	Fractional seconds portion of timestamp is zero [a]	Data [d]	6 [f]	N/A
	Fractional seconds portion of timestamp is non-zero [a]	Truncated data [d], [e]	6 [f]	01S07
SQL_C_TYPE_TIMESTAMP	Fractional seconds portion of timestamp is not truncated [a]	Data [e]	6 [f]	N/A
	Fractional seconds portion of timestamp is truncated [a]	Truncated data [e]	6 [f]	01S07

Table 172. Timestamp SQL Data to ODBC C Data Types (continued)

C Type Identifier	Test	*TargetValuePtr	*StrLen_or_IndPtr	SQLSTATE
Note:				
[a] The value of BufferLength is ignored for this conversion. The driver assumes that the size of *TargetValuePtr is the size of the C data type.				
[b] The fractional seconds of the timestamp are truncated.				
[c] The time portion of the timestamp is truncated.				
[d] The date portion of the timestamp is ignored.				
[e] The fractional seconds portion of the timestamp is truncated.				
[f] This is the size of the corresponding C data type.				

When timestamp SQL data is converted to character C data, the resulting string is in the "yyyy-mm-dd hh:mm:ss [.f ...]"format, where up to nine digits may be used for fractional seconds. The format is not affected by the Windows country setting. (Except for the decimal point and fractional seconds, the entire format must be used, regardless of the precision of the timestamp SQL data type.)

E.12.2 SQL to C data conversion examples

The following examples illustrate how the driver converts SQL data to C data.

Table 173. SQL to C Data Conversion Examples

SQL Type Identifier	SQL Data Values	C Type Identifier	Buffer Length	*TargetValuePtr	SQLSTATES
SQL_CHAR	abcdef	SQL_C_CHAR	7	abcdef\0 [a]	N/A
SQL_CHAR	abcdef	SQL_C_CHAR	6	abcde\0 [a]	01004
SQL_DECIMAL	1234.56	SQL_C_CHAR	8	1234.56\0 [a]	N/A
SQL_DECIMAL	1234.56	SQL_C_CHAR	5	1234\0 [a]	01004
SQL_DECIMAL	1234.56	SQL_C_CHAR	4	----	22003
SQL_DECIMAL	1234.56	SQL_C_FLOAT	ignored	1234.56	N/A
SQL_DECIMAL	1234.56	SQL_C_SSHORT	ignored	1234	01S07
SQL_DECIMAL	1234.56	SQL_C_STINYINT	ignored	----	22003
SQL_DOUBLE	1.2345678	SQL_C_DOUBLE	ignored	1.2345678	N/A
SQL_DOUBLE	1.2345678	SQL_C_FLOAT	ignored	1.234567	N/A
SQL_DOUBLE	1.2345678	SQL_C_STINYINT	ignored	1	N/A

Table 173. SQL to C Data Conversion Examples (continued)

SQL Type Identifier	SQL Data Values	C Type Identifier	Buffer Length	*TargetValuePtr	SQLSTATES
SQL_TYPE_DATE	1992-12-31	SQL_C_CHAR	11	1992-12-31\0 [a]	N/A
SQL_TYPE_DATE	1992-12-31	SQL_C_CHAR	10	----	22003
SQL_TYPE_DATE	1992-12-31	SQL_C_TIMESTAMP	ignored	1992,12,31, 0,0,0,0 [b]	N/A
SQL_TYPE_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	23	1992-12-31 23:45:55.12\0 [a]	N/A
SQL_TYPE_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	22	1992-12-31 23:45:55.1\0 [a]	01004
SQL_TYPE_TIMESTAMP	1992-12-31 23:45:55.12	SQL_C_CHAR	18	----	22003

[a] "\0" represents a null-termination byte. The driver always null-terminates SQL_C_CHAR data.

[b] The numbers in this list are the numbers stored in the fields of the TIMESTAMP_STRUCT structure.

E.13 Converting data from C to SQL data types

This section provides information about converting data from C to SQL data types.

When an application calls `SQLExecute` or `SQLExecDirect`, the driver retrieves the data for any parameters bound with `SQLBindParameter` from storage locations in the application. For data-at-execution parameters, the application sends the parameter data with `SQLPutData`. If necessary, the driver converts the data from the data type specified by the `ValueType` argument in `SQLBindParameter` to the data type specified by the `ParameterType` argument in `SQLBindParameter`. Finally, the driver sends the data to the data source.

The following table shows the supported conversions from ODBC C data types to ODBC SQL data types. A solid circle indicates the default conversion for an SQL data type (the C data type from which the data will be converted when the value of `ValueType` or the `SQL_DESC_CONCISE_TYPE` descriptor field is `SQL_C_DEFAULT`). A hollow circle indicates a supported conversion.

The format of the converted data is not affected by the Windows country setting.

solidDB supports only signed, not unsigned, integer data types (`SQL_TINYINT`, `SQL_SMALLINT`, `SQL_INTEGER`, `SQL_BIGINT`). You may bind an unsigned C variable to a signed SQL column, but you must make sure that the values you store fit within the range supported by both data types.

solidDB does not support the `BIT/SQL_BIT` data type for SQL columns. However, you may bind a numeric SQL column to a `BIT` data type in your C application. For

example, you may use a TINYINT column in your database and bind that column to a C variable of type SQL_C_BIT. The solidDB ODBC driver will try to convert numeric types in the database to BIT data types for the C variables. The numeric data values must be 1 or 0 or NULL; other values cause a data conversion error. The table below does not discuss BIT/SQL_BIT data types.

CAUTION:

Although the table above shows a wide range of ODBC conversions, including conversions involving unsigned data types, solidDB supports only signed integer data types (for example, TINYINT, SMALLINT, INT, and BIGINT).

Table 174. SQL Data Type — SQL_datatype where Datatype Is:

C Data Type	CHAR	CHAR	CHAR	WCHAR	WCHAR	WCHAR	DECIMAL	NUMERIC	TINYINT (signed)	TINYINT (unsigned)	SMALLINT (signed)	SMALLINT (unsigned)	INTEGER (signed)	INTEGER (unsigned)	BIGINT (signed)	BIGINT (unsigned)	REAL	FLOAT	DOUBLE	BINARY	VARBINARY	LONGVARBINARY	DATE	TIME	TIMESTAMP
SQL_C_CHAR	*	*	*	o	o	o	*	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_C_WCHAR	o	o	o	*	*	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o
SQL_C_NUMERIC	*	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o						
SQL_C_STINYINT	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o	o	o	o	o						
SQL_C_UTINYINT	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o	o	o	o						
SQL_C_TINYINT	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o						
SQL_C_SSHORT	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o	o	o						
SQL_C_USHORT	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o	o						
SQL_C_SHORT	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o						
SQL_C_SLONG	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o	o						
SQL_C_ULONG	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o	o	o						
SQL_C_LONG	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o						
SQL_C_SBIGINT	o	o	o	o	o	o	o	o	o	o	o	o	o	o	*	*	o	o	o						
SQL_C_UBIGINT	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o	o						
SQL_C_FLOAT	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	o	*	o	o						

driver does not support the conversion from the specific ODBC C data type to that driver-specific SQL data type, SQLBindParameter returns SQLSTATE HYC00 (Optional feature not implemented).

If the ParameterValuePtr and StrLen_or_IndPtr arguments specified in SQLBindParameter are both null pointers, that function returns SQLSTATE HY009 (Invalid use of null pointer). Although it is not shown in the tables, an application sets the value pointed to by the StrLen_or_IndPtr argument of SQLBindParameter or the value of the StrLen_or_IndPtr argument to SQL_NULL_DATA to specify a NULL SQL data value. (The StrLen_or_IndPtr argument corresponds to the SQL_DESC_OCTET_LENGTH_PTR field of the APD.) The application sets these values to SQL_NTS to specify that the value in *ParameterValuePtr in SQLBindParameter or *DataPtr in SQLPutData (pointed to by the SQL_DESC_DATA_PTR field of the APD) is a null-terminated string.

The following terms are used in the tables:

- *Byte length of data* is the number of bytes of SQL data available to send to the data source, regardless of whether the data will be truncated before it is sent to the data source. For string data, this does not include the null-termination character.
- *Column byte length* is the number of bytes required to store the data at the data source.
- *Character byte length* is the maximum number of bytes needed to display data in character form.
- *Number of digits* is the number of characters used to represent a number, including the minus sign, decimal point, and exponent (if needed).
- Words in italics represent elements of the ODBC SQL grammar. For the syntax of grammar elements, see Appendix D, “Minimum SQL grammar requirements for ODBC,” on page 215.

C to SQL: Character

The character ODBC C data type is:

SQL_C_CHAR SQL_C_WCHAR

The following table shows the ODBC SQL data types to which C character data may be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (C to SQL)” on page 256.

Note: The length of the Unicode data type must be an even number when character C data is converted to Unicode SQL data.

Table 175. C Character Data to ODBC SQL Data Types

SQL Type Identifier	Test	SQLSTATE
SQL_CHAR	Byte length of data <= Column length	N/A
SQL_VARCHAR		22001
SQL_LONGVARCHAR	Byte length of data > Column length	

Table 175. C Character Data to ODBC SQL Data Types (continued)

SQL Type Identifier	Test	SQLSTATE
SQL_WCHAR	Character length of data <= Column length	N/A
SQL_WVARCHAR		22001
SQL_WLONGVARCHAR	Character length of data > Column length	
SQL_DECIMAL	Data converted without truncation	N/A
SQL_NUMERIC	Data converted with truncation of fractional digits [e]	22001
SQL_TINYINT		22001
SQL_SMALLINT	Conversion of data would result in loss of whole (as opposed to fractional) digits [e]	22018
SQL_INTEGER		
SQL_BIGINT	Data value is not a <i>numeric-literal</i>	
SQL_REAL	Data is within the range of the data type to which the number is being converted	N/A
SQL_FLOAT		22003
SQL_DOUBLE	Data is outside the range of the data type to which the number is being converted	22005
	Data value is not a <i>numeric-literal</i>	
SQL_BIT	Data is 0 or 1	N/A
	Data is greater than 0, less than 2, and not equal to 1	22001
		22003
	Data is less than 0 or greater than or equal to 2	22018
	Data is not a <i>numeric-literal</i> .	
	Note: solidDB does not support SQL_BIT.	
SQL_BINARY	(Byte length of data) / 2 <= Column byte length	N/A
SQL_VARBINARY		22001
SQL_LONG-VARBINARY	(Byte length of data) / 2 > Column byte length	22018
	Data value is not a hexadecimal value	

Table 175. C Character Data to ODBC SQL Data Types (continued)

SQL Type Identifier	Test	SQLSTATE
SQL_TYPE_DATE	<p>Data value is a valid <i>ODBC_date_literal</i></p> <p>Data value is a valid <i>ODBC_timestamp_literal</i>; time portion is zero</p> <p>Data value is a valid <i>ODBC_timestamp_literal</i>; time portion is non-zero [a]</p> <p>Data value is not a valid <i>ODBC_date_literal</i> or <i>ODBC_timestamp_literal</i></p>	<p>N/A</p> <p>N/A</p> <p>22008</p> <p>22018</p>
SQL_TYPE_TIME	<p>Data value is a <i>valid ODBC_time_literal</i></p> <p>Data value is a valid <i>ODBC_timestamp_literal</i>; fractional seconds portion is zero [b]</p> <p>Data value is a valid <i>ODBC_timestamp_literal</i>; fractional seconds portion is non-zero [b]</p> <p>Data value is not a valid <i>ODBC_time_literal</i> or <i>ODBC_timestamp_literal</i></p>	<p>N/A</p> <p>N/A</p> <p>22008</p> <p>22018</p>
SQL_TYPE_TIMESTAMP	<p>Data value is a valid <i>ODBC_timestamp_literal</i>; fractional seconds portion not truncated</p> <p>Data value is a valid <i>ODBC-time-literal</i>; fractional seconds portion truncated</p> <p>Data value is a valid <i>ODBC-date-literal</i> [c]</p> <p>Data value is a valid <i>ODBC-time-literal</i> [d]</p> <p>Data value is not a valid <i>ODBC-date-literal</i>, <i>ODBC-time-literal</i>, or <i>ODBC-timestamp-literal</i></p>	<p>N/A</p> <p>22008</p> <p>N/A</p> <p>N/A</p> <p>22018</p>

Table 175. C Character Data to ODBC SQL Data Types (continued)

SQL Type Identifier	Test	SQLSTATE
<p>Note:</p> <p>[a] The time portion of the timestamp is truncated.</p> <p>[b] The date portion of the timestamp is ignored.</p> <p>[c] The time portion of the timestamp is set to zero.</p> <p>[d] The date portion of the timestamp is set to the current date.</p> <p>[e] The driver/data source effectively waits until the entire string has been received (even if the character data is sent in pieces by calls to SQLPutData) before attempting to perform the conversion.</p>		

When character C data is converted to numeric, date, time, or timestamp SQL data, leading and trailing blanks are ignored.

When character C data is converted to binary SQL data, each two bytes of character data are converted to a single byte (8 bits) of binary data. Each two bytes of character data represent a number in hexadecimal form. For example, "01" is converted to a binary 00000001 and "FF" is converted to a binary 11111111.

The driver always converts pairs of hexadecimal digits to individual bytes and ignores the null termination byte. Because of this, if the length of the character string is odd, the last byte of the string (excluding the null termination byte, if any) is not converted.

Note: Because binding character C data to a binary SQL data type is inefficient and slow, refrain from doing this.

C to SQL: Numeric

The numeric ODBC C data types are:

- SQL_C_STINYINT
- SQL_C_SLONG
- SQL_C_UTINYINT
- SQL_C_ULONG
- SQL_C_TINYINT
- SQL_C_LONG
- SQL_C_SSHORT
- SQL_C_FLOAT
- SQL_C_USHORT
- SQL_C_DOUBLE
- SQL_C_SHORT
- SQL_C_NUMERIC
- SQL_C_SBIGINT
- SQL_C_UBIGINT

For more information about the SQL_C_TINYINT, SQL_C_SHORT, and SQL_C_LONG data types, see E.5, "C data types," on page 228. The following

table shows the ODBC SQL data types to which numeric C data may be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (C to SQL)” on page 256.

Table 176. Numeric C Data to ODBC SQL Data Types

ParameterType	Test	SQLSTATE
SQL_CHAR	Number of digits <= Column byte length	N/A
SQL_VARCHAR		22001
SQL_LONGVARCHAR		
SQL_WCHAR	Number of characters <= Column character length	N/A
SQL_WVARCHAR		22001
SQL_WLONGVARCHAR		
SQL_DECIMAL [a]	Data converted without truncation or with truncated of fractional digits	N/A
SQL_NUMERIC [a]		22003
SQL_TINYINT [a]	Data converted with truncation of whole digits	
SQL_SMALLINT [a]		
SQL_INTEGER [a]		
SQL_BIGINT [a]		
SQL_REAL	Data is within the range of the data type to which the number is being converted	N/A
SQL_FLOAT		22003
SQL_DOUBLE	Data is outside the range of the data type to which the number is being converted	
Note:		
[a] For the "n/a" case, a driver may optionally return SQL_SUCCESS_WITH_INFO and 01S07 when there is a fractional truncation.		

The driver ignores the length or indicator value when converting data from the numeric C data types and assumes that the size of the data buffer is the size of the numeric C data type. The length or indicator value is passed in the StrLen_or_IndPtr argument in SQLPutData and in the buffer specified with the StrLen_or_IndPtr argument in SQLBindParameter. The data buffer is specified with the DataPtr argument in SQLPutData and the ParameterValuePtr argument in SQLBindParameter.

C to SQL: Bit

The bit ODBC C data type is:

SQL_C_BIT

The following table shows the ODBC SQL data types to which bit C data may be converted. For an explanation of the columns and terms in the table, see For an explanation of the columns and terms in the table, see “Conversion Table Description (C to SQL)” on page 256.

Table 177. Bit C Data to ODBC SQL Data Types

SQL Type Identifier	Test	SQLSTATE
SQL_CHAR	None	N/A
SQL_VARCHAR		
SQL_LONGVARCHAR		
SQL_WCHAR		
SQL_WVARCHAR		
SQL_WLONGVARCHAR		
SQL_DECIMAL	None	N/A
SQL_NUMERIC		
SQL_TINYINT		
SQL_SMALLINT		
SQL_INTEGER		
SQL_BIGINT		
SQL_REAL		
SQL_FLOAT		
SQL_DOUBLE		

The driver ignores the length or indicator value when converting data from the bit C data types and assumes that the size of the data buffer is the size of the bit C data type. The length or indicator value is passed in the StrLen_or_Ind argument in SQLPutData and in the buffer specified with the StrLen_or_IndPtr argument in SQLBindParameter. The data buffer is specified with the DataPtr argument in SQLPutData and the ParameterValuePtr argument in SQLBindParameter.

C to SQL: Binary

The binary ODBC C data type is:

SQL_C_BINARY

The following table shows the ODBC SQL data types to which binary C data may be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (C to SQL)” on page 256.

Table 178. Binary C Data to ODBC SQL Data Types

SQL Type Identifier	Test	SQLSTATE
SQL_CHAR SQL_VARCHAR SQL_LONGVARCHAR	Byte length of data <= Column byte length Byte length of data > Column length	N/A 22001
SQL_WCHAR SQL_WVARCHAR SQL_WLONGVARCHAR	Character length of data <= Column character length Character length of data > Column character length	N/A 22001
SQL_DECIMAL SQL_NUMERIC SQL_TINYINT SQL_SMALLINT SQL_INTEGER SQL_BIGINT SQL_REAL SQL_FLOAT SQL_DOUBLE SQL_TYPE_DATE SQL_TYPE_TIME SQL_TYPE_TIMESTAMP	Byte length of data = SQL data length Length of data <> SQL data length	N/A 22003
SQL_BINARY SQL_VARBINARY SQL_LONGVARBINARY	Length of data <= Column length Length of data > Column length	N/A 22001

C to SQL: Date

The date ODBC C data type is:

SQL_C_DATE

The following table shows the ODBC SQL data types to which date C data may be converted. For an explanation of the columns and terms in the table, see “Conversion Table Description (C to SQL)” on page 256.

Table 179. Date C Data to ODBC SQL Data Types

SQL Type Identifier	Test	SQLSTATE
SQL_CHAR	Column byte length >= 10	N/A
SQL_VARCHAR	Column byte length < 10	22001
SQL_LONGVARCHAR	Data value is not a valid date	22008
SQL_CHAR	Column character length >= 10	N/A
SQL_VARCHAR	Column character length < 10	22001
SQL_LONGVARCHAR	Data value is not a valid date	22008
SQL_TYPE_DATE	Data value is a valid date	N/A
	Data value is not a valid date	22007
SQL_TYPE_TIMESTAMP	Data value is a valid date [a]	N/A
	Data value is not a valid date	22007
Note: [a] The time portion of the timestamp is set to zero.		

For information about what values are valid in an SQL_C_TYPE_DATE structure, see E.5, "C data types," on page 228.

When date C data is converted to character SQL data, the resulting character data is in the "yyyy-mm-dd" format.

The driver ignores the length or indicator value when converting data from the date C data types and assumes that the size of the data buffer is the size of the date C data type. The length or indicator value is passed in the StrLen_or_Ind argument in SQLPutData and in the buffer specified with the StrLen_or_IndPtr argument in SQLBindParameter. The data buffer is specified with the DataPtr argument in SQLPutData and the ParameterValuePtr argument in SQLBindParameter.

C to SQL: Time

The time ODBC C data type is:

SQL_C_TIME

The following table shows the ODBC SQL data types to which time C data may be converted. For an explanation of the columns and terms in the table, see "Conversion Table Description (C to SQL)" on page 256.

Table 180. Time C Data to ODBC SQL Data Types

SQL Type Identifier	Test	SQLSTATE
SQL_CHAR	Column byte length >= 8	N/A
SQL_VARCHAR SQL_LONGVARCHAR	Column byte length < 8 Data value is not a valid time	22001 22008
SQL_WCHAR SQL_WVARCHAR SQL_WLONGVARCHAR	Column character length >= 8 Column character length < 8 Data value is not a valid time	N/A 22001 22008
SQL_TYPE_TIME	Data value is a valid time Data value is not a valid time	N/A 22007
SQL_TYPE_TIMESTAMP	Data value is a valid time [a] Data value is not a valid time	N/A 22007
<p>Note: [a] The date portion of the timestamp is set to the current date and the fractional seconds portion of the timestamp is set to zero.</p>		

For information about what values are valid in an SQL_C_TYPE_TIME structure, see E.5, "C data types," on page 228.

When time C data is converted to character SQL data, the resulting character data is in the "hh:mm:ss" format.

The driver ignores the length or indicator value when converting data from the time C data types and assumes that the size of the data buffer is the size of the time C data type. The length or indicator value is passed in the StrLen_or_Ind argument in SQLPutData and in the buffer specified with the StrLen_or_IndPtr argument in SQLBindParameter. The data buffer is specified with the DataPtr argument in SQLPutData and the ParameterValuePtr argument in SQLBindParameter.

C to SQL: Timestamp

The timestamp ODBC C data type is:

SQL_C_TIMESTAMP

The following table shows the ODBC SQL data types to which timestamp C data may be converted. For an explanation of the columns and terms in the table, see "Conversion Table Description (C to SQL)" on page 256.

Table 181. Timestamp C Data to ODBC SQL Data Types

SQL Type Identifier	Test	SQLSTATE
SQL_CHAR SQL_VARCHAR SQL_LONGVARCHAR	Column byte length >= Character byte length 19 <= Column byte length < Character byte length Column byte length < 19 Data value is not a valid date	N/A 22001 22001 22008
SQL_WCHAR SQL_WVARCHAR SQL_WLONGVARCHAR	Column character length >= Character length of data 19 <= Column character length < Character length of data Column character length < 19 Data value is not a valid timestamp	N/A 22001 22001 22008
SQL_TYPE_DATE	Time fields are zero Time fields are non-zero Data value does not contain a valid date	N/A 22008 22007
SQL_TYPE_TIME	Fractional seconds fields are zero [a] Fractional seconds fields are non-zero [a] Data value does not contain a valid time	N/A 22008 22007
SQL_TYPE_TIMESTAMP	Fractional seconds fields are not truncated Fractional seconds fields are truncated Data value is not a valid timestamp	N/A 22008 22007
Note: [a] The date fields of the timestamp structure are ignored.		

For information about what values are valid in an SQL_C_TIMESTAMP structure, see E.5, "C data types," on page 228.

When timestamp C data is converted to character SQL data, the resulting character data is in the "yyyy-mm-dd hh:mm:ss [.f. ..]" format.

The driver ignores the length or indicator value when converting data from the timestamp C data types and assumes that the size of the data buffer is the size of the timestamp C data type. The length or indicator value is passed in the StrLen_or_Ind argument in SQLPutData and in the buffer specified with the StrLen_or_IndPtr argument in SQLBindParameter. The data buffer is specified with the DataPtr argument in SQLPutData and the ParameterValuePtr argument in SQLBindParameter.

E.13.2 C to SQL data conversion examples

The following examples illustrate how the driver converts C data to SQL data.

Table 182. C Data to SQL Data

C Data Type	C Data Value	SQL Data Type	Column Length	SQL Data Value	SQLSTATE
SQL_C_CHAR	abcdef\0 [a]	SQL_CHAR	6	abcdef	N/A
SQL_C_CHAR	abcdef\0 [a]	SQL_CHAR	5	abcde	22001
SQL_C_CHAR	1234.56\0 [a]	SQL_DECIMAL	8 [b]	1234.56	N/A
SQL_C_CHAR	1234.56\0 [a]	SQL_DECIMAL	7 [b]	1234.5	22001
SQL_C_CHAR	1234.56\0 [a]	SQL_DECIMAL	4	----	22003
SQL_C_FLOAT	1234.56	SQL_FLOAT	not applicable	1234.56	N/A
SQL_C_FLOAT	1234.56	SQL_INTEGER	not applicable	1234	22001
SQL_C_FLOAT	1234.56	SQL_TINYINT	not applicable	----	22003
SQL_C_TYPE_DATE	1992,12,31 [c]	SQL_CHAR	10	1992-12-31	N/A
SQL_C_TYPE_DATE	1992,12,31 [c]	SQL_CHAR	9	----	22003
SQL_C_TYPE_DATE	1992,12,31 [c]	SQL_TIMESTAMP	not applicable	1992-12-31 00:00:00.0	N/A
SQL_C_TYPE_TIMESTAMP	1992,12,31, 23,45,55, 120000000 [d]	SQL_CHAR	22	1992-12-31 23:45:55.12	N/A
SQL_C_TYPE_TIMESTAMP	1992,12,31, 23,45,55, 120000000 [d]	SQL_CHAR	21	1992-12-31 23:45:55.1	22001
SQL_C_TYPE_TIMESTAMP	1992,12,31, 23,45,55, 120000000 [d]	SQL_CHAR	18	----	22003

Table 182. C Data to SQL Data (continued)

C Data Type	C Data Value	SQL Data Type	Column Length	SQL Data Value	SQLSTATE
<p>Note:</p> <p>[a] "\0" represents a null-termination byte. The null-termination byte is required only if the length of the data is SQL_NTS.</p> <p>[b] In addition to bytes for numbers, one byte is required for a sign and another byte is required for the decimal point.</p> <p>[c] The numbers in this list are the numbers stored in the fields of the SQL_DATE_STRUCT structure.</p> <p>[d] The numbers in this list are the numbers stored in the fields of the SQL_TIMESTAMP_STRUCT structure.</p>					

Appendix F. Scalar functions

This section provides additional information about ODBC scalar functions.

ODBC specifies five types of scalar functions:

- String functions
- Numeric functions
- Time and date functions
- System functions
- Data type conversion functions

A scalar function is a function that returns one value for each row in the query. Functions like SQRT() and ABS() are scalar functions. Functions like SUM() and AVG() are not scalar functions because they return a single value even if they process more than one row.

This section includes tables for each scalar function category. Within each table, functions have been added in ODBC 3.0 to align with SQL-92. Each table also provides the version number when the function was introduced.

F.1 ODBC and SQL-92 scalar functions

This topic provides information about ODBC and SQL-92 scalar functions.

Because functions are often data-source-specific, ODBC does not require a data type for return values from scalar functions. To force data type conversion, applications should use the CONVERT scalar function.

Note:

ODBC and SQL-92 classify functions in different ways. ODBC classifies scalar functions by argument type, whereas SQL-92 classifies them by return value. For example, in ODBC, the EXTRACT function is classified as a timedate function because the extract-field argument is a timedate keyword and the extract_source argument is a timedate or interval expression. In SQL-92, however, the EXTRACT function is classified as a numeric scalar function because the return value is numeric.

Applications need to call SQLGetInfo to determine which scalar functions a driver supports. ODBC and SQL-92 information types are available for scalar function classifications. Because ODBC and SQL-92 use different classifications, the information types for the same function may differ between ODBC and SQL-92. For example, to determine support for the EXTRACT function requires SQL_TIMEDATE_FUNCTIONS information type in ODBC and SQL_SQL92_NUMERIC_VALUE_FUNCTIONS information type in SQL-92.

F.2 String functions

This topic lists string manipulation functions.

Applications can call SQLGetInfo with the SQL_STRING_FUNCTIONS information type to determine which string functions are supported by a driver.

String Function Arguments

Table 183. String Function Arguments

Arguments denoted as...	Definition
<i>string_exp</i>	These arguments can be the name of a column, a string literal, or the result of another scalar function, where the underlying data type can be represented as SQL_CHAR, SQL_VARCHAR, or SQL_LONGVARCHAR.
<i>start, length or count</i>	These arguments can be a numeric literal or the result of another scalar function, where the underlying data type can be represented as SQL_TINYINT, SQL_SMALLINT, or SQL_INTEGER
<i>character_exp</i>	These arguments are a variable-length character string

The following string functions are 1-based, that is, the first character in the string is character 1, not character 0.

Note: BIT_LENGTH, CHAR_LENGTH, CHARACTER_LENGTH, OCTET_LENGTH, and POSITION string scalar functions were added in ODBC 3.0 to align with SQL-92.

List of String Functions

Table 184. List of String Functions

Function	Description
ASCII(<i>string_exp</i>) (ODBC 1.0)	Returns the ASCII code value of the leftmost character of <i>string_exp</i> as an integer.
BIT_LENGTH(<i>string_exp</i>) (ODBC 3.0)	Returns the length in bits of string expression.
CHAR(<i>code</i>) (ODBC 1.0)	Returns the character that has the ASCII code value specified by <i>code</i> . The value of <i>code</i> should be between 0 and 255; otherwise, the return value is data source-dependent.
CHAR_LENGTH(<i>string_exp</i>) (ODBC 3.0)	Returns the length in characters of the string expression, if the string expression is of a character data type; otherwise, returns the length in bytes of the string expression (the smallest integer not less than the number of bits divided by 8). (This function is the same as CHARACTER_LENGTH function.)

Table 184. List of String Functions (continued)

Function	Description
CHARACTER_LENGTH(string_exp) (ODBC 3.0)	Returns the length in characters of the string expression, if the string expression is of a character data type; otherwise, returns the length in bytes of the string expression (the smallest integer not less than the number of bits divided by 8). (This function is the same as the CHAR_LENGTH function.)
CONCAT(string_exp1, string_exp2) (ODBC 1.0)	Returns a character string that is the result of concatenating <i>string_exp2</i> to <i>string_exp1</i> . The resulting string is DBMS-dependent.
DIFFERENCE(string_exp1, string_exp2) (ODBC 2.0)	The function returns the difference between the soundex (see soundex function below) values of two character expressions, as an integer. The integer returned is the number of characters in the soundex values that are the same. The return value ranges from 0 through 4: 0 indicates little or no similarity, and 4 indicates strong similarity or identical values.
INSERT(string_exp1, start, length, string_exp2) (ODBC 1.0)	Returns a character string where <i>length</i> characters have been deleted from <i>string_exp1</i> beginning at <i>start</i> and where <i>string_exp2</i> has been inserted into <i>string_exp</i> , beginning at <i>start</i> .
LCASE(string_exp) (ODBC 1.0)	Returns a string equal to that <i>string_exp</i> , with all uppercase characters converted to lowercase.
LEFT(string_exp, count) (ODBC 1.0)	Returns the leftmost <i>count</i> of characters of <i>string_exp</i> .
LENGTH(string_exp) (ODBC 1.0)	Returns the number of characters in <i>string_exp</i> , excluding trailing blanks.

Table 184. List of String Functions (continued)

Function	Description
LOCATE(string_exp1, string_exp2[, start])	<p>Returns the starting position of the first occurrence of <i>string_exp1</i> within <i>string_exp2</i>. The search for the first occurrence of <i>string_exp1</i> begins with the first character position in <i>string_exp2</i> unless the optional argument, <i>start</i>, is specified. If <i>start</i> is specified, the search begins with the character position indicated by the value of <i>start</i>. The first character position in <i>string_exp2</i> is indicated by the value 1. If <i>string_exp1</i> is not found within <i>string_exp2</i>, the value 0 is returned.</p> <p>If an application can call the LOCATE scalar function with the <i>string_exp1</i>, <i>string_exp2</i>, and <i>start</i> arguments, the driver returns SQL_FN_STR_LOCATE when SQLGetInfo is called with an option of SQL_STRING_FUNCTIONS. If the application can call the LOCATE scalar function with only the <i>string_exp1</i> and <i>string_exp2</i> arguments, the driver returns SQL_FN_STR_LOCATE_2 when SQLGetInfo is called with an option of SQL_STRING_FUNCTIONS. Drivers that support calling the LOCATE function with either two or three arguments return both SQL_FN_STR_LOCATE and SQL_FN_STR_LOCATE_2.</p>
LTRIM(string_exp) (ODBC 1.0)	Returns the characters of <i>string_exp</i> , with leading blanks removed.
OCTET_LENGTH(string_exp) (ODBC 3.0)	Returns the length in bytes of the string expression. The result is the smallest integer not less than the number of bits divided by 8.
POSITION(character_exp IN character_exp) (ODBC 3.0)	Returns the position of the first character expression in the second character expression. The result is an exact numeric with an implementation-defined precision and a scale of 0.
REPEAT(string_exp, count) (ODBC 1.0)	Returns a character string composed of <i>string_exp</i> repeated <i>count</i> times.
REPLACE(string_exp1, string_exp2, string_exp3) (ODBC 1.0)	Search <i>string_exp1</i> for occurrences of <i>string_exp2</i> , and replace with <i>string_exp3</i> .
RIGHT(string_exp, count) (ODBC 1.0)	Returns the rightmost <i>count</i> of characters of <i>string_exp</i> .

Table 184. List of String Functions (continued)

Function	Description
RTRIM(string_exp) (ODBC 1.0)	Returns the characters of <i>string_exp</i> with trailing blanks removed.
SOUNDEX(string_exp1) (ODBC 2.0)	Returns a character string containing the phonetic representation of the argument. This function lets you compare words that are spelled differently, but sound alike in English. If you supply a word to Soundex, it returns a 4-character phonetic code used by the U.S.Census Bureau since 1930s.
SPACE(count) (ODBC 2.0)	Returns a character string consisting of <i>count</i> spaces.
SUBSTRING(string_exp, start, length) (ODBC 1.0)	Returns a character string that is derived from <i>string_exp</i> , beginning at the character position specified by <i>start</i> for <i>length</i> characters.
TRIM(string_exp)	Returns the characters of <i>string_exp</i> with leading blanks and trailing blanks removed.
UCASE(string_exp) (ODBC 1.0)	Returns a string equal to that in <i>string_exp</i> , with all lowercase characters converted to uppercase.

F.3 Numeric functions

This topic describes numeric functions that are included in the ODBC scalar function set.

Applications can call SQLGetInfo with the SQL_NUMERIC_FUNCTIONS information type to determine which numeric functions are supported by a driver.

Except for ABS, ROUND, TRUNCATE, SIGN, FLOOR, and CEILING (which return values of the same data type as the input parameters), all numeric functions return values of data type SQL_FLOAT.

Numeric Function Arguments

Table 185. Numeric Function Arguments

Arguments denoted as...	Definition
<i>numeric_exp</i>	These arguments can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type could be represented as SQL_NUMERIC, SQL_DECIMAL, SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, SQL_BIGINT, SQL_FLOAT, SQL_REAL, or SQL_DOUBLE

Table 185. Numeric Function Arguments (continued)

Arguments denoted as...	Definition
<i>float_exp</i>	These arguments can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type can be represented as SQL_FLOAT.
<i>integer_exp</i>	These arguments can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type can be represented as SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, or SQL_BIGINT

List of Numeric Functions

Table 186. List of Numeric Functions

Function	Description
ABS(<i>numeric_exp</i>) (ODBC 1.0)	Returns the absolute value of <i>numeric_exp</i> .
ACOS(<i>float_exp</i>) (ODBC 1.0)	Returns the arccosine of <i>float_exp</i> as an angle, expressed in radians.
ASIN(<i>float_exp</i>) (ODBC 1.0)	Returns the arcsine of <i>float_exp</i> as an angle, expressed in radians.
ATAN(<i>float_exp</i>) (ODBC 1.0)	Returns the arctangent of <i>float_exp</i> as an angle, expressed in radians.
ATAN2(<i>float_exp1</i> , <i>float_exp2</i>) (ODBC 2.0)	Returns the arctangent of the x and y coordinates, specified by <i>float_exp1</i> and <i>float_exp2</i> , respectively, as an angle, expressed in radians.
CEILING(<i>numeric_exp</i>) (ODBC 1.0)	Returns the smallest integer greater than or equal to <i>numeric_exp</i> . The return value is of the same data type as the input parameter.
COS(<i>float_exp</i>) (ODBC 1.0)	Returns the cosine of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
COT(<i>float_exp</i>) (ODBC 1.0)	Returns the cotangent of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
DEGREES(<i>numeric_exp</i>) (ODBC 2.0)	Returns the number of degrees converted from <i>numeric_exp</i> radians.

Table 186. List of Numeric Functions (continued)

Function	Description
EXP(float_exp) (ODBC 1.0)	Returns the exponential value of <i>float_exp</i> .
FLOOR(numeric_exp) (ODBC 1.0)	Returns largest integer less than or equal to <i>numeric_exp</i> . The return value is of the same data type as the input parameter.
LOG(float_exp) (ODBC 1.0)	Returns the natural logarithm of <i>float_exp</i> .
LOG10(float_exp) (ODBC 2.0)	Returns the base 10 logarithm of <i>float_exp</i> .
MOD(integer_exp1, integer_exp2) (ODBC 1.0)	Returns the remainder (modulus) of <i>integer_exp1</i> divided by <i>integer_exp2</i> .
PI() (ODBC 1.0)	Returns the constant value of pi as a floating point value.
POWER(numeric_exp, integer_exp)	Returns the value of <i>numeric_exp</i> to the power of <i>integer_exp</i> .
RADIANS(numeric_exp) (ODBC 2.0)	Returns the number of radians converted from <i>numeric_exp</i> degrees.
ROUND(numeric_exp, integer_exp) (ODBC 2.0)	Returns <i>numeric_exp</i> rounded to <i>integer_exp</i> places right of the decimal point. If <i>integer_exp</i> is negative, <i>numeric_exp</i> is rounded to $ integer_exp $ places to the left of the decimal point.
SIGN(numeric_exp) (ODBC 1.0)	Returns an indicator or the sign of <i>numeric_exp</i> . If <i>numeric_exp</i> is less than zero, -1 is returned. If <i>numeric_exp</i> equals zero, 0 is returned. If <i>numeric_exp</i> is greater than zero, 1 is returned.
SIN(float_exp) (ODBC 1.0)	Returns the sine of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.
SQRT(float_exp) (ODBC 1.0)	Returns the square root of <i>float_exp</i> .
TAN(float_exp) (ODBC 1.0)	Returns the tangent of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians.

Table 186. List of Numeric Functions (continued)

Function	Description
TRUNCATE(numeric_exp, integer_exp) (ODBC 2.0)	Returns <i>numeric_exp</i> truncated to <i>integer_exp</i> places right of the decimal point. If <i>integer_exp</i> is negative, <i>numeric_exp</i> is truncated to $ integer_exp $ places to the left of the decimal point.

F.4 Time and date functions

This section lists time and date functions that are included in the ODBC scalar function set.

Applications can call SQLGetInfo with the SQL_TIMEDATE_FUNCTIONS information type to determine which time and date functions are supported by a driver.

Time and Data Arguments

Table 187. Time and Data Arguments

Arguments denoted as...	Definition
<i>timestamp_exp</i>	These arguments can be the name of a column, the result of another scalar function, or an <i>ODBC_time_escape</i> , <i>ODBC_date_escape</i> , or <i>ODBC_timestamp_escape</i> , where the underlying data type could be represented as SQL_CHAR, SQL_VARCHAR, SQL_TYPE_TIME, SQL_TYPE_DATE, or SQL_TYPE_TIMESTAMP.
<i>date_exp</i>	These arguments can be the name of a column, the result of another scalar function, or an <i>ODBC_date_escape</i> or <i>ODBC_timestamp_escape</i> , where the underlying data type could be represented as SQL_CHAR, SQL_VARCHAR, SQL_TYPE_DATE, or SQL_TYPE_TIMESTAMP.
<i>time_exp</i>	These arguments can be the name of a column, the result of another scalar function, or an <i>ODBC_time_escape</i> or <i>ODBC_timestamp_escape</i> , where the underlying data type could be represented as SQL_CHAR, SQL_VARCHAR, SQL_TYPE_TIME, or SQL_TYPE_TIMESTAMP.

List of Time and Date Functions

Table 188. List of Time and Date Functions

Function	Description
CURRENTTIME [[<i>time_precision</i>]] (ODBC 3.0)	Returns the current local time as a time value. The <i>time_precision</i> argument (0-6) determines the milliseconds precision of the returned value. Value 0 means no timestamp or time fractions are shown. If the value is not specified, the value 0 is used.

Table 188. List of Time and Date Functions (continued)

Function	Description
CURRENT_TIMESTAMP [(<i>timestamp_precision</i>)] (ODBC 3.0)	Returns the current local date and local time as a timestamp value. The <i>timestamp_precision</i> argument (0-6) determines the milliseconds precision of the returned timestamp. Value 0 means no timestamp or time fractions are shown. If the value is not specified, the value 0 is used.
CURDATE() (ODBC 1.0)	Returns the current date.
CURTIME[(<i>time_precision</i>)] (ODBC 1.0)	Returns the current local time. The <i>time_precision</i> argument (0-6) determines the milliseconds precision of the returned value. Value 0 means no timestamp or time fractions are shown. If the value is not specified, the value 0 is used.
DAYNAME(<i>date_exp</i>) (ODBC 2.0)	Returns a character string containing the data source-specific name of the day (for example, Sunday, through Saturday or Sun. through Sat. for a data source that uses English, or Sonntag through Samstag for a data source that uses German) for the day portion of <i>date_exp</i> .
DAYOFMONTH(<i>date_exp</i>) (ODBC 1.0)	Returns the day of the month in <i>date_exp</i> as an integer value in the range of 1-31.
DAYOFWEEK(<i>date_exp</i>) (ODBC 1.0)	Returns the day of the week based on the week field in <i>date_exp</i> as an integer value in the range of 1-7, where 1 represents Sunday.
DAYOFYEAR(<i>date_exp</i>) (ODBC 1.0)	Returns the day of the year based on the year field in <i>date_exp</i> as an integer value in the range of 1-366.
EXTRACT(<i>extract_field</i> FROM <i>extract_source</i>) (ODBC 3.0)	Returns the <i>extract_field</i> portion of the <i>extract_source</i> . The <i>extract_source</i> argument is a datetime or interval expression. The <i>extract_field</i> argument can be one of the following keywords" YEAR MONTH DAY HOUR MINUTE SECOND The precision of the returned value is implementation-defined. The scale is 0 unless SECOND is specified, in which case the scale is not less than the fractional seconds precision of the <i>extract_source</i> field.

Table 188. List of Time and Date Functions (continued)

Function	Description
HOUR(time_exp) (ODBC 1.0)	Returns the hour based on the hour field in <i>time_exp</i> as an integer value in the range of 0-23.
MINUTE(time_exp) (ODBC 1.0)	Returns the minute based on the minute field in <i>time_exp</i> as an integer value in the range of 0-59.
MONTH(date_exp) (ODBC 1.0)	Returns the month based on the month field in <i>date_exp</i> as an integer value in the range of 1-12.
MONTHNAME(date_exp) (ODBC 2.0)	Returns a character string containing the data source-specific name of the month (for example, January through December or Jan. through Dec. for a data source that uses English, or Januar through Dezember for a data source that uses German) for the month portion of <i>date_exp</i> .
NOW [(timestamp_precision)] (ODBC 1.0)	Returns current date and time as a timestamp value. The <i>timestamp_precision</i> argument (0-6) determines the milliseconds precision of the returned timestamp. Value 0 means no timestamp or time fractions are shown. If the value is not specified, the value 0 is used.
QUARTER(date_exp) (ODBC 1.0)	Returns the quarter in <i>date_exp</i> as an integer value in the range of 1- 4, where 1 represents January 1 through March 31.
SECOND(time_exp) (ODBC 1.0)	Returns the second in <i>time_exp</i> as an integer value in the range of 0-59.

Table 188. List of Time and Date Functions (continued)

Function	Description
<p>TIMESTAMPADD(interval, integer_exp, timestamp_exp)</p> <p>(ODBC 2.0)</p>	<p>Returns the timestamp calculated by adding <i>integer_exp</i> intervals of type interval to <i>timestamp_exp</i>. Valid values of interval are the following keywords:</p> <p>SQL_TSI_FRAC_SECOND SQL_TSI_SECOND SQL_TSI_MINUTE SQL_TSI_HOUR SQL_TSI_DAY SQL_TSI_WEEK SQL_TSI_MONTH SQL_TSI_QUARTER SQL_TSI_YEAR</p> <p>where fractional seconds are expressed in billionths of a second (nanoseconds). For example, the following SQL statement returns the name of each employee and his or her one-year anniversary date:</p> <pre>SELECT NAME, {fn TIMESTAMPADD(SQL_TSI_YEAR, 1, HIRE_DATE)} FROM EMPLOYEES</pre> <p>If <i>timestamp_exp</i> is a time value and interval specifies day, weeks, months, quarters, or years, the date portion of <i>timestamp_exp</i> is set to the current date before calculating the resulting timestamp.</p> <p>If <i>timestamp_exp</i> is a date value and interval specifies fractional seconds, seconds, minutes, or hours, the time portion of <i>timestamp_exp</i> is set to 0 before calculating the resulting timestamp.</p> <p>An application determines which intervals a data source supports by calling SQLGetInfo with the SQL_TIMEDATE_ADD_INTERVALS option.</p>

Table 188. List of Time and Date Functions (continued)

Function	Description
<p>TIMESTAMPDIFF(interval, timestamp_exp1, timestamp_exp2)</p> <p>(ODBC 2.0)</p>	<p>Returns the number of unit intervals (as integers) of type <i>interval</i> between <i>timestamp_exp1</i> and <i>timestamp_exp2</i>.</p> <p>If an application relies on the old TIMESTAMPDIFF semantics, the old behavior can be emulated by the following configuration setting in the SQL section of the <i>solid.ini</i> file.</p> <pre>[SQL] Emulate01dTIMESTAMPDIFF=YES</pre> <p>Note that the old semantics returns the integer number of intervals of type <i>interval</i> by which <i>timestamp_exp2</i> is greater than <i>timestamp_exp1</i>.</p> <p>Valid values of <i>interval</i> are the following keywords:</p> <pre>SQL_TSI_FRAC_SECOND SQL_TSI_SECOND SQL_TSI_MINUTE SQL_TSI_HOUR SQL_TSI_DAY SQL_TSI_WEEK SQL_TSI_MONTH SQL_TSI_QUARTER SQL_TSI_YEAR</pre> <p>where fractional seconds are expressed in billionths of a second (nanoseconds). For example, the following SQL statement returns the name of each employee and the number of years they have been employed:</p> <pre>SELECT NAME, {fn TIMESTAMPDIFF(SQL_TSI_YEAR, {fn CURDATE()} , HIRE_DATE)} FROM EMPLOYEES</pre> <p>If either timestamp expression is a time value and interval specifies days, weeks, months, quarters, or years, the date portion of that timestamp is set to the current date before calculating the difference between the timestamps.</p> <p>If either timestamp expression is a date value and interval specifies fractional seconds, seconds, minutes, or hours, the time portion of that timestamp is set to 0 before calculating the difference between the timestamps.</p> <p>An application determines which intervals a data source supports by calling SQLGetInfo with the SQL_TIMEDATE_DIFF_INTERVALS option.</p>
<p>WEEK(date_exp)</p> <p>(ODBC 1.0)</p>	<p>Returns the week of the year based on the week field in <i>date_exp</i> as an integer value in the range of 1-53.</p>
<p>YEAR(date_exp)</p> <p>(ODBC 1.0)</p>	<p>Returns the year based on the year field in <i>date_exp</i> as an integer value. The range is data source-dependent.</p>

F.5 System functions

This section lists system functions that are included in the ODBC scalar function set.

Applications can call SQLGetInfo with the SQL_SYSTEM_FUNCTIONS information type to determine which string functions are supported by a driver.

System Functions Arguments

Table 189. System Function Arguments

Arguments denoted as...	Definition
<i>exp</i>	These arguments can be the name of a column, the result of another scalar function, or a literal, where the underlying data type could be represented as SQL_NUMERIC, SQL_DECIMAL, SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, SQL_BIGINT, SQL_FLOAT, SQL_REAL, SQL_DOUBLE, SQL_TYPE_DATE, SQL_TYPE_TIME, or SQL_TYPE_TIMESTAMP.
<i>value</i>	These arguments can be a literal constant, where the underlying data type can be represented as SQL_NUMERIC, SQL_DECIMAL, SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, SQL_BIGINT, SQL_FLOAT, SQL_REAL, SQL_DOUBLE, SQL_TYPE_DATE, SQL_TYPE_TIME, or SQL_TYPE_TIMESTAMP.
<i>integer_exp</i>	These arguments can be the name of a column, the result of another scalar function, or a numeric literal, where the underlying data type can be represented as SQL_TINYINT, SQL_SMALLINT, SQL_INTEGER, or SQL_BIGINT

Values returned are represented as ODBC data types

List of System Functions

Table 190. List of System Functions

Function	Description
DATABASE() (ODBC 1.0)	Returns the name of the database corresponding to the connection handle. (The name of the database is also available by calling SQLGetConnectOption with the SQL_CURRENT_QUALIFIER connection option.)
IFNULL(<i>exp</i> , <i>value</i>) (ODBC 1.0)	If <i>exp</i> is null, <i>value</i> is returned. If <i>exp</i> is not null, <i>exp</i> is returned. The possible data type(s) of <i>value</i> must be compatible with the data type of <i>exp</i>
USER() (ODBC 1.0)	Returns the user's name in the DBMS. (The user's authorization name is also available via SQLGetInfo by specifying the information type: SQL_USER_NAME.) This can be different from the login time.

F.6 Explicit data type conversion

Explicit data type conversion is specified in terms of SQL data type definitions.

The ODBC syntax for the explicit data type conversion function does not restrict conversions. The validity of specific conversions of one data type to another data type is dependent on each driver-specific implementation. The driver, as it translates the ODBC syntax into the native syntax, reject those conversions that, although legal in the ODBC syntax, are not supported by the data source. Applications can call the ODBC function `SQLGetInfo` to inquire about conversions supported by the data source.

The format of the `CONVERT` function is:

```
CONVERT(value_exp, data_type)
```

The function returns the value specified by `value_exp` converted to the specified `data_type`, where `data_type` is one of the following keywords:

- `SQL_BIGINT`
- `SQL_SMALLINT`
- `SQL_BINARY`
- `SQL_DATE`
- `SQL_CHAR`
- `SQL_TIME`
- `SQL_DECIMAL`
- `SQL_TIMESTAMP`
- `SQL_DOUBLE`
- `SQL_TINYINT`
- `SQL_FLOAT`
- `SQL_VARBINARY`
- `SQL_INTEGER`
- `SQL_VARCHAR`
- `SQL_LONGVARBINARY`
- `SQL_WCHAR`
- `SQL_LONGVARCHAR`
- `SQL_WLONGVARCHAR`
- `SQL_NUMERIC`
- `SQL_WVARCHAR`
- `SQL_REAL`

The ODBC syntax for the explicit data type conversion function does not support specification of conversion format. If specification of explicit formats is supported by the underlying data source, a driver must specify a default value or implement format specification.

The argument `value_exp` can be a column name, the result of another scalar function, or a numeric or string literal. The following example converts the output of the `CURDATE` scalar function to a character string:

```
{ fn CONVERT( { fn CURDATE() }, SQL_CHAR) }
```

ODBC does not require a data type for return values from scalar functions (because the functions are often data source-specific); applications should use the CONVERT scalar function whenever possible to force data type conversion.

The following two examples illustrate the use of the CONVERT function. These examples assume the existence of a table called EMPLOYEES, with an EMPNO column of type SQL_SMALLINT and an EMPNAME column of type SQL_CHAR.

If an application specifies the following:

```
SELECT EMPNO FROM EMPLOYEES WHERE {fn CONVERT(EMPNO,SQL_CHAR)}LIKE '1%'
```

solidDB ODBC driver translates the request to:

```
SELECT EMPNO FROM EMPLOYEES WHERE CONVERT_CHAR(EMPNO) LIKE '1%'
```

F.7 SQL-92 CAST function

The ODBC CONVERT function has an equivalent function in SQL-92: the CAST function.

The syntax for these equivalent functions is as follows::

```
{ fn CONVERT (value_exp, data_type)} /* ODBC */  
CAST (value_exp AS data_type) /* SQL 92 */
```

Support for the CAST function is at the FIPS Transitional level. For details on data type conversion in the CAST function, see the SQL-92 specification.

To determine application support for the CAST function, call SQLGetInfo with the SQL_SQL_CONFORMANCE information type. The CAST function is supported if the return value for the information type is:

- SQL_SC_FIPS127_2_TRANSITIONAL
- SQL_SC_SQL92_INTERMEDIATE
- SQL_SC_SQL92_FULL

If the return value is SQL_SC_ENTRY or 0, call SQLGetInfo with the SQL_SQL92_VALUE_EXPRESSIONS information type. If the SQL_SVE_CAST bit is set, the CAST function is supported.

Appendix G. Timeout controls

In solidDB, some actions can get timed out. A timeout can be activated by the main server, the client drivers, the Primary or Secondary server, or the Master or Replica server.

Timeouts have factory default values and they can usually be set with different .ini parameters. Some startup defaults can be dynamically changed with different controls, by using SQL, or by using the driver interfaces and connection string parameters.

G.1 Client timeouts

Timeouts related to the database client are introduced in this topic.

Login timeout

This timeout refers to the number of seconds the driver waits for the login (SQLConnect) to succeed. The default value is driver-dependent. If the value (or ValuePtr in ODBC) is 0, the timeout is disabled and a connection attempt will wait indefinitely. If the specified timeout exceeds the maximum login timeout in the data source, the driver substitutes that value and returns SQLSTATE 01S02 (Option value changed).

This timeout applies for the TCP protocol only.

Table 191. Login timeouts

INI parameter	Overridden with SQL	Driver	Connection string
Client-side: [Com] ConnectTimeout= <i>milliseconds</i> or: Com.Connect parameter with option <i>-cmilliseconds</i> For example: [Com] Connect=tcp -c1000 1315		ODBC: SQL_ATTR_LOGIN_TIMEOUT (in seconds) SQL_ATTR_LOGIN_TIMEOUT_MS (in milliseconds, non-standard) JDBC: Method (JDBC 2.0) DriverManager.setLoginTimeout(seconds); Connection property (non-standard) "solid_login_timeout_ms" (milliseconds)	-c <i>milliseconds</i>

Timeout error code and message:

ODBC:

HYT00, Timeout expired

Connection timeout

This timeout refers to the number of seconds (or milliseconds) the driver waits for any request on the connection to complete. This timeout is not associated with the query execution or login. Upon timeout, the driver disconnects from the solidDB server.

If the client can detect reliably that the server is not reachable, it does not wait for the timeout to expire. This can happen, for example, if the server is expected to run on the same host, and the client detects that the server is not running.

The driver returns SQLSTATE HYT00 (Timeout expired) if it is possible to time out in a situation not associated with query execution or login. If the value (or ValuePtr in ODBC) is 0 (the default value), there is no timeout.

This timeout applies to all ODBC functions (ODBC 3.5 specifications) except the following:

- SQLDrivers
- SQLDataSources
- SQLGetEnvAttr
- SQLSetEnvAttr

Table 192. Connection timeout

INI parameter	Overridden with SQL	Driver	Connection string
Client-side: [Com] ClientReadTimeout= <i>milliseconds</i> or: Com.Connect parameter with option <i>-rmilliseconds</i> For example: [Com] Connect=tcp -r1000 1315		ODBC: SQL_ATTR_CONNECTION_TIMEOUT (in seconds) SQL_ATTR_CONNECTION_TIMEOUT_MS (in milliseconds, non-standard) JDBC: Non-standard connection property: "solid_connection_timeout_ms" (milliseconds) or method: SolidConnection.setConnectionTimeout() (milliseconds)	<i>-r milliseconds</i>

Timeout error code and message:

ODBC:

HYT01, Connection timeout expired

See also:

SOLID Server Error 14518:
 Connection to the server is broken, connection lost.

SOLID Communication Error 21328 and SOLID Session Error 20024:
 Timeout while resolving host name.

SOLID Communication Error 21329 and SOLID Session Error 20025:
Timeout while connecting to a remote host.

Query timeout

This timeout refers to the number of seconds the driver waits for an SQL statement to execute. If the value (or ValuePtr in ODBC) is 0 (the default value), there is no timeout.

If the specified timeout exceeds the maximum timeout in the data source, or if the specified timeout is smaller than the minimum timeout, `SQLSetStmtAttr` substitutes that value and returns `SQLSTATE 01S02` (Option value changed).

This timeout applies to the ODBC functions (ODBC 3.5 specifications) as follows:

`SQLBrowseConnect`

`SQLBulkOperations`

`SQLColumnPrivileges`

`SQLColumns`

`SQLConnect`

`SQLDriverConnect`

`SQLExecDirect`

`SQLExecute`

`SQLExtendedFetch`

`SQLForeignKeys`

`SQLGetTypeInfo`

`SQLParamData`

`SQLPrepare`

`SQLPrimaryKeys`

`SQLProcedureColumns`

`SQLProcedures`

`SQLSetPos`

`SQLSpecialColumns`

`SQLStatistics`

`SQLTablePrivileges`

SQLTables

Note: The application need not call `SQLCloseCursor` to reuse the statement if a `SELECT` statement timed out. The query timeout set in this statement attribute is valid in both synchronous and asynchronous modes.

Table 193. Query Timeout

INI parameter	Overridden with SQL	Driver	Connection string
		ODBC: SQL_ATTR_QUERY_TIMEOUT (in seconds) SQL_ATTR_QUERY_TIMEOUT_MS (in milliseconds, non-standard)	

Timeout error code and message:

ODBC:

HYT00, Timeout expired

G.2 Server timeouts

Timeouts related to the database server are introduced in this topic.

SQL Statement Execution Timeout

The server can control the amount of time spent on the execution of one SQL statement. When the time expires, the server terminates the statements and returns a corresponding error code. This timeout applies to the following calls (ODBC 3.5 specifications):

- `SQLExecute()`
- `SQLExecDirect()`
- `SQLPrepare()`
- `SQLForeignKeys()`
- `SQLColumns()`
- `SQLProcedureColumns()`
- `SQLSpecialColumns()`
- `SQLStatistics()`
- `SQLPrimaryKeys()`
- `SQLProcedures()`
- `SQLTables()`
- `SQLTablePrivileges()`
- `SQLColumnPrivileges()`
- `SQLGetTypeInfo()`

The timeout also applies to the corresponding JDBC calls.

Table 194. SQL statement execution timeouts

INI parameter	Overridden with SQL	Driver	Connection string
	SET STATEMENT MAXTIME minutes	ODBC: SQL_ATTR_QUERY_TIMEOUT (in seconds) SQL_ATTR_QUERY_TIMEOUT_MS (in milliseconds, non-standard) JDBC: statement.setQueryTimeout()	

Timeout error code and message:

HYT00, Timeout expired

See also:

SOLID Server Error 14518:
Connection to the server is broken, connection lost.

SOLID Server Error 14529:
The operation timed out.

Lock wait timeout

The *lock wait timeout* specifies the time in seconds (or milliseconds) that the engine waits for a lock to be released. When the timeout interval is reached, solidDB terminates the timed-out transaction. The default value is 30 seconds.

Lock wait timeout is used in deadlock resolution. In that case, the oldest transaction participating in a deadlock is aborted.

Table 195. Lock wait timeout

INI parameter	Overridden with SQL	Driver	Connection string
[General] LockWaitTimeOut= <i>seconds</i>	SET LOCK TIMEOUT {seconds milliseconds MS}		

Timeout error code and message:

SOLID Database Error 10006:
Concurrency conflict, two transactions updated or deleted the same row.

Optimistic lock timeout

The *optimistic lock timeout* specifies the time in seconds (or milliseconds) for optimistic lock timeout. Optimistic lock is an additional lock that can be enacted in order to ensure that SELECT FOR UPDATE will always lead to successful updates in the optimistic concurrency mode. The default is zero; no optimistic lock is used, and a transaction may be aborted after each statement as a result of early

transaction validation. When the timeout is set to a non-zero value, SELECT FOR UPDATE will wait until the lock is obtained, or it is timed-out and aborted. When set, the timeout affects also all DELETE and UPDATE statements.

Table 196. Optimistic lock wait timeout

INI parameter	Overridden with SQL	Driver	Connection string
	SET OPTIMISTIC LOCK TIMEOUT {seconds milliseconds MS}		

Timeout error code and message:

SOLID Database Error 10006:
Concurrency conflict, two transactions updated or deleted the same row.

Table lock wait timeout

Occasionally, the transaction will acquire an exclusive lock to a table. This may be result of a lock escalation, an attempt to execute the ALTER TABLE statement, or as a side effect of some advanced replication commands. If there is a table-level conflict, this setting provides the transaction's wait period until the exclusive or shared lock is released. The unit is seconds, the default value is 30 seconds, and the parameter access mode is read/write.

To be more specific, table level locks are used when the PESSIMISTIC keyword is explicitly provided in the following commands:

```
IMPORT SUBSCRIPTIONMESSAGE message_name EXECUTE
(only with NO EXECUTE option)
MESSAGE message_name FORWARD
MESSAGE message_name GET REPLY
DROP SUBSCRIPTION.
```

Table 197. Table lock wait timeout

INI parameter	Overridden with SQL	Driver	Connection string
[General] TableLockWaitTimeout=seconds			

Timeout error code and message:

SOLID Database Error 10006:
Concurrency conflict, two transactions updated or deleted the same row.

Transaction idle timeout

The *transaction idle timeout* specifies the time in minutes after an idle transaction is aborted; a negative or zero value means infinite. The default value is 120 minutes.

Table 198. Transaction Idle Timeout

INI parameter	Overridden with SQL	Driver	Connection string
[Srv] AbortTimeOut			

Timeout error code and message:

SOLID Database Error 10026:
Transaction is timed out.

Connection idle timeout

The *connection idle timeout* specifies the continuous idle time in minutes (or seconds/milliseconds in a statement) after which a connection is dropped (by the server); negative or zero value indicates an infinite value. The default value is 480 minutes.

Table 199. connection idle timeout

INI parameter	Overridden with SQL	Driver	Connection string
[Srv] ConnectTimeOut= <i>minutes</i>	SET IDLE TIMEOUT {seconds milliseconds MS}	JDBC: Connection property (non-standard): "solid_idle_timeout_min"	

Timeout error code and message:

SOLID Communication Error 21308:
Connection is broken (protocol read/write
operation failed with code internal code).

See also the solmsg.out file.

If the SET IDLE TIMEOUT has been set and the transaction is idle for the given period, the error below is given:

SOLID Database Error 10026:
Transaction is timed out

G.3 HotStandby timeouts

Timeouts related to the HotStandby server are described in this topic.

Connect timeout

By specifying a connect timeout value with the **HotStandby.ConnectTimeout** parameter, you can set the maximum time in milliseconds that a HotStandby connect operation waits for a connection to a remote machine. The **HotStandby.ConnectTimeout** parameter is only used with the following subset of ADMIN COMMANDs:

```
hotstandby connect
hotstandby switch primary
hotstandby switch secondary
```

Table 200. Connect timeout

INI parameter	Overridden with SQL	Driver	Connection string
[HotStandby] ConnectTimeout= <i>milliseconds</i>			

Ping timeout

The **HotStandby.PingTimeout** parameter specifies in milliseconds how long a server waits before concluding that the other server is down or inaccessible.

Table 201. Ping timeout

INI parameter	Overridden with SQL	Driver	Connection string
[HotStandby] PingTimeout= <i>milliseconds</i>			

Transparent connection timeout

The client-side **TransparentFailover.ReconnectTimeout** parameter specifies how long (in milliseconds) the driver should wait until it tries to reconnect to the primary in case of TF switchover or failover.

The client-side **TransparentFailover.WaitTimeout** parameter specifies how long (in milliseconds) the driver should wait for the server to switch state.

Table 202. Transparent connection timeout

INI parameter	Overridden with SQL	Driver	Connection string
[TransparentFailover] ReconnectTimeout= <i>milliseconds</i>		ODBC: SQL_ATTR_TF_RECONNECT_TIMEOUT JDBC: solid_tf1_reconnect_timeout	
[TransparentFailover] WaitTimeout= <i>milliseconds</i>		ODBC: SQL_ATTR_TF_WAIT_TIMEOUT JDBC: solid_tf_wait_timeout	

Appendix H. Client-side configuration parameters

The client-side configuration parameters are stored in the client-side `solid.ini` configuration file and are read when the client starts.

Generally, the factory value settings offer the best performance and operability, but in some special cases modifying a parameter will improve performance. You can change the parameters by editing the `solid.ini` configuration file.

The parameter values set in the client side configuration file come to effect each time an application issues a call to the `SqlConnection` ODBC function. If the values are changed in the file during the program's run time, they affect the connections established thereafter.

H.1 Setting client-side parameters through the `solid.ini` configuration file

This topic provides details about the `solid.ini` configuration file.

When `solidDB` is started, it attempts to open the configuration file `solid.ini`. If the file does not exist, `solidDB` uses the factory values for the parameters. If the file exists, but a value for a particular parameter is not set, `solidDB` will use a factory value for that parameter. The factory values may depend on the operating system you are using.

By default, the client looks for the `solid.ini` file in the current working directory, which is normally the directory from which you started the client. When searching for the file, `solidDB` uses the following precedence (from high to low):

- location specified by the `SOLIDDIR` environment variable (if this environment variable is set)
- current working directory

Rules for formatting the client-side `solid.ini` file

When you format the client-side `solid.ini` file, the same rules apply as for the server-side `solid.ini` file. For more information, refer to section Rules for formatting the `solid.ini` file in *IBM solidDB Administrator Guide*.

Client-side `solid.ini` file

```
[Com]
;use this connect string if no data source given
Listen = tcp host1.acme.com 1315

[Client]
;at SqlConnection, timeout after this time (ms)
ConnectTimeout = 5000

;at any ODBC network request, timeout after this time (ms)
ClientReadTimeout = 10000

[Data Sources]
Primary_Server = tcp irix1 1315, The Primary Server
Secondary_Server = tcp irix2 1315, The Secondary Server
```

H.2 Client section

Table 203. Client parameters

[Client]	Description	Factory Value
ExecRowsPerMessage	<p>This parameter specifies how many result rows are sent (pre-fetched) to the client driver in response to the SQLExecute call with a SELECT statement. The result rows are subsequently returned to the application with the first SQLFetch calls issued by the application. The value 2 allows for prefetching of single-row results. If your SELECT statements usually return larger number of rows, setting this to an appropriate value can improve performance significantly.</p> <p>See also the RowsPerMessage parameter.</p>	decided by the server
NoAssertMessages	<p>If set to yes, the Windows runtime error dialog is not shown.</p> <p>This parameter is relevant to the Windows platform only.</p>	no
ODBCCharBinding	<p>Defines the binding method for character data.</p> <p>The options are:</p> <ul style="list-style-type: none"> • raw (binary) • locale (the current client locale is used) • locale:<locale name> (specific code page is used) <p>The convention for <locale name> depends on the operating system. For example, in Linux environments, the locale name for the code page GB18030 in Chinese/China is zh_CN.gb18030. In Windows environments, the locale name for Latin1 code page in Finnish/Finland is fin_fin.1252.</p> <p>The value 'raw' can be used when you want your database to use the binding used in the 6.3 or earlier versions of solidDB.</p>	locale
RowsPerMessage	<p>Specifies the number of rows returned from the server in one network message when an SQLFetch call is executed (and there are no pre-fetched rows).</p> <p>See also the ExecRowsPerMessage parameter.</p>	decided by the server
StatementCache	<p>Statement cache is an internal memory storing a few previously prepared SQL statements. With this parameter, you can set the number of cached statements per session.</p>	6
UseEncryption	<p>This parameter defines whether passwords are encrypted. If set to no, passwords are not encrypted.</p>	yes

H.3 Com section

Table 204. Client-side communication parameters

[Com]	Description	Factory Value
ClientReadTimeout	This parameter defines the connection (or read) timeout in milliseconds. A network request fails if no response is received during the time specified. The value 0 sets the timeout to infinite. This value can be overridden with the connect string option <code>-r</code> and, further on, with the ODBC attribute <code>SQL_ATTR_CONNECTION_TIMEOUT</code> . Note: This parameter applies only to the TCP protocol.	0 (infinite)
Connect	The Connect parameter defines the default network name (connect string) that the client uses to connect to the solidDB server, if the connect string is not specified in the connection parameters explicitly. This value is used also when the <code>SQLConnect()</code> call is issued with an empty data source name. The format of the standard solidDB connect string is: <code>protocol_name [options] [host_computer_name] server_name</code> where <code>options</code> and <code>server_name</code> depend on the communication protocol. Important: In HotStandby and SMA setups, additional connect string attributes are used to specify further functionality, such as Transparent Connectivity (TC). For more details, see Network name and connect string syntax.	tcp localhost 1964
ConnectTimeout	The ConnectTimeout parameter defines the login timeout in milliseconds. This value can be overridden with the connect string option <code>-c</code> and, further on, with the ODBC attribute <code>SQL_ATTR_LOGIN_TIMEOUT</code> . Note: This parameter applies only to the TCP protocol.	OS-specific
ODBCHandleValidation	This parameter switches ODBC handle validation on or off. See also section <i>ODBC handle validation</i> in <i>IBM solidDB Programmer Guide</i> for more information about the <code>SQL_ATTR_HANDLE_VALIDATION</code> ODBC attribute.	no
Trace	If this parameter is set to <code>yes</code> , trace information about network messages for the established network connection is written to a file specified with the TraceFile parameter.	no
TraceFile	If the Trace parameter is set to <code>yes</code> , trace information about network messages is written to a file specified with this parameter. The trace file is output to the current working directory of the server or client, depending on which end the tracing is started.	soltrace.out

H.4 Data Sources section

Table 205. Data Sources parameters

[Data Sources]	Description	Factory Value	Access Mode
logical name = network name, Description	These parameters can be used to give a logical name to a solidDB server in a <code>solid.ini</code> file of the client application.		N/A

H.5 SharedMemoryAccess section

Table 206. Shared memory access parameters (client-side)

[SharedMemoryAccess]	Description	Factory value	Startup
SignalHandler	<p>The SignalHandler parameter controls the SMA signal handler functionality.</p> <p>When set to yes, the SMA driver signal handler handles the signals defined with the Signals parameter.</p> <p>The SMA driver signal handler enables the SMA system to survive the most common application failures, such as killing or interrupting the applications from outside, or when one of the application threads runs within the server code, and another thread running application code causes application to crash.</p> <p>Upon the capture of certain signals, the signal handler closes the SMA connections safely and exits the SMA application. This means that in most cases, the SMA server continues to run despite abnormal application exits.</p> <p>The SMA driver signal handler installs itself when the first SMA connection is established and uninstalls itself when the last SMA connection is closed. Previously installed signal handlers are retained.</p>	yes	NA
Signals	<p>This parameter defines the signals that can break the SMA connection and should be handled by the SMA driver.</p> <p>The signals are defined as integers or with the following mnemonics: SIGSTOP, SIGKILL, SIGINT, SIGTERM, SIGQUIT, SIGABORT.</p> <p>Note: If the SMA application loops outside of the SMA driver (for example, does not call any functions), the signal can fail to terminate the application. In such a case:</p> <ol style="list-style-type: none"> 1. Throw out the connections at the server. admin command 'throwout <userid>' 2. Use SIGKILL signal to force the SMA application to exit. kill -SIGKILL <pid> 	<p>Linux and UNIX: SIGINT, SIGTERM</p> <p>Windows: SIGINT</p>	NA

H.6 TransparentFailover section

Table 207. TransparentFailover parameters

[TransparentFailover]	Description	Factory value
ReconnectTimeout	This parameter specifies how long (in milliseconds) the driver should wait until it tries to reconnect to the primary in case of TF switchover or failover. If the driver cannot find the new primary (reconnect), an error is returned and the TF connection becomes broken.	10000
WaitTimeout	This parameter specifies how long (in milliseconds) the driver should wait for the server to switch state. When the driver tries to reconnect to the servers, it might connect to the server being in an intermediate (switching or uncertain) state.	10000

Index

A

- ABS (function) 274
- ACOS (function) 274
- ad hoc query example 41
- APD (Application Parameter Descriptor) 228
- APIs
 - JDBC Driver 51
- application development
 - HotStandby
 - creation 41
 - testing and debugging 50
- Application Parameter Descriptor (APD) 228
- Application Row Descriptor (ARD) 228
- ARD (Application Row Descriptor) 228
- Array interface 56
- ASCII (function) 270
- ASIN (function) 274
- ATAN (function) 274
- ATAN2 (function) 274
- autocommit
 - warnings about SELECT statements 25
- autocommit mode
 - cursors 25
 - JDBC Driver 54
 - transactions 25

B

- binding
 - assigning storage for rowsets 32
 - column-wise 32
 - row-wise 32
 - Unicode 164
- BIT
 - SQL_BIT 241, 254
- BIT_LENGTH (function) 270
- BLOBs (Binary Large Objects)
 - interface 56
- block cursor 32
- bookmarks
 - description 39
 - using 39

C

- CallableStatement interface 56
- calling procedures 27
- CAST (function)
 - description 283
- CEILING (function) 274
- CHAR (function) 270
- CHAR_LENGTH (function) 270
- CHARACTER_LENGTH (function) 271
- Client timeouts 285
- client-side configuration parameters 293
- ClientReadTimeout (parameter) 295
- CLOB data type
 - JDBC interface 56
- CONCAT (function) 271

- configuration file
 - client-side 5, 293
- configuring
 - client-side configuration file 5
 - default settings 5
 - factory values 5
 - parameter settings 5
 - solid.ini 5
- Connect (parameter) 295
- connect string 3
 - using 16
- Connect Timeout 291
- connecting to solidDB
 - with sample application 92
- connection idle timeout 291
- connection interface 56
- connection timeout 286
- ConnectionPoolDataSource API Functions
 - Constructor 68
 - getConnectionURL 68
 - getDescription 68
 - getLoginTimeout 68
 - getLogWriter 68
 - getPassword 68
 - getPooledConnection 68
 - getURL 68
 - getUser 68
 - setConnectionURL 68
 - setDescription 68
 - setLoginTimeout 68
 - setLogWriter 68
 - setPassword 68
 - setURL 68
 - setUser 68
- connections
 - terminating 41
- ConnectTimeOut (parameter) 295
- constraints
 - Gregorian calendar 240
- conversion
 - explicit data type 282
- CONVERT (function)
 - description 282
- converting data
 - from C to SQL data types 254
 - from SQL to C data types 241
- COS (function) 274
- COT (function) 274
- CURDATE (function) 277
- CURRENT_CATALOG() scalar function 15
- CURRENT_SCHEMA() scalar function 15
- CURRENT_TIMESTAMP (function) 277
- CURRENTTIME (function) 276
- cursors
 - autocommit 25
 - block cursor 32
 - dynamic 33, 34
 - forward 34
 - scrollable 33
 - specifying the type 34
 - static 34

cursors (*continued*)
 types supported 33, 34
 using 32
CURTIME (function) 277

D

Data Sources
 configuring for Windows 20
 connecting to 16
 defining in solid.ini 19
 empty data source name 20
 retrieving catalog information 26
data types 223
 explicit conversion 282
DATABASE (function) 281
DatabaseMetaData interface
 methods 56
DAYNAME (function) 277
DAYOFMONTH (function) 277
DAYOFWEEK (function) 277
DAYOFYEAR (function) 277
debugging
 applications 50
DEGREES (function) 274
DIFFERENCE (function) 271
driver interface (JDBC) 56
dynamic libraries 12

E

END LOOP 217
errors
 JDBC Driver 54
 processing messages 40
 SA functions 93
 sample messages 39
ExecRowsPerMessage (parameter) 294
EXP (function) 275
EXTRACT (function) 277

F

FLOOR (function) 275
fn
 usage in {fn func_name} 14, 279
forward cursor 34
functions
 ABS 274
 ACOS 274
 additional extensions to SQL 29
 ASCII 270
 ASIN 274
 ATAN 274
 ATAN2 274
 BIT_LENGTH 270
 CAST 283
 CEILING 274
 CHAR 270
 CHAR_LENGTH 270
 CHARACTER_LENGTH 271
 CONCAT 271
 CONVERT 282
 COS 274
 COT 274
 CURDATE 277

functions (*continued*)

CURRENT_TIMESTAMP 277
CURRENTTIME 276
CURTIME 277
DATABASE 281
DAYNAME 277
DAYOFMONTH 277
DAYOFWEEK 277
DAYOFYEAR 277
DEGREES 274
DIFFERENCE 271
executing asynchronously 26
EXP 275
EXTRACT 277
FLOOR 275
guidelines for calling 14
HOUR 278
IFNULL 281
INSERT 271
LCASE 271
LEFT 271
LENGTH 271
LOCATE 272
LOG 275
LOG10 275
LTRIM 272
MINUTE 278
MOD 275
MONTH 278
MONTHNAME 278
NOW 278
OCTET_LENGTH 272
PI 275
POSITION 272
POWER 275
prototypes 14
QUARTER 278
RADIANS 275
REPEAT 272
REPLACE 272
return codes 15
RIGHT 272
ROUND 275
RTRIM 273
scalar 14
Scalar 269
SECOND 278
SIGN 275
SIN 275
SOUNDEX 273
SPACE 273
SQLAllocConnect 175
SQLAllocEnv 175
SQLAllocHandle 175
SQLAllocStmt 177
SQLBindCol 179
SQLBindParameter 177
SQLBrowseConnect 175
SQLBulkOperations 180
SQLCancel 182
SQLCloseCursor 182
SQLColAttribute 179
SQLColAttributes 179
SQLColumnPrivileges 181
SQLColumns 181
SQLConnect 175
SQLCopyDesc 177

functions (*continued*)

- SQLDataSources 176
- SQLDescribeCol 179
- SQLDescribeParam 178
- SQLDisconnect 182
- SQLDriverConnect 175
- SQLDrivers 176
- SQLEndTran 182
- SQLError 180
- SQLExecDirect 178
- SQLExecute 178
- SQLExtendedFetch 180
- SQLFetch 179
- SQLFetchScroll 180
- SQLForeignKeys 181
- SQLFreeConnect 182
- SQLFreeEnv 182
- SQLFreeHandle 182
- SQLFreeStmt 182
- SQLGetConnectAttr 176
- SQLGetConnectOption 177
- SQLGetCursorName 178
- SQLGetData 180
- SQLGetDescField 177
- SQLGetDescRec 177
- SQLGetDiagField 180
- SQLGetDiagRec 180
- SQLGetEnvAttr 177
- SQLGetFunctions 176
- SQLGetInfo 175
- SQLGetStmtAttr 177
- SQLGetStmtOption 177
- SQLGetTypeInfo 176
- SQLMoreResults 180
- SQLNativeSQL 178
- SQLNumParams 178
- SQLNumResultCols 179
- SQLParamData 178
- SQLParamOptions 178
- SQLPrepare 177
- SQLPrimaryKeys 181
- SQLProcedureColumns 181
- SQLProcedures 181
- SQLPutData 178
- SQLRowCount 179
- SQLSetConnectAttr 176
- SQLSetConnectOption 177
- SQLSetCursorName 178
- SQLSetDescField 177
- SQLSetDescRec 177
- SQLSetEnvAttr 177
- SQLSetParam 178
- SQLSetPos 180
- SQLSetScrollOptions 178
- SQLSetStmtAttr 177
- SQLSetStmtOption 177
- SQLSpecialColumns 181
- SQLStatistics 181
- SQLTablePrivileges 182
- SQLTables 182
- SQLTransact 182
- SQRT 275
- SUBSTRING 273
- system functions 281
- TAN 275
- time and date 276
- TIMESTAMPADD 279

functions (*continued*)

- TIMESTAMPDIFF 280
- TRIM 273
- TRUNCATE 276
- UCASE 273
- Unicode strings 164
- USER 281
- WEEK 280
- YEAR 280

H

- header files 14
- hints 28
- HotStandby
 - timeouts 291
- HOUR (function) 278

I

- IFNULL (system function) 281
- INSERT (string function) 271

J

- Java interfaces
 - Array 56
 - Blob 56
 - CallableStatement 56
 - Clob 56
 - Connection 56
 - database access 51
 - DatabaseMetaData 56
 - Driver 56
 - Naming and directory interface 78
 - PreparedStatement 56
 - Ref 56
 - ResultSet 56
 - ResultSet class 56
 - ResultSetMetaData 56
 - SQLData 56
 - SQLInput 56
 - SQLOutput 56
 - Statement 56
 - Struct 56
- Java Transaction API (JTA) 63
- JDBC Connection Pooling 68
 - ConnectionPoolDataSource 68
 - PooledConnection 68
- JDBC Driver 78
- JNDI 78
- JTA (Java Transaction API) 63

L

- LCASE (function) 271
- LEFT (function) 271
- LENGTH (function) 271
- listen name 16
- LOCATE (function) 272
- lock wait timeout 289
- LOG (function) 275
- LOG10 (function) 275
- login
 - timeout 285

LOGIN_CATALOG() scalar function 15
LOOP 217
LTRIM (function) 272

M

MaxSpace (parameter) 169
MINUTE (function) 278
MOD (function) 275
MONTH (function) 278
MONTHNAME (function) 278

N

native scalar functions 14
network names 16
NoAssertMessages (parameter) 294
non-standard behavior
 ODBC 5
NOW (function) 278
numeric functions
 ODBC 273

O

octet length 238
OCTET_LENGTH (function) 272
ODBC
 additional functions to SQL 29
 Driver 14
 extensions 27
 function support 175
 non-standard behavior 5
 solidDB extensions for ODBC API 30
ODBC handles 24
ODBCCharBinding (parameter) 294
ODBCHandleValidation (parameter) 295
optimistic lock timeout 289
optimizer hints 28

P

parameters
 client-side 293
PI (function) 275
Ping Timeout 291
PooledConnection API Functions
 addConnectionEventListener 68
 close 68
 getConnection 68
 removeConnectionEventListener 68
POSITION (function) 272
POWER (function) 275
PreparedStatement interface
 methods 56
procedures
 calling in ODBC 27

Q

QUARTER (function) 278
query timeout 287

R

RADIANS (function) 275
Ref interface
 methods 56
REPEAT (function) 272
REPLACE (function) 272
ResultSet interface
 methods 56
ResultSetMetaData interface
 methods 56
return code
 for functions 15
RIGHT (function) 272
ROUND (function) 275
rowset 32
RowsPerMessage (parameter) 294
RTRIM (function) 273

S

SA
 sample program 92
SaErrorInfo
 solidDB SA 93
scalar functions 14
 native 14
 ODBC 269
 SQL-92 269
scrollable cursors 33
SECOND (function) 278
server timeouts 288
SET LOGREADER BATCH (statement) 174
SIGN (function) 275
SIN (function) 275
solid.ini 293
 configuration parameters 293
solid.jdbc.SolidBaseRowSet 76
solidDB Data Dictionary
 Unicode 161
solidDB Export
 Unicode 161
solidDB JDBC Driver
 classes and methods 56
 connection to the database 54
 conversion matrix 89
 DatabaseMetaData interface 56
 description 6, 51
 Driver class 56
 getting started 51
 PreparedStatement interface 56
 Ref interface 56
 registering 53
 ResultSet interface 56
 ResultSetMetaData interface 56
 SQLData interface 56
 SQLInput interface 56
 SQLOutput interface 56
 Statement interface 56
 Struct interface 56
 Unicode 164
solidDB ODBC API
 Unicode 164
solidDB ODBC Driver
 description 13
 driver manager 14
 files 13

- solidDB ODBC Driver (*continued*)
 - on Microsoft Windows 13
 - Unicode 164
 - using 11
- solidDB ODBC functions 175
- solidDB SA 101
 - building a sample program 92
 - connection to the database 92
 - delete 93
 - description 91
 - getting started 92
 - handling database errors 93
 - reading data without SQL 93
 - running SQL statements 93
 - setting up the development environment 92
 - transactions and autocommit mode 93
 - update 93
 - writing data without SQL 93
- solidDB SQL Editor 161
- SOUNDEX (function) 273
- SPACE (function) 273
- Speed Loader
 - Unicode 161
- SQL statements
 - execution timeout 288
- SQL_C_BIT
 - binding C variable of type SQL_C_BIT 241, 254
- SQL_C_DEFAULT
 - avoid use of 232
- SQL_CLOSE
 - option in SQLFreeStmt() function call 32
- SQL_DELETE
 - option in SQLSetPos() function call 32
- SQL_NTS
 - null-terminated string 267
- SQL_POSITION
 - option in SQLSetPos() function call 32
- SQL_ROWSET_SIZE
 - option in SQLSetStmtAttr() function call 32
- SQL_UPDATE
 - option in SQLSetPos() function call 32
- SQLAllocConnect (function) 175
- SQLAllocEnv (function) 175
- SQLAllocHandle (function) 175
- SQLAllocStmt (function) 177
- SQLBindCol
 - function description 32
- SQLBindCol (function) 179
- SQLBindParameter (function) 177
- SQLBrowseConnect (function) 175
- SQLBulkOperations (function) 180
- SQLCancel (function) 182
- SQLCloseCursor (function) 182
- SQLColAttribute (function) 179
- SQLColAttributes (function) 179
- SQLColumnPrivileges (function) 181
- SQLColumns (function) 181
- SQLConnect (function) 175
- SQLCopyDesc (function) 177
- SQLData interface
 - methods 56
- SQLDataSources (function) 176
- SQLDescribeCol (function) 179
- SQLDescribeParam (function) 178
- SQLDisconnect (function) 182
- SQLDriverConnect (function) 175
- SQLDrivers (function) 176
- SQLEndTran (function) 182
- SQLError (function) 180
- SQLExecDirect (function) 178
- SQLExecute (function) 178
- SQLExtendedFetch (function) 32, 34, 180
- SQLFetch (function) 32, 179
- SQLFetchScroll (function) 32, 34, 180
- SQLForeignKeys (function) 181
- SQLFreeConnect (function) 182
- SQLFreeEnv (function) 182
- SQLFreeHandle (function) 182
- SQLFreeStmt (function) 32, 182
- SQLGetConnectAttr (function) 176
- SQLGetConnectOption (function) 177
- SQLGetCursorName (function) 178
- SQLGetData (function) 180
- SQLGetDescField (function) 177
- SQLGetDescRec (function) 177
- SQLGetDiagField (function) 180
- SQLGetDiagRec (function) 180
- SQLGetEnvAttr (function) 177
- SQLGetFunctions (function) 176
- SQLGetInfo (function) 175
- SQLGetStmtAttr (function) 177
- SQLGetStmtOption (function) 177
- SQLGetTypeInfo (function) 176
- SQLInput interface
 - methods 56
- SQLMoreResults (function) 180
- SQLNativeSQL (function) 178
- SQLNumParams (function) 178
- SQLNumResultCols (function) 179
- SQLOutput interface
 - methods 56
- SQLParamData (function) 178
- SQLParamOptions (function) 178
- SQLPrepare (function) 177
- SQLPrimaryKeys (function) 181
- SQLProcedureColumns (function) 181
- SQLProcedures (function) 181
- SQLPutData (function) 178
- SQLRowCount (function) 179
- SQLSetConnectAttr (function) 176
- SQLSetConnectOption (function) 177
- SQLSetCursorName (function) 178
- SQLSetDescField (function) 177
- SQLSetDescRec (function) 177
- SQLSetEnvAttr (function) 177
- SQLSetParam (function) 178
- SQLSetPos (function) 32, 180
- SQLSetScrollOptions (function) 178
- SQLSetStmtAttr (function) 32, 177
 - dynamic cursors 33, 34
- SQLSetStmtOption (function) 177
- SQLSpecialColumns (function) 181
- SQLStatistics (function) 181
- SQLTablePrivileges (function) 182
- SQLTables (function) 182
- SQLTransact (function) 182
- SQRT (function) 275
- Statement interface
 - methods 56
- StatementCache (parameter) 294
- static cursor 34
- static library 12
- static SQL
 - code example 41

- stored procedures
 - JDBC Driver 56
- string functions
 - ODBC 270
- struct interface
 - solidDB JDBC Driver 56
- SUBSTRING (function) 273

T

- table lock wait timeout 290
- TAN (function) 275
- TC Info 16
- testing
 - applications 50
- Timeout controls 285
- TIMESTAMPADD (function) 279
- TIMESTAMPDIFF (function) 280
- Trace (parameter) 295
- TraceFile (parameter) 295
- transaction idle timeout 290
- transactions
 - autocommit mode 25
 - committing read-only 25
 - JDBC Driver 54
 - terminating 41
- Transfer Octet Length 238
- translation
 - effect on Unicode columns 164
- TRIM (function) 273
- TRUNCATE (function) 276

U

- UCASE (function) 273
- Unicode
 - character translation 164
 - compliance 157
 - converting 162
 - creating columns for storing data 159
 - description 157
 - encoding forms 158
 - loading data 159, 161
 - setting up 159, 164
 - solidDB Data Dictionary 161
 - solidDB Export 161
 - solidDB JDBC Driver 164
 - solidDB ODBC API 164
 - solidDB ODBC Driver 164
 - solidDB Remote Control 161
 - solidDB SQL Editor 161
 - Speed Loader 161
 - standard 158
 - string functions 164
 - user names and passwords 159
 - using in database entity names 159
 - variables and binding 164
- unixODBC 21
- UseEncryption (parameter) 294
- USER (function) 281
- UTF-16
 - description 158
- UTF-8
 - description 158

V

- Variables
 - Unicode 164

W

- WebSphere
 - compatibility 63
- WEEK (function) 280

Y

- YEAR (function) 280

Notices

© Copyright Oy International Business Machines Ab 1993, 2011.

All rights reserved.

No portion of this product may be used in any way except as expressly authorized in writing by Oy International Business Machines Ab.

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 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.



Printed in USA

SC27-3840-00

