

IBM solidDB
IBM solidDB Universal Cache
Version 6.3

Linked Library Access User Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 57.

First edition, third revision

This edition applies to version 6, release 3 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 1993, 2011

Contents

Figures v

Tables vii

Summary of changes. ix

About this manual xi

Typographic conventions xi

Syntax notation conventions. xii

1 Introducing linked library access 1

Linked library access library 2

 Disk-based vs. diskless servers 3

 Library contents 3

 Application types used with linked library access 4

solidDB client APIs and drivers for linked library

access. 5

 solidDB SA API 5

 solidDB ODBC API 6

 solidDB JDBC API 6

 solidDB Server Control API (SSC API). 7

2 Creating and running an linked library access application 9

Accessing the linked library access library 9

 Libraries for remote applications. 9

 Sample C applications. 10

 Using data synchronization 10

Linking applications for the linked library access 11

 Preparing user applications for the linked library

 access 12

 Establishing a local or remote connection to

 solidDB with the linked library access 14

Starting and shutting down solidDB linked library

access 15

 Explicit startup with the Control API function

 SSCStartServer 16

 Implicit startup with ODBC API function call

 SQLConnect 18

 Implicit startup with SA API function call

 SaConnect. 20

 Shutting down solidDB linked library access 20

 Implicit start configuration parameter 21

3 Description of control API 23

Retrieving task information 23

Notifying functions of a special event 23

 Obtaining solidDB status and server information 23

Summary of control API functions. 24

Control API and equivalent ADMIN COMMANDs 24

Control API reference 25

 Function synopsis 25

 Return value 26

 Control API error codes and messages 26

SSCGetServerHandle 27

 Synopsis 27

 Comments. 27

 Return value 27

SSCGetStatusNum 27

 Synopsis 28

 Comments. 28

 Return value 28

SSCIsRunning 28

 Synopsis 28

 Return value 28

 Comments. 29

SSCIsThisLocalServer 29

 Synopsis 29

 Return value 29

 Comments. 29

SSCRegisterThread 29

 Synopsis 29

 Return value 29

 Comments. 29

 See also. 30

SSCSetNotifier 30

 Synopsis 30

 Return value 31

 Comments. 32

 Example 32

SSCSetState 33

 Synopsis 33

 Return value 34

 Comments. 34

SSCStartDisklessServer 34

 Synopsis 34

 SSCStartDisklessServer parameter options 35

 Return value 35

 Comments. 35

 Example 36

 See also. 36

SSCStartServer 36

 Synopsis 36

 Return value 37

 Comments. 38

 See also. 38

SSCStopServer 38

 Synopsis 38

 Return value 39

 Comments. 39

 See also. 39

SSCUnregisterThread 39

 Synopsis 39

 Return value 40

 Comments. 40

 See also. 40

4 Using the diskless capability 41

Configuration parameters for a diskless server. 41

Parameters used in diskless servers	41
Configuration parameters that do not apply to diskless engines	43
5 Using solidDB linked library access with Java.	45
Overview of solidDB JDBC Accelerator (SJA)	45
How the Accelerator works	45
System requirements	47
Basic usage	47
Installation	47
Compiling and running a program	47
Making JDBC connections	48

Limitations	49
solidDB Server Control (SSC) API	49

Appendix. Linked library access parameters 53

Linked library access parameters	53
Accelerator section	53

Index 55

Notices 57

Figures

1. solidDB with Linked Library Access.	2	3. solidDB with Linked Library Access - APIs	7
2. Linking to solidDB	4		

Tables

1.	Typographic conventions	xi	12.	SCCRegisterThread parameters	29
2.	Syntax notation conventions	xii	13.	SSCSetNotifier parameters	30
3.	Linked library access system libraries	11	14.	SSCSetState parameters	33
4.	Library files	12	15.	SSCStartDisklessServer parameters.	34
5.	SSCStartServer parameters	16	16.	Command line options for the <i>argv</i> parameter	35
6.	SSCStartServer argv options	17	17.	SSCStartServer parameters	36
7.	Summary of control API functions	24	18.	SSCStopServer parameters	38
8.	Control API parameter usage types	25	19.	SCCUnregisterThread parameters	39
9.	Error codes and messages for control API functions	26	20.	Configuration parameters not applicable to diskless engines	43
10.	SSCGetStatusNum parameters	28	21.	Accelerator parameters.	53
11.	SSCIsRunning parameters.	28			

Summary of changes

Changes for revision 03

- Editorial corrections.

Changes for revision 02

- Editorial corrections.

Changes for revision 01

- Editorial corrections.

About this manual

The IBM® solidDB® Linked Library Access is a higher performance version of solidDB® data management solution. To avoid network delays, the solidDB executable and the user application are linked in the same program space to produce a single executable. By replacing the network connection and Remote Procedure Calls (RPCs) with local function calls, performance is improved significantly.

This guide contains information specific to the linked library access. This guide supplements the information contained in the *IBM solidDB Administrator Guide*, which contains details on administration and maintenance of solidDB.

This guide assumes a working knowledge of the C programming language, general DBMS knowledge, familiarity with SQL, and knowledge of a solidDB data management product, such as solidDB in-memory database, or solidDB disk-based engine. If you are going to work with the solidDB Java™ Accelerator, then this manual also assumes a working knowledge of Java.

Typographic conventions

solidDB documentation uses the following typographic conventions:

Table 1. Typographic conventions

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

Table 1. *Typographic conventions (continued)*

Format	Used for
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 Introducing linked library access

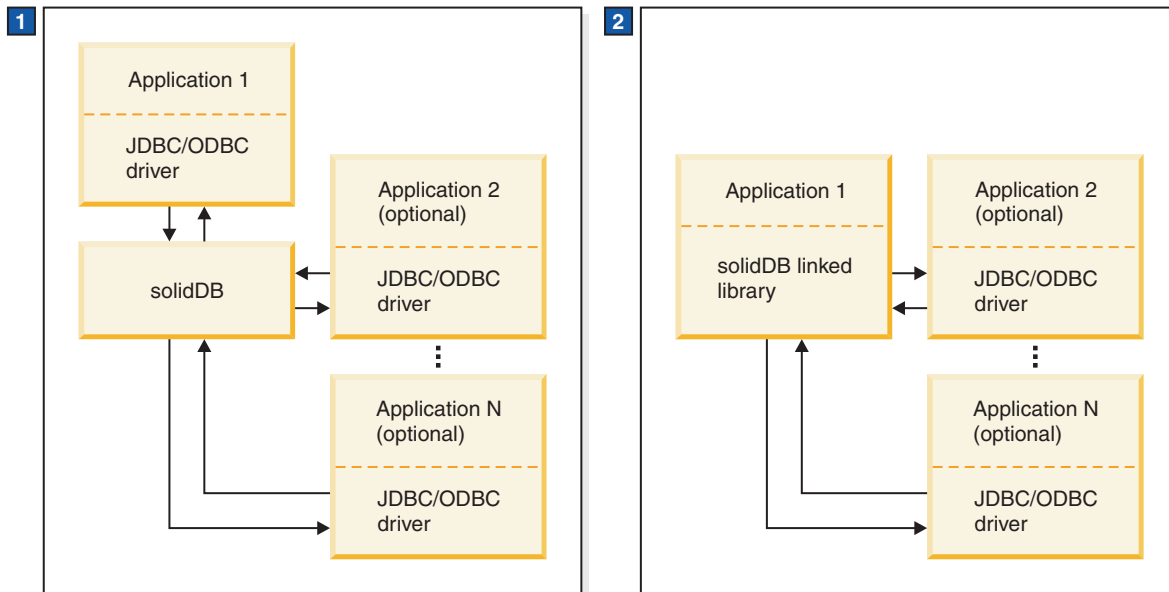
The IBM® solidDB Linked Library Access is a function library that provides the same functionality and interfaces available with the solidDB. A user application may be linked to this library. The linked application communicates with the server by using direct function calls, thus skipping the overhead required when the client and server communicate through network protocols such as TCP/IP. Linking the application and server into a single executable provides higher performance.

Your application does not have to be re-written to use the linked library access library. For example, you do not need to call proprietary functions (except a few to start and stop the database server). Instead, your application may continue to use the same ODBC function calls that it has always used. When the linked library access library is linked to your application, these ODBC function calls go directly to the server, bypassing the network.

Your application also has access to some additional linked library access function calls to do things such as scheduling tasks within the server. However, you are not required to use these function calls unless you want to.

The fact that your server is linked to your application does not mean that your linked application is the only client that can use the server. A solidDB server that is executing as an linked library access function library is accessible not only to the "local" client application (the application that is linked directly to the library), but also to "remote" client applications (which connect to the server through communications protocols such as TCP/IP). Your remote clients see the linked library access server as similar to any other solidDB server, while your local client sees a faster, more precisely controllable version of the solidDB server.

Note: Although "remote" applications usually run on a different computer from the one that the server is running on, an application is also considered "remote" if it uses a network communication protocol to communicate with the server, even if that client runs on the same computer as the database server runs on.



1. In a standard solidDB database configuration, the applications and the server are separate programs.
2. solidDB linked library is a library of subroutines that are linked into an application. Other applications may also communicate with the server.

Figure 1. solidDB with Linked Library Access

The figure above shows a sample solidDB that uses the linked library access library.

Note:

Local application requests are handled through solidDB SA API or ODBC API direct function calls. For the local application, linked library access also provides a Control API which handles local requests for controlling solidDB background processes and client tasks. You may also use JDBC calls with the linked library access. See 5, "Using solidDB linked library access with Java," on page 45.

As you can see in the illustration, remote clients communicate through an ODBC or JDBC driver that is linked to the client application, while the local client application does not need any remote communication driver.

Linked library access library

In a standard (non-linked library access) solidDB configuration, the application (the "client") and the database engine (the "server") are separate processes that communicate through a network protocol. The client must link to a communications driver (such as an ODBC or JDBC driver) that communicates with the database server through the network.

With the linked library access, an application links to a static library (for example, .lib or .a for UNIX) that contains the full database server functionality. This means solidDB runs in the same executable with the application, eliminating the need to transfer data through the network. The application that links to the linked library access library can also have multiple connections, using both ODBC API and SA API. Both of these APIs are reentrant, allowing simultaneous connections from separate threads.

A user application that links directly to the linked library access library can also create remote connections to other database servers. The connect string that is passed to the ODBC API or SA API connect function defines whether the connection type is local or remote.

For details on linking an application, read “Linking applications for the linked library access” on page 11.

When you start your application, only the code in your application starts running automatically. The server code is largely independent of your application code, and you must explicitly start the server by calling a function. (In most or all implementations, the server runs on threads that are separate from the thread(s) used by the application. Calling the function to start a server will perform any initialization steps required by the server code, create the appropriate additional threads if necessary, and start the server running on those threads.)

Disk-based vs. diskless servers

Linked library access library contains two different function calls to start the server. One of the function calls starts a normal (that is, disk-based) server, while the other starts a server that does not use the disk drive. For more information, see 4, “Using the diskless capability,” on page 41 and the descriptions of the `SSCStartServer` and `SSCStartDisklessServer` functions.


Library contents

Linked library access library includes functions for three separate APIs:

- solidDB Control API (SSC API) library that contains functions to control task scheduling.
- solidDB ODBC Driver functions that allows for direct communication with the server library, without going through the network.
- solidDB SA API library which may be required for additional functionality using the linked library access. For example, this library allows you to insert, delete, and select records from a table.

Because your application gets linked to a library with all three of these APIs (SSC, SA, and ODBC), your application program may call functions from any combination of these APIs. For details on each of these APIs, read “solidDB client APIs and drivers for linked library access” on page 5.

Note: Remote applications have access to the same three APIs (SSC, SA, and ODBC). However, the functions for these three APIs are not all in the same file for remote applications. For details on remote and dual role applications, read “Application types used with linked library access” on page 4. For information on API files for remote applications, read “solidDB client APIs and drivers for linked library access” on page 5.



1. In a standard solidDB database configuration, the applications and the server are separate programs.
2. solidDB linked library is a library of subroutines that are linked into an application. Other applications may also communicate with the server.

Figure 2. Linking to solidDB

Application types used with linked library access

The linked library access application is "local" to the server; the server and the application are combined into a single program. Calls to ODBC functions actually go directly to the server, rather than going through an ODBC driver and the communications protocol (such as TCP/IP).

In addition to handling requests from the local application that is linked to the linked library access library, the server also handles requests from remote applications.

A remote application is not linked to the linked library access library. It is a separate executable that must communicate with the server using a network connection (such as TCP/IP) or other connection (for example shared memory). Remote applications are usually, but not always, run on a different computer from the one that is running the server. However, a single computer can run an linked library access local application, while running one or more remote applications as separate processes.

Most applications are either local (that is linked to the linked library access library in a single executable) or remote (never linked to the linked library access library). However, it is also possible to write an application that can be either local and remote; it switches modes, depending upon how it is compiled and linked. Such a *dual mode application* uses, for example, the same C-language application code in either local or remote mode; but it is linked to a different library when in local mode than when in remote mode.

Using dual-mode applications with linked library access

In the case of linked library access, for example, a dual-mode application must be linked to the local linked library access library when it is run locally. However, when it is run remotely, the dual-mode application must be linked to the linked library access control API stub library (for example, `solidctrlstub.lib` in Windows), so that it can be compiled, linked, and executed without link-time errors.

The "Control API stub library" is required for remote applications because the linked library access's own Control API (which is provided in the local linked library access library) cannot be used with remote applications. For example, assume you have a local application (containing Control API functions) that links to a standard ODBC library. You want to run the same application remotely. By linking to the Control API stub library, you avoid having to remove the Control API function calls from your code. In this way, you can easily turn your linked library access local application into a normal remote client application.

Note: The Control API stub library contains "do-nothing" functions; if you call them in a remote application, they have no effect on the server.

A dual-mode application is useful for other reasons as well:

- You may want to test your local application first before linking it with the linked library access library.
- You may want all users/processes to have the same application logic whether they are local or remote.

Dual-mode application

Assume there are two users who are running the same application. User1 runs the application locally (benefiting from higher performance). User2 runs the same application remotely.

User1 (local user) compiles and links with the server library (`solidac.a`, for example) and is responsible for starting and stopping the server and performing other scheduling tasks using the linked library access's Control API. User2 (remote user) runs the same application, but is not able to connect to the server until User1 has started the server. Thus, only User1 is able to control the tasking system.

solidDB client APIs and drivers for linked library access

Below is a brief description of the APIs available for use with the linked library access.

Note:

These descriptions use the term "local" and "remote" applications as defined in "Application types used with linked library access" on page 4.

solidDB SA API

solidDB SA defined SA API is a low-level proprietary C-language API to solidDB data management services. It is included in the linked library access library (for example, `ssolidacxx.dll` for Windows or `solidac.a` for UNIX). The linked library access library includes the SA-API library that provides support for local applications using SA API function calls.

The SA API 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. solidDB SA API can enhance performance significantly. You can use SA API to optimize the performance of batch insert operations, for example.

For remote applications, the linked library access library also provides support for the SA API function calls. However, you must link to a separate SA API library file (for example, `solidimpsa.lib` for Windows).

For details on the solidDB SA API, see *solidDB Programmer Guide*.

solidDB ODBC API

solidDB ODBC API provides a standards-compliant way to access data of a local or remote solidDB database through SQL. It provides functions for controlling database connections, executing SQL statements, retrieving result sets, committing transactions, and other data management functionality.

ODBC API, a Call Level Interface (CLI) for solidDB databases, is compliant with ANSI X3H2 SQL CLI, and is included in the linked library access library (for example, `ssolidacxx.dll` for Windows or `solidac.a` for UNIX).

Linked library access supports the ODBC 3.51 standard. The linked library access library includes solidDB ODBC 3.x, which provides support for local applications that require direct function calls to the server.

For local applications, the linked library access library provides support for ODBC function calls. For remote applications (or for a dual-mode application that is to be run remotely), you must link the ODBC Driver to get the same functionality.

If your application is a dual mode application (i.e. can be run either locally or remotely), and if it uses linked library access's Control API and ODBC, then you will need two different executables, one to be run locally and one to be run remotely. When you link your application to run it locally, you will link it to the linked library access library, which provides support for both the ODBC functions and the Control API library. When you link your application to run it remotely, you must link it to both the ODBC driver and to the Control API stub library (for example, `solidctrlstub.lib` for Windows). This stub library does not actually give your remote application any control over the server; it simply allows you to compile and link your program without getting errors about "unresolved symbols".

Note:

When ODBC functions (in a dual mode application) are called remotely, then the calls go through the network from the client to the server. When ODBC functions are called locally (in accelerated applications), then the ODBC subroutine library bypasses the network and directly connects the local application to the server.

Read *IBM solidDB Programmer Guide* for more details on ODBC API.

solidDB JDBC API

JDBC API is used by remote applications only. As the core API for JDK 1.2, it defines Java classes to represent database connections, SQL statements, result sets, database metadata, etc. It allows you to issue SQL statements and process the

results. JDBC is the primary API for database access in Java. Linked Library Access supports both JDBC 1.x and 2.x. Read *IBM solidDB Programmer Guide* for more details.

solidDB Server Control API (SSC API)

solidDB Server Control API (SSC API) is a C-language, thread-safe interface to control the server behavior in solidDB database products.

The Control API is included in the linked library access library (for example, `ssolidacxx.dll` for Windows or `solidac.a` for UNIX). The linked library access library provides support for local applications using Control API function calls and a separate library is available for remote-only applications.

If your application will run remotely and contains Control API function calls, then you must link the Control API Stub library (for example, `solidctrlstub.lib` for Windows). This library does not actually give your remote application control of the server; it merely allows you to compile and link your application as a remote application without getting link-time errors from solidDB with linked library access.

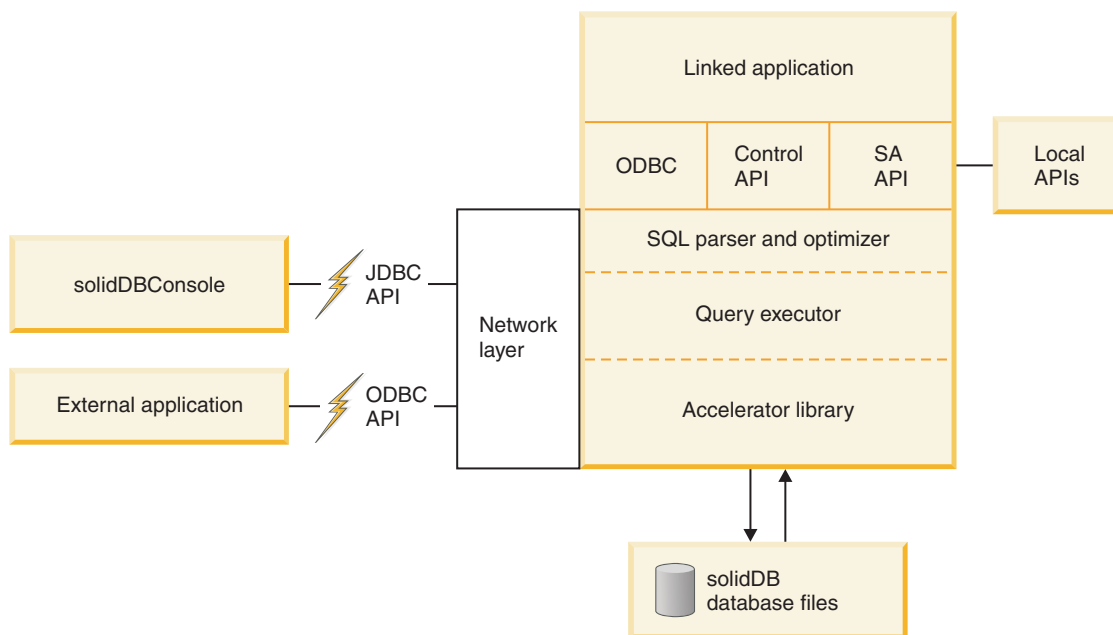


Figure 3. solidDB with Linked Library Access - APIs

2 Creating and running an linked library access application

This chapter describes how to create and run the linked library access application. It includes the following topics:

- Accessing the linked library access library
- Linking Applications to the linked library access library
- Creating or using an existing database
- Starting and stopping solidDB with the linked library access

Note: This chapter provides linked library access-specific additions, supplements, and linked library access usage differences from solidDB without the linked library access. For information on solidDB SQL, solidDB data management tools, general solidDB administration and maintenance, and database error codes, refer to the *IBM solidDB Administrator Guide*. Read 3, "Description of control API," on page 23 and *IBM solidDB Programmer Guide* for details on developing applications with an linked library access supported API.

Accessing the linked library access library

The solidDB with linked library access is a library file that is included in the solidDB Development Kit.

For example, if you are using solidDB with HP-UX, the linked library access library file is `solidac.a`. Refer to "Linking applications for the linked library access" on page 11 for a list of platform-specific libraries.

The linked library access library for all platforms contains the following:

- solidDB data management functionality
- SA API header (`sa.h`) for local user applications
- solidDB Control API interface header (`sscap.h`) for local user applications

For details on linking a user application to the linked library access library, read "Linking applications for the linked library access" on page 11.

Libraries for remote applications

For the purposes of this linked library access guide, a "remote" application is any application that is not linked to the server - that is, any application that is not using the linked library access library. Thus an application that is running on the same node as the database server, but that is not linked to it, is considered to be a "remote" application. A remote application communicates with the server through a network communications protocol such as TCP/IP. A "local" application, on the other hand, is linked to the linked library access library, and can call functions in that library directly, without going through a network protocol.

Because a remote application goes through the network communications protocol, the linked library access does not improve performance for remote applications. Only the local application (the one that is directly linked to the accelerator library) has higher performance.

In some cases, however, remote applications can benefit from improved performance by using the SA API, which allows low-level operations to read from and write to the database.

If you are using a remote application, you may need to link to the following libraries in the solidDB SDK into your application.

- Link to the solidDB Control API stub library (`solidctrlstub.lib` for Windows platforms), when you have Control API function calls in your application and you want to run your application remotely. (Note that if your application is a local rather than remote application —i.e. if it is directly linked to the linked library access library - then you do not need `solidctrlstub.lib`.)

For more details on the Control API Stub library (`solidctrlstub.lib`), read “Using dual-mode applications with linked library access” on page 5.

- Link to solidDB SA API (`solidimpsa.lib` for Windows platforms) if you are running a remote solidDB SA API application (without linked library access).

If you are using ODBC, SA API, or JDBC as remote applications only (that do not use Control API function calls), then you do not need to link to `solidctrlstub.lib`.

Sample C applications

For Accelerator Control API usage samples (available in C programming language), refer to `samples/aclib`, `samples/aclib_replication` and `samples/aclib_control_api` under the installation directory. These C samples reflect linked applications that use ODBC API functions to connect to solidDB servers.

Using data synchronization

If you are new to solidDB data synchronization, *IBM solidDB Advanced Replication User Guide* contains introductions on how to use the sample scripts provided with solidDB.

Before you run the sample C application `acsnet.c` (under directory `samples/aclib_replication`), it is recommended that you become familiar with solidDB functionality by doing at least one of the following:

- Using solidDB (without the linked library access) to run the SQL scripts contained in *IBM solidDB Advanced Replication User Guide*. These scripts are found in `samples/replication`.
- Running the SQL scripts locally, using the solidDB linked library access. As a prerequisite, you are required to set up an application to start the server according to the instructions in this chapter. For details, read “Linking applications for the linked library access” on page 11 and “Starting and shutting down solidDB linked library access” on page 15.

Note: You cannot use the SA API to run synchronization commands.

- Running the implementation sample file `aclibstandalone.c`, which with the linked library access library, emulates a normal server. The sample file is located in directory `samples/aclib`.

After using any of these methods, it is possible to run all the steps in *IBM solidDB Advanced Replication User Guide's* chapter titled *Getting Started with Data Synchronization* using solidDB SQL Editor (`solsql`).

Setting up your ODBC application with the advanced replication scripts

You can build an ODBC application, similar to the sample C application `acsNet.c`, to execute all statements required to set up, configure, and run a synchronizing environment. You can find `acsNet.c` under directory `samples/aclib_replication`.

To set up sample databases for use with an ODBC client application, you can execute sample scripts `replica3.sql`, `replica4.sql`, `replica5.sql`, and `replica6.sql`, all of which you can find in the `samples/replication/eval_setup` directory. These sample scripts contain SQL statements that write new data to replica(s) and control the execution of synchronization messages. These scripts may be run independently through the solidDB SQL Editor (`solsql`).

Alternatively, you can embed the SQL statements into a C/ODBC application, compile, and link it directly to the linked library access library. When linked with the linked library access, the sample scripts allow you to get the performance benefit inherent in linked library access's architecture.

The sample program `embed.c` in the `samples/odbc` directory illustrates how to set up databases with an ODBC client application using linked library access. You can insert the SQL commands from the sample scripts, such as `replica3.sql`, etc., into the `embed.c` application.

Linking applications for the linked library access

The solidDB linked library access is a library that must be linked to a user application. As long as the application is running, local and remote application requests for solidDB data management services are available through the library.

Note: If you are writing remote user applications that use solidDB Control API, you will need to link your remote application to the solidDB Control API stub library (for example, `solidctrlstub.lib` for Windows). If you are using solidDB SA API remotely (without linked library access) then you need to link to a separate solidDB API library (`solidimpsa.lib` for Windows). If you are only using ODBC, SA API, or JDBC remotely, without Control API, then there is no need to link to the solidDB Control API stub library.

You link only one application directly to the linked library access library at one time. However, once the linked application is up and running, and the server started, any network client can connect to the server using any of the protocols supported by the server, which depends on the operating system. These are for example, TCP/IP, shared memory and named pipes. Remote clients cannot use direct function calls.

When linking an application to solidDB with the linked library access, use one of the following libraries required for your operating system. Refer to your operating system documentation.

Table 3. Linked library access system libraries

Platforms	solidDB with linked library access library
Windows	<code>solidimpac.lib</code> (this is an import library file that gives you access to the real library file, which is <code>ssolidacxx.dll</code>)
Solaris	<code>solidac.a</code>

Table 3. Linked library access system libraries (continued)

Platforms	solidDB with linked library access library
HP-UX	solidac.a
Linux®	solidac.a
VxWorks	solidac.a

Preparing user applications for the linked library access

To allow your application to use the solidDB with linked library access, be sure to:

- Link to the linked library access library instead of to the driver libraries.
If you are using remote applications, you may need to link to other libraries. For details, read “Libraries for remote applications” on page 9.
- Change the connect string to the local or remote server name. For details, read “Establishing a local or remote connection to solidDB with the linked library access” on page 14.
- If needed, add calls to SSCStartServer and SSCStopServer or other Control API calls. For details, read “Control API reference” on page 25.

Signal handlers

Signal handlers are used to report the occurrence of an exceptional event to the application, for example division by zero. You must not set signal handlers in user applications because they would override the signal handlers that are set by the linked library access. For example, if the user application sets a signal handler for floating point exceptions, that setting overrides the handler set by the linked library access. Thus the server is unable to catch, for example, division by zero.

Dynamic link library

solidDB provides both a "static" and a "dynamic" version of the linked library access library. The names of the dynamic link library files are shown below for some major platforms. (For the names of the static libraries, see “Linking applications for the linked library access” on page 11.)

Table 4. Library files

Platforms	solidDB with Linked Library Access Library
Windows	ssolidacxx.dll
Solaris	ssolidacxx.so
HP-UX	ssolidacxx.sl
Linux	ssolidacxx.so

Both the static and dynamic library files contain a complete copy of the solidDB server, in library format. When you use a static library file (e.g. lib/solidac.a), you link your program directly to it, and both your code and the library code are written to the resulting executable file. If you link to a dynamic library file, the

code from the library is not included in the output file that contains your executable program. Instead, the code is loaded from the dynamic link library separately when your program runs.

Other than changing the size of your executable, there is very little difference between linking to the static library file vs. the dynamic library file. The total amount of code in memory at any one time is similar (assuming that you are executing a single client and a single server on your computer). Performance is also similar, although there is a slight amount of extra overhead if you use the dynamic library.

The main advantage of using the dynamic link library file is that you can save memory IF you execute more than one copy of the server in the same computer. For example, if you are doing development work on a single computer and you want to have both a advanced replication Master and advanced replication Replica on the computer at the same time, or you'd like to have a HotStandby Primary and a HotStandby Secondary at the same time, then you may prefer to use the dynamic library so that you don't have multiple copies of the linked library access library in memory at the same time.

On Microsoft Windows, the solidDB linked library access includes the additional file `lib/solidimpac.lib`. On Microsoft Windows, if you want to use a dynamic link library, you do not link directly to the `ssolidacxx.dll` dynamic link library itself; instead you link to `solidimpac.lib`, which is an import library. This links only a small amount of code to your client executable. At the time that your client program actually executes, the `ssolidacxx.dll` file will automatically be loaded by the Microsoft Windows operating system, and your client will be able to call the usual linked library access functions in that `.dll` file. The `.dll` file must be in your load path when you run the program that references that `.dll`.

Note: Using the dynamic link library file does not mean that you can have multiple "local" clients linked to the solidDB server. Even with the dynamic library approach, you are still limited to a single local client; all other clients must be remote clients, which means that they will communicate with the solidDB server by using TCP or some other network protocol, rather than the direct function calls available to the local client.

Makefile examples

Following are examples for providing the library name in Windows and Vxworks.

Microsoft Windows MakeFile example

For the Microsoft Windows makefile example below, the solidDB library name for the linked library access is used, `solidimpac.lib`.

```
# compiler
CC          = cl
# compiler flags
CFLAGS     = -I. -DSS_WINDOWS -DSS_WINNT
# linker flags and directives
SYSLIBS = libcmt.lib kernel32.lib advapi32.lib netapi32.lib wsock32.lib
user32.lib oldnames.lib gdi32.lib
LFLAGS    = ..\solidimpac.lib
OUTFILE   = -Fe

# MyApp building
all: myapp
```

```
myapp: myapp.c
      $(CC) $(CFLAGS) $(OUTFILE)myapp myapp.c /link$(LFLAGS)
      /NODEFAULTLIB:libc.lib
```

VxWorks MakeFile example

For the VxWorks makefile example below, the solidDB library name for the linked library access is used, `solidac.a`. Note that the example uses backslashes. If your makefile program does not support backslashes in pathnames, then change the backslashes to slashes.

```
CC      = ccppc
CFLAGS = -DSS_UNIX -DSS_VXW -I. -I..\..\include -I$(WIND_BASE)
        \target\h \
        -DCPU=PPC603 -DMV2600
LFLAGS = -nostartfiles -s -r ....\lib\solidac.a
OUTFILE = -o

# solidDB with AcceleratorLib samples building

all:   acsNet acsrv

acsNet: acsNet.c
        $(CC) $(CFLAGS) $(OUTFILE)acsNet acsNet.c $(LFLAGS)

acsrv:  acsrv.c
        $(CC) $(CFLAGS) $(OUTFILE)acsrv acsrv.c $(LFLAGS)
```

Establishing a local or remote connection to solidDB with the linked library access

Once an application is linked to the linked library access library, it can use ODBC API or SA API to establish a local or remote connection directly to the local server. An application can also establish remote connections to other solidDB servers, including others using the linked library access.

Establishing a local connection

When you establish a local connection, the client's calls to the server are direct function calls to the linked library access library; they do not go through the network.

In the ODBC API, to establish a connection to a local server (i.e. to the server that was linked to the application), the user application calls the `SQLConnect` function with the literal string "localserver". Note that for the local server connection you can also specify an empty source name "". You can also specify a local server name, but this will cause linked library access to use a "remote" connection (to go through the network rather than to use the direct function calls to the linked library access library).

The following ODBC API code examples connect directly to a local solidDB server with username `dba` and password `dba` :

```
rc = SQLConnect(hdbc, "localserver", (WORD)SQL_NTS, "dba", 3, "dba", 3);
```

or

```
rc = SQLConnect(hdbc, "", (WORD)SQL_NTS, "dba", 3, "dba", 3);
```

In the SA API, to establish a connection, the user application calls the `SaConnect` function with the literal string "localserver" (not the server name). Note that for the local server connection you can also specify an empty source name "". You can also

specify a local server name, but this will cause linked library access to use a "remote" connection (to go through the network rather than to use the direct function calls to the linked library access library).

The following SA API example code connects directly to a solidDB server with username dba and password dba :

```
SaConnectT* sc = SaConnect("localserver", "dba", "dba");
```

or

```
SaConnectT* sc = SaConnect("", "dba", "dba");
```

Establishing a remote connection

When you establish a remote connection, the client's calls to the server will go through the network rather than use the direct function calls to the linked library access library.

In the ODBC API, to establish a remote connection, the user application calls the SQLConnect function with the name of the remote server. The following ODBC API code example connects to a remote solidDB server with username dba and password dba. In this example, the network protocol that the client and server use is "tcp" (TCP/IP). The server is named "remote_server1" and the port that it listens on is 1313.

```
rc = SQLConnect(hdbc, "tcp remote_server1 1313",  
(SWORD)SQL_NTS, "dba", 3, "dba", 3);
```

In the SA API, to establish a remote connection, the user application calls the SaConnect function with the name of the remote server. In this example, the network protocol that the client and server use is "tcp" (TCP/IP). The server is named "remote_server1" and the port that it listens on is 1313.

```
SaConnectT* sc = SaConnect("tcp remote_server1 1313", "dba", "dba");
```

Starting and shutting down solidDB linked library access

You can start up, restart, and shut down the solidDB server from the following APIs:

- Explicitly, from local (linked) user application by calling the Control API function SSCStartServer to start solidDB and SSCStopServer to shut it down.

When you start a new solidDB server that does not already have a database, you must explicitly specify that solidDB create a new database with the function SSCStartServer() with the

```
-Uusername  
-Ppassword  
-Ccatalogname (the default database catalog name)
```

parameters. For details, read "Explicit startup with the Control API function SSCStartServer" on page 16.

- Implicitly, when connecting locally to solidDB for the first time, either using ODBC API function SQLConnect or SA API function SaConnect. In this case, shut down occurs when the last local connection disconnects from solidDB using either function SQLDisconnect or SaDisconnect.

When solidDB engine/server is started implicitly from the application, it checks if a database already exists in the solidDB directory. If a database file is found, solidDB will automatically open that database. If a database file is not found, then solidDB will give an error. (solidDB will not create a new database during

implicit startup. To create a new database, you must use an explicit startup function, such as `SSCStartServer`, and pass the appropriate parameters.)

For details, read “Implicit startup with ODBC API function call `SQLConnect`” on page 18 and “Implicit startup with SA API function call `SaConnect`” on page 20.

Note:

1. At server start up, recovery is performed if needed before control returns to the application. Therefore, if the server is successfully started, it is ready to serve application requests. For the duration of the application process, the server can be started or stopped as needed.
2. If you want to start a diskless server, you must start the server with Control API function `SSCStartDisklessServer`.

Explicit startup with the Control API function `SSCStartServer`

To start solidDB explicitly, have the user application call the following Control API function:

```
SSCStartServer (int argc, char* argv [ ],
               SscServerT* h, SscStateT runflags)
```

where parameters are:

Table 5. SSCStartServer parameters

Parameter	Description
argc	The number of command line arguments.
argv	Array of command line arguments that are used during the function call. The argument <code>argv[0]</code> is reserved for the path and filename of the user application only and must be present. For valid options, see <code>SSCStartServer</code> options below.
h	Each server has a "handle" (a pointer to a data structure) that identifies that server and indicates where information about that server is stored. This handle is required when referencing the server with other Control API functions. The handle of the server is provided to you when you call the <code>SSCStartServer</code> function. To get the handle of the server, you create a variable that is of type pointer-to-server-handle (i.e. you create an <code>SSCServerT *</code> , which is a pointer to a handle — essentially a pointer to a pointer) and you pass that when you call <code>SSCStartServer</code> . If the server is created successfully, then the <code>SSCStartServer</code> function will write the handle (pointer) of the new server into the variable whose address you passed.
runflags	The options for this parameter are <code>SSC_STATE_OPEN</code> (remote connections are allowed) and <code>SSC_STATE_PREFETCH</code> (the server performs a prefetch if needed). Prefetch refers to the memory and/or disk cache that provides read-ahead capability for table content. See below for a runflags parameter entry: <code>runflags = SSC_STATE_OPEN SSC_STATE_PREFETCH;</code>

When you start the server for the first time, solidDB creates a new database only if you have specified the database administrator's username, password, and a name for the default database catalog. For details on the database catalog, read the

section "Managing Database Objects" in chapter "Using solidDB SQL for Data Administration" in *IBM solidDB Administrator Guide*.

For example:

```
SscServerT h; char* argv[4];
argv[0] = "appname"; /* path and filename of the user app. */
argv[1] = "-UDBA"; /* user name */
argv[2] = "-PDBA"; /* user's password */
argv[3] = "-CDBA"; /* catalog name */
/* Start the server */
rc = SSCStartServer(argc, argv, &h, run_flags);
```

If you start the server without an existing database and do not specify a database catalog name, solidDB returns an error that the database is not found.

Note: If you already have an existing database, you do not need to specify the username and password, or the catalog name.

By default, the database will be created as one file (with the default name, `solid.db`, or the name you specified in the `solid.ini` file) in the solidDB working directory, where the current working directory is located. An empty database containing only the system tables and views uses approximately 850 KB of disk space. The time it takes to create the database depends on the hardware platform you are using.

After the database has been created, solidDB starts listening to the network for remote client connection requests.

SSCStartServer argv parameter options

Following are the command line options for the argv parameter. Note that all options are case sensitive.

Table 6. SSCStartServer argv options

Option	Description
-c <i>dir</i>	Changes the working directory.
-m	Monitors users' messages and SQL statements.
-n <i>name</i>	Sets the server name.
-U <i>username</i>	Specifies the user name of the administrator for the database being created. The user name is case insensitive and it requires at least two characters. For user name, the maximum number of characters is 80. A user name must begin with a letter or an underscore. Use lower case letters from a to z, upper case letters from A to Z, the underscore character '_', and digits from 0 to 9. CAUTION: You must remember your user name to be able to connect to solidDB. There are no default user names; the user name you enter when creating the database is the only user name available for connecting to the new database.

Table 6. SSCStartServer argv options (continued)

Option	Description
-P <i>password</i>	Specifies the password of the administrator for the database being created. The password is case insensitive and it requires at least three characters. Passwords can begin with a letter, an underscore, or a digit. Use lower case letters from a to z, upper case letters from A to Z, the underscore character '_', and digits from 0 to 9.
-C <i>catalogname</i>	Specifies the name of the default catalog of the database, which is required if you are starting the server for the first time. For details on catalogs, read the section "Managing Database Objects" in chapter "Using solidDB SQL for Data Management" in <i>IBM solidDB Administrator Guide</i> .
-xautoconvert	Converts database format to current version and starts server process.
-xforcerecovery	Does a forced roll-forward recovery.
-xignoreerrors	Ignores index errors.
-xtestblocks	Tests database blocks.
-xtestindex	Tests database index.

Starting up SSCStartServer

Start up SSCStartServer with the servername, the catalog name, and the administrator's username and password:

```
SscStateT runflags = SSC_STATE_OPEN; SscServerT h; char* argv[5];
argv[0] = "appname"; /* path and filename of the user app. */
argv[1] = "-nsolid1"; argv[2] = "-UDBA" argv[3] = "-PDDBA";
argv[4] = "-CDBA"; /* Start the server */ rc =
SSCStartServer(argc, argv, &h, run_flags);
```

Note: If you already have an existing database, you do not need to specify the username and password, or the catalog name.

Shut down with SSCStopServer

If the server is started by SSCStartServer, then it must be shut down with the following function call in the embedded application:

```
SSCStopServer()
```

For example:

```
/* Stop the server */
SSCStopServer (h, TRUE);
```

Implicit startup with ODBC API function call SQLConnect

When function SQLConnect is called for the first time, the server is implicitly started. The server is shut down implicitly when the user application calls function SQLDisconnect and this is the last open local connection. Note that the server will shut down regardless of currently existing remote connections.

Note: When you start the server for the first time, you must create a solidDB database by using function `SSCStartServer()` and specifying the default database catalog, along with the administrator's username and password. For a description and example, read "Explicit startup with the Control API function `SSCStartServer`" on page 16.

Following is an example of implicit start up and shut down with `SQLConnect` and `SQLDisconnect`:

```

/* Connection #1 */
rc = SQLConnect (hdbc1, "", SQL_NTS, "dba", SQL_NTS, "dba",
SQL_NTS); //Server Started Here
... odbc calls

/* Disconnect #1 */
SQLDisconnect (hdbc1); //Server Shut Down Here

/* Connection #2 */
rc = SQLConnect (hdbc2, "", SQL_NTS, "dba", SQL_NTS, "dba",
SQL_NTS); //Server Started Here
... odbc calls

/* Disconnect #2 */
SQLDisconnect (hdbc2); //Server Shut Down Here

```

OR

```

/* Connection #1*/
rc = SQLConnect (hdbc1, "", SQL_NTS, "dba",
SQL_NTS, "dba", SQL_NTS); // Server Started Here

/* Connection #2*/
rc = SQLConnect (hdbc2, "", SQL_NTS, "dba", SQL_NTS, "dba", SQL_NTS);

... odbc calls

/* Disconnect #1 */
SQLDisconnect (hdbc1);
/* Disconnect #2 */
SQLDisconnect (hdbc2); // Server Shut Down Here

```

Note: If the server is started with an `SSCStartServer` function call, then `SQLDisconnect` does not do implicit shut down. The server must be shut down explicitly, either by `SSCStopServer`, `ADMIN COMMAND 'shutdown'`, or other explicit shutdown methods.

```

SscStateT runflags = SSC_STATE_OPEN;
SscServerT server;
SQLHDBC hdbc;
SQLHENV henv;
SQLHSTMT hstmt;

```

```

/* Start the server */
SSCStartServer (argc, argv, &server, runflags); // Server Started Here

/* Alloc environment */
rc = SQLAAllocEnv (&henv);

/* Connect to the database */
rc = SQLAAllocConnect (henv, &hdbc);
rc = SQLConnect (hdbc, "", SQL_NTS, "dba", SQL_NTS, "dba", SQL_NTS);

/* Delete all the rows from table foo */
rc = SQLAAllocStmt (hdbc, &hstmt);
rc = SQLExecDirect (hstmt, (SQLCHAR *) "DELETE FROM FOO", SQL_NTS);

```

```

/* Commit */
rc = SQLTransact (henv, hdbc, SQL_COMMIT);
rc = SQLFreeStmt (hstmt, SQL_DROP);

/* Disconnect */
SQLDisconnect (hdbc);
SQLFreeConnect (hdbc);

/* Free the environment */
SQLFreeEnv(henv);

/* Stop the server */
SSCStopServer (server, TRUE); // Server Shut Down Here

```

Implicit startup with SA API function call SaConnect

When function SaConnect is called for the first time, the server is implicitly started. The server is shut down implicitly when the user application calls function SaDisconnect and there are no more subsequent connections.

Note: When you start the server for the first time, you must create a solidDB database by using function SSCStartServer() and specifying the default database catalog, along with the username and password. For a description and example, read “Explicit startup with the Control API function SSCStartServer” on page 16.

Following is an example of implicit start up and shut down with SaConnect and SaDisconnect:

```

/* Open Connection */
SaConnect(...);

Server Started Here
... sa calls

/* Close Connection */
SaDisconnect(...);

Server Shut Down Here

```

Note: If the server is started with an SSCStartServer function call, then it must be shut down only with an SSCStopServer function call.

Shutting down solidDB linked library access

From solidDB client interfaces and even from another remote solidDB connection, you can shut down the solidDB server as long as you have SYS_ADMIN_ROLE privileges.

Programmatically, you can perform the shut down from an application such as solidDB SQL Editor (solsql), or solidDB Remote Control¹.

To do this, perform the following steps:

1. To prevent new connections to solidDB, close the database(s) by entering the following command:
ADMIN COMMAND 'close'
2. Exit all solidDB users by entering the following command:
ADMIN COMMAND 'throwout all'
3. Stop solidDB by entering the following command:

1. When using solidDB Remote Control for steps 1-3, you enter the command name only without quotes (for example, close).

ADMIN COMMAND 'shutdown'

All the shutdown mechanisms will start the same routine, which writes all buffered data to the database file, frees cache memory, and finally terminates the server program. Shutting down a server may take awhile since the server must write all buffered data from main memory to the disk.

Note: You can use explicit methods (e.g. `SSCStopServer`) to shut down a server that was started with implicit methods (e.g. `SQLConnect`). The converse is not true; for example, you cannot use `SQLDisconnect` to stop a server that was started with `SSCStartServer`.

Implicit start configuration parameter

`solidDB` implicitly starts up the server only when a local connection is established. In the Accelerator section of the `solid.ini` configuration file, the parameter `ImplicitStart`, by default, is set to `Yes`. This default setting starts the server automatically when you use the function `SQLConnect` which is required for any ODBC connection. The function `SaConnect` behaves similarly. When it is called for the first time, the server is implicitly started.

3 Description of control API

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

Retrieving task information

To retrieve a list of all active tasks, use the `SSCGetActiveTaskClass` function. To retrieve a list of all suspended tasks, use the `SSCGetSuspendedTaskClass` function. To get the priority of a task class, use the `SSCGetTaskClassPrio` function.

Notifying functions of a special event

The linked library access provides fine tuning of priority tasks. You can use the `SSCSetNotifier()` function to establish that solidDB calls a specified user-defined function whenever a special event occurs. Special events that the function detects are:

- solidDB server shutdown
- Bonsai merge from the index to the storage tree
- Bonsai merge interval maximum
- Backup or checkpoint request
- Idle server state
- Netcopy request (which is a request to send a network copy of the Primary database to the Secondary server) received from the Primary server.
- Completion of a netcopy request, which occurs when the server is started up with the new database received through the network copy (netcopy).

Obtaining solidDB status and server information

You can use the function `SSCGetStatusNum` to view current status information of the solidDB database server. The following information is displayed:

- Number of rows that are not merged from the Bonsai Tree to the Storage Tree

The `SSCGetServerHandle` function returns the solidDB server handle if the server is running.

You can also use the function `SSCIsRunning` to verify if the server is running and the function `SSCIsThisLocalServer` to verify whether an application is linked to the local linked library access server library (for example, `ssolidacxx.dll` for Windows platforms) or a "dummy" server library (for example, `solidctrlstub.lib` for Windows platforms) that are used to test remote applications that are using Control API.

Summary of control API functions

The following is a brief summary of Control API functions and where the function is described in the Control API Function Reference section.

Table 7. Summary of control API functions

Function	Description	For more details, see
SSCStartServer	Starts a solidDB linked library access server.	See "SSCStartServer" on page 36.
SSCStartDisklessServer	Starts a solidDB linked library access diskless server.	See "SSCStartDisklessServer" on page 34.
SSCSetState	Sets the state of a solidDB server (for example, SSC_STATE_OPEN indicates if subsequent connections are allowed). Setting the state to ~SSC_STATE_OPEN will block local, as well as remote, connections.	See "SSCSetState" on page 33.
SSCRegisterThread	Registers an linked library access application thread for the server. Registration is required in every thread in the user application before any Accelerator API function can be called.	See "SSCRegisterThread" on page 29.
SSCUnregisterThread	Unregisters an linked library access application thread for the server. Registration removal is required in every thread that is registered before terminating.	See "SSCUnregisterThread" on page 39.
SSCStopServer	Stops solidDB server.	See "SSCStopServer" on page 38.
SSCSetNotifier	Specifies a user-defined function which solidDB calls at a specified event, such as merge, backup, shutdown, etc.	See "SSCSetNotifier" on page 30.
SSCIsRunning	Returns non-zero if the server is running.	See "SSCIsRunning" on page 28.
SSCIsThisLocalServer	Indicates whether the application is linked to the solidDB server with the linked library access or the "dummy" (solidctrlstub) library to test solidDB remote applications using the linked library access's Control API.	See "SSCIsThisLocalServer" on page 29.
SSCGetServerHandle	Returns the solidDB server handle if the server is running.	See "SSCGetServerHandle" on page 27.
SSCGetStatusNum	Gets solidDB status information.	See "SSCGetStatusNum" on page 27.

Control API and equivalent ADMIN COMMANDS

Control API functions have equivalent solidDB SQL extension ADMIN COMMANDS. You can execute these commands from both remote and local sites through solidDB tools, such as solidDB Remote Control (solcon), and solidDB SQL Editor (solsql).

Refer to "Linked library access parameters" on page 53 for details on Control API equivalent ADMIN Commands.

Control API reference

The following pages describe each Control API function in alphabetic order. Each description includes the purpose, synopsis, parameters, return value, and comments.

Function synopsis

The declaration synopsis for the function is:

```
ReturnType SSC_CALL function(modifier parameter[...]);
```

The ReturnType varies, but is usually a value that indicates success or failure of the call. Return values are described in more detail later in this section.

SSC_CALL is required for portability. SSC_CALL specifies the calling convention of the function. It is defined appropriately for each platform in the `sscapi.h` file.

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 below), 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 Control API 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 8. Control API 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
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.
take	Applies only to a pointer parameter. It means that the parameter value is taken by the function. The caller cannot reference 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.

Table 8. Control API parameter usage types (continued)

Usage Type	Meaning
hold	<p>Applies only to a pointer parameter. It means that the function holds the parameter value even after the function call. The caller can continue to reference the parameter value after the function call and is responsible for releasing the parameter.</p> <p>Attention:</p> <p>Because this parameter is shared by the user and the server, you must not release it until the server is finished with it. In general, you can free the held object after you free the object that is holding it. For example:</p> <pre>conn = SaConnect("", "dba", "dba"); /* Connection is held until cursor is freed */ scur = SaCursorCreate(conn, "mytable"); ... SaCursorFree(scur); /* After we free the cursor, it is safe to free */ /* the connection (or, as in this case, call a */ /* server function that frees the connection). */ SaDisconnect(conn);</pre>

Return value

Each function description indicates if the function returns a value and the type of value that is returned.

SscTaskSetT

When functions return a value of type SscTaskSetT, this definition is used as a bit mask. SscTaskSetT is defined in sscapi.h with the following possible values:

```
SSC_TASK_NONE
SSC_TASK_CHECKPOINT
SSC_TASK_BACKUP
SSC_TASK_MERGE
SSC_TASK_LOCALUSERS
SSC_TASK_REMOTEUSERS
SSC_TASK_SYNC_HISTCLEAN
SSC_TASK_SYNC_MESSAGE
SSC_TASK_HOTSTANDBY
SSC_TASK_HOTSTANDBY_CATCHUP
SSC_TASK_ALL (all of the above tasks)
```

Note that the HotStandby "netcopy" and HotStandby "copy" operations are performed by the task "SSC_TASK_BACKUP"; there is no separate task "SSC_TASK_NETCOPY".

Control API error codes and messages

Control API functions may return the following error codes and messages:

Table 9. Error codes and messages for control API functions

Error Code/Message	Description
SSC_SUCCESS	Operation is successful.

Table 9. Error codes and messages for control API functions (continued)

Error Code/Message	Description
SSC_ERROR	Generic error.
SSC_ABORT	Operation aborted.
SSC_FINISHED	SSCAdvanceTasks returns this message if all tasks are executed.
SSC_CONT	SSCAdvanceTasks returns this message if there are still more tasks to execute.
SSC_CONNECTIONS_EXIST	There are open connections.
SSC_UNFINISHED_TASKS	There are unfinished tasks.
SSC_INFO_SERVER_RUNNING	The server is already running.
SSC_INVALID_HANDLE	Invalid local server handle given. This server does not match the one started through SSCStartServer.
SSC_INVALID_LICENSE	No license or invalid license file found.
SSC_NODATABASEFILE	No database file found.
SSC_SERVER_NOTRUNNING	The server is not running.
SSC_SERVER_INNETCOPYMODE	The server is in netcopy mode (applies only with High Availability/HotStandby).

These constants (SSC_SUCCESS, etc.) are defined in the sscapi.h file.

SSCGetServerHandle

SSCGetServerHandle returns the solidDB server handle if the server is running.

Synopsis

```
SscServerT SSC_CALL SSCGetServerHandle(void)
```

Comments

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

Return value

- NULL if the server is not running.
- The server handle if the server is running.

SSCGetStatusNum

SSCGetStatusNum gets the status information of solidDB.

Synopsis

```
SscRetT SSC_CALL SSCGetStatusNum(SscServerT h, SscStatusT stat,  
    long * num)
```

The SSCGetStatusNum function accepts the following parameters:

Table 10. SSCGetStatusNum parameters

Parameters	Usage Type	Description
h	in, use	Handle to server.
stat	in	Specifies the status identifier for retrieval:
num	out	If the function was successful, then when it returns this parameter's value will be set to either the number of writes not merged, or the number of server threads, depending upon which information was requested.

Comments

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

If you call SSCGetStatusNum and pass it an unrecognized value for the stat parameter, then the function will return SSC_SUCCESS.

Return value

- SSC_SUCCESS - Operation is successful. This value is also returned if you pass an invalid value for the stat parameter.
- SSC_ERROR - Operation failed.
- SSC_SERVER_INNETCOPYMODE - The server is in netcopy mode (HotStandby only)
- SSC_SERVER_NOTRUNNING - The server is not running.

SSCIsRunning

SSCIsRunning returns non-zero if the server is running.

Synopsis

```
int SSC_CALL SSCIsRunning(SscServerT h)
```

The SSCIsRunning function accepts the following parameters:

Table 11. SSCIsRunning parameters

Parameters	Usage Type	Description
h	in, use	Handle to server

Return value

- 0 - The server is not running.
- nonzero - The server is running.

Comments

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

SSCIsThisLocalServer

SSCIsThisLocalServer indicates whether the application is linked to a solidDB server or the "dummy" (solidctrlstub) library. The solidctrlstub library allows developers to test solidDB remote applications using Control API without linking the linked library access library and modifying the source code.

Synopsis

```
int SSC_CALL SSCIsThisLocalServer(void)
```

Return value

- 0 - The application is not linked to the solidDB server.
- 1 - The application is linked to the solidDB server.

Comments

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

SSCRegisterThread

SSCRegisterThread registers a solidDB application thread for the server. Every thread that uses Control API, ODBC API, or SA API must be registered. The SSCRegisterThread function must be called by the thread before any other linked library access API function can be used.

If the application has only one (main) thread, that is, if the application creates no threads itself, then registration is not required.

Before a thread terminates, it must unregister itself by calling the function SSCUnregisterThread.

Synopsis

```
SscRetT SSC_CALL SSCRegisterThread(SscServerT h)
```

The SCCRegisterThread function accepts the following parameters:

Table 12. SCCRegisterThread parameters

Parameters	Usage Type	Description
h	In, Use	Handle to server

Return value

- SSC_SUCCESS
- SSC_INVALID_HANDLE

Comments

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

See also

SSCUnregisterThread

SSCSetNotifier

SSCSetNotifier sets the callback functions that an linked library access server calls when it is started or stopped. The function does not have a corresponding ADMIN COMMAND.

Synopsis

```
SscRetT SSC_CALL SSCSetNotifier(SscServerT h, SscNotFunT what,  
                                notify_fun handler, void* userdata  
                                )
```

The SSCSetNotifier function accepts the following parameters:

Table 13. SSCSetNotifier parameters

Parameters	Usage Type	Description
<i>h</i>	in	Handle to the server.

Table 13. *SSCSetNotifier* parameters (continued)

Parameters	Usage Type	Description
<i>what</i>	in	<p>Specifies the event for notification. Options are:</p> <ul style="list-style-type: none"> • <code>SSC_NOTIFY_EMERGENCY_EXIT</code> This function is called if a server crashes after it has been activated with <code>SSCStartServer()</code>. The notifier call <code>SSCSetNotifier()</code> has to be issued before <code>SSCStartServer()</code> • <code>SSC_NOTIFY_SHUTDOWN</code> Function is called at shutdown. • <code>SSC_NOTIFY_SHUTDOWN_REQUEST</code> Function is called when the server receives the shutdown request and may shut down if the user-defined function accepts the request. You can refuse the shut down by returning <code>SSC_ABORT</code> from the notified function, or proceed with the request by returning <code>SSC_CONT</code>. • <code>SSC_NOTIFY_ROWSTOMERGE</code> Function is called when there is data in the bonsai index tree that needs to be merged to the storage server. • <code>SSC_NOTIFY_MERGE_REQUEST</code> Function is called when the <code>MergeInterval</code> parameter setting in the <code>solid.ini</code> configuration file is exceeded and the merge has to start. • <code>SSC_NOTIFY_BACKUP_REQUEST</code> Function is called when a backup is requested. You can refuse the backup by returning <code>SSC_ABORT</code> from the notified function. • <code>SSC_NOTIFY_CHECKPOINT_REQUEST</code> Function is called when a checkpoint is requested. You can refuse the checkpoint by returning <code>SSC_ABORT</code> from the notified function. • <code>SSC_NOTIFY_IDLE</code> Function is called when the server switches to the idle state. • <code>SSC_NOTIFY_NETCOPY_REQUEST</code> This callback function applies to the HotStandby component only. The function is called when a netcopy request (which is a request to send a network copy of the Primary database to the Secondary server) is received from the Primary server. For details on the netcopy command, refer to <i>IBM solidDB High Availability User Guide</i>. • <code>SSC_NOTIFY_NETCOPY_FINISHED</code> This callback function applies to the HotStandby component only. The function is called when a netcopy request is finished. When finished, the server is started up with the new database received through the network copy (netcopy) and <code>SSC_NOTIFY_FINISHED</code> is called to inform the application that the server is again available.
<i>notify_fun_handler</i>	in, hold	User function to call.
<i>userdata</i>	in, hold	<p>User data to be passed to the notify function.</p> <p>Be sure to read the warning on releasing a parameter of usage type <i>hold</i> under "Parameter description" on page 25.</p>

Return value

- `SSC_SUCCESS` - Request from the server accepted.

HotStandby only:

If `SSC_NOTIFY_NETCOPY_FINISHED` returns `SSC_SUCCESS`, then all other application connections are terminated and the server is set to "netcopy listening mode". In this mode the server accepts the connection from the Primary server

and the only possible operation for the Secondary server is to receive the data from the hotstandby netcopy command. For more details on "netcopy listening mode", read *IBM solidDB High Availability User Guide*. (Note that in the past, "netcopy listening mode" was also called "backup listening mode".)

- SSC_ABORT - Request from the server denied.
HotStandby only:
If the SSC_NOTIFY_NETCOPY_REQUEST returns SSC_ABORT, then the netcopy is not started and an error code (SRV_ERR_OPERATIONREFUSED) is returned to the Primary server.
- SSC_INNETCOPYMODE - The server is in netcopy mode (HotStandby only).
- SSC_SERVER_NOTRUNNING - The server is not running.

Comments

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

Releasing a parameter of usage type *hold* should be done with caution. Read the warning for hold "Parameter description" on page 25.

The user-defined notifier function should not call any SA, SSC, or ODBC function.

When creating a user-defined notifier function, you must conform to the following prototype:

```
int SSC_CALL mynotifyfun(SscServerT h, SscNotFunT what, void* userdata);
```

Once you have used SSC_CALL to explicitly define the convention for your user function, then you use the SSCSetNotifier function to register the function so that it is called during the specified event; for example:

```
SscRetT SSCSetNotifier(h, SSC_NOTIFY_IDLE, mynotifyfun, NULL);
```

Example

Calling a function upon shutdown

Assume a user creates the function `user_own_shutdownrequest`, which is called every time a shut down is requested:

```
int SSC_CALL user_own_shutdownrequest(SscServerT h, SscNotFunT what, void
*userdata);
{
  if (shutdown not needed) {
    return SSC_ABORT;
  }
  return SSC_CONT; /*Proceed with shutdown*/
}
```

The SSCSetNotifier function can then be called as follows to specify that `user_own_shutdownrequest` gets called before the server is shut down.

```
SSCSetNotifier(handle, SSC_NOTIFY_SHUTDOWN, user_own_shutdownrequest, NULL);
```

Note:

If function `user_own_shutdownrequest` returns SSC_ABORT, the shut down is not allowed and if the function returns SSC_CONT, the shut down can proceed.

SSCSetState

SSCSetState sets the state of an linked library access server. This allows you to control whether the server accepts subsequent connections, and whether the server uses prefetch.

If the server is set to "open", then the server will accept connections. If the server is set to "closed", then it will not accept any further connections (this applies to both local connections and remote connections); however, any connections that have already been made are allowed to continue.

Turning on prefetch tells the server to "read ahead" to fetch data that is likely to be referenced soon. Prefetch requires more memory or disk cache space. When prefetch is on, performance is generally higher. When prefetch is off, less memory is required. Turning on prefetch is most useful if you have queries that involve large sequential scans of the server. For example, if you use reports or aggregate functions to get values for the entire database (or large portions of it), then prefetch may help. Prefetch is generally not useful if all your queries involve only one or a few records. Because prefetch uses up memory, prefetch may actually reduce performance in systems with little available memory.

The following guidelines may help you decide when to use prefetch.

DO use prefetch when: you have a lot of available memory (or disk cache space) and your queries require large sequential scans.

DO NOT use prefetch when: you have little available memory and your queries generally read unrelated records one at a time.

Synopsis

```
SscRetT SSC_CALL SSCSetState(SscServerT h,SscStateT runflags)
```

The SSCSetState function accepts the following parameters:

Table 14. SSCSetState parameters

Parameter	Usage Type	Description
<i>h</i>	in, use	Handle to the server.
<i>runflags</i>	in	Options can be a combination of the flags SSC_STATE_OPEN, which means new remote connections are allowed and SSC_STATE_PREFETCH, which means the user allows the server to do a prefetch if needed. Following is an example of the possible combinations: <ul style="list-style-type: none">• set server open: state = state SSC_STATE_OPEN;• set server closed: state = state & ~SSC_STATE_OPEN;• set prefetch on: state = state SSC_STATE_PREFETCH;• set prefetch off: state = state & ~SSC_STATE_PREFETCH;

Return value

- SSC_SUCCESS - Operation is successful.
- SSC_ERROR - Operation failed.
- SSC_SERVER_INNETCOPYMODE - The server is in netcopy mode (HotStandby only).
- SSC_SERVER_NOTRUNNING - The server is not running.

Comments

This function has a corresponding solidDB SQL extension ADMIN COMMAND. The command is:

```
ADMIN COMMAND 'close';
```

SSCStartDisklessServer

SSCStartDisklessServer starts a diskless server using the linked library access.

Synopsis

```
SscRetT SSC_CALL SSCStartDisklessServer (int argc, char* argv[ ],  
    SscServerT * h, SscStateT runflags, char* lic_string, char* ini_string);
```

The SSCStartDisklessServer function accepts the following parameters:

Table 15. SSCStartDisklessServer parameters

Parameters	Usage Type	Description
<i>argc</i>	in	The number of command line arguments.
<i>argv</i>	in, use	Array of command line arguments that are used during the function call. The argument argv[0] is reserved only for the path and filename of the user application and must be present. For a list of valid arguments, refer to the SSCStartDisklessServer parameter options listed below.
<i>h</i>	out	Returns a handle to the started server. This handle is needed when referencing the server with other Control API functions.
<i>runflags</i>	in	The only option for this parameter is: SSC_STATE_OPEN - Remote connections are allowed. runflags = SSC_STATE_OPEN
<i>lic_string</i>	in	Specifies the string containing the solidDB license file.
<i>ini_string</i>	in	Specifies the string containing the solidDB configuration file.

SSCStartDisklessServer parameter options

Following are the command line options for the *argv* parameter.

Table 16. Command line options for the *argv* parameter

Option	Description
-h	Displays help.
-n <i>name</i>	Sets server name.
-U <i>username</i>	Specifies the username for the data. The username is case insensitive. The username requires at least two characters. For username, the maximum number of characters is 80. A user name must begin with a letter or an underscore. Use lower case letters from a to z, upper case letters from A to Z and the underscore character '_', and digits from 0 to 9. Note: You must remember your username to be able to connect to solidDB. There are no default usernames ; the username you enter when creating the database is the only username available for connecting to the new database.
-P <i>password</i>	Specifies the given password for the data. The password is case insensitive. The password requires at least three characters. Passwords can begin with a letter, underscore, or a number. Use lower case letters from a to z, upper case letters from A to Z and the underscore character '_', and digits from 0 to 9.
-C <i>catalogname</i>	Specifies the catalog name for the data, required if you are starting the server for the first time. When specifying this parameter, be sure to use uppercase C. For details on catalogs, read the section "Managing Database Objects" in chapter "Using solidDB SQL for Data Administration" in <i>IBM solidDB SQL Guide</i> .
-x ignoreerrors	Ignores index errors.

Return value

- SSC_SUCCESS - The server is started.
- SSC_ERROR - The server failed to start.
- SSC_SERVER_INNETCOPYMODE - The server is netcopy mode (HotStandby only).
- SSC_INFO_SERVER_RUNNING - The server is already running.
- SSC_INVALID_HANDLE - Invalid local server handle given.
- SSC_INVALID_LICENSE - No license or invalid license file found.

Comments

By default, the state is set to SSC_STATE_OPEN.

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

Example

SSCStartDisklessServer

```
SscStateT runflags = SSC_STATE_OPEN;
SscServerT h;
char* argv[4]; /* pointers to four parameter strings */
int argc = 4;
char* lic = get_lic(); /* get the license */
char* ini = get_ini(); /* get the solid.ini */
SscRetT rc;
argv[0] = "appname"; /* path and filename of the user app. */
argv[1] = "-Udba"; /* user name */
argv[2] = "-Pdba"; /* user's password */
argv[3] = "-Cdba"; /* catalog name */
/* Start the diskless server */
rc = SSCStartDisklessServer(argc, argv, &h, runflags, lic, ini);
```

Note:

In the example, `get_ini()` and `get_lic()` are functions that a user must write. Each must return a string that contains the `solid.ini` file text or the `solid.lic` license file.

If you do not specify a catalog name, `solidDB` returns an error.

See also

`SSCStopServer`

See also 4, "Using the diskless capability," on page 41.

SSCStartServer

`SSCStartServer` starts the linked library access. In multi-thread environments, the server runs in a separate thread(s) from the client. For the duration of the application, the application can start or stop the server subroutines as needed.

Note that the third parameter is an "out" parameter. If the server is started successfully, then the `SSCStartServer` routine will set this parameter to point to the handle for this server.

Note:

If you are starting a diskless server, you must start the server with Control API function `SSCStartDisklessServer`. Read "SSCStartDisklessServer" on page 34.

Synopsis

```
SscRetT SSC_CALL SSCStartServer(int argc, char* argv[], SscServerT* h
    SscStateT runflags)
```

The `SSCStartServer` function accepts the following parameters:

Table 17. *SSCStartServer* parameters

Parameters	Usage Type	Description
<i>argc</i>	in	Number of command line arguments.

Table 17. *SSCStartServer* parameters (continued)

Parameters	Usage Type	Description
<i>argv</i>	in, use	Array of command line arguments. For a list of valid arguments, refer to “SSCStartServer” on page 36.
<i>h</i>	out	Returns a handle to the started server. This handle is needed when referencing the server with other Control API functions.
<i>runflags</i>	in	Options can be one or both of the following: <ul style="list-style-type: none"> • SSC_STATE_OPEN - Remote connections are allowed. • SSC_STATE_PREFETCH - Server will do a prefetch if needed. For example: runflags = SSC_STATE_OPEN & SSC_STATE_PREFETCH

Return value

- SSC_SUCCESS - The server started.
- SSC_ERROR - The server failed to start.
- SSC_ABORT
- SSC_BROKENNETCOPY - Database corrupted because of incomplete netcopy.
- SSC_FINISHED
- SSC_CONT
- SSC_CONNECTIONS_EXIST
- SSC_UNFINISHED_TASKS
- SSC_INVALID_HANDLE - Invalid local server handle given.
- SSC_INVALID_LICENSE - No license or invalid license file found.
- SSC_NODATABASEFILE - No database file found.
- SSC_SERVER_NOTRUNNING
- SSC_INFO_SERVER_RUNNING - The server is already running.
- SSC_SERVER_INNETCOPYMODE - The server is in netcopy mode (HotStandby only).
- SSC_DBOPENFAIL - Failed to open database.
- SSC_DBCONNFAIL - Failed to connect to database.
- SSC_DBTESTFAIL - Database test failed.
- SSC_DBFIXFAIL - Database fix failed.
- SSC_MUSTCONVERT - Database must be converted.
- SSC_DBEXIST - Database exists.
- SSC_DBNOTCREATED - Database not created.
- SSC_DBCREATEFAIL - Database create failed.
- SSC_COMINITFAIL - Communication init failed.
- SSC_COMLISTENFAIL - Communication listen failed.
- SSC_SERVICEFAIL - Service operation failed.

- SSC_ILLARGUMENT - Illegal command line argument.
- SSC_CHDIRFAIL - Failed to change directory.
- SSC_INFILEOPENFAIL - Input file open failed.
- SSC_OUTFILEOPENFAIL - Output file open failed.
- SSC_SRVCONNFAIL - Server connect failed.
- SSC_INITERROR - Operation init failed.
- SSC_CORRUPTED_DBFILE - Assert or other fatal error.
- SSC_CORRUPTED_LOGFILE - Assert or other fatal error.

Comments

By default, the state is set to SSC_STATE_OPEN.

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

When you start a new solidDB server, you must explicitly specify that solidDB create a new database with the function SSCStartServer() with the `-U username -P password -C catalogname` (the default database catalog name) parameters. For details, read “Explicit startup with the Control API function SSCStartServer” on page 16.

If you are restarting a database server (i.e. a database already exists in the directory), then SSCStartServer will use the existing database.

The SSCStartServer function may spawn multiple threads to run the server tasks. The server tasks include processing local and remote client requests, as well as running various background tasks, such as checkpoints, merges, etc.

See also

SSCStopServer

SSCStopServer

SSCStopServer stops an linked library access server.

Note that you can use explicit methods (e.g. SSCStopServer) to shut down a server that was started with implicit methods (e.g. SQLConnect). The converse is not true; for example, you cannot use SQLDisconnect to stop a server that was started with SSCStartServer.

An application is not limited to starting and stopping the server once each time that the application is run. After the server has been stopped, the application can re-start the server by using SSCStartServer.

Synopsis

```
SscRetT SSC_CALL SSCStopServer(SscServerT h, bool force)
```

The SSCStopServer function accepts the following parameters:

Table 18. SSCStopServer parameters

Parameter	Usage Type	Description
<i>h</i>	in, use	Handle to server

Table 18. *SSCStopServer* parameters (continued)

Parameter	Usage Type	Description
<i>force</i>	in	Options are: <ul style="list-style-type: none"> • TRUE - stop server in all cases. • FALSE - stop server if there are no open connections. Otherwise, stop fails.

Return value

- SSC_SUCCESS - The server is stopped.
- SSC_CONNECTIONS_EXIT - There are open connections.
- SSC_UNFINISHED_TASKS - Tasks that are executing.
- SSC_ABORT
- SSC_ERROR

Comments

Remote users can stop solidDB by using ADMIN COMMAND 'shutdown'. Refer to "Linked library access parameters" on page 53 for details.

The FALSE option does not permit shut down if there are open connections to the database or existing users. This option is equivalent to solidDB SQL extension ADMIN COMMAND 'shutdown'.

The SSCSetState() with the &~SSC_STATE_OPEN option prevents new connections to solidDB.

See also

SSCStartServer

SSCSetState

SSCUnregisterThread

SSCUnregisterThread unregisters a solidDB application thread for the server. The SSCUnregisterThread function must be called by every thread that has registered itself with the function SSCRegisterThread. The function is called before the thread terminates.

Synopsis

```
SscRetT SSC_CALL SSCUnregisterThread(SscServerT h)
```

The SSCUnregisterThread function accepts the following parameters:

Table 19. *SSCUnregisterThread* parameters

Parameter	Usage Type	Description
<i>h</i>	in, use	Handle to server

Return value

- SSC_SUCCESS
- SSC_INVALID_HANDLE

Comments

SSC_CALL is required to explicitly define the calling convention of your user function. It is defined in the sscapi.h file appropriately for each platform.

This function has no corresponding solidDB SQL extension ADMIN COMMAND.

See also

SSCRegisterThread

4 Using the diskless capability

solidDB linked library access allows you to create a database engine that runs without any disk storage space. This is useful in embedded systems that do not have hard disks, such as line cards in a network router or switch.

There are two main ways to run a diskless server: alone, and as a replica in a advanced replication system. In each case, you will start the server by using the linked library access function call `SSCStartDisklessServer()`.

Diskless Server Alone

If you run a diskless server alone, it has no way to read data when it starts up and no way to write data when it shuts down. This means that each time the server starts, it starts without any previous data.

Furthermore, since the server has no way to write data to disk, if the server is shut down abnormally (due to a power failure, for example), then any data in the server is lost and cannot be recovered. You can reduce the risk of data loss by using the solidDB HotStandby component to create a "hot standby" machine that contains a copy of the data. For more information about this hot standby capability, see *IBM solidDB High Availability User Guide*.

Diskless Server as Part of a Advanced Replication System

A diskless server may be a replica in a advanced replication system. In this situation, the replica may send data to the master server and may download data from that master server. Thus, even though the replica has no disk storage or other permanent storage of its own, it may make some or all of its data persistent within the advanced replication system.

Configuration parameters for a diskless server

This section describes the parameter settings for implementing and maintaining a diskless server.

Parameters used in diskless servers

The following sections of the configuration file contain parameters that have specific settings for diskless servers.

Index file section

Following are the configuration parameters that affect the index file.

Filespec_[1...n] parameter

The FileSpec parameter describes the name and the maximum size of the database file. To define the maximum size in bytes for the main memory engine, the FileSpec parameter accepts the following arguments:

- database file name - Since the diskless server does not create a physical database file, this parameter is not used; however, a dummy value must be provided for this argument.

- maximum file size - This setting is required. You need to specify the size in bytes that is large enough to store all the data in the diskless server. Note that the maximum file size must be smaller than the cache size, which is set with the **CacheSize** parameter.

The default value for the FileSpec parameter is solidr.db, 5000000 bytes. For example:

```
FileSpec_1=SOLIDR.db 5000000
```

Note: If you specify multiple files, then the maximum file size setting must be the sum of all the FileSpec parameter settings.

Not surprisingly, the maximum size is limited by the physical memory available, since a diskless machine has no disk to use as swap space for virtual memory. Note that on some platforms, the amount of physical memory available to the applications may be less than the amount of physical memory in the machine. For example, in some versions of Linux on 32-bit systems, the amount of memory available to applications is limited to one half or one quarter of the theoretical address space (4GB) because Linux reserves the 1 or 2 most significant bits of the address for its own memory manager.

If the data in memory exceeds the maximum file size, the error message 11003 is displayed:

```
File write failed, configuration exceeded
```

CacheSize

The **CacheSize** parameter defines the amount of main memory in bytes that the server allocates for the buffer cache. For example:

```
CacheSize=10000000
```

The setting for this value depends on the following criteria for diskless servers:

- For disk based tables, the cache size (in bytes) should be at least 20% larger than the maximum file size (that is, the amount of data) set with the FileSpec parameter since this data is held in the buffer cache. The 20% overhead is an estimate that may vary depending on the usage of the database. For example:

```
[IndexFile]
FileSpec_1=solid.db 10MB
CacheSize=12MB
```

- Even if no disk-based tables are used (the database is created by using in-memory tables), the cache is necessary to hold system tables. In that case, the minimum cache size is 1-2 MB. The space occupied by the system tables depends of the number and complexity of database objects and whether advanced replication is used or not.
- The cache size must be less than the physical memory available for running the diskless server.

Total memory used by the diskless server can be estimated as follows. (Note that the TOTAL of all of these must fit within the amount of physical memory available, which means that the cache size must in fact be significantly smaller than the amount of physical memory available to the server:)

```
CacheSize
+ 5MB
+ (100K * number of users * number of active statements per user)
+ in-memory table space
+ (HSB operations to be sent to the Secondary) [1][2]
```

[1] This term of the equation applies to HotStandby users only. An HSB Primary server needs some memory to store HotStandby operations that are to be sent to the Secondary server. During a temporary network failure between the Primary server and the Secondary diskless server, the Primary may continue to accept transactions from an application. When the network connection is restored between the servers, updates from the Primary server are sent to the Secondary server. (HotStandby uses the transaction log to store these operations. A diskless server cannot write the transaction log to disk; the information must be stored in memory.) This memory is separate from the Cache.

[2] For this term of the equation, the maximum limit is currently 1 MB or 512 operations, whichever is lower. Unlike on a disk-based server, the transaction log is not allowed to keep growing until it uses up all available space.

The exact amount required also depends on other factors, including the nature of the queries executed against the server. Naturally, the amount of memory available to the server is less than the total physical memory, since the operating system etc. will use up some of the physical memory.

Com section

Following are the configuration parameters that affect communication between the master and the diskless replica server (if you are using the diskless server as a advanced replication replica server).

Listen parameter [Com]

This is the protocol and name that the diskless server uses when it starts listening to the network. Its default is Operating System dependent. Refer to "Managing Network Connections" in *IBM solidDB Administrator Guide*.

Configuration parameters that do not apply to diskless engines

The following configuration file parameters (grouped by section) are disabled or inoperable for diskless servers. These parameters affect behaviors that do not apply to diskless engines.

Table 20. Configuration parameters not applicable to diskless engines

Parameter	Description
<i>[General] Section</i>	
CheckpointInterval	This parameter is disabled since checkpoints do not apply to diskless servers.
<i>[IndexFile] Section</i>	
ReadAhead	No physical read from the database file, so this parameter is inoperable
PreFlushPercent	No physical write to the database file, so this parameter is inoperable
<i>[Logging] Section</i>	

Table 20. Configuration parameters not applicable to diskless engines (continued)

Parameter	Description
LogEnabled	<p>This parameter is disabled since transaction logging is always disabled for diskless servers.</p> <p>Note: Diskless mode supports transaction rollback only. Transaction rollbacks are typically used when some failure interrupts a half-completed transaction. The diskless mode does not support rollforward recovery.</p>

5 Using solidDB linked library access with Java

Note: This chapter assumes that you are already familiar with the material in the preceding chapters. If you jumped straight to this chapter because you are interested only in Java/JDBC, not C/ODBC, you will have missed too much material to understand this entire chapter.

Overview of solidDB JDBC Accelerator (SJA)

A Java/JDBC program, like a C/ODBC program, may use the solidDB linked library access to get higher performance and greater control over the server. SJA enables a Java application to start a local solidDB server, which will be loaded into the Java Virtual Machine context from a dynamic library called 'ssolidacxx'. The Java application will then be able to connect to the solidDB server and use the services solidDB DBMS provides through a standard JDBC API.

The client application program will get higher performance because it is directly linked to the server library, so calls to server functions do not have the overhead of network (RPC) calls. The application will have greater control because it can call functions (methods) in the solidDB Server Control (SSC) library to do things such as assign priorities to certain types of tasks. For example, the application might give itself a high priority and might give remote client applications a low priority.

solidDB JDBC Accelerator (SJA) can only be used when the server and client are linked together; thus, if the Java application and the solidDB server are to be run in separate hosts, SJA cannot be used.

Only the "local" client (the one that is linked to the linked library access library) can bypass the network and get the higher performance of the linked library access. Other client programs may also use the server, but they must connect through the network, and are treated as "remote" programs even if they are running on the same computer as the solidDB server. You may only have one "local" client; the rest are "remote". The remote programs may be a mix of C and Java programs.

The language in which the local client is written does not restrict which languages the remote clients can be written in. For example, if you use JDBC Accelerator, the remote client programs may use C, Java, or both.

How the Accelerator works

As with C programs, Java/JDBC programs that want to use the linked library access must link to the solidDB linked library access library (ssolidacxx). This library contains the entire solidDB server, except that it is in the form of a callable library instead of a standalone executable program. The ssolidacxx used with Java/JDBC is the same as the ssolidacxx that was explained in previous chapters; there are not separate versions for Java and C clients. Linking to the library allows a client program to avoid the overhead of RPC (Remote Procedure Calls) through the network.

When you use the linked library access with Java/JDBC, you link the following into a single executable process:

- solidDB linked library access library,

- your Java-language client program, and
- the JVM.

The "layers" in the executable process are, from top to bottom:

- Local Java/JDBC client application
- JVM (Java Virtual Machine)
- solidDB Accelerator Library (ssolidacxx)

Java commands in your client are executed by the JVM. If the command is a JDBC function call, then the JVM calls the appropriate function in `ssolidacxx`. The function call is "direct", rather than going through the network (through RPC). The calls are made using JNI (Java Native Interface). Note that you do not need to know about these low-level details. You do not need to write any JNI code yourself; you simply have to call the same JDBC functions that you would call if you were a remote client program.

Accessing a solidDB database from Java Accelerator is identical to accessing a solidDB database through RPC — with one exception: in order to access the database services, the application using Java Accelerator must first start the solidDB linked library access server. This is done with a proprietary API called SolidServerControl (SSC). SSC API calls are used to start, as well as to stop, the solidDB DBMS. The actual database connections are done with normal JDBC API. Both the SolidServerControl API and solidDB's JDBC driver can be found in a .jar file named `SolidDriver2.0.jar`.

When the local solidDB server is started, it will be loaded into the Java Virtual Machine context from a dynamic library called `ssolidacxx`. The Java application will then be able to connect to the solidDB server and use the services solidDB DBMS provides through a standard JDBC API.

Every local client program that uses solidDB Java Accelerator follows the same basic three-step pattern:

1. Start the accelerator server with SolidServerControl
2. Access the database by using normal JDBC API
3. When database processing is done, stop the accelerator server again with SolidServerControl

The SolidServerControl classes for accessing solidDB accelerator server have been embedded inside solidDB JDBC driver file, inside the `solid.ssc` package. The solidDB JDBC driver jar file (`SolidDriver2.0.jar`) contains the following packages:

- `solid.jdbc.*` solidDB JDBC driver classes
- `solid.ssc.*` solidDB Server Control classes (proprietary interface)

The classes inside the solidDB Server Control (`solid.ssc`) package are:

- SolidServerControl (for starting and stopping solidDB server from Java)
- SolidServerControlInitializationError (for reporting errors)

For detailed information on SolidServerControl (SSC) class interface, see "solidDB Server Control (SSC) API" on page 49.

To start a solidDB server from a Java application, you must instantiate the class SolidServerControl in the beginning of your application and call the `startServer`

method with correct parameters (examples are given below). After you've started the server, you should be ready to make a JDBC connection to the server.

System requirements

You need the following to use the solidDB Java Accelerator:

- The solidDB linked library access library itself. This is a file named `ssolidacxx`. The filename extension varies depending upon the platform; some common names and platforms are listed below:
 - Microsoft Windows: `ssolidacxx.dll` and the import library `solidimpac.lib`
 - Solaris and Linux: `ssolidacxx.so`
 - HP-UX: `ssolidacxx.sl`
- A valid license file for using the solidDB server and the linked library access
- solidDB JDBC2 driver file (`SolidDriver2.0.jar`)
- solidDB communication libraries for your platform (these are normally installed when you install the solidDB Development Kit).
- Java Development Kit (JDK) 1.4.2 or newer

Basic usage

Installation

If you have installed a Java Development Kit, then you do not need to do any further installation. When solidDB is installed, it includes the library(s) that are needed when using the solidDB Java Accelerator.

Note: You may need to set `PATH` and `CLASSPATH` environment variables to appropriate values so that you can access the Java compiler.

Compiling and running a program

In order for the server startup to succeed, you need to have at least a valid license for using solidDB and linked library access.

The `ssolidacxx` dynamic link library must be in the system search path. Proceed as follows:

1. Set the paths (examples from Microsoft Windows command prompt)
`set PATH=<path to your ssolidacxx DLL>;%PATH%`
Make sure you have the directory containing solidDB communication libraries in your path too.
2. Set your path environment variable to include JDK's HOTSPOT runtime environment in (SJA has only been tested in hotspot JRE's). For example,
`set PATH=<your JDK directory>\jre\bin\hotspot;%PATH%`
3. Compile the sample `SJASample.java` file (located in the `samples/aclib_java` directory) with the following command:
`javac -classpath <IBM solidDB JDBC driver directory>/SolidDriver2.0.jar; \ SJASample.java`
4. Run your application with a command line resembling the next one:
`java -Djava.library.path=<path to ssolidacxx DLL> \ -classpath <IBM solidDB JDBC driver directory>/SolidDriver2.0.jar; \ <your application name>`

For example, on Microsoft Windows, if you installed the server to C:\soliddb and would like to run the SJASample program, then your command line would look like:

```
java -Djava.library.path=C:\soliddb\bin
-classpath C:\soliddb\jdbc\SolidDriver2.0.jar;. SJASample
```

(On Microsoft Windows, the `ssolidacxx.dll` dynamic library is in the `bin` subdirectory of the solidDB root installation directory.)

As in the example class `SJASample`, you must pass the solidDB server at least the following parameters with `SolidServerControl`'s `startServer` method:

```
-c<directory containing solidDB license file>
-U<username>
-P<password>
-C<catalog>
```

Note that upper and lower case "C" are both used, and they mean different things.

Assuming you have all the necessary files (`ssolidacxx` library, communication libraries, JDBC driver and `solid.lic`) in your current working directory, you can start `SJASample` with a command line like the following one:

```
java -Djava.library.path=. -classpath SolidDriver2.0.jar. <your application>
```

If all things went as they were supposed to go, you should now have a solidDB accelerator server up and running.

Making JDBC connections

solidDB Java accelerator supports both local database connections as well as RPC based connections.

In order to make a local (non RPC-based) JDBC connection, you need to specify the JDBC driver that you are using 'localserver' at port 0. Thus, if you are making the database connection by using, for example, JDBC class `DriverManager`, connect by using the following statement (as also presented in the example code `SJASample` further below)

```
DriverManager.getConnection("jdbc:solid://localserver:0", myLogin, myPwd);
```

As you can see, the `DriverManager` uses the URL `"jdbc:solid://localserver:0"` for making a connection to the local server. If the `getConnection` subroutine is given another URL, the driver will probably try to connect with RPC.

So remember the URL -

```
jdbc:solid://localserver:0
```

when making Java accelerator connections.

Note:

If you are using multiple threads (`java.lang.Thread` objects) that access solidDB linked library access server inside your Java application, you must register each thread separately with the solidDB linked library access server before you start any JDBC-related activities using that thread. The thread registration is done by calling `SolidServerControl` API's `registerThread` method in the thread's context. The thread registration must be done explicitly for each user thread (except the main thread) using solidDB's JDBC driver.

The user must also explicitly unregister each thread that has been registered to the solidDB linked library access server. To unregister a thread, call SolidServerControl API's unregisterThread function.

Limitations

- solidDB 'admin commands' do not work in the Java accelerator context.
- Java does not behave consistently if something fails outside the VM context (for example, inside a native method call). If something should assert (or even crash) in the solidDB server native code, Java either exits (when it notices an unexpected exception) or hangs up completely. In the latter case, you may have to kill the dangling Java process manually.
- To minimize memory consumption, we recommend that users explicitly drop all allocated statements; that is, all allocated JDBC Statement objects must be explicitly freed by calling the close() method.
- The server can crash if you access the same statement object from multiple Java threads. You must open a separate JDBC connection (and statement) for each thread that needs to use JDBC.

solidDB Server Control (SSC) API

Below is the complete public interface for the SolidServerControl class. For an example of a program that uses some of the methods in this class, see the file samples/aclib_java/SJASample.java

```
/**
 * See solidDB Linked Library Access User Guide
 * for the following constants
 */
    public final static int SSC_SUCCESS = 0;
    public final static int SSC_ERROR = 1;
    public final static int SSC_ABORT = 2;
    public final static int SSC_FINISHED = 3;
    public final static int SSC_CONT = 4;
    public final static int SSC_CONNECTIONS_EXIST = 5;
    public final static int SSC_UNFINISHED_TASKS = 6;
    public final static int SSC_INVALID_HANDLE = 7;
    public final static int SSC_INVALID_LICENSE = 8;
    public final static int SSC_NODATABASEFILE = 9;
    public final static int SSC_SERVER_NOTRUNNING = 10;
    public final static int SSC_INFO_SERVER_RUNNING = 11;
    public final static int SSC_SERVER_INNETCOPYMODE = 12;

    public final static int SSC_STATE_OPEN      = (1 << 0);
    public final static int SSC_STATE_PREFETCH = (1 << 1);

/**
 * Initiates a SolidServerControl class. Output is not directed to any
 * PrintStream.
 *
 * @return      SolidServerControl instance
 */
public static SolidServerControl instance()
    throws SolidServerInitializationError;

/**
 * Initiates a SolidServerControl class. Output is being directed
 * to a PrintStream object given in parameter 'os'.
 *
 * @param os    the PrintStream for output
```

```

* @return      SolidServerControl instance
*
*/
public static SolidServerControl instance( PrintStream os )
    throws SolidServerInitializationError;

/**
* setOutputStream method sets the output to the given PrintStream
*
* @param os      the PrintStream for output
*/
public void setOutputStream( PrintStream os );

/**
* getOutputStream returns the stream used for output in class
* SolidServerControl
*
* @return      returns the outputstream of this object
*/
public PrintStream getOutputStream();

/**
* startServer starts the solidSB Linked Library Access server
*
* @param argv      parameter vector for the accelerator server
* ( be sure to give the working directory containing
* solidDB license file (f.ex. -c\tmp) first, in front
* of other parameters. ) See solidDB Linked Library
* Access User Guide for details of parameters that can
* be passed to the Linked Library Access server.
*
* @param runflags  Options for this parameter are SSC_STATE_OPEN
* (remote connections are allowed) and
* SSC_STATE_PREFETCH (server will do a "prefetch"
* if needed). Prefetch refers to the memory
* and/or disk cache that provides read-ahead
* capability for table content. Following is
* a runflags parameter entry:
* runflags |= SSC_STATE_OPEN & SSC_STATE_PREFETCH
*
* @return  the return value from the server :
* SSC_SUCCESS
* SSC_ERROR
* SSC_INVALID_LICENSE - No license or invalid license file found.
* SSC_NODATABASEFILE - No database file found.
*/
public long startServer( String[] argv, long runflags );

/**
* stopServer stops the solidDB Linked Library Access server
*
* @param runflags  Runflags for stopping solidDB Linked Library Access server.
* See solidDB Linked Library Access User Guide for more
* details.
*
* @return  the return value from the server
* SSC_SUCCESS if server is stopped.
* SSC_CONNECTIONS_EXIT if there are open connections.
* SSC_UNFINISHED_TASKS if there are still tasks that are
* executing.
* SSC_SERVER_NOTRUNNING if the server is not running.
*/
public long stopServer( int runflags );

```

```

/**
 * returns the state of the server, i.e. is the server running or not
 *
 * @return SSC_STATE_OPEN if server is up and running
 */
public int getState();

/**
 * registerThread registers this user thread to solidDB Linked Library Access server
 *
 *
 * @return the return value from the server
 * SSC_SUCCESS Registration succeeded.
 * SSC_ERROR Registration failed.
 * SSC_INVALID_HANDLE Invalid local server handle
 given.
 * SSC_SERVER_NOTRUNNING Server is not running.
 */
public long registerThread();

/**
 * unregisterThread unregisters this user thread from the
 * solidDB Linked Library Access server
 *
 *
 * @return the return value from the server
 * SSC_SUCCESS Registration succeeded.
 * SSC_ERROR Registration failed.
 * SSC_INVALID_HANDLE Invalid local server handle given.
 * SSC_SERVER_NOTRUNNING Server is not running.
 */
public long unregisterThread();

```

Appendix. Linked library access parameters

Linked library access parameters

This appendix provides a list of all parameters for the linked library access. Linked library access parameters appear in the [Accelerator] section of the solidDB configuration file (`solid.ini`).

For a description of all other solidDB parameters, refer to the appropriate Appendix in the *IBM solidDB Administrator Guide*.

Note that you can change solidDB parameters in the following ways:

- Entering the ADMIN COMMAND 'parameter' command in solidDB `solsql`.
- Manually editing the `solid.ini` configuration file.

Note that any changes to the `solid.ini` file using the methods above do not take effect until the next time that the server starts.

Accelerator section

Table 21. Accelerator parameters

[Accelerator]	Description	Factory Value
ImplicitStart	If set to yes, this parameter starts solidDB automatically as soon as the ODBC API function <code>SQLConnect</code> is called in a user application. If set to no, solidDB must be explicitly started with a call to the Control API function <code>SSCStartServer</code> .	yes

Index

A

- accessing linked library access
 - description 9
- administering diskless servers
 - defining solidDB configuration file options 41
- applications
 - preparing for the linked library access 12

B

- backup listening mode 31

C

- C applications
 - samples 10
- CacheSize (parameter)
 - configuring for diskless 42
- Com section
 - configuring for diskless 43
- configuration file
 - CacheSize (parameter) 42
 - configuring 41
- connections
 - establishing for linked library access 14
 - ODBC remote without server startup 21
- Control API
 - ADMIN COMMAND equivalents 24
 - SSCGetActiveTaskClass (function) 23
 - SSCGetServerHandle (function) 23
 - SSCGetStatusNum (function) 23
 - SSCGetTaskClassState (function) 23
 - SSCIsRunning (function) 23
 - SSCIsThisLocalServer (function) 23
 - SSCSetNotifier (function) 23
 - summary of scheduling functions 24

D

- database
 - Index file section 41
 - size 16
- diskless
 - parameter settings 41
- dual mode application
 - description 4

E

- events
 - notifying function of 23

F

- FileSpec (parameter) 41
 - configuring for diskless 41

I

- implicit startup 21
- ImplicitStart (parameter) 21, 53
- IndexFile section
 - configuring for diskless 41

J

- JDBC API
 - description 6

L

- library
 - contents of linked library access 9
 - for remote user applications 9
 - solidimpac 12
- linked library access
 - accessing 9
 - Components 1
 - Described 1
 - linking applications 11
 - shutting down 20
 - starting 15
- linking applications
 - linked library access 11
- Linux
 - memory limitations with 41
- Listen (parameter)
 - configuring for diskless 43
- local application
 - description 4

M

- makefile examples 13
- memory
 - CacheSize (for diskless server) 42
 - total used by diskless server 42

N

- netcopy listening mode 31

O

- ODBC
 - API 6
 - applications
 - building with advanced replication scripts 11

P

- parameters
 - FileSpec 41
- passwords
 - criteria 18, 35

R

remote application 4

S

SaConnect
implicit startup with 20
server information
retrieving 23
shutting down
linked library access 20
solidctrlstub 5, 6, 9, 24
solidDB client APIs and drivers 5
solidDB configuration file
FileSpec (parameter) 41
Listen (parameter) 43
parameter settings 41
solidDB Control API
description 7
solidctrlstub 7
solidDB drivers and client APIs 5
solidDB SA 5
solidDB Server Control (SSC) API 49
solidimpac 12
SQLConnect (function)
implicit startup 18
SSC API 49
SSC_ABORT 27
SSC_CALL 25
SSC_CONNECTIONS_EXIST 27
SSC_CONT 27
SSC_ERROR 27
SSC_FINISHED 27
SSC_INFO_SERVER_RUNNING 27
SSC_INVALID_HANDLE 27
SSC_INVALID_LICENSE 27
SSC_NODATABASEFILE 27
SSC_SERVER_INNETCOPYMODE 27
SSC_SERVER_NOTRUNNING 27
SSC_STATE_OPEN 33, 34, 37
SSC_STATE_PREFETCH 33, 37
SSC_SUCCESS 26
SSC_TASK_ALL 26
SSC_TASK_BACKUP 26
SSC_TASK_CHECKPOINT 26
SSC_TASK_HOTSTANDBY 26
SSC_TASK_HOTSTANDBY_CATCHUP 26
SSC_TASK_LOCALUSERS 26
SSC_TASK_MERGE 26
SSC_TASK_NONE 26
SSC_TASK_REMOTEUSERS 26
SSC_TASK_SYNC_HISTCLEAN 26
SSC_TASK_SYNC_MESSAGE 26
SSC_UNFINISHED_TASKS 27
sscap.h 26
SSCGetServerHandle
function description 27
SSCGetStatusNum
function description 27
SSCsRunning
function description 28
SSCsThisLocalServer
function description 29
SSCRegisterThread
function description 29
SSCServerT 16

SSCSetNotifier
function description 30
SSCSetState
function description 33
SSCStartDisklessServer
function description 34
SSCStartServer
explicit startup with 16
function description 36
SSCStopServer
function description 38
shut down with 18
SscTaskSetT 26
SSCUnregisterThread
function description 39
starting solidDB
with linked library access 15
status information
retrieving 23
synchronization
using 10

T

task information
retrieving 23

U

usernames
default 17, 35

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

SC23-9831-03

