

IBM solidDB
Version 7.0

*Shared Memory Access and Linked
Library Access User Guide*



Note

Before using this information and the product it supports, read the information in "Notices" on page 95.

First edition, fifth revision

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

© Oy IBM Finland Ab 1993, 2013

Contents

Figures v

Tables vii

Summary of changes. ix

About this manual xi

Typographic conventions xi

Syntax notation conventions. xii

1 Overview of shared memory access and linked library access 1

1.1 Shared memory access (SMA) 2

1.1.1 System requirements for SMA. 4

1.1.2 SMA components and packaging. 4

1.2 Linked library access (LLA) 6

1.2.1 System requirements for LLA 6

1.2.2 LLA components and packaging 6

1.2.3 Static and dynamic link libraries for LLA 8

1.3 solidDB APIs and drivers for SMA and LLA 9

1.3.1 solidDB SA API 9

1.3.2 solidDB ODBC API 9

1.3.3 solidDB JDBC API 10

1.3.4 solidDB Server Control API (SSC API). 10

1.3.5 solidDB Server Control API (SSC API) for Java 10

1.4 Configurations with local and remote applications types 11

2 Creating and running SMA applications 13

2.1 Creating SMA applications - overview 13

2.1.1 Modifying shared memory kernel parameters - overview. 14

2.1.2 Preparing applications for SMA use with driver manager 20

2.1.3 Preparing applications for SMA use without driver manager 21

2.1.4 Establishing local connections for SMA 23

2.2 Starting and shutting down SMA server 24

2.2.1 Starting SMA server. 24

2.2.2 Shutting down SMA server 24

2.2.3 Starting SMA server as a service (Windows) 25

2.3 Monitoring SMA 26

2.4 Troubleshooting SMA. 27

3 Creating and running SMA applications with Java 31

3.1 Overview of using SMA with Java 31

3.2 Configuring your environment for SMA use with Java 31

3.3 Starting and shutting down SMA server 33

3.3.1 Starting SMA server. 33

3.3.2 Shutting down SMA server 33

3.4 Making JDBC connections for SMA 33

4 SMA with HotStandby 35

4.1 Configuring SMA TC with HotStandby 36

5 Creating and running LLA applications 39

5.1 Configuring your environment for LLA use 39

5.2 Establishing a local connection for LLA. 40

5.3 Starting and shutting down LLA server. 41

5.3.1 Explicit startup with SSC API function

SSCStartServer 42

5.3.2 Implicit startup with ODBC API function

call SQLConnect. 44

5.3.3 Implicit startup with SA API function call

SaConnect. 45

5.3.4 Shutting down LLA server 45

5.4 Sample C applications for LLA. 46

5.4.1 Samples for LLA with advanced replication 46

6 Creating and running LLA applications with Java 49

6.1 Overview of using LLA with Java. 49

6.1.1 Limitations. 50

6.2 Configuring your environment for LLA use with Java 50

6.3 Starting and stopping LLA server with SSC API for Java. 52

6.4 Making JDBC connections for LLA 52

6.5 Compiling and running a sample LLA program 52

7 Using the diskless capability 55

8 Creating and running remote or dual-mode applications 57

8.1 Example: Creating a dual-mode LLA application with ODBC and SSC API function calls 57

8.2 Establishing remote connections 57

Appendix A. Shared memory access parameters 59

Appendix B. Linked library access parameters 61

Appendix C. Configuration parameters for a diskless server. 63

C.1 IndexFile.Filespec_[1...n] parameter in diskless servers 63

C.2 IndexFile.CacheSize parameter in diskless servers 64

C.3 Com.Listen in diskless servers	65
C.4 Configuration parameters that do not apply to diskless engines	65

Appendix D. solidDB Server Control

API (SSC API) 67

D.1 Summary of SSC API functions	67
D.2 SSC API reference	69
D.2.1 SSCGetServerHandle	71
D.2.2 SSCGetStatusNum	72
D.2.3 SSCIsRunning	72
D.2.4 SSCIsThisLocalServer	73
D.2.5 SSCRegisterThread	73
D.2.6 SSCSetNotifier	74
D.2.7 SSCSetState	76
D.2.8 SSCStartDisklessServer	77

D.2.9 SSCStartServer	79
D.2.10 SSCStartSMADisklessServer	82
D.2.11 SSCStartSMAServer	83
D.2.12 SSCStopServer	85
D.2.13 SSCUnregisterThread	86
D.2.14 Retrieving task information.	87
D.2.15 Obtaining solidDB status and server information	87

Appendix E. SolidServerControl class

interface 89

Index 93

Notices 95

Figures

1. Configurations with SMA, LLA, and network-connection-based solidDB server . . . 2
2. Example: SMA and LLA APIs for C/C++ programs. 9
3. Architecture of SMA Transparent Connectivity with HotStandby. 35
4. Example: HotStandby with SMA configuration 37

Tables

1. Typographic conventions	xi	18. Linked library access (LLA) system libraries	50
2. Syntax notation conventions	xii	19. Shared memory access parameters	59
3. SMA drivers (libraries)	5	20. Shared memory access parameters (client-side)	60
4. SMA server applications	5	21. Accelerator parameters.	61
5. Linked library access (LLA) system libraries	7	22. Configuration parameters not applicable to diskless engines	65
6. SSC API and SA API libraries for remote applications	12	23. Summary of control API functions	67
7. Minimum requirements of shared memory kernel parameters for SMA (HP-UX)	17	24. SSC API parameter usage types.	70
8. Minimum requirements of shared memory kernel parameters for SMA (Linux)	18	25. Error codes and messages for SSC API functions	71
9. Minimum requirements of shared memory kernel parameters for SMA (Solaris)	20	26. SSCGetStatusNum parameters	72
10. SMA drivers (libraries).	20	27. SSCIsRunning parameters.	73
11. SMA drivers (libraries).	22	28. SCCRegisterThread parameters	74
12. Starting the SMA server	24	29. SSCSetNotifier parameters	74
13. SMA default address spaces	28	30. SSCSetState parameters	77
14. SMA drivers (libraries).	32	31. SSCStartDisklessServer parameters.	78
15. Starting the SMA server	33	32. SSCStartServer parameters	79
16. Linked library access (LLA) system libraries	39	33. SSCStartSMADisklessServer parameters	82
17. SSCStartServer parameters	42	34. SSCStartSMAServer parameters.	83
		35. SSCStopServer parameters	85
		36. SCCUnregisterThread parameters	87

Summary of changes

Changes for revision 05

- Information about configuring the cache size in diskless servers updated in sections Using the diskless capability and Configuration parameters for a diskless server.
- Information about the SSC_NOTIFY_SHUTDOWN_REQUEST event in the SSCSetNotifier function updated in section SSCSetNotifier: to accept the shutdown request, SSC_SUCCESS must be returned, not SSC_CONT.

Changes for revision 04

- Editorial corrections.

Changes for revision 03

- Editorial corrections.

Changes for revision 02

- Section Troubleshooting SMA updated.

Changes for revision 01

- Section Configuring SMA TC with HotStandby updated with information about specifying load balancing method when using SMA with Transparent Connectivity. When using SMA with TC, the load balancing method must be set to LOCAL_READ. If READ_MOSTLY or WRITE_MOSTLY is specified, a network connection is used instead of the SMA connection.

About this manual

The IBM® solidDB® shared memory access (SMA) and linked library access (LLA) enable applications to link to solidDB server directly, without the need to communicate through network protocols such as TCP/IP. SMA allows to link multiple applications while LLA allows to link one application. By replacing the network connection with local function calls, performance is improved significantly.

This guide contains information specific to SMA and LLA. 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 use SMA or LLA with Java™, 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.
<i>solid.ini</i>	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 Overview of shared memory access and linked library access

The solidDB *shared memory access* (SMA) and *linked library access* (LLA) allow applications to link to solidDB server directly, without the need to communicate through performance-consuming network protocols such as TCP/IP. SMA allows to link multiple applications while LLA allows to link one application.

SMA and LLA are implemented as library files that contain a complete copy of the solidDB server in library format.

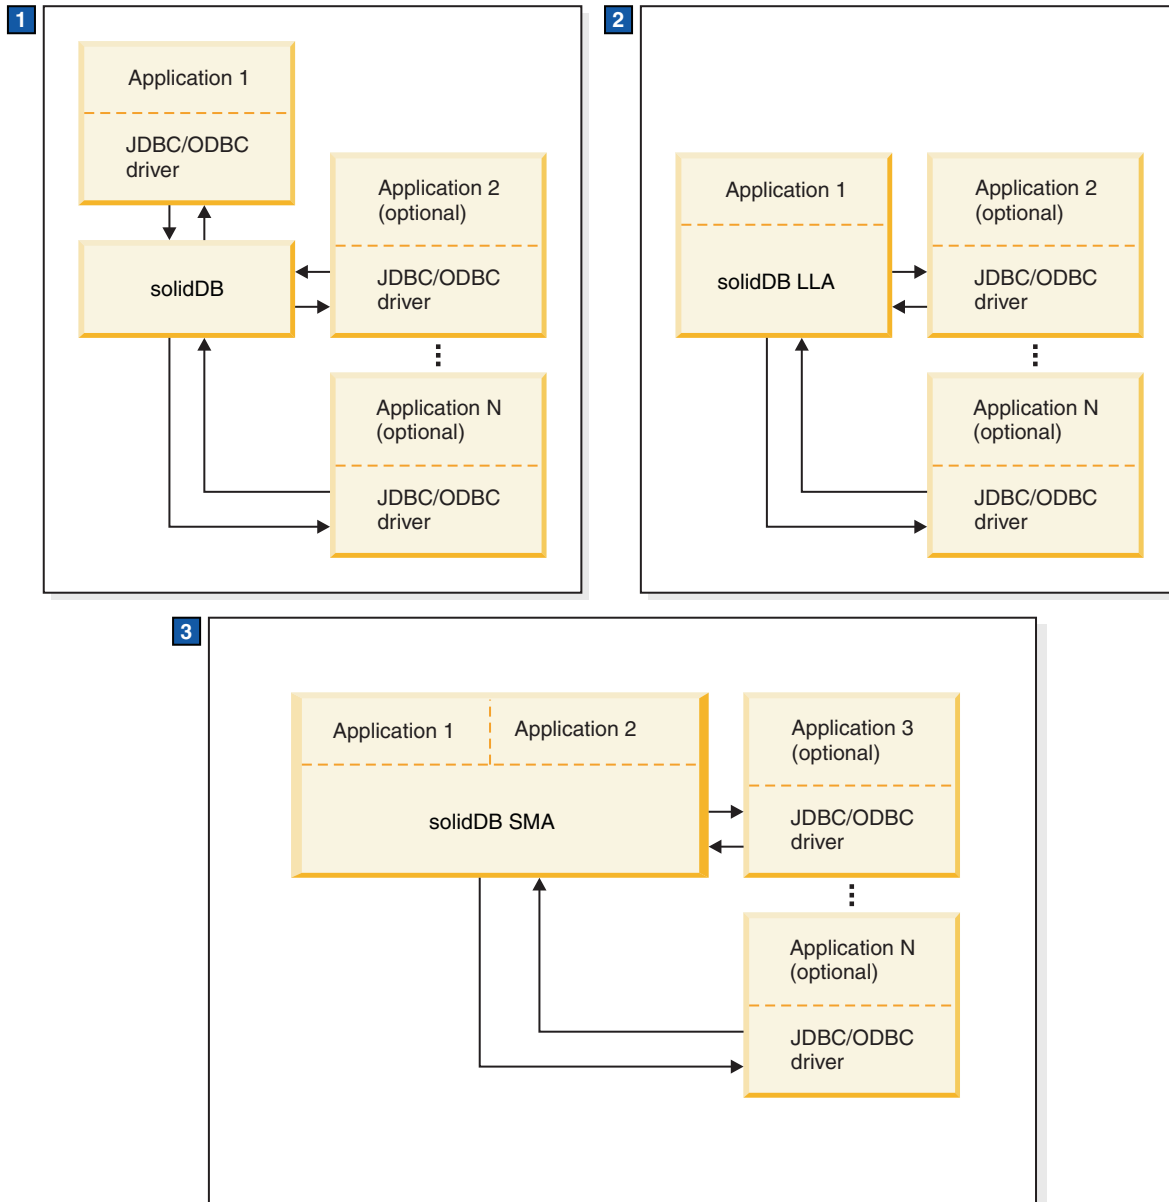
- With SMA, the library that the applications link to can be seen as a driver. Before the linked application is started, the solidDB server is started in a SMA mode, which loads the SMA driver dynamically and allocates and initializes a shared memory segment that the applications use to access the database.
- With LLA, the application links to the LLA library and the application and server are built into a single executable program.

Your application does not have to be rewritten to use SMA or LLA. The applications communicate with the solidDB server using ODBC or JDBC calls, or the solidDB proprietary SA API.

The SMA and LLA servers can also handle requests from remote applications which connect to the server through communications protocols such as TCP/IP. The remote applications see the SMA or LLA server as similar to any other solidDB server, while the local SMA and LLA applications see a faster, more precisely controllable version of the solidDB server.

Also, similarly to network-based servers, multiple SMA and LLA servers can be run on the same node.

The solidDB server used with SMA and LLA can be disk-based or diskless. Both in-memory tables and disk-based tables are also supported.



1. In a standard solidDB database configuration the applications and the server are separate programs.
2. In LLA configuration, LLA is a library that is linked into an application. Other applications may also communicate with the LLA server.
3. In SMA configuration, SMA is a driver library that multiple applications can link to. Other network connections based applications may also communicate with the SMA server.

Figure 1. Configurations with SMA, LLA, and network-connection-based solidDB server

1.1 Shared memory access (SMA)

With shared memory access (SMA), multiple applications can be linked to a dynamic driver library that contains the full database server functionality. This means that the applications ODBC or JDBC requests are processed almost fully in the application process space, without a need for a context switch among

processes. To facilitate the processing of a common database, the driver has access to a shared memory segment initialized by the server.

The application that is linked to the solidDB server using the SMA driver is called an *SMA application*; the server is called *SMA server*.

SMA server

Before the first application is started with SMA, the solidDB server is initialized in a SMA mode by starting a small application (`solidsma`) that loads the SMA driver library dynamically. This SMA server application starts the server code internally and allocates and initializes a shared memory segment that the applications use to access the database.

The SMA server has the full functionality of the network server:

- The SMA server process takes care of all client-independent tasks: startup and recovery, checkpointing and logging, making backups, and so on.
- The solidDB configuration parameters, admin commands, and command-line parameters can be used.
- Both in-memory and disk-based tables can be accessed equally.
- The SMA server can be used in solidDB Universal Cache, and with High Availability, Advanced Replication, and InfoSphere® CDC Replication configurations.
- The SMA server can also be used as a regular network-connection-based server.

When the server is started in the SMA mode, it accepts connection requests from the SMA driver at the normal listening port. By assigning different port numbers to different SMA servers, it is possible to run several SMA servers on a single system simultaneously.

When the server is shut down, or all the users are thrown out, the applications receive the error `Connection lost` on the next request. If the application is waiting for a response during a forceful shutdown, it receives a shutdown notification.

The SMA server can be disk-based or diskless. Diskless server can be useful in embedded systems that do not have hard disks, such as line cards in a network router or switch.

The solidDB data management tools can be used with network-based connections to the SMA server.

SMA application

An existing ODBC or JDBC application does not need to be modified in order to use SMA, with the exception of a data source name or a connect string. For example, in an ODBC application, the connection is requested with the regular ODBC `SQLConnect()` call.

An existing LLA application can be changed into an SMA application or vice versa. An application can also be changed from a SMA application to a network-based application.

SMA driver

The SMA driver is a dynamic library that contains a complete copy of the solidDB server in library format. The applications can link to the SMA driver directly or using a driver manager.

The footprint of the driver's binary file corresponds to the full solidDB server, which is 3-6 MB, depending on the platform. However, because all applications link to the same driver, the in-memory footprint is not multiplied when additional applications are added. The total memory footprint of the whole application system (applications, drivers, and the server) is comparable to the one of the client-server model.

SMA connection

Once the SMA server is running, applications can establish either SMA or network connections. For SMA connections, the applications have to be located on the same node as the server. The connection type is defined within the connect string. In the case of ODBC applications, if a driver manager is used, the SMA data source can be configured in the same way as an ODBC driver.

The connection request is sent over a network connection (handshake connection) using any locally available protocol (tcpip, named pipes, unix pipes). During the connect handshake, the shared-memory segment handle is passed to the driver so that it can access the server's shared memory segment.

1.1.1 System requirements for SMA

SMA is available on 64-bit platforms as well as 32-bit Linux. For SMA with Java, Java Runtime Environment (JRE) 1.4.2 or Java Development Kit (JDK) 1.4.2 or newer is required.

Supported platforms for SMA

- AIX®
- HP-UX
- Linux 64-bit
- Linux 32-bit
- Solaris
- Windows 64-bit

For more details on the supported platforms, see section *solidDB supported platforms* in the *IBM solidDB Getting Started Guide* .

1.1.2 SMA components and packaging

The SMA driver library and SMA server application are included in the solidDB software package. For SMA with Java, the solidDB JDBC Driver is needed.

SMA driver (library)

The SMA driver libraries for most common platforms are shown in the table below:

Table 3. SMA drivers (libraries)

Platform	SMA driver library	Default installation location
Windows	ssolidsm xx .dll Note: If you link to the SMA driver directly (without driver manager), you link to the solid ma .lib import library file that gives you access to the actual .dll library file.	Library: <solidDB installation directory>\bin Import library: <solidDB installation directory>\lib
Linux	ssolidsm xx .so	<solidDB installation directory>/bin
Solaris	ssolidsm xx .so	<solidDB installation directory>/bin
HP-UX	ssolidsm xx .so	<solidDB installation directory>/bin
AIX	ssolidsm xx .so	<solidDB installation directory>/bin

xx is the version number of the driver library, for example, ssolid $ma70$.so.

The SMA driver library for all platforms contains the following:

- full solidDB server functionality
- functions for three APIs
 - solidDB ODBC driver functions that allow for direct communication with the server library, without going through the network.
 - solidDB Control API (SSC API) library that contains functions to control starting and shutting down the SMA server.
 - 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.

Since the library that your application links to contains three APIs, your application program may call functions from any combination of these APIs. For details on each of these APIs, see 1.3, “solidDB APIs and drivers for SMA and LLA,” on page 9.

For SMA use with Java, the solidDB JDBC Driver is needed; the solidDB JDBC Driver (SolidDriver2.0.jar) is installed during solidDB server installation into the jdbc directory in the solidDB installation directory.

SMA server application

Table 4. SMA server applications

Platform	SMA application
Windows	solid ma .exe
Linux	solid ma
Solaris	solid ma
HP-UX	solid ma
AIX	solid ma

1.2 Linked library access (LLA)

With the linked library access (LLA), an application links to a static library or a dynamic library that contains the full database server functionality. This means solidDB runs in the same executable program with the application, eliminating the need to transfer data through the network.

The application that is linked to the solidDB server using the LLA library is called an *LLA application*; the server is called *LLA server*.

The application that links to the LLA library can also have multiple connections, using ODBC API, SA API, and JDBC API. All these APIs are reentrant, allowing simultaneous connections from separate threads.

The application that links directly to LLA library can also create remote connections to other database servers. The connection type (local or remote) is defined in the connect string that is passed to the ODBC API or SA API connect function or defined in the JDBC connection properties.

Principles of operation

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 implementations, the server runs on threads that are separate from the thread or threads used by the application. Calling the function to start a server performs any initialization steps required by the server code, creates the appropriate additional threads if necessary, and starts the server running on those threads.

Disk-based and diskless server

The solidDB server used with LLA can be disk-based or diskless. The LLA library contains two different function calls to start the server. The `SSCStartServer` function call starts a normal disk-based server, while `SSCStartDisklessServer` starts a server that does not use the disk drive.

1.2.1 System requirements for LLA

LLA is available on all platforms that solidDB supports. For LLA with Java, Java Runtime Environment (JRE) 1.4.2 or Java Development Kit (JDK) 1.4.2 or newer is required.

For more details on the supported platforms, see section *solidDB supported platforms* in the *IBM solidDB Getting Started Guide*.

1.2.2 LLA components and packaging

The LLA library is included in the solidDB software package. For LLA with Java, the JDBC driver and solidDB proprietary control classes are embedded in the solidDB JDBC Driver.

The LLA libraries for most common platforms are shown in the table below:

Table 5. Linked library access (LLA) system libraries

Platform	Static LLA library	Dynamic/Shared LLA Library
Windows	bin\ssolidacxx.dll	lib\solidimpac.lib This is an import library file that gives you access to the actual library file bin\ssolidacxx.dll.
AIX	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so.
HP-UX	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so.
Linux	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so
Solaris	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so

The LLA library for all platforms contains the following:

- full solidDB server functionality
- 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.

The library that your application links to contains all three APIs (SSC, SA, and ODBC); your application program can call functions from any combination of these APIs. For details on each of these APIs, see 1.3, “solidDB APIs and drivers for SMA and LLA,” on page 9.

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

1.4, “Configurations with local and remote applications types,” on page 11. For information about API files for remote applications, read 1.3, “solidDB APIs and drivers for SMA and LLA,” on page 9.

For LLA use with Java, the solidDB JDBC driver is needed; 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 SSC API for Java interface)

1.2.3 Static and dynamic link libraries for LLA

solidDB provides both a *static* and a *dynamic* version of the linked library access library.

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 (for example, `lib/solidac.a`), you link your program directly to it. 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 no difference between linking to the static library file or the dynamic library file. For example, the total amount of code in memory at any one time is similar. 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 if you execute more than one copy of the server in the same computer, you can save memory. For example, if you are doing development work on a single computer and you want to have both an advanced replication master and replica on the computer at the same time, or you would like to have a HotStandby Primary and a HotStandby Secondary at the same time, then you might prefer to use the dynamic library so that you do not have multiple copies of LLA in memory at the same time.

On Windows environments, the solidDB linked library access includes the additional file `lib/solidimpac.lib`. 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 type of linking links only a small amount of code to your client executable program. When your client program executes, the Windows operating system loads the `ssolidacxx.dll` file will automatically. Your client can then call the usual linked library access functions in the `.dll` file. The `.dll` file must be in your load path when you run the program that references it.

Note: Using the dynamic link library file does not mean that you can have multiple LLA applications clients linked to a single solidDB server. Even with the dynamic library approach, you are still limited to a single local client; all other clients must be remote clients. The remote clients communicate with the solidDB server by using TCP or some other network protocol, rather than the direct function calls available to the local client. To enable multiple local applications to access solidDB, design your environment using shared memory access (SMA).

1.3 solidDB APIs and drivers for SMA and LLA

The SMA and LLA application requests are typically handled through ODBC API direct function calls or JDBC calls. The solidDB proprietary solidDB API is available also. The solidDB Server Control API (SSC API and SSC API for Java) is included in the LLA library for handling local requests to control solidDB background processes and client tasks. SMA includes limited support for SSC API; only calls for starting or stopping the SMA server are included.

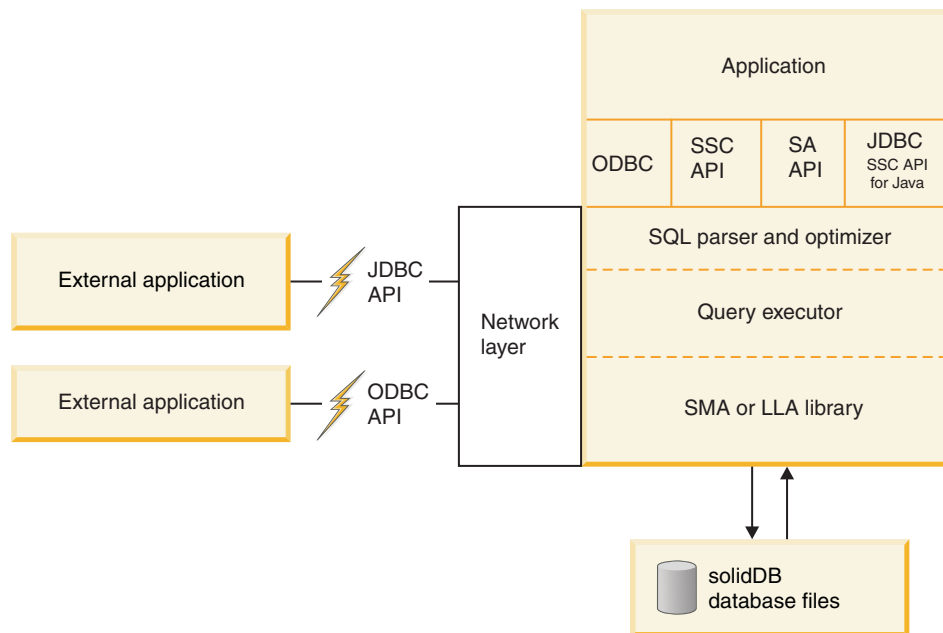


Figure 2. Example: SMA and LLA APIs for C/C++ programs

1.3.1 solidDB SA API

solidDB SA API is a low-level proprietary C-language API for data management services. The SA API provides support for local applications using SA API function calls.

The SA API library is used internally in the solidDB server. It provides access to data in database tables. The library contains about 90 functions providing low-level mechanisms for connecting the database and running cursor-based operations. The SA API can enhance performance significantly. You can use SA API to optimize the performance of batch insert operations, for example.

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

Remote applications can use the SA API function calls also. Remote applications must link to a separate SA API library file (for example, `solidimpsa.lib` for Windows).

1.3.2 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.

solidDB ODBC API is a Call Level Interface (CLI) for solidDB databases. It is compliant with ANSI X3H2 SQL CLI.

SMA and LLA support the ODBC 3.51 standard. The SMA and LLA libraries include solidDB ODBC 3.x, which provides support for local applications that require direct function calls to the server.

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

1.3.3 solidDB JDBC API

The JDBC API defines Java classes to represent database connections, SQL statements, result sets, database metadata, and so on. It allows you to issue SQL statements and process the results. JDBC is the primary API for database access in Java.

SMA and LLA support both JDBC 1.x and 2.x.

The JDBC interface and the solidDB JDBC Driver, including descriptions of the solidDB specific enhancements are documented in *IBM solidDB Programmer Guide*.

1.3.4 solidDB Server Control API (SSC API)

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

The SSC API functions are included in the SMA driver and LLA library files. However, for SMA, only the functions for starting and stopping the server are supported.

LLA provides support for local applications using SSC API function calls and a separate library is available for remote-only applications.

If your LLA application runs remotely and contains SSC API function calls, then you must link to the SSC 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 LLA.

1.3.5 solidDB Server Control API (SSC API) for Java

The solidDB Server Control API (SSC API) for Java is a proprietary API, named after `SolidServerControl` class. The SSC API for Java calls are used to start and stop the LLA server. The actual database connections are done with normal solidDB JDBC API. Both the SSC API for Java classes and solidDB JDBC driver classes are included in the solidDB JDBC driver (`SolidDriver2.0.jar`).

The `SolidServerControl` classes for accessing solidDB 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 API interface)

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

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

1.4 Configurations with local and remote applications types

With SMA and LLA, the applications always connect to a local solidDB server (SMA server or LLA server); the application and solidDB server are located on the same node. In addition to handling requests from the local SMA or LLA application, the SMA or LLA server can also handle requests from remote applications, which connect to the server through communications protocols such as TCP/IP. A dual-mode application can also be written; it switches modes between local and remote, depending upon how it is compiled and linked.

The SMA or LLA application is a *local application*; the server and the application are located on the same node. For example, calls to ODBC functions go directly to the server, rather than going through an ODBC driver and the communications protocol (such as TCP/IP).

A *remote application* is not linked to the SMA driver or LLA library. It is a separate executable that must communicate with the server using a network connection (such as TCP/IP) or other connection. Remote applications are usually run on a different node from the one that is running the server, but the application is also considered remote if it uses a network communication protocol to communicate with the server. A single node can run a local SMA or LLA application, while running one or more remote applications as separate processes.

The remote applications see the SMA and LLA server as similar to any other solidDB server, while the local application sees a faster, more precisely controllable version of the solidDB server.

Most applications are either local (linked to the SMA driver or LLA library) or remote (never linked). However, it is also possible to write a *dual-mode application* that uses both local and network-based connections. For example, the application can use the same C-language application code in either local or remote mode. The application is linked to a different library when in local mode than when in remote mode.

Dual-mode applications may be useful for example in the following cases:

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

The remote applications 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 LLA with Java, the remote client programs may use C, Java, or both.

SSC API and SA API libraries for remote applications

Remote applications that contain SSC API or SA API function calls must link to separate libraries.

Table 6. SSC API and SA API libraries for remote applications

Platform	SSC API stub library	SA API library	Default location
Windows	solidctrlstub.lib	solidimpsa.lib	<solidDB installation directory>\lib
Other platforms	solidctrlstub.a	solidsa.a	<solidDB installation directory>/bin

The SSC API stub library is required for remote applications because the SSC API functions included in the SMA and LLA libraries cannot be used with remote applications. For example, assume that you have a local application (containing SSC API functions) that links to a standard ODBC library. You want to run the same application remotely. By linking to the SSC API stub library, you avoid having to remove the SSC API function calls from your code. In this way, you can easily turn your LLA or SMA application into a normal remote client application.

Note: The SSC API stub library contains "do-nothing" functions; if you call them in a remote application, they have no effect on the server. The SSC API stub library does not actually give your remote application control of the server. Instead, it allows you to compile and link your application as a remote application, without getting link-time errors from solidDB with LLA or SMA.

2 Creating and running SMA applications

To create SMA applications, you configure solidDB as necessary, link your applications to the SMA driver, start an SMA server, and establish a local connection between the applications and the server. After you have created your application, you can monitor the SMA performance using the monitoring features provided with solidDB.

Important: The instructions for creating and running SMA applications provide SMA-specific additions, supplements, and usage differences in comparison to the solidDB without SMA.

For information about solidDB SQL, solidDB data management tools, general solidDB administration and maintenance, and database error codes, refer to the *IBM solidDB Administrator Guide*.

For detailed information about the APIs and solidDB JDBC and ODBC Drivers, see *IBM solidDB Programmer Guide*.

2.1 Creating SMA applications - overview

To create an application that uses SMA, you must prepare your system for SMA use, configure solidDB, set your application to use the SMA driver, start the SMA server, and connect your application to it.

About this task

This procedure provides an overview on how to create SMA applications. The SMA applications for C/ODBC environments are created in a similar way to applications that are not using SMA.

Note: When developing your application, it is recommended to use a network-based connection. Once your application is ready, switch to using the SMA connection.

For an example of a SMA application written in C, see the SMA sample in the `samples/sma` directory in your solidDB installation directory.

Procedure

1. Check the system settings for shared memory use in your environment.

The default values for shared memory use in your environment may not be sufficient for using SMA. For details on viewing and setting the shared memory system parameters on your system, see 2.1.1, "Modifying shared memory kernel parameters - overview," on page 14.

2. Set up your database environment by creating a working directory, your solidDB database, and user accounts.

For instructions, see *Creating a new database* in the *IBM solidDB Administrator Guide*.

Note: The application and the SMA server processes must have identical file access permissions (database files, log files, and so on). The file access

permissions are not checked at startup; subsequently, insufficient file access permissions may cause the SMA server to crash at a later point.

3. **Configure solidDB to meet your environment, performance, and operation needs.**

Use the `solid.ini` configuration file to define basic configuration settings such as database file names and locations, database block size, and so on.

- In normal setups, it is not necessary to modify the SMA-specific parameters in the `[SharedMemoryAccess]` section of the `solid.ini` file. The factory values are applicable to most use cases.
- Do not set the `Srv.ProcessMemoryLimit` parameter when using SMA. If you need to limit the memory the SMA server uses, use the `SharedMemoryAccess.MaxSharedMemorySize` parameter.

If there is no configuration file, the factory values are used.

4. **Prepare your application for SMA use.**

You can set up your application to use SMA with or without a driver manager.

- 2.1.2, “Preparing applications for SMA use with driver manager,” on page 20
- 2.1.3, “Preparing applications for SMA use without driver manager,” on page 21

5. **Start the SMA server.**

For instructions, see 2.2.1, “Starting SMA server,” on page 24.

6. **Start the application.**

2.1.1 Modifying shared memory kernel parameters - overview

Shared memory is allocated in segments. The shared memory system parameters control the maximum size and number of segments allowed on your system.

Typically solidDB uses 32 MB segment sizes.

The shared memory parameters and their management mechanisms depend on the system. In Linux and UNIX environments, you may need to address the type of kernel parameters described below.

Important: This section and the sections below discuss only the requirements set by solidDB. Other processes running on the same system may require higher limit values.

- **Maximum size of a shared memory segment**

Typically you do not need to modify the default system setting. This is because the solidDB segment size of 32 MB is considerably small.

- **Maximum number of shared memory segments in a system/process**

- Because solidDB allocates most of the segments in 32 MB, you may need more segments than allowed by your system by default, especially if you have a large database.

The maximum number of shared memory segments must be at least the solidDB process size in MB divided by 32.

For example, for a process size of 1 GB (1024 MB), at least 32 segments are needed.

- Set the maximum number of segments to a clearly higher value than required by your database size. A higher value has no side effects.

- solidDB uses only one process; if your environment requires you to set the maximum number of segments for a process and for the system separately, you can use the same value for both.
- Maximum total size of all shared memory segments
The total combined size of all shared memory segments depends on the size of your database and availability of disk space.

Note: In addition to this kernel parameter, the maximum total size of shared memory used by solidDB is controlled with the solidDB parameter **MaxSharedMemorySize** (in the [SharedMemoryAccess] section of the solid.ini file) as follows:

- The value set with the **MaxSharedMemorySize** parameter takes precedence over the value set with the kernel parameter. Thus, the value set with the **MaxSharedMemorySize** parameter must never be higher than the value set with the kernel parameter.
- By default, solidDB is set to use the maximum size of the physical memory of the computer (**MaxSharedMemorySize=0**). Thus, the default value set with the kernel parameter may be too low.

Example 1

If the system has 2 GB of memory and **MaxSharedMemorySize** is set to 0, solidDB uses maximum of 2 GB of memory. If the kernel parameter for maximum total size of all shared memory segments is then set to 1 GB, solidDB runs out of memory when the 1 GB is reached.

Example 2

If the system has 2 GB of memory and **MaxSharedMemorySize** is set to 500M, solidDB never uses more than 500 MB of memory. As long as the kernel parameter for maximum total size of all shared memory segments is set to 500 MB or higher, solidDB never runs out of memory. It is recommended to set the value to a higher value than the memory required by solidDB.

Shared memory kernel parameters for SMA on AIX

On AIX systems, the shared memory kernel parameters do not need to be modified. The upper limits are defined for the AIX IPC mechanisms, which are not configurable. The shared memory limits are allocated and deallocated dynamically as needed, so that the memory requirements always depend on the current system usage.

Important: Do not set the page space allocation policy to early page space allocation. Instead, use the deferred (default) or late space allocation policy.

For more details, see the following sections in the IBM Systems Information Center (<http://publib16.boulder.ibm.com/pseries/index.htm>):

- Inter-Process Communication (IPC) Limits – Shared memory default limits and the IPC mechanisms
- Page space allocation – Page space allocation policies

Modifying shared memory kernel parameters for SMA on HP-UX

The default values for shared memory kernel parameters on HP-UX may not be sufficient for running a SMA application. The kernel parameter values can be changed dynamically using the kctune command.

Before you begin

You must have root authority to modify shared memory kernel parameters.

About this task

The following steps show how to set shared memory kernel parameters on HP-UX. The minimum values that are shown according to the requirements set by solidDB. Other processes running on the same system may require higher limit values.

In HP-UX environments, you may need to modify the following shared memory kernel parameters:

- **shmmni** — Maximum number of shared memory segments on the system
- **shmseg** — Maximum number of shared memory segments attached to a process
- **shmmax** — Maximum size of a single shared memory segment (bytes)

Procedure

1. **View the shared memory kernel parameters to determine if there are any necessary changes required for your system.**

View the **shmmni** parameter:

```
kctune -v shmmni
Tunable          shmmni
Description      Maximum number of shared memory segments on the system
Module          vm_asi
Current Value    400 [Default]
Value at Next Boot 400 [Default]
Value at Last Boot 400
Default Value    400
Constraints      shmmni >= 3
                 shmmni <= 32768
                 shmmni >= shmseg
Can Change       Immediately or at Next Boot
```

View the **shmseg** parameter:

```
kctune -v shmseg
Tunable          shmseg
Description      Maximum number of shared memory segments attached to a process
Module          vm_asi
Current Value    300 [Default]
Value at Next Boot 300 [Default]
Value at Last Boot 300
Default Value    300
Constraints      shmseg >= 1
                 shmseg <= shmmni
Can Change       Immediately or at Next Boot
```

View the **shmmax** parameter:

```
kctune -v shmmax
Tunable          shmmax
Description      Maximum size of a shared memory segment (bytes)
Module          vm_asi
Current Value    1073741824 [Default]
Value at Next Boot 1073741824 [Default]
Value at Last Boot 1073741824
Default Value    1073741824
Constraints      shmmax >= 2048
                 shmmax <= 4398046511104
Can Change       Immediately or at Next Boot
```

Table 7. Minimum requirements of shared memory kernel parameters for SMA (HP-UX)

Parameter	Description	When to modify	Notes
shmmni	Maximum number of shared memory segments on the system	If the value is smaller than the solidDB process size (MB) divided by 32 For example, for a process size of 1 GB (1024 MB), at least 32 segments are needed.	Set this parameter to a clearly higher value than required by your database size; a higher value has no side effects.
shmseg	Maximum number of shared memory segments attached to a process	If the value is smaller than the solidDB process size (MB) divided by 32 For example, for a process size of 1 GB (1024 MB), at least 32 segments are needed.	Because solidDB uses only 1 process, the value of shmmni and shmseg can be the same.
shmmax	Maximum size of a single shared memory segment (bytes)	If the value is smaller than 32768 KB (32 MB)	Setting this parameter to a higher value has no side effects.

2. **To modify the parameters, use the kctune command.** For example, to set the maximum number of shared memory segments to 1024, use the following command:

```
kctune shmmni=1024
```

The parameter value change becomes effective immediately and stays effective after reboot.

What to do next

If you modified the shared memory parameters after getting an out of memory error, you may need to clear hanging shared memory segments with the `ipcrm` command. For more details, see 2.4, “Troubleshooting SMA,” on page 27.

Modifying shared memory kernel parameters for SMA on Linux

The default values for shared memory kernel parameters on Linux may not be sufficient for running a SMA application. To modify the shared memory kernel parameters on Linux, edit the `/etc/sysctl.conf` file.

Before you begin

You must have root authority to modify kernel parameters.

About this task

The following steps show how to update kernel parameters on Red Hat and SUSE Linux using the shared memory requirements set by solidDB. Other processes running on the same system may require higher limit values.

In Linux environments, you may need to modify the following shared memory parameters:

- **SHMMNI** — Maximum number of shared memory segments on a system
- **SHMMAX** — Maximum size of a single shared memory segment on a system
- **SHMALL** — Maximum allocation of shared memory pages on a system

Procedure

1. Run the `ipcs -l` command.

For example:

Note: Comments have been added following the `//` to show what the parameter names are.

```
# ipcs -l

----- Shared Memory Limits -----
max number of segments = 4096           // SHMMNI
max seg size (kbytes) = 32768           // SHMMAX
max total shared memory (kbytes) = 8388608 // SHMALL
```

2. Analyze the output to determine if there are any necessary changes required for your system.

Table 8. Minimum requirements of shared memory kernel parameters for SMA (Linux)

Kernel parameter	Description	When to modify	Notes
SHMMNI	Maximum number of shared memory segments on a system	If the value is smaller than the solidDB process size (MB) divided by 32 For example, for a process size of 1 GB (1024 MB), at least 32 segments are needed.	Set this parameter to a clearly higher value than required by your database size; a higher value has no side effects.
SHMMAX	Maximum size of a single shared memory segment on a system	If the value is smaller than 32768 KB (32 MB)	Setting this parameter to a higher value has no side effects. Note: The <code>ipcs</code> output has converted SHMMAX kilobytes. The kernel requires the SHMMAX value in bytes.
SHMALL	Maximum allocation of shared memory pages on a system	If <code>MaxSharedMemorySize=0</code> and the value of this parameter is smaller than the maximum size of the physical memory size of your computer in KB divided by 4. or If the value of this parameter is smaller than the value (in KB divided by 4) you have set with the parameter <code>MaxSharedMemorySize</code> .	<ul style="list-style-type: none"> The value set with the <code>MaxSharedMemorySize</code> parameter takes precedence over the value set with the kernel parameter. Thus, the value set with the <code>MaxSharedMemorySize</code> parameter must never be higher than the value set with the kernel parameter. By default, solidDB is set to use the maximum size of the physical memory of the computer (<code>MaxSharedMemorySize=0</code>). Thus, the default value set with the kernel parameter may be too low. <p>Note: The <code>ipcs</code> output has converted SHMALL into kilobytes. The kernel requires the SHMALL value as a number of pages.</p>

3. To modify these kernel parameters, edit the `/etc/sysctl.conf` file.

If the file does not exist, create it. For example, create a file with the following lines:

```
#Example shmmni for a 1 GB database
kernel.shmmni=400
#Example shmmax for a 64-bit system
kernel.shmmax=1073741824
#Example shmall for 16 GB memory
kernel.shmall=3774873
```

4. Run `sysctl` with `-p` parameter to load in `sysctl` settings from the default file `/etc/sysctl.conf`.

```
sysctl -p
```

5. Make the changes effective after every reboot.

- In SUSE Linux: make `boot.sysctl` active.
- In Red Hat Linux: the `rc.sysinit` initialization script reads the `/etc/sysctl.conf` file automatically.

What to do next

If you modified the shared memory parameters after getting an out of memory error, you may need to clear hanging shared memory segments with the `ipcrm` command. For more details, see 2.4, “Troubleshooting SMA,” on page 27.

Modifying shared memory kernel parameters for SMA on Solaris

The default values for shared memory kernel parameters on Solaris 10 may not be sufficient for running a SMA application. In Solaris 10, the shared memory kernel parameter values can be changed dynamically with the Solaris resource control facilities.

Before you begin

You must have root authority to modify shared memory parameters.

About this task

The following steps show how to set shared memory kernel parameters on Solaris 10. The minimum values that are shown according to the requirements set by `solidDB`. Other processes running on the same system may require higher limit values.

In Solaris environments, you may need to modify the following shared memory parameters:

- **max-shm-ids** — Maximum number of shared memory segments on a system
- **max-shm-memory** — Maximum size of all shared memory segments on a system (MB)

In the examples below, the operating system default project is used.

Procedure

1. **View the shared memory parameters to determine if there are any necessary changes required for your system.**

View the **project.max-shm-ids** parameter:

```
prctl -n project.max-shm-ids -i project default
project: 3: default
NAME      PRIVILEGE      VALUE      FLAG      ACTION      RECIPIENT
project.max-shm-ids
  privileged      128        -      deny      -
  system          16.8M      max      deny      -
```

View the **project.max-shm-memory** parameter:

```
prctl -n project.max-shm-memory -i project default
project: 3: default
NAME      PRIVILEGE      VALUE      FLAG      ACTION      RECIPIENT
project.max-shm-memory
  privileged      62.7GB      -      deny      -
  system          16.0EB      max      deny
```


Table 9. Minimum requirements of shared memory kernel parameters for SMA (Solaris)

Parameter	Description	When to modify	Notes
max-shm-ids	Maximum number of shared memory segments on a system	If the value is smaller than the solidDB process size (MB) divided by 32 For example, for a process size of 1 GB (1024 MB), at least 32 segments are needed.	Set this parameter to a clearly higher value than required by your database size; a higher value has no side effects.
max-shm-memory	Maximum size of all shared memory segments on a system	If MaxSharedMemorySize=0 and the memory size set with this parameter is smaller than the maximum size of the physical memory size of your computer. or If the memory size set with this parameter is smaller than the memory size you have set with the parameter MaxSharedMemorySize .	Setting this parameter to a higher value has no side effects. <ul style="list-style-type: none"> The value set with the MaxSharedMemorySize parameter takes precedence over the value set with the kernel parameter. Thus, the value set with the MaxSharedMemorySize parameter must never be higher than the value set with the kernel parameter. By default, solidDB is set to use the maximum size of the physical memory of the computer (MaxSharedMemorySize=0). Thus, the default value set with the kernel parameter may be too low.

2. To modify the parameters, use the **prctl** command.

For example, to set the maximum number of shared memory segments to 1024, use the following command:

```
prctl -n project.max-shm-ids -r -v 1024 -i project default
```

3. Make the changes effective after every reboot.

```
/usr/sbin/projmod -sK "project.max-shm-ids=(privileged,1024,deny)" default
```

What to do next

If you modified the shared memory parameters after getting an out of memory error, you may need to clear hanging shared memory segments with the **ipcrm** command. For more details, see 2.4, “Troubleshooting SMA,” on page 27.

2.1.2 Preparing applications for SMA use with driver manager

When using SMA with a driver manager, you connect to a SMA data source in a similar way as when connecting to a regular solidDB ODBC data source.

About this task

The SMA driver library file is installed during the solidDB server installation. The table below lists the file names and their default installation locations for the most common platforms.

Table 10. SMA drivers (libraries)

Platform	SMA driver library	Default installation location
Windows	ssolidsmmax.dll Note: If you link to the SMA driver directly (without driver manager), you link to the solidma.lib import library file that gives you access to the actual .dll library file.	Library: <solidDB installation directory>\bin Import library: <solidDB installation directory>\lib

Table 10. SMA drivers (libraries) (continued)

Platform	SMA driver library	Default installation location
Linux	ssolidsm xx .so	<solidDB installation directory>/bin
Solaris	ssolidsm xx .so	<solidDB installation directory>/bin
HP-UX	ssolidsm xx .so	<solidDB installation directory>/bin
AIX	ssolidsm xx .so	<solidDB installation directory>/bin
xx is the version number of the driver library, for example, ssolidsm 70 .so.		

Procedure

1. **Configure the SMA data source according to the instructions provided with your driver manager.**

2. **Connect to the SMA data source.**

Use the SMA-specific connect string when defining the data source connection information.

The connect string syntax for a SMA connection is:

```
sma <protocol name> <port number or pipe name>
```

For example:

```
sma tcp 2315
```

3. **Check the use of signal handlers.**

Signal handlers are used to report the occurrence of an exceptional event to the application. The SMA driver installs by default its own signal handler that can help the SMA system to survive the most common application failures, such as killing or interrupting the applications from outside. Upon the capture of certain signals, the signal handler closes the SMA connections safely and exits the SMA application. This means that in most cases, the SMA server continues to run despite abnormal application exits.

By default, the SMA driver handles the following signals that can cause the SMA connection to break:

- Linux and UNIX: SIGINT, SIGTERM
- Windows: SIGINT

You can modify the set of signals that the SMA driver handles with the client-side parameter **SharedMemoryAccess.Signals**. You can also disable the SMA driver signal handler by setting the client-side parameter **SharedMemoryAccess.SignalHandler** to no.

If the **SharedMemoryAccess.SignalHandler** parameter is set to 'yes' (default), do not set signal handlers in your application for those signals that are handled by the SMA driver; the application setting will override the SMA driver settings.

Related reference:

Appendix A, "Shared memory access parameters," on page 59

2.1.3 Preparing applications for SMA use without driver manager

When you are using SMA without a driver manager, you link your application to the SMA driver library directly. You link to the SMA driver in a similar way as when linking directly to an solidDB ODBC driver library.

Procedure

1. **Link your application to the SMA driver library.**

The SMA driver library files are installed during the solidDB installation. The table below lists the file names and their default installation locations for the most common platforms.

Table 11. SMA drivers (libraries)

Platform	SMA driver library	Default installation location
Windows	ssolidsmmax.dll Note: If you link to the SMA driver directly (without driver manager), you link to the solidsma.lib import library file that gives you access to the actual .dll library file.	Library: <solidDB installation directory>\bin Import library: <solidDB installation directory>\lib
Linux	ssolidsmmax.so	<solidDB installation directory>/bin
Solaris	ssolidsmmax.so	<solidDB installation directory>/bin
HP-UX	ssolidsmmax.so	<solidDB installation directory>/bin
AIX	ssolidsmmax.so	<solidDB installation directory>/bin
xx is the version number of the driver library, for example, ssolidisma70.so.		

2. Change the connect string to the local SMA server name.

The connect string syntax for a SMA connection is:

```
sma <protocol name> <port number or pipe name>
```

For example:

```
sma tcp 2315
```

For examples of the connect string when using ODBC API or SA API, see 2.1.4, "Establishing local connections for SMA," on page 23.

3. Check the use of signal handlers.

Signal handlers are used to report the occurrence of an exceptional event to the application. The SMA driver installs by default its own signal handler that can help the SMA system to survive the most common application failures, such as killing or interrupting the applications from outside. Upon the capture of certain signals, the signal handler closes the SMA connections safely and exits the SMA application. This means that in most cases, the SMA server continues to run despite abnormal application exits.

By default, the SMA driver handles the following signals that can cause the SMA connection to break:

- Linux and UNIX: SIGINT, SIGTERM
- Windows: SIGINT

You can modify the set of signals that the SMA driver handles with the client-side parameter **SharedMemoryAccess.Signals**. You can also disable the SMA driver signal handler by setting the client-side parameter **SharedMemoryAccess.SignalHandler** to no.

If the **SharedMemoryAccess.SignalHandler** parameter is set to 'yes' (default), do not set signal handlers in your application for those signals that are handled by the SMA driver; the application setting will override the SMA driver settings.

Related reference:

Appendix A, "Shared memory access parameters," on page 59

2.1.4 Establishing local connections for SMA

For SMA use, the application needs to establish a local SMA connection with the SMA server. The connection type is defined with a SMA-specific connect string.

For SMA, the connection request is sent over the network connection (handshake connection) using any locally available protocol (tcpip, named pipes, unix pipes). During the connect handshake, the shared-memory segment handle is passed to the driver so that it can access the server's shared memory.

The connect string syntax for a SMA connection is:

```
sma <protocol name> <port number or pipe name>
```

For example:

```
sma tcp 2315
```

If an SMA connect request is made to the server that is not an SMA Server, a connect error is returned.

Important: A single application can connect to only one SMA server. However, the SMA application may make regular network-based connections to any local or remote server.

ODBC API

When using the ODBC API, define the SMA-specific connect string in the SQLConnect function call.

Examples

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

```
rc = SQLConnect(hdbc, "sma tcp 1315", (SWORD)SQL_NTS, "dba", 3, "dba", 3);
```

or

```
rc = SQLConnect(hdbc, "sma upipe SOLID", (SWORD)SQL_NTS, "dba", 3, "dba", 3);
```

SA API

When using the SA API, define the SMA-specific connect string in the SaConnect function call.

The following code examples connect directly to the local SMA solidDB server with username dba and password dba:

```
SaConnectT* sc = SaConnect("sma tcp 1315", "dba", "dba");
```

or

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

Driver manager

When using a Driver manager, define the SMA-specific connect string in the SMA Data Source.

2.2 Starting and shutting down SMA server

The SMA server is started, restarted, and shutdown in a similar way as the normal network-connection-based solidDB server.

2.2.1 Starting SMA server

The SMA server is started using the command prompt in the same way as with a normal network-connection-based solidDB server. When SMA server is started, it checks if a database exists. The server first looks for a `solid.ini` configuration file and reads the value of `FileSpec` parameter. If a database file is found with the names and paths specified in the `FileSpec` parameter, that database is opened automatically. If no database is found, the server prompts you to create a database.

Procedure

To start the SMA server:

Table 12. Starting the SMA server

Operating system	To start the SMA server:
Linux and UNIX	Enter the command <code>solidsma</code> at the command prompt. When you start the server for the first time, enter the command <code>solidsma -f</code> at the command prompt to force the server to run in the foreground.
Windows	Enter the command <code>solidsma</code> at the command prompt. To start the SMA server as a service:

Results

When the server is started in the SMA mode, it loads the SMA driver library dynamically, accepting connection requests from the SMA driver at the normal listening port. By assigning different port numbers to different SMA servers, it is possible to run several SMA servers on a single system at the same time.

Tip: It is also possible to start the solidDB server in SMA mode by having the application call an SSC API function `SSCStartSMAServer`. However, in such a setup, only one application can start (and stop) the SMA server. For details on the SSC API calls, see D.2.11, “`SSCStartSMAServer`,” on page 83.

2.2.2 Shutting down SMA server

SMA server is shut down using the solidDB ADMIN COMMANDS.

Procedure

1. To prevent new connections to solidDB, close the database 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:
ADMIN COMMAND 'shutdown'

Results

All the shutdown mechanisms 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.

Tip: It is also possible to stop the solidDB server in SMA mode by having the application call an SSC API function `SSCStopServer()`. Only one application can start and stop the SMA server. The same application that started the SMA server must also perform the shutdown. For details, see D.2.12, “SSCStopServer,” on page 85.

2.2.3 Starting SMA server as a service (Windows)

solidDB with SMA can be run as a service in Windows. The first time you want to run the SMA server as a service, you must install the service, that is, allow Windows to run the SMA server as a service. After that, you can start and stop the services with the Windows Service dialog or command prompt, or remove the services using solidDB command-line options.

Before you begin

To be able to install and start services in some Windows environments (for example, Windows 2008 Server), you need to run the Windows command prompt with administrator rights.

1. In the **Start** menu, right-click **Command Prompt**.
2. Select **Run as administrator**.
3. Log in with an administrator account.

About this task

The first time you want to run the SMA server as a service, you must first install the service, and then start the service with the Windows Service dialog or command prompt.

Procedure

1. **Allow (install) Windows to run the SMA server as a service.**

In the solidDB working directory, issue the following command:

```
solidsma -s"install,<name>,<fullexepath> -c<working directory>[,autostart]"
```

where

<name> is the service name

<fullexepath> is the full path for solidsma.exe

<working directory> is the full path for solidDB working directory (where your `solid.ini` configuration file and license file are located)

[autostart] is an optional parameter that sets the Startup Type of the service to Automatic, that is, the SMA server runs automatically as a service when Windows is started.

Note:

Regardless of the [autostart] parameter, the service is not started automatically at the time of installation. For the first time, the service has to be

started manually in the Windows Services dialog or command prompt. (See step 2 below.)

Example 1

The following command installs a service named SOLIDSMA (with Startup Type Manual) when the SMA server is installed into the directory C:\soliddb and the working directory is C:\soliddb.

```
solidsma -s"install,SOLIDSMA,C:\soliddb\bin\solidsma.exe -cC:\soliddb"
```

Example 2

The following command installs a service named SOLIDSMA (with Startup Type Automatic) when the SMA server is installed into the directory C:\soliddb and the working directory is C:\soliddb. The next time Windows is started, the SMA server will automatically run as a service.

```
solidsma -s"install,SOLIDSMA,C:\soliddb\bin\solidsma.exe -cC:\soliddb,autostart"
```

Tip:

Alternatively, you can create the service using the Windows command-line utility `sc.exe`. In that case, to start the SMA server in a services mode, you must include the `solidDB -sstart` command-line option in the command. For example:

```
sc create SOLIDSMA binPath= "c:\soliddb\bin\solidsma.exe -cC:\soliddb -sstart"
```

The `-sstart` command-line option is required to remove the GUI-based interactions between the SMA server and the user. Programs running as a Windows service cannot use the GUI operations.

2. Start the service manually in the Windows Services dialog or command prompt.

- You can access the Windows Services dialog through Control Panel: **Control Panel > Administrative Tools > Services**.

- In the command prompt, issue the following command:

```
sc start <name>
```

For example:

```
sc start SOLIDSMA
```

Results

When running as an Windows service, `solidDB` with SMA will log warning and error messages to the Windows event log. These messages can be viewed from Windows by using the Event Viewer, available through Control Panel: **Control Panel > Administrative Tools > Event Viewer**. Messages are also logged to the `solmsg.out` file.

2.3 Monitoring SMA

`solidDB` includes means for monitoring and collecting data on the type and number of SMA connections as well as the SMA memory segment sizes.

- Use ADMIN COMMAND 'userlist' to print a list of the type of user connections (network client or SMA client).
- Use ADMIN COMMAND 'report' to print a list of the connections by the type.
- Check the `solmsg.out` file login entries for the type of the connections made.
- Use the performance counter SMA connection count to collect data on the number of SMA connections.

- Use the performance counter SMA shared memory used to collect data on the SMA memory segment size.

For details on using the ADMIN COMMANDs, the `solmsg.out`, and the performance counters, see *IBM solidDB Administrator Guide*.

Measuring processor utilization in SMA-based databases

The SMA server process (`solidsma`) cannot be used to measure the processor utilization of the database used by SMA applications. With SMA, most of the solidDB server code executes in the application's address space. On the other hand, the processor utilization reported for an SMA application reflects both the load inflicted by the application code itself and the solidDB code invoked to process the application's request.

If the applications are not processor-intensive, the best approximation of the processor utilization of the total database system is the sum of the processor utilization values of all the SMA applications and the SMA server process.

The SMA server process consumes very little of the processor capacity in its own address space. The server process consists of only housekeeping and asynchronous tasks like checkpointing, backup, and so on.

2.4 Troubleshooting SMA

This section provides instructions and guidelines on how to prevent or troubleshoot common problems while configuring or using SMA.

Error: Server could not allocate shared memory segment by id -1

Symptoms

When trying to start a SMA server, the following type of error is displayed, and the SMA server cannot be started.

```
IBM solidDB process has encountered an internal error and is unable to
continue normally. Report the following information to technical support.
SOLID Fatal error: Out of central memory when allocating buffer memory (size = 33554432)
Date: 2012-04-24 15:39:44
Product: IBM solidDB
Version: 7.0.0.2 Build 2012-04-20
```

```
[solid1]~ ./solidsma -f -c .
Server could not allocate shared memory segment by id -1
```

Causes

The SMA server startup fails because there is no memory available. This situation can occur if:

- When a SMA application or SMA server terminates abnormally, they can leave shared memory allocated. Even if you shut down all SMA processes, the shared memory is still left reserved.
- You have allocated too little memory for SMA use.

This leads to a situation where all memory is used and you cannot start a SMA server any more.

Resolving the problem

In Linux and UNIX environments, clear the hanging shared memory segments with the `ipcrm` command.

For example in Linux environments, use the following script to identify and remove the unused shared memory segments.

```
#!/bin/sh

if [ $# -ne 1 ]
then
    echo "$0 user"
    exit 1
fi

for shm_id in $(ipcs -m|grep $1|awk -v owner=$1 ' { if ( owner == $3 ) {print $2} }')
do
    ipcrm -m $shm_id
done
```

For more details on the ipcrm command, see your operating system documentation.

Cannot map shared memory area

Symptoms

When trying to connect to a SMA server, the following type of error is displayed, and the connection fails.

- Linux and UNIX operating systems
cannot map shared memory area 1288077395 to 0x2b0029800000
Cannot connect to target database.
- Windows operating systems
SQL State "08004"; Native Error Code "25215";
Error Text "SMA failed in MapViewOfFileExt,
desired addr: 0000000800000000, got addr: 0000000000000000, error: 6.

Causes

When started, the SMA starts attaching shared memory segments to an address space that is used by another process.

Resolving the problem

In general, the earlier your application connects to the SMA server, the less likely it is that the address space requested by solidDB is in use.

The SMA server uses the following address spaces by default:

Table 13. SMA default address spaces

Operating system	Default start address space*
AIX	0x700000010000000ul
Linux 64-bit	0x2c0000000000
Linux 32-bit	0x50000000
Solaris Intel	0x2b0000000000
Solaris Sparc	0xffffffff60000000
Windows	0x0000000080000000

*The start address space is the value of the parameter **shmaddr** in the shmat() system call.

1. Force the start address space for the SMA server to a different address space using the environment variable SOLSMASRT.
 - Linux and UNIX operating systems:
export SOLSMASRT=<start_address_space>
For example:
export SOLSMASRT=0x2b0000000000

- Windows operating systems:
set SOLSMASRT=<start_address_space>
For example:
set SOLSMASRT=0x0000000800000000
2. Restart the SMA server.

Error 21300: Protocol 'sma' is not supported

Symptoms

When trying to connect to a SMA server, the following type of error is displayed:

```
Error HY000: SOLID Communication Error 21300:  
Protocol 'sma' is not supported  
SQLConnect failed
```

Causes

The application has been linked both to the solidDB ODBC library and the SMA library (ssolidmaxx).

Resolving the problem

Check your application code and remove any references to the solidDB ODBC libraries (for example, sac12x70.so or socw6470.dll).

3 Creating and running SMA applications with Java

Java applications are linked to the SMA driver library. The actual database connections are done with the normal JDBC API.

3.1 Overview of using SMA with Java

A Java application that uses SMA is created in the same way as an application that uses a regular solidDB server, with the exception that you start an SMA server instead of a regular solidDB server. The Java application connects to the SMA server and uses the services solidDB server provides through a standard JDBC API. Linking to the dynamic library allows the application to avoid the overhead of RPC (Remote Procedure Calls) through the network.

Java/JDBC programs that want to use SMA link to the SMA driver library (`ssolidmaxx`). This library contains the entire solidDB server, except that it is in the form of a callable library instead of a stand-alone executable program. The libraries used with Java/JDBC are the same as the ones used with C/C++ applications; there are no separate versions for Java.

When you use SMA with Java/JDBC, you link the following components into a single executable process:

- SMA driver library
- your Java-language client program
- the JVM

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

- Local Java (JDBC) client application
- JVM (Java Virtual Machine)
- SMA driver library

The Java commands in your client are executed by the JVM. If the command is a JDBC function call, the JVM calls the appropriate function in the SMA driver library. The function call is direct, it does not go through the network (through RPC). The calls are made using Java Native Interface (JNI). 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 writing a remote client program.

Every application that uses SMA follows the same basic four-step pattern:

1. Configure the solidDB server and connection settings.
2. Start the SMA server.
3. Access the database by using normal JDBC API.
4. When database processing is done, stop the SMA server.

3.2 Configuring your environment for SMA use with Java

When using SMA with Java, your `LD_LIBRARY_PATH` or `LIBPATH` (Linux and UNIX) or `PATH` (Windows) environment variable must include the location of the SMA driver library.

Before you begin

It is assumed that you have a working installation of the solidDB JDBC Driver.

About this task

The SMA driver library file is installed during the solidDB installation. The following table lists the file names and their default installation locations for the most common platforms.

Table 14. SMA drivers (libraries)

Platform	SMA driver library	Default installation location
Windows	ssolidmaxx.dll Note: If you link to the SMA driver directly (without driver manager), you link to the solidma.lib import library file that gives you access to the actual .dll library file.	Library: <solidDB installation directory>\bin Import library: <solidDB installation directory>\lib
Linux	ssolidmaxx.so	<solidDB installation directory>/bin
Solaris	ssolidmaxx.so	<solidDB installation directory>/bin
HP-UX	ssolidmaxx.so	<solidDB installation directory>/bin
AIX	ssolidmaxx.so	<solidDB installation directory>/bin

xx is the version number of the driver library, for example, ssolidma70.so.

Procedure

1. **Add the location of the SMA driver library to the LD_LIBRARY_PATH or LIBPATH (Linux and UNIX) or PATH (Windows) environment variable.**

- In Linux and UNIX environments, use the following syntax:
export LD_LIBRARY_PATH=<path to SMA library>:\$LD_LIBRARY_PATH
or
in AIX environments:
export LIBPATH=<path to SMA library>:\$LIBPATH
- In Windows environments, use the following syntax:
set PATH=<path to SMA library>;%PATH%

2. **Set up your database environment by creating a working directory, your solidDB database, and user accounts.**

For instructions, see *Creating a new database* in the *IBM solidDB Administrator Guide*.

Note: The application and the SMA server processes must have identical file access permissions (database files, log files, and so on). The file access permissions are not checked at startup; subsequently, insufficient file access permissions may cause the SMA server to crash at a later point.

3. **Configure solidDB to meet your environment, performance, and operation needs.**

Use the solid.ini configuration file to define basic configuration settings such as database file names and locations, database block size, and so on.

- In normal setups, it is not necessary to modify the SMA-specific parameters in the [SharedMemoryAccess] section of the solid.ini file. The factory values are applicable to most use cases.

- Do not set the **Srv.ProcessMemoryLimit** parameter when using SMA. If you need to limit the memory the SMA server uses, use the **SharedMemoryAccess.MaxSharedMemorySize** parameter. If there is no configuration file, the factory values are used.

3.3 Starting and shutting down SMA server

The SMA server is started, restarted, and shutdown in a similar way as the normal network-connection-based solidDB server.

3.3.1 Starting SMA server

The SMA server is started using the command prompt in the same way as with a normal network-connection-based solidDB server. When SMA server is started, it checks if a database exists. The server first looks for a `solid.ini` configuration file and reads the value of `FileSpec` parameter. If a database file is found with the names and paths specified in the `FileSpec` parameter, that database is opened automatically. If no database is found, the server prompts you to create a database.

Procedure

To start the SMA server:

Table 15. Starting the SMA server

Operating system	To start the SMA server:
Linux and UNIX	Enter the command <code>solidsma</code> at the command prompt. When you start the server for the first time, enter the command <code>solidsma -f</code> at the command prompt to force the server to run in the foreground.
Windows	Enter the command <code>solidsma</code> at the command prompt. To start the SMA server as a service:

3.3.2 Shutting down SMA server

SMA server is shut down using the solidDB ADMIN COMMANDs.

Procedure

1. To prevent new connections to solidDB, close the database 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:
`ADMIN COMMAND 'shutdown'`

3.4 Making JDBC connections for SMA

To make a local (non-RPC-based) JDBC connection to the SMA server, you need to connect to a SMA server using the non-standard connection property `solid_shared_memory`, and use a local server at a given port in the JDBC URL (connection string).

Connecting with Driver Manager

1. Set the non-standard connection property `solid_shared_memory` to `yes`.
2. Set the connection string to use a local server (`localhost`) and define an available port number.

For example:

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

Defining connection property in the connect string

Include the connection property `solid_shared_memory=yes` in the connect string, and use a local server (`localhost`), defining an available port number.

For example:

```
Connection c = DriverManager.getConnection
("jdbc:solid://localhost:1315?solid_shared_memory=yes");
```

Note: In addition to the `DriverManager` class, a similar syntax is available for classes `SolidDataSource` and `SolidConnectionPoolDataSource`.

4 SMA with HotStandby

The SMA server node can be made highly available with the solidDB HotStandby component.

In an SMA with HotStandby setup, there can be one or more SMA applications on each node. The application connection to the database can be configured as a regular SMA connection (SMA Basic Connectivity) or as a Transparent Connectivity SMA connection (SMA TC). With both connectivity types, the application on the Primary node uses an SMA connection to execute reads and writes locally, and the application on the Secondary uses an SMA connection to execute reads locally. Additionally, with SMA TC connection, write transactions from the application on the Secondary can be executed on the Primary server using a network connection. Furthermore, if the load balancing option is enabled with the SMA TC connection, the applications can operate in an active-active manner; on each node, the full functionality of database access is available.

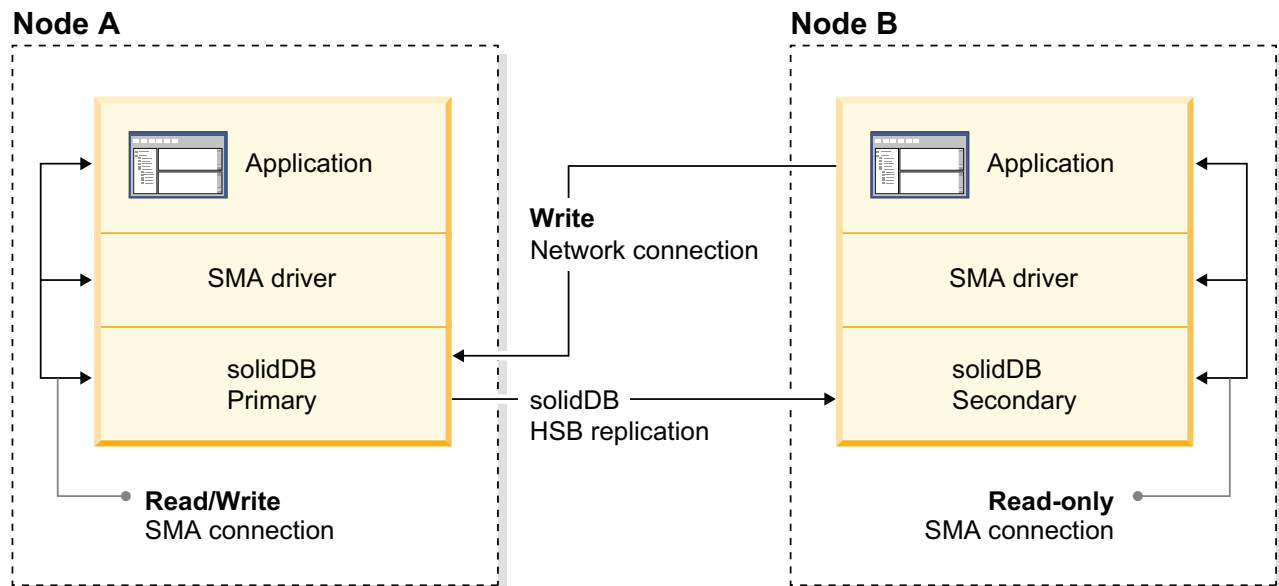


Figure 3. Architecture of SMA Transparent Connectivity with HotStandby

With SMA TC, the application on each node must be able to connect to the local server with a SMA connection and to the remote server with a network-based connection.

Failover and switchover handling

- The connection handle is maintained over switchovers and failovers for as long as one of the servers is in the PRIMARY ACTIVE, PRIMARY ALONE, or STANDALONE state.
- If the SMA server fails, the application might fail also. To ensure high availability in such a failure scenario, your system needs to include an application-level failover mechanism that moves the service offered by the application from the failed application instance to another one.

4.1 Configuring SMA TC with HotStandby

When using SMA with Transparent Connectivity (TC), the applications on the Primary and the Secondary must connect to the databases using a SMA-specific TC connect info syntax.

About this task

With SMA TC, the application on each node must be able to connect to the local server with a SMA connection and to the remote server with a network-based connection.

The format of the TC connect target list for SMA with HotStandby is the following:

```
connect_target_list::=[SERVERS:]sma_connect_string, network_connect_string
```

where

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

```
network_connect_string:= protocol_name IP_address | host_computer_name  
                        port_number | pipe_name
```

Additionally, you need to set the load balancing method to LOCAL_READ (PREFERRED_ACCESS=LOCAL_READ).

Important: When using SMA with TC, if you set the load balancing method to READ_MOSTLY or WRITE_MOSTLY (default), a network connection is used instead of the SMA connection. Thus, when using SMA with TC, always set the load balancing method to LOCAL_READ.

Procedure

1. Set up the two HotStandby servers.
2. Set up SMA on both servers.
3. For both applications, define the TC connection using the SMA-specific connect target list syntax and the load balancing attribute PREFERRED_ACCESS=LOCAL_READ.
4. Compile and start the applications.

Example

Connect info of the application on host1 where solidDB is listening at port 1964:
PREFERRED_ACCESS=LOCAL_READ SERVERS=sma tcp 1964, tcp host2 2315

Connect string of the application on host2 where solidDB is listening at port 2315:
PREFERRED_ACCESS=LOCAL_READ SERVERS=sma tcp 2315, tcp host1 1964




Figure 4. Example: HotStandby with SMA configuration

5 Creating and running LLA applications

Creating LLA applications includes linking your application to the library, starting the server, and establishing a local connection between the application and the server. You can start and stop the server with the SSC API, ODBC API, and SA API.

The instructions contain LLA-specific additions, supplements, and usage differences in comparison to the solidDB without LLA.

For information about solidDB SQL, solidDB data management tools, general solidDB administration and maintenance, and database error codes, refer to the *IBM solidDB Administrator Guide*.

For detailed information about the APIs and solidDB JDBC and ODBC Drivers, see *IBM solidDB Programmer Guide*.

5.1 Configuring your environment for LLA use

When using LLA, you must link your application to the LLA library file.

Procedure

1. **Link your application to the LLA library file specific to your operating system.**

Table 16. Linked library access (LLA) system libraries

Platform	Static LLA library	Dynamic/Shared LLA Library
Windows	bin\ssolidacxx.dll	lib\solidimpac.lib This is an import library file that gives you access to the actual library file bin\ssolidacxx.dll.
AIX	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so.
HP-UX	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so.
Linux	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so

Table 16. Linked library access (LLA) system libraries (continued)

Platform	Static LLA library	Dynamic/Shared LLA Library
Solaris	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so

Note: In Linux and UNIX environments, you need to link to the symbolic link library `libssolidacxx` that is available in the `<solidDB_installation_directory/lib>` directory. Alternatively, rename the `ssolidacxx` library in the `<solidDB_installation_directory/bin>` directory as `libssolidacxx`.

Example: Makefile for providing the LLA library name in Windows

In the makefile example below, the `solidimpac.lib` library is used.

```
# 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
```

2. If you do not plan to use the implicit start method for starting the solidDB server with SSC API, set the `ImplicitStart` parameter to no.

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 in the same way. When the `SQLConnect` or `SaConnect` function is called for the first time, the server is implicitly started.

3. Disable 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.

5.2 Establishing a local connection for LLA

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.

In the ODBC API, to establish a connection to a local server (the server that was linked to the application), the user application calls the `SQLConnect` function with the literal string "localserver". You can also specify an empty source name "" for the local server connection. Furthermore, you can also specify a local server name. However, specifying a local server name causes the linked library access to use a "remote" connection, that is, 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). You can also specify an empty source name "" for the local server connection. Furthermore, you can also specify a local server name. However, specifying a local server name causes the linked library access to use a "remote" connection, that is, 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");
```

5.3 Starting and shutting down LLA server

You can start, restart, and shut down the LLA server using the SSC API, ODBC API, or SA API function calls. The ODBC API and SA API function calls can be used to start the server only if a database exists already. The SSC API can be used to create a database at the startup.

At server startup, 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.

Explicit startup and shutdown with SSC API

The SSC API is used to start and shut down the LLA server *explicitly*. The application calls the SSC API function `SSCStartServer` to start the server and `SSCStopServer` to shut it down.

When you start a new LLA server that does not already have a database, you must specify explicitly that a new database is created. To create a database, include the following parameters with the `SSCStartServer()` function:

```
-Username  
-Ppassword  
-Catalogname (the default database catalog name)
```

Note: If you want to start a diskless server, you must start the server with SSC API function `SSCStartDisklessServer`.
For details, see 5.3.1, “Explicit startup with SSC API function `SSCStartServer`.”

Implicit startup and shutdown with ODBC API and SA API

The ODBC API and SA API can only be used to start and shut down the LLA server *implicitly*. When the application connects locally to LLA server for the first time, it calls the ODBC API function `SQLConnect` or SA API function `SaConnect`. In this case, shut down occurs when the last local connection disconnects from the server using either function `SQLDisconnect` or `SaDisconnect`.

When the LLA server is started implicitly from the application, it checks if a database exists in the working directory. If a database file is found, the server opens that database automatically. If a database file is not found, the server returns an error.

The server does not create a database during implicit startup. To create a database, you must use an explicit startup function, such as `SSCStartServer` with the appropriate parameters, or create a database as for a non-linked server.

For details, see 5.3.2, “Implicit startup with ODBC API function call `SQLConnect`,” on page 44 and 5.3.3, “Implicit startup with SA API function call `SaConnect`,” on page 45.

For instructions on how to create a database in a non-linked server setup, see section *Creating a new database* in the *IBM solidDB Administrator Guide*.

5.3.1 Explicit startup with SSC API function `SSCStartServer`

To start solidDB explicitly, have the user application call the solidDB Server Control API function `SSCStartServer()`.

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

where the parameters are:

Table 17. `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. 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, the <code>SSCStartServer</code> function writes the handle (pointer) of the new server into the variable whose address you passed.

Table 17. *SSCStartServer* parameters (continued)

Parameter	Description
runflags	<p>The value for this parameter is a combination of 2 flags: the open flag and the netcopy disabling flag. The following flag symbols can be used:</p> <ul style="list-style-type: none"> • <code>SSC_STATE_OPEN</code> – the open flag is set to 1: new connections are allowed. • <code>SSC_STATE_CLOSED</code> – the open flag is set to 0: all new network and LLA connections are rejected, except for connections from solidDB Remote Control (solcon) program. • <code>SSC_DISABLE_NETCOPY</code> – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which the <code>SSC_DISABLE_NETCOPY</code> is set. <p>The flag does not prevent the server to act as a source of netcopy. If only <code>SSC_DISABLE_NETCOPY</code> flag is set, the server is in the closed state. To enable netcopy, use the SSC API function <code>SSCSetState()</code> with the runflag value <code>SSC_STATE_OPEN</code> or <code>SSC_STATE_CLOSED</code>.</p> <pre>runflags = SSC_STATE_OPEN SSC_STATE_CLOSED SSC_DISABLE_NETCOPY</pre> <p>Tip: The flags can be used in combinations, for example:</p> <pre>... rc = SSCStartServer(g_argc, g_argv, &hh, SSC_STATE_OPEN SSC_DISABLE_NETCOPY); ...</pre> <p>If the server is started as closed, it can be opened with ADMIN COMMAND 'open', or with the solcon command open. The same effect can be achieved with the SSC API function <code>SSCSetState()</code>.</p>

Starting LLA server without an existing database

When you start the server for the first time, solidDB creates a database only if you have specified the database administrator username, password, and a name for the default database catalog.

For example:

```
SscServerT h; char* argv[4];
argv[0] = "apname"; /* 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.

By default, the database is 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. 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.

Starting LLA server with an existing database

If you already have an existing database, you do not need to specify the username and password, or the catalog name in the `SSCStartServer` function call.

5.3.2 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, which is the last open local connection.

Note: The server shuts 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 username and password. For a description and example, read 5.3.1, "Explicit startup with SSC API function SSCStartServer," on page 42.

Following is an example of implicit startup and shutdown 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 the SSCStartServer function call, the SQLDisconnect function call does not do implicit shutdown. The server must be shut down explicitly, either by SSCStopServer function call, the **ADMIN COMMAND 'shutdown'** command, 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 = SQLAllocConnect (henv, &hdbc);
rc = SQLConnect (hdbc, "", SQL_NTS, "dba", SQL_NTS, "dba", SQL_NTS);

/* Delete all the rows from table foo */
rc = SQLAllocStmt (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

```

5.3.3 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 5.3.1, “Explicit startup with SSC API function SSCStartServer,” on page 42.

Following is an example of implicit startup and shutdown 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 the SSCStartServer function call, it can only be shut down with the SSCStopServer function call.

5.3.4 Shutting down LLA server

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

Programmatically, you can perform the shutdown from an application such as solidDB SQL Editor (**solsql**), or solidDB Remote Control (**solcon**).

To shutdown solidDB:

1. To prevent new connections to solidDB, close the database 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:

ADMIN COMMAND 'shutdown'

All the shutdown mechanisms 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 the explicit method `SSCStopServer` to shut down a server that was started with implicit methods (`SQLConnect`). The converse is not true; for example, you cannot use `SQLDisconnect` to stop a server that was started with `SSCStartServer`.

Shutting down LLA server 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);
```

5.4 Sample C applications for LLA

The solidDB package includes samples of LLA applications written in C that use ODBC API functions to connect to solidDB servers.

The samples are located in the following directories in the solidDB installation directory:

- `samples/aclib`: sample of an LLA application using a single solidDB
- `samples/aclib_control_api`: sample of an LLA application that uses the SSC API functions.
- `samples/aclib_diskless`: sample of an LLA application using a diskless solidDB server
- `samples/aclib_replication`: sample of an LLA application that combines LLA and advanced replication

See 5.4.1, “Samples for LLA with advanced replication” for information about how to use the replication samples.

Each directory contains a `readme.txt` file that provides instructions on how to set up your system and run the samples.

5.4.1 Samples for LLA with advanced replication

If you are new to solidDB data synchronization, *IBM solidDB Advanced Replication User Guide* contains information about 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:

- Use solidDB without SMA or LLA to run the SQL scripts contained in *IBM solidDB Advanced Replication User Guide*. These scripts are found in `samples/replication`.
- Run the SQL scripts locally, using the solidDB SMA or LLA. As a prerequisite, you are required to set up an application to start the server according to the instructions in this document.

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 the section *Getting started with data synchronization* in the *IBM solidDB Advanced Replication User Guide* using solidDB SQL Editor (`solsql`).

Setting up your ODBC application with the advanced replication sample 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 replicas and control the execution of synchronization messages. These scripts can 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`, into the `embed.c` application.

6 Creating and running LLA applications with Java

Java applications are linked to the LLA library and they use solidDB SSC API calls to start and stop the solidDB. The actual database connections are done with the normal JDBC API. Both the SolidServerControl API calls and JDBC driver are included in the solidDB JDBC Driver .jar file (SolidDriver2.0.jar).

6.1 Overview of using LLA with Java

LLA enables Java applications to start a local solidDB server, which will be loaded into the Java Virtual Machine context from a dynamic library. The Java application is then able to connect to the solidDB server and use the services solidDB server provides through a standard JDBC API. Linking to the dynamic library allows the application to avoid the overhead of RPC (Remote Procedure Calls) through the network.

Java/JDBC programs that want to use LLA link to the LLA library (ssolidacxx). The LLA library contains the entire solidDB server, except that it is in the form of a callable library instead of a stand-alone executable program. The libraries used with Java/JDBC are the same as the ones used with C/C++ applications; there are not separate versions for Java.

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

- LLA 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)
- LLA library

Java commands in your client are executed by the JVM. If the command is a JDBC function call, the JVM calls the appropriate function in ssolidacxx. The function call is direct, it does not go through the network (through RPC). The calls are made using Java Native Interface (JNI). 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 with LLA is identical to accessing a solidDB database through RPC — with one exception: to access the database services, the application using LLA must first start the LLA server. You can start the LLA server with a proprietary API called solidDB Server Control (SSC) API for Java (named after SolidServerControl class. The actual database connections are done with normal solidDB JDBC API. Both the SSC API for Java and solidDB JDBC driver can be found in the .jar file named SolidDriver2.0.jar.

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

Every application that uses LLA follows the same basic four-step pattern:

1. Configure the solidDB server and the connection settings.
2. Start the LLA server with SolidServerControl class.
3. Access the database by using normal JDBC API.
4. When database processing is done, stop the LLA server again with SolidServerControl class.

6.1.1 Limitations

- All solidDB 'admin commands' cannot be used when using LLA with Java.
- 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 abnormally or hangs up completely. In the latter case, you may have to kill the dangling Java process manually.
- To minimize memory consumption, drop all allocated statements explicitly; all allocated JDBC statement objects must be explicitly freed by calling the close() method. Freeing objects is important especially if your setup use Transparent Connectivity (TC).

6.2 Configuring your environment for LLA use with Java

When using LLA with Java, your LD_LIBRARY_PATH or LIBPATH (Linux and UNIX) or PATH (Windows) environment variable must include the location of the LLA library. The LLA link libraries are available in the <solidDB_installation_directory/bin> and <solidDB_installation_directory/lib> directories.

Before you begin

It is assumed that you have installed and registered the solidDB JDBC Driver.

About this task

Table 18. Linked library access (LLA) system libraries

Platform	Static LLA library	Dynamic/Shared LLA Library
Windows	bin\ssolidacxx.dll	lib\solidimpac.lib This is an import library file that gives you access to the actual library file bin\ssolidacxx.dll.
AIX	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so.

Table 18. Linked library access (LLA) system libraries (continued)

Platform	Static LLA library	Dynamic/Shared LLA Library
HP-UX	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so.
Linux	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so
Solaris	bin/solidac.a	lib/libssolidacxx.so This is a symbolic link that gives you access to the actual library file bin/ssolidacxx.so

xx is the version number of the driver library, for example, ssolidac70.so.

Note: In Linux and UNIX environments, you need to link to the symbolic link library libssolidacxx that is available in the <solidDB_installation_directory/lib> directory. Alternatively, rename the ssolidacxx library in the <solidDB_installation_directory/bin> directory as libssolidacxx.

Procedure

1. Add the location of the LLA library to the LD_LIBRARY_PATH or LIBPATH (Linux and UNIX) or PATH (Windows) environment variable.

- In Linux and UNIX environments, use the following syntax:
export LD_LIBRARY_PATH=<path to LLA library>:\$LD_LIBRARY_PATH
or
in AIX environments:
export LIBPATH=<path to LLA library>:\$LIBPATH

Example:

```
export LD_LIBRARY_PATH=/opt/solidDB/soliddb-7.0/lib
```

- In Windows environments, use the following syntax:
set PATH=<path to LLA library>;%PATH%

2. Set up your database environment by creating a working directory, your solidDB database, and user accounts.

For instructions, see *Creating a new database* in the *IBM solidDB Administrator Guide*.

Note: The application and the SMA server processes must have identical file access permissions (database files, log files, and so on). The file access permissions are not checked at startup; subsequently, insufficient file access permissions may cause the SMA server to crash at a later point.

6.3 Starting and stopping LLA server with SSC API for Java

To start the solidDB server from a Java application, you must instantiate the class `SolidServerControl` in the beginning of your application and call the `ssc.startServer` method with correct parameters. After starting the server, you are ready to make a JDBC connection to the server. Similarly, the server is stopped with the call `ssc.stopServer`.

Procedure

1. Starting the server

- **LLA server:** `ssc.startServer`

When starting the server, you must pass the solidDB server at least the following parameters:

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

Tip: Uppercase C and lowercase c are not interchangeable; they provide different functionality.

2. Stopping the server

- **LLA server:** `ssc.stopServer`

6.4 Making JDBC connections for LLA

Linked library access (LLA) with Java 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.

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

If you are making the database connection by using, for example, JDBC class `DriverManager`, connect by using the following statement:

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

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 try to connect with RPC.

6.5 Compiling and running a sample LLA program

About this task

The examples in this procedure are given for Windows command prompt.

Procedure

1. Set the paths.

```
set PATH=<path to your ssolidacxx DLL>;%PATH%
```

Make sure that you have the directory containing solidDB communication libraries in your path too.

2. Set your path environment variable to include JDK HOTSPOT runtime environment in (SJA has only been tested in hotspot JRE).

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 the sample 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;. \ SJASample
```

For example, 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 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 the SolidServerControl startServer method:

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

Note: Uppercase and lowercase "C" are both used, and they mean different things.

5. If you have all the necessary files (ssolidacxx library, communication libraries, JDBC driver, and license file 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;. SJASample
```

Results

Your LLA server is up and running.

7 Using the diskless capability

SMA and LLA servers can be used to create a database engine that runs without any disk storage space. Diskless server is useful in embedded systems that do not have hard disks, such as line cards in a network router or switch.

Diskless server can contain both disk-based tables (D-tables) and in-memory tables (M-tables). D-tables typically take less space than M-tables. Because a diskless server does not use any disk storage space, it maintains all D-tables in the database cache. The database cache size (**IndexFile.CacheSize**) has to be configured to contain the whole database.

There are two ways to run a diskless server: as a single server (alone) or as a replica in an advanced replication system. In both cases, you need to start the server by using the SSC API or SSC API for Java function calls.

SSC API

- Use `SSCStartSMADisklessServer` to start a diskless SMA server.
- Use `SSCStartDisklessServer` to start a diskless LLA server.

SSC API for Java

- Use `startDisklessServer` to start a diskless LLA server.

Diskless server alone

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

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 HotStandby feature to create a 'hot standby' machine that contains a copy of the data. For more information about the HotStandby capabilities, see *IBM solidDB High Availability User Guide*.

Diskless server as part of advanced replication systems

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

Related information:

C.2, "IndexFile.CacheSize parameter in diskless servers," on page 64

8 Creating and running remote or dual-mode applications

The SMA and LLA server can be accessed by remote applications. If your remote application includes SSC API or SA API function calls, you must link to separate SSC API and SA API libraries. The SSC API and SA API libraries are needed because the functions included in the SMA and LLA libraries cannot be accessed by remote applications. If your remote application uses only ODBC or JDBC, you can build you applications normally using the ODBC and JDBC interfaces. The remote connection type is defined in the connect string.

8.1 Example: Creating a dual-mode LLA application with ODBC and SSC API function calls

If your application is a dual-mode application that uses SSC API and ODBC function calls, you need two different executable programs, one to be run locally and one to be run remotely.

Procedure

1. Create the application version that runs in local mode.
 - a. Link the application to the LLA library (for example, `solidimpac.lib` for Windows).

The LLA library provides support for both the ODBC functions and the SSC API functions.
 - b. Modify the connection string to use local connection.
2. Create the application version that runs in remote mode.
 - a. Link the application to both the `solidDB` ODBC driver and to the SSC API stub library (for example, `solidctrlstub.lib` for Windows).

The 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".
 - b. Modify the connection string to use remote connection.

8.2 Establishing remote connections

When you establish a remote connection, the application calls to the server go through the network rather than use the direct function calls to the SMA or LLA library.

ODBC API

With the ODBC API, to establish a remote connection, the application calls the `SQLConnect` function with the name of the remote server.

Example

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);
```

SA API

With the SA API, to establish a remote connection, the application calls the SaConnect function with the name of the remote server.

Example

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");
```

JDBC

With JDBC, to establish a remote connection, the name of the remote server is defined in the connect string.

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

Appendix A. Shared memory access parameters

Server-side parameters

Table 19. Shared memory access parameters

[SharedMemoryAccess]	Description	Factory value	Startup
MaxSharedMemorySize	<p>This parameter sets the maximum total size of the shared memory area used by solidDB.</p> <p>If the SMA server tries to allocate more, an "out of memory" error occurs. With value "0", the maximum value is set automatically to be the size of the physical memory of the computer (platform specific).</p> <p>Note: The value set with the SharedMemoryAccess.MaxSharedMemorySize parameter takes precedence over the value set with any corresponding kernel parameter (for example, SHMALL in Linux environments). Thus, the value set with the SharedMemoryAccess.MaxSharedMemorySize parameter must not be higher than the value set with the corresponding kernel parameter.</p> <p>If you set the SharedMemoryAccess.MaxSharedMemorySize parameter, do not use the Srv.ProcessMemoryLimit parameter.</p>	<p>0 (automatic)</p> <p>Unit: 1 byte, G=GB, M=MB, K=KB</p>	RW
SharedMemoryAccessRights	<p>This parameter sets a validation context for the user access to the shared memory area.</p> <p>The validation context is modeled after a traditional file validation mask. The possible values are:</p> <ul style="list-style-type: none"> • user – access is granted only to the same user as the one that started the SMA server • group – access is granted to any user belonging to the same group as the one that started the SMA server • all – access is granted to all users 	group	RW

Client-side parameters

Table 20. Shared memory access parameters (client-side)

[SharedMemoryAccess]	Description	Factory value	Startup
SignalHandler	<p>The SignalHandler parameter controls the SMA signal handler functionality.</p> <p>When set to yes, the SMA driver signal handler handles the signals defined with the Signals parameter.</p> <p>The SMA driver signal handler enables the SMA system to survive the most common application failures, such as killing or interrupting the applications from outside, or when one of the application threads runs within the server code, and another thread running application code causes application to crash.</p> <p>Upon the capture of certain signals, the signal handler closes the SMA connections safely and exits the SMA application. In most cases, the SMA server continues to run despite abnormal application exits.</p> <p>The SMA driver signal handler installs itself when the first SMA connection is established and uninstalls itself when the last SMA connection is closed. Previously installed signal handlers are retained.</p>	yes	NA
Signals	<p>This parameter defines the signals that can break the SMA connection and is handled by the SMA driver.</p> <p>The signals are defined as integers or with the following mnemonics: SIGSTOP, SIGKILL, SIGINT, SIGTERM, SIGQUIT, SIGABORT.</p> <p>Note: If the SMA application loops outside of the SMA driver (for example, does not call any functions), the signal can fail to terminate the application. In such a case:</p> <ol style="list-style-type: none"> 1. Throw out the connections at the server. admin command 'throwout <userid>' 2. Use SIGKILL signal to force the SMA application to exit. kill -SIGKILL <pid> 	<p>Linux and UNIX: SIGINT, SIGTERM</p> <p>Windows: SIGINT</p>	NA

Appendix B. Linked library access parameters

Linked library access (LLA) parameters appear in the [Accelerator] section of the solid.ini configuration file.

Table 21. Accelerator parameters

[Accelerator]	Description	Factory Value	Access Mode
ImplicitStart	If set to yes, solidDB starts automatically as soon as the ODBC API function SQLConnect is called in a user application. If set to no, solidDB must be explicitly started with a call to the SSC API function SSCStartServer.	yes	RW/ Startup
ReturnListenErrors	If this parameter is set to yes and network listening fails, the SSCStartServer function returns an error. If this parameter is set to no and network listening fails, the SSCStartServer starts the LLA server but network connections are not possible.	no	RW/ Startup

Related reference:

D.2.9, "SSCStartServer," on page 79

The SSCStartServer functions starts the linked library access server.

Appendix C. Configuration parameters for a diskless server

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

In the **IndexFile** section, the **FileSpec** and **CacheSize** parameters have specific settings for a diskless server.

If you are using the diskless server as an advanced replication replica server, the **Listen** parameter in the **Com** section affects the communication between the master and the diskless replica server.

C.1 IndexFile.FileSpec_[1...n] parameter in diskless servers

The **IndexFile.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 **IndexFile.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. The maximum file size must be smaller than the cache size, which is set with the **IndexFile.CacheSize** parameter.

The default value for the **IndexFile.FileSpec** parameter is `solid.db, 2147483647` bytes (2 GB -1). For example:

```
FileSpec_1=solid.db 2147483647
```

Note: If you specify multiple files, the maximum file size setting is the sum of all the **IndexFile.FileSpec** parameter settings.

The maximum size is limited by the physical memory available. A diskless machine has no disk to use as swap space for virtual memory.

Note: On some platforms, the amount of physical memory available to the applications can be less than the amount of physical memory in the machine.

For example, in some versions of Linux 32-bit systems, the amount of memory available to applications is limited to one half or one quarter of the theoretical address space (4 GB) 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:

```
11003,System,Fatal Error,File write failed, configuration exceeded.  
Writing to the database file failed because the maximum database file size  
set with IndexFile.FileSpec parameter has been exceeded. Increase the maximum  
file size limit or divide the database into multiple files.
```

C.2 IndexFile.CacheSize parameter in diskless servers

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

For example:

```
[IndexFile]
CacheSize=10000000
```

With diskless servers, the cache size depends on the following criteria:

- Because a diskless server does not use any disk storage space, it maintains all D-tables in the database cache. The database cache size (**IndexFile.CacheSize**) has to be configured to contain the whole database. Thus, if the database on a diskless server contains mainly disk-based tables (D-tables), the cache size (in bytes) needs to be at least 20% larger than the maximum file size (that is, the amount of data) set with the **IndexFile.FileSpec** parameter. The 20% buffer is an estimate that can vary depending on the usage of the database.

For example:

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

- If the database on a diskless server contains mainly in-memory tables (M-tables), a database cache still exists: it holds the system tables. With M-tables, 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.

You can estimate the total memory used by the diskless server by using the following formula. The total memory must fit within the amount of physical memory available, which means that the cache size must be 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. The amount of memory available to the server is less than the total physical memory, since, for example, the operating system uses some of the physical memory.

C.3 Com.Listen in diskless servers

If you are using the diskless server as an advanced replication replica server, the **Listen** parameter in the **Com** section affects the communication between the master and the diskless replica server. The **Com.Listen** parameter defines the protocol and the name that the diskless server uses when it starts listening to the network.

For example:

```
[Com]
Listen=tcpip 2315
```

The default value of the **Com.Listen** parameter is tcp 1964.

For more information about the network names and protocols, see section *Managing network connections* in *IBM solidDB Administrator Guide*.

C.4 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 22. 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>	
LogEnabled	This parameter is disabled since transaction logging is always disabled for diskless servers. 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 roll-forward recovery.

Appendix D. solidDB Server Control API (SSC API)

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

Note: Some information about the functions applies also to SSC API for Java. For more information about SSC API for Java, see section Appendix E, “SolidServerControl class interface,” on page 89.

D.1 Summary of SSC API functions

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

Table 23. Summary of control API functions

Function	Description	Supported in	For more details, see
SSCStartSmaServer	Starts a SMA server.	SMA	See D.2.11, “SSCStartSMAServer,” on page 83.
SSCStartSMADisklessServer	Starts a diskless SMA server.	SMA	See D.2.10, “SSCStartSMADisklessServer,” on page 82.
SSCStartServer	Starts an LLA server.	LLA and SSC API stub library	See D.2.9, “SSCStartServer,” on page 79.
SSCStartDisklessServer	Starts a diskless LLA server.	LLA and SSC API stub library	See D.2.8, “SSCStartDisklessServer,” on page 77.

Table 23. Summary of control API functions (continued)

Function	Description	Supported in	For more details, see
SSCSetState	<p>Sets the state of a solidDB server (for example, SSC_STATE_OPEN indicates that subsequent connections are allowed). Setting the state to ~SSC_STATE_OPEN will block LLA connections and remote network connections.</p> <p>The following flag symbols can be used:</p> <ul style="list-style-type: none"> • SSC_STATE_OPEN – the open flag is set to 1: new connections are allowed. • SSC_STATE_CLOSED – the open flag is set to 0: all new network and LLA connections are rejected, except for connections from solidDB Remote Control (solcon) program. • SSC_DISABLE_NETCOPY – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which SSC_DISABLE_NETCOPY is set. <p>The flag does not prevent the server to act as a source of netcopy. If only SSC_DISABLE_NETCOPY flag is set, the server is in the closed state. To enable netcopy, use the SSC API function SSCSetState() with the runflag value SSC_STATE_OPEN or SSC_STATE_CLOSED.</p>	LLA and SSC API stub library	See D.2.7, "SSCSetState," on page 76.
SSCRegisterThread - deprecated as of 6.5 FP1	Registers a linked library access application thread for the server. Registration is required in every thread in the user application before any LLA API function can be called.	LLA and SSC API stub library	See D.2.5, "SSCRegisterThread," on page 73.
SSCUnregisterThread - deprecated as of 6.5 FP1	Unregisters an LLA application thread for the server. Registration removal is required in every thread that is registered before terminating.	LLA and SSC API stub library	See D.2.13, "SSCUnregisterThread," on page 86.
SSCStopServer	Stops SMA or LLA server.	SMA, LLA, and SSC API stub library	See D.2.12, "SSCStopServer," on page 85.
SSCSetNotifier	Specifies a user-defined function which solidDB calls at a specified event, such as merge, backup, or shutdown.	LLA and SSC API stub library	See D.2.6, "SSCSetNotifier," on page 74.

Table 23. Summary of control API functions (continued)

Function	Description	Supported in	For more details, see
SSCIsRunning	Returns non-zero if the server is running.	LLA and SSC API stub library	See D.2.3, "SSCIsRunning," on page 72.
SSCIsThisLocalServer	Indicates whether the application is linked to the solidDB server with the LLA or the "dummy" (<code>solidctrlstub</code>) library to test solidDB remote applications using the SSC API.	LLA and SSC API stub library	See D.2.4, "SSCIsThisLocalServer," on page 73.
SSCGetServerHandle	Returns the solidDB server handle if the server is running.	LLA and SSC API stub library	See D.2.1, "SSCGetServerHandle," on page 71.
SSCGetStatusNum	Gets solidDB status information.	LLA and SSC API stub library	See D.2.2, "SSCGetStatusNum," on page 72.

D.2 SSC API reference

The SSC API reference describes each SSC API function in alphabetic order. Each description includes the purpose, synopsis, parameters, return value, and comments.

- "Function synopsis"
- "Parameters"
- "Return values" on page 70
- "SSC API error codes and messages" on page 71

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.

Parameters

In each function description, parameters are described in a table format. The table includes 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 SSC 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 24. SSC 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 used during the function call. The caller can do whatever it wants with the parameter after the function call. use 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.
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 values

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 `sscap.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)

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

SSC API error codes and messages

SSC API functions may return the error codes and messages listed in the table below.

These constants are defined in the `sscap.h` file.

Table 25. Error codes and messages for SSC API functions

Error Code/Message	Description
SSC_SUCCESS	Operation is successful.
SSC_ERROR	Generic error.
SSC_ABORT	Operation cancelled.
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 HotStandby).

D.2.1 SSCGetServerHandle

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

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.

D.2.2 SSCGetStatusNum

SSCGetStatusNum returns the status information of the solidDB server.

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

The SSCGetStatusNum function accepts the following parameters:

Table 26. SSCGetStatusNum parameters

Parameters	Usage Type	Description
h	in, use	Handle to server.
stat	in	Specifies the status identifier for retrieval.
num	out	Upon successful return of the function, the value of the parameter value is 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 an unrecognized value for the stat parameter, the function returns 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.

D.2.3 SSCIsRunning

If the server is running, SSCIsRunning returns a non-zero value.

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

The SSCIsRunning function accepts the following parameters:

Table 27. *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.

D.2.4 SSCIsThisLocalServer

The `SSCIsThisLocalServer` function indicates whether the application is linked to a solidDB server or the `solidctrlstub` library. You can use the `solidctrlstub` library to test remote applications using SSC API without linking the linked library access library and modifying the source code.

```
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.

D.2.5 SSCRegisterThread

Note: `SSCRegisterThread` is deprecated as of 6.5 FP1; it is no longer necessary to register and unregister threads explicitly when using solidDB with linked library access (LLA). As of 6.5 FP1, thread registration is handled implicitly.

`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 `SSCRegisterThread` function accepts the following parameters:

Table 28. *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

D.2.13, “SSCUnregisterThread,” on page 86

D.2.6 SSCSetNotifier

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

You can use the SSCSetNotifier function to verify that the solidDB server calls a specified user-defined function whenever a special event occurs. The function detects the following events:

- Server shutdown
- Bonsai merge from the index to the storage tree
- Bonsai merge interval maximum
- Backup or checkpoint request
- Idle server state
- Netcopy request received from the Primary server
- Completion of a netcopy request, which occurs when the server is started with the new database received through netcopy

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

The SSCSetNotifier function accepts the following parameters:

Table 29. *SSCSetNotifier* parameters

Parameters	Usage Type	Description
h	in	Handle to the server

Table 29. *SSCSetNotifier* parameters (continued)

Parameters	Usage Type	Description
<i>what</i>	in	<p>Specifies the event for notification.</p> <p>The options are:</p> <ul style="list-style-type: none"> • <code>SSC_NOTIFY_EMERGENCY_EXIT</code> This function is called if a server shuts down abnormally after it has been activated with <code>SSCStartServer()</code>. The notifier call <code>SSCSetNotifier()</code> must 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 might shut down if the user-defined function accepts the request. You can refuse the shutdown by returning <code>SSC_ABORT</code> from the notified function, or proceed with the request by returning <code>SSC_SUCCESS</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 General.MergeInterval 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 is received from the Primary server. • <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 with the new database received through the network copy (netcopy). <code>SSC_NOTIFY_NETCOPY_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>Note: Check the warning on releasing parameters of usage type <i>hold</i> in section "Parameters" on page 69.</p>

Return value

- `SSC_SUCCESS` - Request from the server accepted.

HotStandby only:

If `SSC_NOTIFY_NETCOPY_FINISHED` returns `SSC_SUCCESS`, all other application connections are terminated and the server is set to netcopy listening 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.

- `SSC_ABORT` - Request from the server denied.
HotStandby only:
If the `SSC_NOTIFY_NETCOPY_REQUEST` returns `SSC_ABORT`, the netcopy is not started and the error `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.

The user-defined notifier functions cannot call any SA, SSC, or ODBC functions.

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

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

After you have used `SSC_CALL` to explicitly define the convention for your user function, 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 that a user creates the function `user_own_shutdownrequest`, which is called every time a shutdown is requested:

```
int SSC_CALL user_own_shutdownrequest(SscServerT h, SscNotFunT what, void
*userdata);
{
  if (shutdown not needed) {
    return SSC_ABORT;
  }
  return SSC_SUCCESS; /*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 shutdown is not allowed. If the function returns `SSC_SUCCESS`, the shutdown can proceed.

D.2.7 SSCSetState

The `SSCSetState` functions controls whether the server accepts subsequent connections.

The `SSCSetState` functions sets the state of an LLA or SMA server to open or closed. If the server is set to open", the server accepts connections. If the server is set to closed, it does not accept any further connections (local or remote). However, any connections that have already been made are allowed to continue.

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

The `SSCSetState` function accepts the following parameters:

Table 30. SSCSetState parameters

Parameter	Usage Type	Description
<i>h</i>	in, use	Handle to the server.
<i>runflags</i>	in	<p>The value for this parameter is a combination of 2 flags: the open flag and the netcopy disabling flag. The following flag symbols can be used:</p> <ul style="list-style-type: none"> SSC_STATE_OPEN – the open flag is set to 1: new connections are allowed. SSC_STATE_CLOSED – the open flag is set to 0: all new network, LLA, and SMA connections are rejected, except for connections from solidDB Remote Control (solcon) program. SSC_DISABLE_NETCOPY – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which SSC_DISABLE_NETCOPY is set. <p>The flag does not prevent the server to act as a source of netcopy. If only SSC_DISABLE_NETCOPY flag is set, the server is in the closed state. To enable netcopy, use the SSC API function SSCSetState() with the runflag value SSC_STATE_OPEN or SSC_STATE_CLOSED.</p> <pre>runflags = SSC_STATE_OPEN SSC_STATE_CLOSED SSC_DISABLE_NETCOPY</pre> <p>Tip: The flags can be used in combinations, for example:</p> <pre>... rc = SSCStartServer(g_argc, g_argv, &hh, SSC_STATE_OPEN SSC_DISABLE_NETCOPY); ...</pre> <p>If the server is started as closed, it can be opened with ADMIN COMMAND 'open', or with the solcon command open. The same effect can be achieved with the SSC API function SSCSetState().</p>

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'

D.2.8 SSCStartDisklessServer

The SSCStartDisklessServer function starts a diskless server using the linked library access.

```
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 31. *SSCStartDisklessServer* parameters

Parameters	Usage Type	Description
<i>argc</i>	in	The number of command-line arguments.
<i>argv</i>	in, use	<p>Array of command-line arguments that are used during the function call. The argument <code>argv[0]</code> is reserved only for the path and filename of the user application and must be present.</p> <p>For a list of available arguments, see section <i>solidDB command-line options</i> in the <i>IBM solidDB Administrator Guide</i>.</p>
<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	<p>The value for this parameter is a combination of 2 flags: the open flag and the netcopy disabling flag. The following flag symbols can be used:</p> <ul style="list-style-type: none"> • <code>SSC_STATE_OPEN</code> – the open flag is set to 1: new connections are allowed. • <code>SSC_STATE_CLOSED</code> – the open flag is set to 0: all new network and LLA connections are rejected, except for connections from solidDB Remote Control (solcon) program. • <code>SSC_DISABLE_NETCOPY</code> – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which the <code>SSC_DISABLE_NETCOPY</code> is set. <p>The flag does not prevent the server to act as a source of netcopy. If only <code>SSC_DISABLE_NETCOPY</code> flag is set, the server is in the closed state. To enable netcopy, use the SSC API function <code>SSCSetState()</code> with the runflag value <code>SSC_STATE_OPEN</code> or <code>SSC_STATE_CLOSED</code>.</p> <p><code>runflags = SSC_STATE_OPEN SSC_STATE_CLOSED SSC_DISABLE_NETCOPY</code></p> <p>Tip: The flags can be used in combinations, for example:</p> <pre>... rc = SSCStartServer(g_argc, g_argv, &hh, SSC_STATE_OPEN SSC_DISABLE_NETCOPY); ...</pre> <p>If the server is started as closed, it can be opened with ADMIN COMMAND 'open', or with the solcon command open. The same effect can be achieved with the SSC API function <code>SSCSetState()</code>.</p>
<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.

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 the 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, the server returns an error.

Related concepts:

7, "Using the diskless capability," on page 55

SMA and LLA servers can be used to create a database engine that runs without any disk storage space. Diskless server is useful in embedded systems that do not have hard disks, such as line cards in a network router or switch.

Related reference:

D.2.12, "SSCStopServer," on page 85

The `SSCStopServer` function stops a linked library access (LLA) server or a shared memory access (SMA) server.

D.2.9 `SSCStartServer`

The `SSCStartServer` functions starts the linked library access server.

In multithreaded environments, the server runs in a separate threads from the client. For the duration of the application, the application can start or stop the server subroutines as needed.

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

The `SSCStartServer` function accepts the parameters described in the following table.

Table 32. `SSCStartServer` parameters

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

Table 32. SSCStartServer parameters (continued)

Parameters	Usage Type	Description
<i>argv</i>	in, use	Array of command-line arguments. For a list of available arguments, see section <i>solidDB command-line options</i> in the <i>IBM solidDB Administrator Guide</i> .
<i>h</i>	out	Returns a handle to the started server. This handle is needed when referencing the server with other Control API functions. If the server is started successfully, the SSCStartServer routine sets this parameter to point to the handle for this server.
<i>runflags</i>	in	<p>The value for this parameter is a combination of 2 flags: the open flag and the netcopy disabling flag. The following flag symbols can be used:</p> <ul style="list-style-type: none"> SSC_STATE_OPEN – the open flag is set to 1: new connections are allowed. SSC_STATE_CLOSED – the open flag is set to 0: all new network and LLA connections are rejected, except for connections from solidDB Remote Control (solcon) program. SSC_DISABLE_NETCOPY – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which the SSC_DISABLE_NETCOPY is set. <p>The flag does not prevent the server to act as a source of netcopy. If only SSC_DISABLE_NETCOPY flag is set, the server is in the closed state. To enable netcopy, use the SSC API function SSCSetState() with the runflag value SSC_STATE_OPEN or SSC_STATE_CLOSED.</p> <pre>runflags = SSC_STATE_OPEN SSC_STATE_CLOSED SSC_DISABLE_NETCOPY</pre> <p>Tip: The flags can be used in combinations, for example:</p> <pre>... rc = SSCStartServer(g_argc, g_argv, &hh, SSC_STATE_OPEN SSC_DISABLE_NETCOPY); ...</pre> <p>If the server is started as closed, it can be opened with ADMIN COMMAND 'open', or with the solcon command open. The same effect can be achieved with the SSC API function SSCSetState().</p>

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 the solidDB server creates a database with the function SSCStartServer() by using the **-U** *username* **-P** *password* **-C** *catalogname* (the default database catalog name) parameters. For details, read 5.3.1, “Explicit startup with SSC API function SSCStartServer,” on page 42.

If you are restarting a database server (a database exists in the directory), the server uses the existing database.

The SSCStartServer function can 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, and so on.

Note: If you want to start a diskless server, use the function D.2.8, “SSCStartDisklessServer,” on page 77.

Example: Starting up SSCStartServer

Start SSCStartServer with the server name, the catalog name, and the administrator username and password:

```

SscStateT runflags = SSC_STATE_OPEN;
SscServerT h;
char* argv[5];
argv[0] = "apname"; /* path and filename of the user app. */
argv[1] = "-nsolid1";
argv[2] = "-Udba";
argv[3] = "-Pdba";
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.

Related reference:

D.2.12, “SSCStopServer,” on page 85

The SSCStopServer function stops a linked library access (LLA) server or a shared memory access (SMA) server.

D.2.10 SSCStartSMADisklessServer

The SSCStartSMADisklessServer function starts a diskless server using SMA.

Synopsis

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

The SSCStartSMADisklessServer function accepts the following parameters:

Table 33. SSCStartSMADisklessServer 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 <i>argv</i> [0] is reserved only for the path and filename of the user application and must be present. For a list of available arguments, see section <i>solidDB command-line options</i> in the <i>IBM solidDB Administrator Guide</i> .
<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 value for this parameter is a combination of 2 flags: the open flag and the netcopy disabling flag. The following flag symbols can be used: <ul style="list-style-type: none"> SSC_STATE_OPEN – the open flag is set to 1: new connections are allowed. SSC_STATE_CLOSED – the open flag is set to 0: all new network and SMA connections are rejected, except for connections from solidDB Remote Control (solcon) program. SSC_DISABLE_NETCOPY – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which SSC_DISABLE_NETCOPY is set. <p>The flag does not prevent the server to act as a source of netcopy. If only SSC_DISABLE_NETCOPY flag is set, the server is in the closed state. To enable netcopy, use the SSC API function SSCSetState() with the runflag value SSC_STATE_OPEN or SSC_STATE_CLOSED.</p> <pre>runflags = SSC_STATE_OPEN SSC_STATE_CLOSED SSC_DISABLE_NETCOPY</pre> <p>Tip: The flags can be used in combinations, for example:</p> <pre>... rc = SSCStartSMAServer(g_argc, g_argv, &hh, SSC_STATE_OPEN SSC_DISABLE_NETCOPY); ...</pre> <p>If the server is started as closed, it can be opened with ADMIN COMMAND 'open', or with the solcon command open. The same effect can be achieved with the SSC API function SSCSetState().</p>

Table 33. *SSCStartSMADisklessServer* parameters (continued)

Parameters	Usage Type	Description
<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.

Return values

- 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.

See also

D.2.12, “SSCStopServer,” on page 85

D.2.11 SSCStartSMAServer

The SSCStartSMAServer function starts a server using SMA.

Synopsis

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

The SSCStartSMAServer function accepts the following parameters:

Table 34. *SSCStartSMAServer* 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 available arguments, see section <i>solidDB command-line options</i> in the <i>IBM solidDB Administrator Guide</i> .
<i>h</i>	out	Returns a handle to the started server. This handle is needed when referencing the server with other Control API functions.

Table 34. SSCStartSMAServer parameters (continued)

Parameters	Usage Type	Description
<i>runflags</i>	in	<p>The value for this parameter is a combination of 2 flags: the open flag and the netcopy disabling flag. The following flag symbols can be used:</p> <ul style="list-style-type: none"> • SSC_STATE_OPEN – the open flag is set to 1: new connections are allowed. • SSC_STATE_CLOSED – the open flag is set to 0: all new network and SMA connections are rejected, except for connections from solidDB Remote Control (solcon) program. • SSC_DISABLE_NETCOPY – the netcopy disabling flag is set to 1: in HotStandby configuration, no netcopy can be received by the server for which SSC_DISABLE_NETCOPY is set. <p>The flag does not prevent the server to act as a source of netcopy. If only SSC_DISABLE_NETCOPY flag is set, the server is in the closed state. To enable netcopy, use the SSC API function SSCSetState() with the runflag value SSC_STATE_OPEN or SSC_STATE_CLOSED.</p> <pre>runflags = SSC_STATE_OPEN SSC_STATE_CLOSED SSC_DISABLE_NETCOPY</pre> <p>Tip: The flags can be used in combinations, for example:</p> <pre>... rc = SSCStartSMAServer(g_argc, g_argv, &hh, SSC_STATE_OPEN SSC_DISABLE_NETCOPY); ...</pre> <p>If the server is started as closed, it can be opened with ADMIN COMMAND 'open', or with the solcon command open. The same effect can be achieved with the SSC API function SSCSetState().</p>

Return values

- 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 database with the function SSCStartSMAServer() with the **-U username -P password -C catalogname** (the default database catalog name) parameters.

If you are restarting a database server (a database exists in the directory), then SSCStartSMAServer uses the existing database.

The SSCStartSMAServer 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, and so on.

See also

D.2.12, “SSCStopServer”

D.2.12 SSCStopServer

The SSCStopServer function stops a linked library access (LLA) server or a shared memory access (SMA) server.

You can use explicit methods (for example, SSCStopServer) to shut down a server that was started with implicit methods (for example, 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 restart the server by using, for example, SSCStartServer.

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

The SSCStopServer function accepts the following parameters:

Table 35. SSCStopServer parameters

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

Table 35. 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'**. For more information, see Appendix B, “Linked library access parameters,” on page 61.

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

The SSCSetState() function with the SSC_STATE_OPEN option prevents new connections to the solidDB server.

Related reference:

D.2.7, “SSCSetState,” on page 76

The SSCSetState functions controls whether the server accepts subsequent connections.

D.2.9, “SSCStartServer,” on page 79

The SSCStartServer functions starts the linked library access server.

D.2.11, “SSCStartSMAServer,” on page 83

The SSCStartSMAServer function starts a server using SMA.

D.2.13 SSCUnregisterThread

Note: SSCUnregisterThread is deprecated as of 6.5 FP1; it is no longer necessary to register and unregister threads explicitly when using solidDB with linked library access (LLA). As of 6.5 FP1, thread registration is handled implicitly.

The SSCUnregisterThread function 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 36. *SCCUnregisterThread* 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 ADMIN COMMAND.

See also

D.2.5, "SSCRegisterThread," on page 73

D.2.14 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.

D.2.15 Obtaining solidDB status and server information

You can use the function `SSCGetStatusNum` to view 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. Use 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.

Appendix E. SolidServerControl class interface

The complete public interface for the SolidServerControl class is described below.

Details on parameters and corresponding ADMIN COMMANDs are included in the SSC function descriptions in section D.2, "SSC API reference," on page 69.

For an example of an LLA program that uses some of the methods in this class, see the LLA for Java sample in the samples/aclib_java directory.

Return value 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;

/**
 * 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();
```

```

/**
 * startDisklessServer starts the Linked Library Access server
 * in a diskless mode
 *
 * @param argv      parameter vector for the accelerator server.
 *
 * @param runflags  Options for this parameter are SSC_STATE_OPEN,
 *                  SSC_STATE_CLOSED, and SSC_DISABLE_NETCOPY.
 *
 * @param lic_file  The contents of the license file as a string,
 *                  since the diskless version cannot read the
 *                  information from the disk
 *
 * @param ini_file  The contents of the solid.ini configuration file as a string,
 *                  because the diskless version cannot read the information
 *                  from the disk
 *
 * @return the return value from the server :
 *         SSC_SUCCESS
 *         SSC_ERROR
 *         SSC_INVALID_LICENSE - No license or license file found.
 *         SSC_NODATABASEFILE - No database file found.
 */
public static long startDisklessServer( String[] argv, long runflags,
String lic_string, String ini_string )

/**
 * startServer starts the Linked Library Access server
 *
 * @param argv      parameter vector for the LLA server
 *                  Note: You must give the working directory containing
 *                  license file (f.ex. -c\tmp) first, in front
 *                  of other parameters.
 *
 * @param runflags  Options for this parameter are SSC_STATE_OPEN,
 *                  SSC_STATE_CLOSED, and SSC_DISABLE_NETCOPY.
 * *
 * @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 LLA server
 *
 * @param runflags  Runflags for stopping LLA server.
 *                  See section SSCStopServer 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: server is running or not

```

```

*
* @return SSC_STATE_OPEN if server is up and running
*/
public int getState();

/**
* registerThread registers this user thread to Linked Library Access server
* - deprecated as of V6.5 Fix Pack 1
*
*
* @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
* Linked Library Access server - deprecated as of V6.5 Fix Pack 1
*
*
* @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();

```

Index

A

- administering
 - diskless server configuration file options 63
- advanced replication
 - linked library access (LLA) 46

B

- backup
 - listening mode 74

C

- C applications
 - samples 46
- CacheSize (parameter) 64
 - diskless server 64
- configuration file
 - diskless 63
- Control API
 - SSCGetActiveTaskClass (function) 87
 - SSCGetServerHandle (function) 87
 - SSCGetStatusNum (function) 87
 - SSCGetTaskClassState (function) 87
 - SSCIsRunning (function) 87
 - SSCIsThisLocalServer (function) 87

D

- database
 - size 42
- diskless server mode
 - overview 55
- dual-mode applications 11, 57

I

- ImplicitStart (parameter) 61

J

- JDBC API 10

L

- linked library access (LLA) 1
 - components 6
 - definition 6
 - shutting down 45
 - starting 41
 - supported platforms 4, 6
- Linux
 - memory limitations 63
- Listen (parameter)
 - configuring for diskless 65
- local application
 - definition 11

M

- MaxSharedMemorySize (parameter) 59
- memory
 - CacheSize (for diskless server) 64
 - total used by diskless server 64

N

- netcopy
 - listening mode 74

O

- ODBC
 - overview 10

R

- remote applications 57
 - definition 11
- ReturnListenErrors (parameter) 61

S

- SA API
 - definition 9
- SaConnect
 - implicit startup with 45
- server information
 - retrieving 87
- shared memory access (SMA) 1
 - components 4
 - configuring 13
 - definition 3
 - monitoring 26
 - troubleshooting 27
- SharedMemoryAccessRights (parameter) 59
- shutdown
 - linked library access 45
- SignalHandler (parameter) 59
- Signals (parameter) 59
- SMA server
 - starting 24, 33
- SMA system parameters
 - AIX 15
 - overview 14
- solidctrlstub 10, 69
- solidDB client APIs and drivers 9
- solidDB configuration file
 - Listen (parameter) 65
- solidDB drivers and client APIs 9
- solidDB SA 9
- solidDB Server Control (SSC) API for Java 10, 89
- solidDB Server Control API (SSC API)
 - definition 10
 - solidctrlstub 10
- solidimpac 8
- SolidServerControl class 89

- SQLConnect (function)
 - implicit startup with 44
- SSC API (Control API) 67
 - definition 10
 - summary of scheduling functions 67
- SSC API for Java 10, 89
- SSC_ABORT 71
- SSC_CALL 69
- SSC_CONNECTIONS_EXIST 71
- SSC_CONT 71
- SSC_ERROR 71
- SSC_FINISHED 71
- SSC_INFO_SERVER_RUNNING 71
- SSC_INVALID_HANDLE 71
- SSC_INVALID_LICENSE 71
- SSC_NODATABASEFILE 71
- SSC_SERVER_INNETCOPYMODE 71
- SSC_SERVER_NOTRUNNING 71
- SSC_STATE_OPEN 78, 82, 84
- SSC_SUCCESS 71
- SSC_TASK_ALL 69
- SSC_TASK_BACKUP 69
- SSC_TASK_CHECKPOINT 69
- SSC_TASK_HOTSTANDBY 69
- SSC_TASK_HOTSTANDBY_CATCHUP 69
- SSC_TASK_LOCALUSERS 69
- SSC_TASK_MERGE 69
- SSC_TASK_NONE 69
- SSC_TASK_REMOTEUSERS 69
- SSC_TASK_SYNC_HISTCLEAN 69
- SSC_TASK_SYNC_MESSAGE 69
- SSC_UNFINISHED_TASKS 71
- sscapi.h 69
- SSCGetServerHandle
 - function description 71
- SSCGetStatusNum
 - function description 72
- SSCIsRunning
 - function description 72
- SSCIsThisLocalServer
 - function description 73
- SSCRegisterThread
 - function description 73
- SSCServerT 42
- SSCSetNotifier
 - function description 74
- SSCSetState
 - function description 76
- SSCStartDisklessServer
 - function description 77
- SSCStartDisklessSMAServer 83
- SSCStartServer
 - explicit startup with 42
 - function description 79
- SSCStartSMADisklessServer 82
- SSCStopServer
 - function description 85
 - shut down with 45
- SscTaskSetT 69
- SSCUnregisterThread
 - function description 86
- starting
 - solidDB
 - with linked library access 41
- status
 - retrieving 87

T
 task information
 retrieving 87

Notices

© Copyright Oy IBM Finland Ab 1993, 2013.

All rights reserved.

No portion of this product may be used in any way except as expressly authorized in writing by IBM.

This product is protected by U.S. patents 6144941, 7136912, 6970876, 7139775, 6978396, 7266702, 7406489, 7502796, and 7587429.

This product is assigned the U.S. Export Control Classification Number ECCN=5D992b.

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, ibm.com[®], Solid, solidDB, InfoSphere, DB2[®], Informix[®], and WebSphere[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.



SC27-3846-05

