



IBM solidDB Getting Started Guide

Version 6.1 | June 2008

IBM solidDB Getting Started Guide

Copyright © Solid Information Technology Ltd. 1993, 2008

Document number: SEGS61

Product version: 06.10.0014

Date: 2008-06-13

All rights reserved. No portion of this product may be used in any way except as expressly authorized in writing by Solid Information Technology Ltd. or International Business Machines Corporation.

"IBM", the IBM logo, "DB2", "Informix", "Solid" and "solidDB" are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

All other products, services, companies and publications are trademarks or registered trademarks of their respective owners.

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

This product contains lexical analyzer Flex. Copyright (c) 1990 The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Vern Paxson. The United States Government has rights in this work pursuant to contract no. DE-AC03-76SF00098 between the United States Department of Energy and the University of California. Redistribution and use in source and binary forms are permitted provided that: (1) source distributions retain this entire copyright notice and comment, and (2) distributions including binaries display the following acknowledgement: "This product includes software developed by the University of California, Berkeley and its contributors" in the documentation or other materials provided with the distribution and in all advertising materials mentioning features or use of this software. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

This product contains zlib general purpose compression library version 1.1.4, March 11th, 2002. Copyright (C) 1995-2002 Jean-loup Gailly and Mark Adler.

This software is provided "as-is", without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software. Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions: 1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required. 2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software. 3. This notice may not be removed or altered from any source distribution.

This product contains the Qsort routine in the external sorter, Copyright (c) 1980, 1983, 1990 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
-

-
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by the University of California, Berkeley and its contributors.
 4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product contains the DES cipher (in ECB mode), parts of this code are Copyright (C) 1996 Geoffrey Keating. All rights reserved.

Its use is FREE FOR COMMERCIAL AND NON-COMMERCIAL USE as long as the following conditions are adhered to.

Copyright remains Geoffrey Keating's, and as such any Copyright notices in the code are not to be removed. If this code is used in a product, Geoffrey Keating should be given attribution as the author of the parts used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by Eric Young (eay@mincom.oz.au)

THIS SOFTWARE IS PROVIDED BY GEOFFREY KEATING ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Parts of this code (in particular, the string representing SPtrans below) are Copyright (C) 1995 Eric Young (eay@mincom.oz.au). All rights reserved.

Its use is FREE FOR COMMERCIAL AND NON-COMMERCIAL USE as long as the following conditions are adhered to.

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this code is used in a product, Eric Young should be given attribution as the author of the parts used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by Eric Young (eay@mincom.oz.au)

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

Table of Contents

1 Welcome	1
1.1 About This Guide	1
1.1.1 Organization	1
1.1.2 Audience	2
1.2 Conventions	2
1.2.1 Typographic Conventions	2
1.2.2 Syntax Notation	3
1.3 IBM solidDB Documentation	4
1.4 Accessibility	5
1.4.1 Keyboard Input and Navigation	5
1.4.2 Accessible Display	6
1.4.3 Compatibility with Assistive Technologies	6
1.4.4 Accessible Documentation	6
1.5 Technical Support	6
2 Introduction	7
2.1 System Architecture	7
2.2 IBM solidDB Directory Structure	8
2.3 Prerequisites	9
2.4 What Is in Your IBM solidDB Development Kit	9
2.4.1 Software	9
2.4.2 Samples	9
2.4.3 Documentation	9
3 Preparing for the Evaluation	11
3.1 Running the License Script	11
3.2 IBM solidDB Configuration File	11
4 Starting IBM solidDB and Creating Your First Database	13
4.1 Starting IBM solidDB	13
4.2 Creating a Database Manually	13
4.3 Automatic Creation of a Database (Windows)	14
4.4 Verifying the Database Installation	14
5 Connecting to IBM solidDB	15
5.1 Using IBM solidDB SQL Editor	15
5.1.1 Starting IBM solidDB SQL Editor	15
5.1.2 Connecting to Database	15
5.1.3 Shutting Down IBM solidDB SQL Editor	16
5.1.4 Using Help	16
5.2 Entering Queries	17
5.3 Checking the Database Status	17
6 Stopping and Restarting the Database	19

6.1 Stopping the Database with IBM solidDB SQL Editor	19
6.2 Restarting the Database	19
7 Running Samples	21
8 IBM solidDB Cache	23
9 What Next?	25
Glossary	27
Index	31

List of Figures

5.1 IBM solidDB SQL Editor Connected	16
5.2 Example Output from Solsql	17

List of Tables

1.1 Typographic Conventions	2
1.2 Syntax Notation Conventions	3
2.1 IBM solidDB Directory Structure	8
4.1 Starting IBM solidDB on Different Operating Systems	13

Chapter 1. Welcome

IBM solidDB provides the features you would expect to find in any industrial-strength database server — multithreaded architecture, stored procedures, and so on. IBM solidDB also provides other features, such as "hot standby" capability to provide protection against server failure, advanced replication technology to distribute data among multiple nodes, and in-memory database capability to increase performance.

A solution is also available for using IBM solidDB as a high-speed and low-latency front-end to IBM data servers like DB2 and Informix Dynamic Server (IDS).

IBM solidDB is available for a very wide range of platforms, including platforms such as Windows,¹ Linux, Solaris, HP-UX and other UNIX platforms as well as Real-Time Operating Systems such as VxWorks and QNX.

1.1 About This Guide

IBM solidDB Getting Started Guide introduces IBM solidDB to you. This manual also gives you an overview of the other manuals in this set. We assume that when you start reading this manual you have already successfully installed the IBM solidDB software package containing the database files, database drivers, and the product documentation. If you have not, do so now. For instructions on installation, see the Evaluation Setup Guide.

1.1.1 Organization

This guide contains the following chapters:

- Chapter 2, *Introduction* describes the situation immediately after the IBM solidDB installation. It also provides you with the background information necessary to administer and maintain IBM solidDB in your network environment.
- Chapter 3, *Preparing for the Evaluation* gives you instructions on running the license script for evaluating IBM solidDB. It also briefly discusses the role and contents of the IBM solidDB configuration file, `solid.ini`.
- Chapter 4, *Starting IBM solidDB and Creating Your First Database* explains how to start IBM solidDB and how to create databases with it.
- Chapter 5, *Connecting to IBM solidDB* explains how to use the IBM solidDB SQL editor to connect to the database.

¹For a detailed list of supported Windows™ versions, please refer to IBM solidDB Release Notes.

- Chapter 6, *Stopping and Restarting the Database* explains how to stop and restart IBM solidDB.
- Chapter 7, *Running Samples* provides an introduction on running samples provided in the IBM solidDB distribution package.
- Chapter 8, *IBM solidDB Cache* provides an introduction to IBM solidDB Cache, the high-performance, low-latency database front-end solution for IBM Data Servers.
- Chapter 9, *What Next?* explains how to continue evaluating IBM solidDB.

Glossary

Glossary provides definitions of terms.

1.1.2 Audience

This guide is for Database Administrators commissioning IBM solidDB for evaluation purposes.

1.2 Conventions

1.2.1 Typographic Conventions

This manual uses the following typographic conventions:

Table 1.1. 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.
<code>solid.ini</code>	These fonts indicate file names and path expressions.
<code>SET SYNC MASTER YES; COMMIT WORK;</code>	This font is used for program code and program output. Example SQL statements also use this font.
<code>run.sh</code>	This font is used for sample command lines.
<code>TRIG_COUNT()</code>	This font is used for function names.
<code>java.sql.Connection</code>	This font is used for interface names.
<code>LockHashSize</code>	This font is used for parameter names, function arguments, and Windows registry entries.

Format	Used for
<i>argument</i>	Words emphasised like this indicate information that the user or the application must provide.
<i>IBM solidDB Administration Guide</i>	This style is used for references to other documents, or chapters in the same document. New terms and emphasised issues are also written like this.
File path presentation	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.

1.2.2 Syntax Notation

This manual uses the following syntax notation conventions:

Table 1.2. Syntax Notation Conventions

Format	Used for
<code>INSERT INTO <i>table_name</i></code>	Syntax descriptions are on this font. Replaceable sections are on <i>this</i> font.
<code>solid.ini</code>	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.

1.3 IBM solidDB Documentation

Below is a complete list of documents available for IBM solidDB. IBM solidDB documentation is distributed in an electronic format, usually PDF files and web pages.

- *Release Notes*. This file contains installation instructions and the most up-to-date information about the specific product version. This file (`releasenotes.txt`) is copied onto your system when you install the software.
- *IBM solidDB Getting Started Guide*. This manual gives you an introduction to IBM solidDB.
- *IBM solidDB SQL Guide*. This manual describes the SQL commands that IBM solidDB supports. This manual also describes some of the system tables, system views, system stored procedures, etc. that the engine makes available to you. This manual contains some basic tutorial material on SQL for those readers who are not already familiar with SQL. Note that some specialized material is covered in other manuals. For example, the IBM solidDB "administrative commands" related to the High Availability (HotStandby) component are described in the *IBM solidDB High Availability User Guide*, not the *IBM solidDB SQL Guide*.
- *IBM solidDB Administration Guide*. This guide describes administrative procedures for IBM solidDB servers. This manual includes configuration information. Note that some administrative commands use an SQL-like syntax and are documented in the *IBM solidDB SQL Guide*.
- *IBM solidDB Programmer Guide*. This guide explains in detail how to use features such as IBM solidDB Stored Procedure Language, triggers, events, and sequences. It also describes the interfaces (APIs and drivers) available for accessing IBM solidDB and how to use them with a IBM solidDB database.
- *IBM solidDB In-Memory Database User Guide*. This manual describes how to use the IBM solidDB in-memory database and main memory engine (MME).
- *IBM solidDB Advanced Replication Guide*. This guide describes how to use the IBM solidDB advanced replication technology to synchronize data across multiple database servers.
- *IBM solidDB Linked Library Access User Guide*. Linking the client application directly to the server improves performance by eliminating network communication overhead. This guide describes how to use the linked library access, a database engine library that can be linked directly to the client application.

This manual also explains how to use two proprietary Application Programming Interfaces (APIs). The first API is the IBM solidDB SA interface, a low-level C-language interface that allows you to perform simple single-table operations (such as inserting a row in a table) quickly. The second API is SSC API, which allows your C-language program can control the behavior of the embedded (linked) database server

This manual also explains how to set up a IBM solidDB to run without a disk drive.

- *IBM solidDB High Availability User Guide.* IBM solidDB HotStandby allows your system to maintain an identical copy of the database in a backup server or "secondary server". This secondary database server can continue working if the primary database server fails.
- *IBM solidDB Connector Guide.* This guide explains in detail how to use the IBM solidDB Cache solution. IBM solidDB Cache provides a high-performance, low-latency database front-end solution for IBM Data Servers, namely DB2™ and Informix™. IBM solidDB Cache solution uses a number of in-memory front-end databases to handle high-volume traffic from the applications. The connectors are applications that manage data between the back-end and the front-ends in IBM solidDB Cache.

1.4 Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. The following list specifies the major accessibility features in IBM solidDB products:

- All IBM solidDB functionality is available using the keyboard for navigation instead of the mouse. For more information, see Section 1.4.1, "Keyboard Input and Navigation".
- You can customize the size and color of the fonts on IBM solidDB user interfaces. For more information, see Section 1.4.2, "Accessible Display".
- IBM solidDB products support accessibility applications. For more information, see Section 1.4.3, "Compatibility with Assistive Technologies".
- IBM solidDB documentation is provided in an accessible format. For more information, see Section 1.4.4, "Accessible Documentation".

1.4.1 Keyboard Input and Navigation

Keyboard Input

You can operate IBM solidDB using only the keyboard. You can use keys or key combinations to perform all operations. Standard operating system keystrokes are used for standard operating system operations. Standard operating system keyboard accessibility features are supported by IBM solidDB products.

For more information about using keys or key combinations to perform operations, see the appropriate operating system documentation.

Keyboard Navigation

You can navigate the IBM solidDB user interface using keys or key combinations only.

For more information about using keys or key combinations to navigate the command line interfaces, see the appropriate operating system documentation.

1.4.2 Accessible Display

IBM solidDB products support standard operating system display settings, such as high contrast and font settings.

Font Settings

You can select the color, size, and font for the text using standard operating system operations.

For more information about specifying font settings, see the appropriate operating system documentation.

Non-Dependence on Color

You do not need to distinguish between colors in order to use any of the functions in this product.

1.4.3 Compatibility with Assistive Technologies

IBM solidDB interacts with the operating system through standard APIs which support interaction with assistive technologies, which enables you to use screen readers and other accessibility tools.

1.4.4 Accessible Documentation

Documentation for IBM solidDB is provided in HTML format, which is viewable in most Web browsers. HTML allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.

1.5 Technical Support

If you need technical support in using IBM solidDB, see <http://www.ibm.com/software/data/soliddb/support/>.

Chapter 2. Introduction

This chapter describes the situation immediately after the IBM solidDB installation. It also provides you with the background information necessary to administer and maintain IBM solidDB in your network environment.

2.1 System Architecture

IBM solidDB uses a client/server model. In practice, the IBM solidDB session consists of the following co-operating processes:

- The server process, which manages the database files, accepts connections to the database from client applications, and carries out actions on the database as requested by the clients.
- The client process, which is used to pass the required tasks (through the server process) to the database. There can be several client types: a client could be a command-line tool, a graphical application, or a database maintenance tool. Typically, different applications act as clients to connect to IBM solidDB.

IBM solidDB can also run within the application process. This is provided by IBM solidDB linked library access. The linked library access is a function library that provides the same functionality and interfaces available with IBM solidDB. A user application may be linked to this library. The linked application communicates with the server by using direct function calls, thus skipping the overhead required when the client and server communicate through network protocols such as the TCP/IP. Linking the application and server into a single executable provides higher performance. For more information, refer to *IBM solidDB Linked Library Access User Guide*.

The client and the server can be on different hosts, in which case they communicate over a network. IBM solidDB provides simultaneous support for multiple network protocols and connection types. Both the database server and the client applications can be simultaneously connected to multiple sites using multiple different network protocols.

To submit a query (an SQL statement) to a database server, a client must be able to communicate with that database server. IBM solidDB, like many other database servers, uses "drivers" to enable this communication. Client applications call functions in the driver, and the driver then handles the communications and other details with the server. For example, you might write a C program that calls functions in the Open Database Connectivity (ODBC) driver, or you might write a Java program that calls functions in the Java Database Connectivity (JDBC) driver.

For more information about the ODBC and JDBC drivers, and how to use them with your client applications, see the *IBM solidDB Programmer Guide*.

2.2 IBM solidDB Directory Structure

After the IBM solidDB installation, you will have a new directory on your server called `IBM solidDB 6.1`. The files and subdirectories in that directory are explained in the table below. For more detailed information, refer to the in-package documentation.

Table 2.1. IBM solidDB Directory Structure

Location	Explanation
Root folder	The root folder contains, for example: <ul style="list-style-type: none"> • A script used to facilitate running samples in the database evaluation phase • The evaluation licence file • Open the <code>welcome.html</code> file in a browser for more information on IBM solidDB.
<code>bin</code>	IBM solidDB binary files.
<code>connector</code>	The connector binary file and related sample scripts.
<code>doc_html</code> <code>doc_txt</code>	Package documentation in HTML and in text.
<code>eval_kit/standalone</code>	IBM solidDB evaluation kit licence and initiation files. It will also hold your database once it is created.
<code>include</code>	C program headers.
<code>jdbc</code>	Java Database Connectivity (JDBC) API for IBM solidDB.
<code>lib</code>	Program libraries.
<code>licence</code>	Licence files.
<code>manuals</code>	IBM solidDB documentation.
<code>odbc</code>	Open DataBase Connectivity (ODBC) API for IBM solidDB.
<code>samples</code>	Samples that can be used in the database evaluation phase and future application development.

2.3 Prerequisites

Make sure that you have a valid IBM solidDB licence certificate in the installation directory. The evaluation licence certificate file, `solideval.lic`, is provided with the product image. The evaluation period expires ninety days after the product is first installed. To enable the product, you must purchase the full product license certificate and copy the file `solid.lic`, `soliddb2.lic` or `solidids.lic` to the installation directory.

2.4 What Is in Your IBM solidDB Development Kit

Your IBM solidDB Development Kit (SDK) contains more than just a IBM solidDB program. It also contains the IBM solidDB linked library access program, sample client programs, and documentation.

2.4.1 Software

The IBM solidDB Development Kit contains a complete set of IBM solidDB software, including the linked library access linkable library, and various "utility" programs.

2.4.2 Samples

The SDK comes with several sample programs written in C, SQL, and Java to help you get started using the features in your IBM solidDB.

2.4.3 Documentation

For a description of the manuals that come with IBM solidDB, see Section 1.3, "IBM solidDB Documentation".

Chapter 3. Preparing for the Evaluation

3.1 Running the License Script

The evaluation license enables you to evaluate IBM solidDB for a limited time, each time you create a new database. To use the evaluation product for a longer period of time, contact your IBM Corporation salesperson for a different license.

In order to be able to run samples, you must put a copy of the license file into each of the sample run directories. The easiest way to do this is to execute the script called

```
copy_licenses (copy_licenses.bat in Microsoft Windows)
```

This script copies the license file from the installation root directory to all appropriate directories.

3.2 IBM solidDB Configuration File

When you start IBM solidDB, it reads configuration parameters from the `solid.ini` configuration file. You can also use IBM solidDB without the configuration file, in which case the factory settings are used. The main difference between the factory settings and the configuration file is the database port number. In factory settings, the port number is `1964` whereas in the configuration file it is `1315`.

The IBM solidDB configuration file will serve most evaluation needs as long as IBM solidDB is concerned. If there is a need to use other components such as the High Availability component, the configuration file must be modified accordingly. Samples of the corresponding components are included with the samples. In other words, you do not have to modify the `solid.ini` file (in the `/eval_kit/standalone` directory) to evaluate other IBM solidDB components.

The `solid.ini` configuration file specifies parameters that help customize and optimize IBM solidDB. For example, the `FileSpec` parameter in the `solid.ini` configuration file specifies the directory and files names of the data fields in which the server stores the user data. Another parameter specifies the block size for the database.

You can view all parameter setting with the `solsql` command below. for instructions, see Section 5.1, “Using IBM solidDB SQL Editor”.

ADMIN COMMAND 'par';

You can view section specific parameter settings with the solsql command:

```
ADMIN COMMAND 'par section_name' ;
```

For more information, please refer to *IBM solidDB Administration Guide*.

Chapter 4. Starting IBM solidDB and Creating Your First Database

4.1 Starting IBM solidDB

You can start IBM solidDB as explained in the table below:

Table 4.1. Starting IBM solidDB on Different Operating Systems

Operating System	To Start the Server...
Unix, Linux	Enter command solid at the command prompt. When you start the server for the first time, enter the command solid -f at the command prompt to force the server to run in the foreground.
Microsoft Windows	Click the icon labeled <i>IBM solidDB Server</i> in the Start menu.

When you start IBM solidDB, it checks if a database already exists. If no database is found, the engine automatically creates a new database in line with the `solid.ini` configuration file. However, if you want to create a database manually to a specific location, proceed as explained below.

4.2 Creating a Database Manually

To create a new database in the directory below:

```
\eval_kit\standalone
```

Open a command prompt window, go to the IBM solidDB root directory and enter the following command:

```
bin\soliddb -c eval_kit\standalone
```

A window with the title "Creating a new database" will open. You will be asked to enter the system catalog name, username, and password.



Caution

There are no defaults for the username and password. You must remember the username and password to be able to access the database again.

Traditionally, IBM solidDB has used "dba" for the username and password of evaluation databases, since it is simple, easy to remember, and can be found in this document in case you forget it. To use it:

1. type any name you prefer as the catalog name
2. type "dba" as username and password
3. retype "dba" as the password
4. press OK

A new database is created and it is running.

4.3 Automatic Creation of a Database (Windows)

Start IBM solidDB through the Start→Programs menu path. If no database is found, the engine automatically creates a new database in line with the `solid.ini` configuration file. In this case, the default name "dba" is used for the catalog, username and password.

A new database is created and it is running.



Note

Do not use these simple strings as username and password in the production environment.

4.4 Verifying the Database Installation

If you successfully created the database, your IBM solidDB process is now running. The process name is `solid` (`solid.exe` in Microsoft Windows).

Chapter 5. Connecting to IBM solidDB

After starting IBM solidDB and creating your first database, you can test the configuration by connecting to the server from your workstation by using the IBM solidDB SQL Editor.

If you have problems in connecting to the database with these instructions or if any errors should occur, please refer to *IBM solidDB Administration Guide*.

5.1 Using IBM solidDB SQL Editor

With IBM solidDB SQL Editor, SQL statements (including the SQL ADMIN COMMANDS) can be issued at the command line, command prompt, or by executing a script file that contains the SQL statements.

5.1.1 Starting IBM solidDB SQL Editor

To start the IBM solidDB SQL Editor, enter the start command at your operating system prompt. The command syntax is as follows:

```
solsql "networkname" [userid [password]]
```

For example:

```
solsql "tcp hobbes 1315" dba dba
```

Or, if the server runs in the same computer:

```
solsql "tcp 1315" dba dba
```

When using Microsoft Windows, start the IBM solidDB SQL Editor from the icon in the Start menu

5.1.2 Connecting to Database

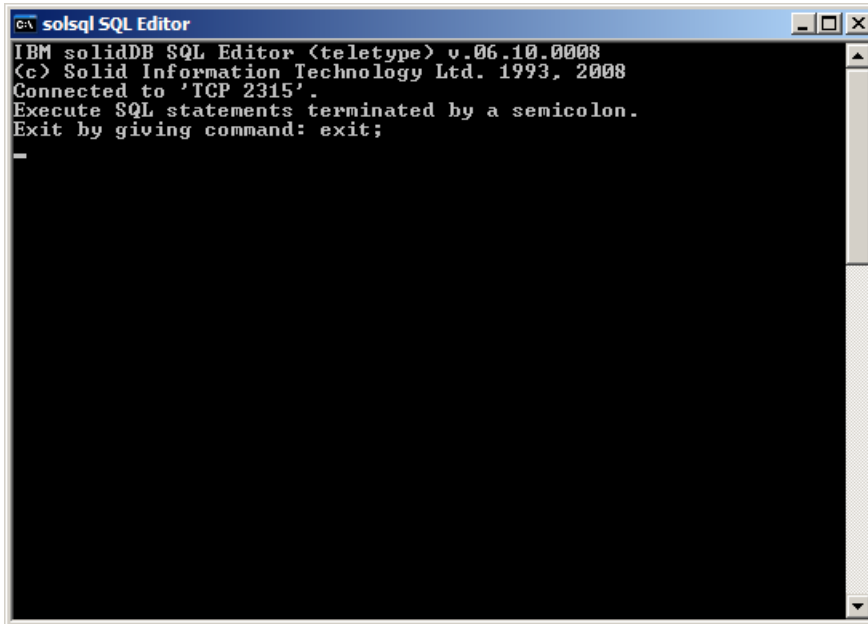
If you entered a valid user name and password when starting the IBM solidDB SQL Editor, you are connected to the database. If you did not give your user name and password when starting the IBM solidDB SQL Editor, you will be prompted for them. Enter your username and password (for example "dba", "dba") in the IBM solidDB SQL Editor.

The figure below shows the IBM solidDB SQL Editor after a successful connection to the database.

 **Note**

It can take a while before the connection can be established. Even though you can use the editor during the connecting phase, do not give any commands before you see a screen such as the one in the figure below.

Figure 5.1. IBM solidDB SQL Editor Connected



5.1.3 Shutting Down IBM solidDB SQL Editor

In all environments, you can close the solsql editor by entering the command:

exit;

5.1.4 Using Help

The IBM solidDB SQL Editor contains more detailed instructions for using the SQL commands. For a list of available commands, enter the command below:

ADMIN COMMAND 'help';

5.2 Entering Queries

Enter an SQL query, for example:

```
select table_name from tables;
```

In solsql, press Enter to execute the query.

Note

Notice the use of semicolon at the end of the SQL command.

For any IBM solidDB, you should receive a resultset of 79 rows.

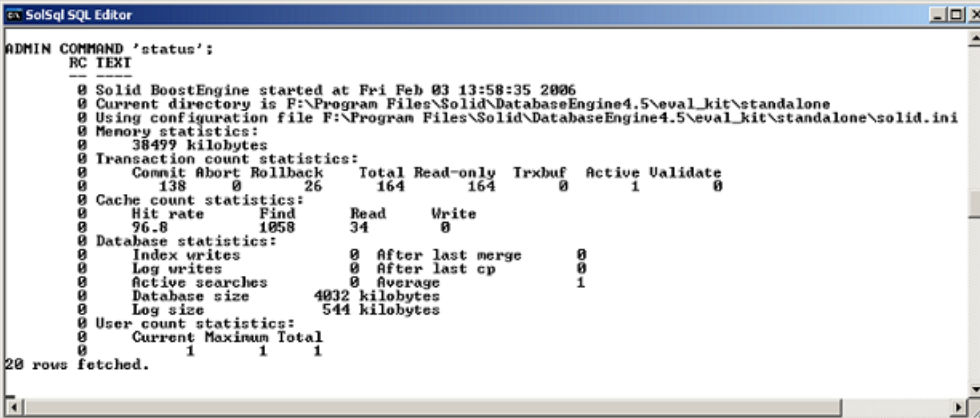
5.3 Checking the Database Status

The general server status may be retrieved by using the following command in the IBM solidDB SQL Editor:

```
ADMIN COMMAND 'status';
```

An example output from the solsql is shown below:

Figure 5.2. Example Output from Solsql



```

SolSql SQL Editor
ADMIN COMMAND 'status':
RC TEXT
-----
0 Solid BoostEngine started at Fri Feb 03 13:58:35 2006
0 Current directory is F:\Program Files\Solid\DatabaseEngine4.5\eval_kit\standalone
0 Using configuration file F:\Program Files\Solid\DatabaseEngine4.5\eval_kit\standalone\solid.ini
0 Memory statistics:
0 38497 kilobytes
0 Transaction count statistics:
0 Commit Abort Rollback Total Read-only Trxbuf Active Validate
0 138 0 26 164 164 0 1 0
0 Cache count statistics:
0 Hit rate Find Read Write
0 96.8 1058 34 0
0 Database statistics:
0 Index writes 0 After last merge 0
0 Log writes 0 After last cp 0
0 Active searches 0 Average 1
0 Database size 4032 kilobytes
0 Log size 544 kilobytes
0 User count statistics:
0 Current Maximum Total
0 1 1 1
20 rows fetched.

```

Chapter 6. Stopping and Restarting the Database

This section gives you instructions on stopping and restarting IBM solidDB.

6.1 Stopping the Database with IBM solidDB SQL Editor

You can stop the database by using the IBM solidDB SQL Editor. Give the commands below. Use the semicolon at the end of the commands, and note that you must use single quotes, not double quotes.

admin command 'close';

admin command 'throwout all';

admin command 'shutdown';

exit;

The first command prevents any additional users from connecting. The second command throws out all connected users (except the one who issued the command). The third command shuts down IBM solidDB. The fourth command exits from the solsql tool.

The first three commands may be replaced with this one:

admin command 'shutdown force';



Caution

When you shut down the server, it breaks the connection to solsql, and your solsql may show an error message such as: 14519: The user was thrown out from the server; connection lost

6.2 Restarting the Database

You can restart IBM solidDB by moving to the IBM solidDB root directory and entering the following command:

./bin/soliddb -c ./eval_kit/standalone

If you have the IBM solidDB "bin" directory in your path, you can go to the `eval_kit/standalone` directory to make it your current working directory, and start IBM solidDB just by executing the command

solid

In Windows, to restart the evaluation database, use the Start menu

Chapter 7. Running Samples

You will find sample programs and scripts in the samples directory below the IBM solidDB installation directory. There is a readme file in each samples subdirectory. Note that some sample scripts may remove and re-create database files residing in sample subdirectories. Only the database in the "standalone" directory is always left intact.

For more information, refer to the sample documentation provided in the IBM solidDB distribution package.

Chapter 8. IBM solidDB Cache

IBM solidDB Cache 6.1 provides a high-performance, low-latency database front-end solution for IBM Data Servers. IBM solidDB Cache solution uses a number of in-memory front-end databases to cache traffic from the applications. Typically, each front-end contains writable partition replica(s) of the back-end database and read-only partition replica(s) of the back-end database. The replication takes place asynchronously. In this way, the read and write load can be distributed in the system extremely efficiently to achieve extreme speed.

IBM solidDB Cache uses connector instances to handle traffic between the front-ends and the back-end.

IBM solidDB package contains a IBM solidDB Cache sample called *solconnector*.

Chapter 9. What Next?

We recommend that you read *IBM solidDB Administration Guide*. This will help you configure, start, and maintain your database server.

We recommend that you try one of the sample programs provided with the IBM solidDB Development Kit.

After you have read *IBM solidDB Administration Guide*, you may want to read the guides that apply to any specific components that you plan to use:

- *IBM solidDB In-Memory Database User Guide*
- *IBM solidDB High Availability User Guide*
- *IBM solidDB Linked Library Access User Guide*
- *IBM solidDB Advanced Replication Guide*
- *IBM solidDB Connector Guide*

Most users will probably want to read the parts of the IBM solidDB Programmer Guide that are relevant to them. If you will be writing client programs that query the server, you will need to read at least part of this manual.

Glossary

This glossary gives you a description of the terminology used in this guide.

A

Application Programming Interface (API)

An API is a way for one piece of code to use another piece of code. In the case of IBM solidDB, client programs use an API (such as ODBC or JDBC) to communicate with the server and make requests of it. For example, your client application program may compose an SQL query and then use an ODBC or JDBC function call to send that query to the server for processing.

IBM solidDB also provides a proprietary database API called SA that can be accessed by using a special library.

C

Client/server computing

Client/server computing divides a large piece of software into modules that need not all be executed within the same memory space nor on the same processor. The calling module becomes the “client” that requests services, and the called module becomes the “server” that provides services. Client and server processes exchange information by sending messages through a computer network. They may run on different hardware and software platforms as appropriate for their special functions.

Two basic client/server architecture types are called two-tier and three-tier application architectures.

Communication protocol

A communication protocol is a set of rules and conventions used in the communication between servers and clients. The server and client have to use the same communication protocol in order to establish a connection. TCP/IP is an example of a common communication protocol.

D

Database administrator

The database administrator is a person responsible for tasks such as:

- managing users, tables, and indices
- backing up data

-
- allocating disk space for the database files

Database Management System (DBMS)

A DBMS is a system that stores information in and retrieves information from a database. A DBMS typically consists of a database server, administration utilities, an application interface, and development tools.

Database procedures

See stored procedures.

I

Index

An index of records has an entry for each key field (for example, employee name, identification number, etc.) and the location of the record. Indexes are used to speed up access to tables. IBM solidDB uses indexes to access the rows in a table directly. Without indexes, the engine would have to search the whole contents of a table to find the desired row. A single table can have more than one index; however, adding indexes does slow down write operations, such as inserts, deletes, and updates on that table. There are two kinds of indexes: non unique indexes and unique indexes. A unique index is an index where all key values are unique.

J

Java Database Connectivity (JDBC)

JDBC is an application programming interface for the Java programming language. It defines how a client may access a database. It provides methods for querying and updating data in a database.

O

Open Database Connectivity (ODBC)

ODBC is a programming interface standard for SQL database programs. IBM solidDB offers a native ODBC programming interface.

Optimizer Hints

Optimizer hints (which are an extension of SQL) are directives specified through embedded pseudo comments within query statements. The Optimizer detects these directives or hints and bases its query execution plan accordingly. Optimizer hints allow applications to be optimized under various conditions to the data, query type, and the database. They not only provide solutions to performance problems occasionally encountered with queries, but shift control of response times from the system to the user.

R

Relational database management system (RDBMS)

IBM solidDB is an RDBMS, which stores and retrieves information that is organized into two-dimensional tables. This name derives from the relational theory that formalizes the data manipulation requests as set operations and allows mathematical analysis of these sets. RDBMSs typically support the SQL language for data manipulation requests.

S

Structured Query Language (SQL)

SQL is a standardized query language designed for handling database requests and administration. The SQL syntax used in IBM solidDB is based on the ANSI X3H2-1989 Level 2 standard including important ANSI X3H2-1992 (SQL-92) extensions. Refer to *IBM solidDB Administration Guide* for a more formal definition of the syntax.

SQL Access Group Call Level Interface (SAG CLI)

SAG CLI is a programming interface standard that defines the functions that are used to submit dynamic SQL clauses to a database server for execution. The ODBC interface is also based on SAG CLI. The IBM solidDB SQL API conforms to the SAG CLI standard.

Stored procedures

Stored procedures allow programmers to split the application logic between the client and the server. These procedures are stored in the database, and they accept parameters in the activation call from the client application. This arrangement is used by intelligent transactions that are implemented with calls to stored procedures.

T

Triggers

Triggers are pieces of logic that IBM solidDB automatically executes when a user attempts to change the data in a table. When a user modifies data within the table, the trigger that corresponds to the command (such as insert, delete, or update) is activated.

Index

C

Client/server model, 7
 client, 7
 server, 7

G

Getting started with IBM solidDB, 7

I

IBM solidDB Configuration File, 11
IBM solidDB Directory Structure, 8

J

JDBC, 7

L

Licence script, 11
Licencing, 11
 Licence script, 11
linked library access, 7, 9

O

ODBC, 7

P

Preparing for the Evaluation, 11

S

solid.ini, 11

T

TCP/IP, 7
