

IBM solidDB
IBM solidDB Universal Cache
Version 6.3

Getting Started Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 43.

First edition, third revision

This edition applies to version 6, release 3 of IBM solidDB (product number 5724-V17) and IBM solidDB Universal Cache (product number 5724-W91) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Oy International Business Machines Ab 1993, 2011

Contents

Figures	v	Post-installation tasks	22
Tables	vii	Uninstalling solidDB	22
Summary of changes	ix	Installing solidDB Documentation package	23
About this manual	xi	Installing fix packs	23
Typographic conventions	xi	5 Starting IBM solidDB and creating your first database	27
Syntax notation conventions	xii	solidDB configuration file	27
Accessibility	xii	Starting solidDB	27
1 Product overview	1	Creating a database manually (example)	27
IBM solidDB product family	1	Creating a database automatically (example)	28
Architectural overview of solidDB	1	6 Connecting to solidDB for the first time	29
Architectural overview of solidDB Universal Cache	3	Using solidDB SQL Editor (solsql)	29
Main features and functionality	4	Starting and shutting down the solidDB SQL Editor	29
2 Product and packaging information	9	Connecting to a database in the solidDB SQL Editor	29
solidDB server package	9	Entering queries in the solidDB SQL Editor	30
Directory structure	9	Checking the database status	30
Library file names	10	7 Stopping and restarting the database 33	
Samples	11	Stopping the database with solidDB SQL Editor	33
InfoSphere CDC components	11	Restarting the database (example)	33
Documentation	11	8 Running samples	35
solidDB documentation	11	9 Using solidDB with graphical SQL clients	37
InfoSphere CDC documentation	12	10 Upgrading solidDB to a new release level	39
3 Pre-requisites	15	Index	41
System requirements	15	Notices	43
Platform support	15		
Available interfaces for client application development	15		
solidDB Java Accelerator	17		
4 Installation	19		
Installing solidDB	19		
Graphical interface installation	19		
Console installation	20		
Silent installation	21		

Figures

1.	solidDB 6.3 products	1
2.	solidDB Universal Cache architecture	3
3.	solidDB SQL Editor Connected	30
4.	Example output from solidDB SQL Editor (solsql)	31

Tables

1.	Typographic conventions	xi	4.	'solidDB6.3' directory structure	9
2.	Syntax notation conventions	xii	5.	Documentation for InfoSphere CDC components	13
3.	Components of solidDB and solidDB Universal Cache product packages.	9	6.	solidDB point release versions for fix packs	23

Summary of changes

Changes for revision 03

- Editorial corrections.

Changes for revision 02

- Section “Installing fix packs” on page 23 updated: To install a fix pack for the InfoSphere™ CDC engines (InfoSphere CDC for solidDB® or InfoSphere CDC for *back-end data server*), you must already have an installation of the InfoSphere CDC engine.
- New section added: 9, “Using solidDB with graphical SQL clients,” on page 37

Changes for revision 01

- New section added: 10, “Upgrading solidDB to a new release level,” on page 39.
- Installation instructions for solidDB updated in section “Installing solidDB” on page 19; console installation added.

About this manual

This guide provides an overview of the IBM® solidDB® product family. It also includes instructions for how to set up solidDB for the first time and how to run samples.

This guide does not provide instructions for how to install and configure solidDB Universal Cache or how to start using the different features and functionality available with the solidDB products. For information on how the documentation is structured, see section Documentation in this guide.

Typographic conventions

solidDB documentation uses the following typographic conventions:

Table 1. Typographic conventions

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

Table 1. *Typographic conventions (continued)*

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

Syntax notation conventions

solidDB documentation uses the following syntax notation conventions:

Table 2. *Syntax notation conventions*

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

Accessibility

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

Keyboard input and navigation

Keyboard input

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

Keyboard navigation

You can navigate the 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.

Accessible display

solidDB supports 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 solidDB.

Compatibility with assistive technologies

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.

Accessible documentation

Documentation for solidDB is provided in HTML format via the Information Center, which is viewable in most Web browsers. Information Center 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.

See the solidDB Information Center opening page for more details on the Information Center accessibility.

1 Product overview

IBM solidDB product family

IBM® solidDB product family features relational, in-memory database technology that delivers extreme speed, performing up to ten times faster than conventional, disk-based databases. IBM solidDB, or simply, solidDB, uses the familiar SQL language to allow applications to achieve tens of thousands of transactions per second with response times measured in microseconds.

IBM solidDB 6.3 product family is composed of two products:

IBM solidDB Universal Cache is the industry's first relational, in-memory caching software that accelerates traditional disk-based relational database servers by caching performance critical data into one or more solidDB in-memory database instances.

IBM solidDB is a fully-featured, relational, in-memory database that delivers extreme speed and extreme availability to meet performance and reliability demands of real-time applications. It provides the ability to have both in-memory tables and on-disk tables within a single solidDB instance. Other features include a High Availability implementation and several different replication topologies.

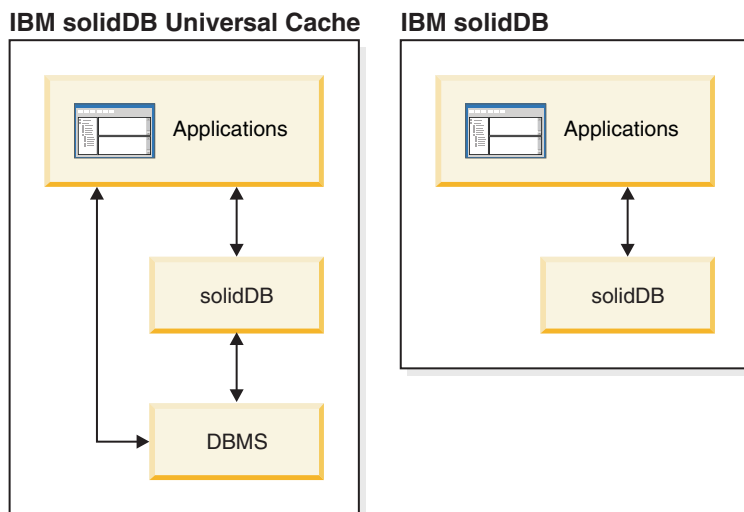


Figure 1. solidDB 6.3 products

Architectural overview of solidDB

This section describes the system architecture of solidDB.

solidDB uses a client/server model. In practice, the solidDB session consists of cooperating server and client processes. The server process 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 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 solidDB.

The client and the server can be located on different hosts (nodes), in which case they communicate over a network. 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.

solidDB can also run within the application process. This is provided by solidDB *linked library access*. In this case, a user application is linked to a function library that is provided with the product. 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, see the *IBM solidDB Linked Library Access User Guide*.

To submit a query (an SQL statement) to a database server, a client must be able to communicate with that database server. 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 (ODBC) driver, or you might write a Java™ program that calls functions in the (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*.

Architectural overview of solidDB Universal Cache

The following diagram illustrates the architecture and key components of a typical configuration of the solidDB Universal Cache.

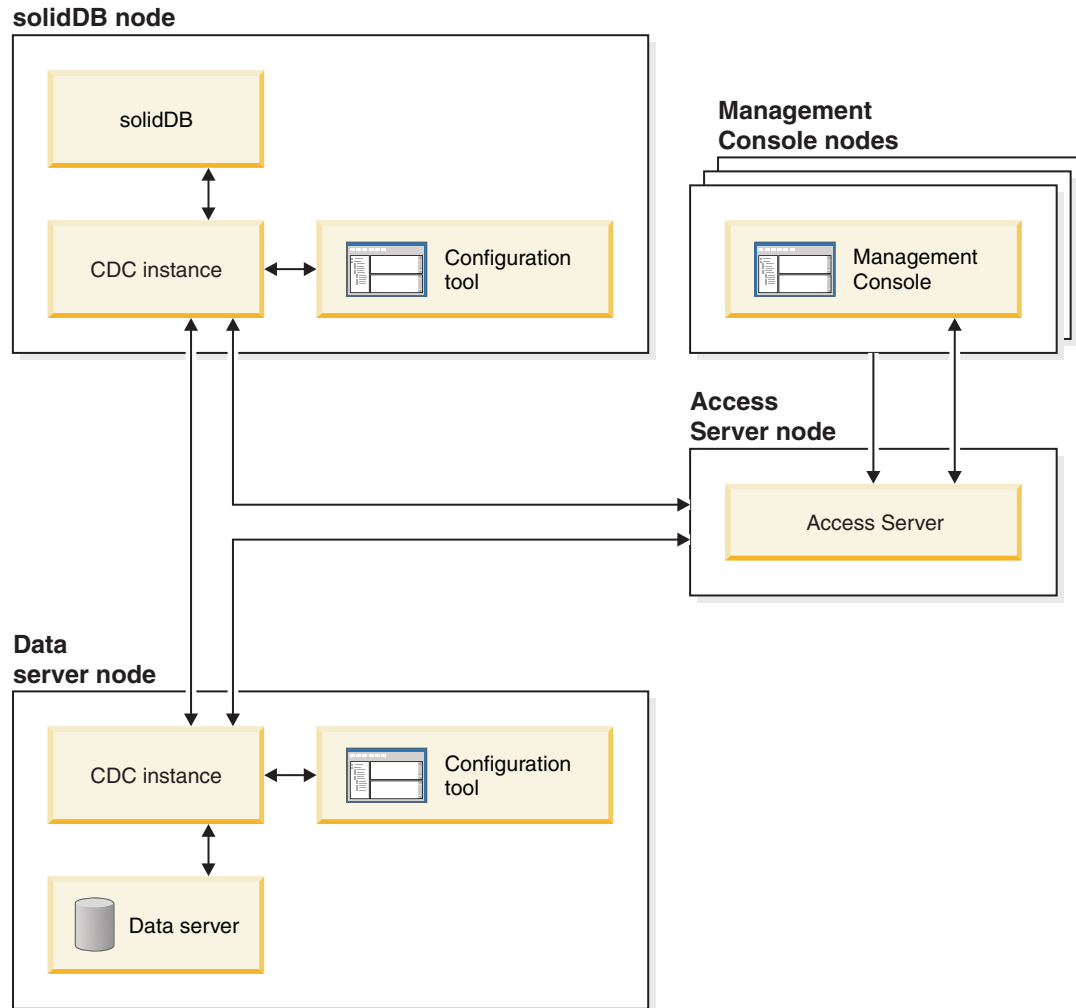


Figure 2. solidDB Universal Cache architecture

The roles and functions of the components are described below.

solidDB — the front-end database which you replicate data to and from. Depending on the replication model, solidDB can be the source or target database, or both.

Data server — the back-end database which you replicate data to and from. Depending on the replication mode, the back-end database can be the source or target database, or both.

InfoSphere CDC instance — a run-time instance of the InfoSphere CDC engine for a given DBMS. To set up InfoSphere CDC instances, you have to have the corresponding InfoSphere CDC agent software (like InfoSphere CDC for DB2®, or InfoSphere CDC for solidDB) installed on the same node where the DBMS in

question is running. Exceptionally for solidDB, the InfoSphere CDC for solidDB agent can be installed and set up on any node connected to the node running solidDB.

Configuration tool — a GUI or console based tool used to configure and create InfoSphere CDC instances. While configuring the instance, you supply the port number for communication with the rest of the InfoSphere CDC system as well as login information to connect to the database.

Access Server — a process (running typically as a service or daemon) allowing the Management Console users to access the InfoSphere CDC instances and configure them. Different users may have access to different instances. During the Access Server installation you are requested to provide a port number and Administrator login information, to be used by the Management Console.

Management Console — an interactive application with a GUI that you can use to configure and monitor replication. It allows you to manage replication on various servers, specify replication parameters, and initiate refresh and mirroring operations from a client workstation.

You start using the Management Console by creating datastores in the Access Manager perspective. A *datastore* is a logical entity of a database and the related InfoSphere CDC instance. Using the datastores defined, you can set up *subscriptions* that embody data replication from one datastore to another.

There are three types of datastores: source, target, and dual. Dual datastore can participate in subscriptions as both a source and target. Typically, dual datastores are used with solidDB Universal Cache, and symmetric mapping pairs are defined for two-way replication. If a replica is meant to be read-only, in the front-end, a single upload subscription is defined.

After you have finished setting up replication, the Management Console can be closed on the client workstation without affecting active data replication activities between source and target servers. Management Console also includes an event log and a monitor. The event log allows you to examine generated InfoSphere CDC event messages. The monitor provides the necessary support to continuously monitor replication operations and latency. Diagrams depicting components of your replication configuration are constructed through direct manipulation of graphical objects. The monitor in Management Console is intended for time-critical working environments that require continuous analysis of data movement.

Main features and functionality

In-memory tables

In-memory tables store all their data in main memory in data-structures which are specifically optimized for main-memory access. The benefit of the in-memory residency is low query latency and high throughput. Depending on the application needs, the in-memory data tables may be configured to be persistent or transient. In addition, flexible logging capabilities allow to define the level of transactional durability required for persistent table, including full durability.

For more information, see the *IBM solidDB In-Memory Database User Guide*.

Disk-based tables

Disk-based tables are traditional relational database tables, equivalent to those found in other disk-based database products. In addition to in-memory tables, normal disk-based tables may be used in the database too. Typically only a part of the data from disk-based tables is resident in main-memory at any point in time. Queries can span in-memory and disk-based tables. For example, an SQL join operation may join an in-memory table with a disk-based table; the actual location of the table is transparent to the user.

The disk based tables, combined with a small footprint and unattended operation, makes solidDB ideal to be used in embedded systems. Despite the small size, the product supports many advanced features like triggers, events, and stored procedures.

Linked library access

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

For more information, see the *IBM solidDB Linked Library Access User Guide*.

Replication technologies

In solidDB product family, data replication can be implemented with three different technologies, namely, Advanced Replication, InfoSphere CDC Replication, and solidDB High Availability (HotStandby).

Replication can be synchronous or asynchronous. In synchronous replication, the data is durably delivered to the receiving end within the transaction executing at the originating end. In asynchronous replication, the data is delivered to the receiving end after the original transaction has been committed. Replication can be based on a push or pull operational model. In the push model, it is the originating end that activates the replication. In the pull model, the receiving end takes the initiative.

solidDB High Availability (HotStandby)

solidDB High Availability, or simply HA or HotStandby, increases the availability of data. Servers are paired up in such a way that, if one fails, the other can take over. Data from the primary server is replicated to the secondary server, which is used as a 'hot standby' unit. If the primary server is unavailable, for example, because of hardware failure or scheduled maintenance, applications can connect to the secondary server and continue immediately without any loss of committed transactions.

In solidDB HA, all the data changes in the primary are propagated to the secondary using a push-based replication protocol. The protocol may set to synchronous (*2-Safe*) or asynchronous (*1-Safe*). In principle, the transaction load is served at the primary. Should the primary fail, the secondary takes up the job (executes a failover) and continues to serve the load as the new primary. With the synchronous replication protocol, there is no risk of data loss during the failover.

Transparent Connectivity, a special connectivity mode of solidDB ODBC and JDBC drivers, offers the applications failover transparency and transparent load balancing between the primary and secondary.

Both relaxed and strict durability can be used with HotStandby. There is also a durability level called *adaptive*, which uses relaxed durability when both primary and secondary servers are active, and which switches to strict durability when the primary runs alone. With adaptive durability, the transactions are secured against any single failure, at any time.

In comparison to Advanced Replication and InfoSphere CDC Replication, with solidDB HA, all the data in the server's database is always replicated. There are a few dynamic controls available, in the form of specialized SQL commands.

For more information, see the *IBM solidDB High Availability User Guide*.

Advanced Replication

Advanced Replication technology represents an asynchronous pull-based approach. It allows users to distribute and synchronize, occasionally, data across multiple database servers. Advanced Replication uses a master/replica model in which a single node holds the master copy of the data. One or more replica nodes can also have a copy of part or all of the master data. Furthermore, each replica can also have data that is specific to that replica only.

Replication is bidirectional; replica nodes can upload data to the master as well as download from it. If replicas submit conflicting data, the master can reject the data or alter the data before making it available to all replicas. The management of replicated data is based on a flexible publish/subscribe model. The management interface is in the form of proprietary extensions to the SQL language.

Advanced Replication's asynchronous approach means that the system is inherently flexible; if some nodes are down for any reason, those nodes can resynchronize when they come back up or reconnect to the system. This allows mobile devices, such as PDAs or laptops, to request updated data, disconnect from the network, and then reconnect later. Users can choose how frequently they want to synchronize data.

For more information, see the *IBM solidDB Advanced Replication User Guide*.

InfoSphere CDC technology

IBM InfoSphere Change Data Capture (InfoSphere CDC) technology can be used to implement asynchronous replication among various databases. InfoSphere CDC technology is based on an asynchronous push model. Unidirectional subscriptions can be created for real-time propagation of data changes from the source side to the target side. Bidirectional capability is achieved by setting up two subscriptions with mirrored source and target definitions.

InfoSphere CDC technology is included both in solidDB and solidDB Universal Cache products.

- **In solidDB**, the InfoSphere CDC technology can be used for replicating data between any pair of solidDB server instances. For example, it enables geographic redundancy, or, it allows two solidDB instances to be configured in an active/active setup whereby two copies of the same data can be equally processed at either side.

For more information, see the *IBM solidDB InfoSphere CDC Replication User Guide*.

- **In solidDB Universal Cache**, the InfoSphere CDC technology can be used to setup a solidDB front-end cache, to speed up access to performance critical data stored in back-end data servers.

For more information, see the *IBM solidDB Universal Cache User Guide*.

In both setups, each solidDB server instance can also be a one solidDB HotStandby pair.

InfoSphere CDC technology includes database-specific components that run as separate processes called InfoSphere CDC instances. A Management Console for managing the replication subscriptions and data mapping is also included.

2 Product and packaging information

The solidDB and solidDB Universal Cache products are composed of solidDB and InfoSphere CDC components. The table below shows the components included in the solidDB and solidDB Universal Cache product packages. Each listed component must be installed separately.

Table 3. Components of solidDB and solidDB Universal Cache product packages

Component	solidDB	solidDB Universal Cache
IBM solidDB 6.3.x	X	X
IBM solidDB 6.3.x Documentation	X	X
InfoSphere Change Data Capture v6.3.x solidDB	X ¹	X
InfoSphere Change Data Capture v6.3.x Management Console	X ¹	X
InfoSphere Change Data Capture v6.3.x Access Server	X ¹	X
InfoSphere Change Data Capture v6.3.x Documentation	X ¹	X
InfoSphere Change Data Capture v6.3.x for a back-end data server		X

¹ Needed only in configurations deploying InfoSphere CDC technology for solidDB-to-solidDB replication (InfoSphere CDC Replication).

solidDB server package

The solidDB package contains a complete set of solidDB software, including the JDBC and ODBC drivers as well as various utility programs.

The solidDB package is delivered with an evaluation license certificate file, *slideval.lic*. The evaluation license enables you to evaluate solidDB for 90 days. For acquiring a permanent license, please contact IBM Corporation.

Directory structure

The default installation of solidDB 6.3 creates a directory called 'solidDB6.3'.

The files and subdirectories in the 'solidDB6.3' installation directory are explained in the table below.

Table 4. 'solidDB6.3' directory structure

Location	Explanation
Root directory	The root directory 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• the welcome.html file for accessing the package documentation
bin	solidDB binary files

Table 4. 'solidDB6.3' directory structure (continued)

Location	Explanation
doc_html, doc_txt	Package documentation in HTML and text format
eval_kit/standalone	solidDB evaluation kit license and initiation files. It will also hold your evaluation database once it is created.
include	C program headers
jdbc	Java Database Connectivity (JDBC) Driver for solidDB Data store helper archive for use with WebSphere
lib	Program libraries
licence	License files
manuals	The English version manuals can be downloaded to this folder and then accessed through the Manuals link on the Welcome-page
samples	Samples that can be used in the database evaluation phase and future application development

Library file names

solidDB provides many files as linkable libraries.

Most of these library files fall into one of the following categories:

- ODBC Drivers
- solidDB Linked Library Access files
- Communication library files
- SA (Server API) library file.

Not all platforms have every file. For example, some communication library files are available on Windows environments only.

Some library files are static, that is, they are linked to your client application's executable program when you do a compile-and-link operation. Other library files are dynamic: these files are stored separately from your executable and are loaded into memory when your program executes. For many libraries, solidDB provides both a static and a dynamic version on some or all platforms.

Library files are generally found in one of two directories:

- 'bin'
- 'lib'

As a rule, the 'bin' directory contains dynamic libraries (in addition to executables), while the 'lib' directory contains static libraries. The 'lib' directory also contains the import libraries.

For more detailed information on the library file names, see the *releasenotes.html* or *releasenotes.txt* in your solidDB installation directory.

Samples

The solidDB package includes several sample programs written in C, SQL, and Java to help you get started using the features of solidDB.

The samples and scripts are located in the 'samples' directory below the solidDB installation directory. Each sample directory includes also a *readme.txt* file that provides instructions for how to use the samples.

InfoSphere CDC components

The InfoSphere CDC components are delivered as separately deployable packages.

- InfoSphere Change Data Capture v6.3 solidDB (InfoSphere CDC for solidDB)

The InfoSphere CDC for solidDB package contains the software for the configuration tool and the InfoSphere CDC instance for solidDB.

- InfoSphere Change Data Capture v6.3 Management Console (InfoSphere CDC) and InfoSphere Change Data Capture v6.3 Access Server (InfoSphere CDC Access Server)

The InfoSphere CDC Management Console and InfoSphere CDC Access Server are delivered and installed as separate packages. They contain the software that you can use to configure and monitor InfoSphere CDC user access and replication subscriptions in the solidDB Universal Cache and InfoSphere CDC Replication configurations.

- InfoSphere Change Data Capture v6.3 *for a back-end database* for other data servers (solidDB Universal Cache only)

The solidDB Universal Cache product includes the InfoSphere CDC component for the back-end data server. The package contains the software for the configuration tool and the InfoSphere CDC instance for the data server in question.

Documentation

Documentation for solidDB 6.3 is composed of an *IBM solidDB 6.3 Documentation* package and an *InfoSphere Change Data Capture v6.3 Documentation* package.

The solidDB Documentation package, which includes the documentation for InfoSphere CDC for solidDB component, is available as an online information center and in PDF format. The InfoSphere CDC Documentation package contains documentation for the InfoSphere CDC Management Console, the InfoSphere CDC Access Server, and the InfoSphere CDC component for the other data server.

solidDB documentation

solidDB documentation is available online in the solidDB 6.3 and solidDB Universal Cache 6.3 Information Center, as well as in PDF format. Most up-to-date information is always available in the Information Center.

Delivery of solidDB documentation

solidDB 6.3 and solidDB Universal Cache 6.3 Information Center

The most up-to-date solidDB documentation is available in the information center format at <http://publib.boulder.ibm.com/infocenter/soliddb/v6r3/>.

solidDB manuals in PDF format

The PDF manuals are available for download at the following locations:

- solidDB Support Web pages <ftp://ftp.software.ibm.com/software/data/soliddb/info/6.3/man/>.
- IBM Publications Center at <http://www.elink.ibm.com/publications/servlet/pbi.wss>

In addition, the PDF format manuals are available as the *IBM solidDB 6.3 Documentation* package. This package is delivered together with the software packages in IBM Passport Advantage®, or in the Quick Start DVD in physical media deliveries.

Tip: If you download the English version PDF files to the 'manuals' directory in your solidDB installation directory, you can access the manuals also through the Welcome page of your solidDB software package. For detailed instructions, see section "Installing solidDB Documentation package" on page 23.

Structure of solidDB documentation

Common functionality of solidDB as a stand-alone server or as part of the Universal Cache is documented in:

- *IBM solidDB Getting Started Guide*
- *IBM solidDB Administrator Guide*
- *IBM solidDB Programmer Guide*
- *IBM solidDB SQL Guide*
- *IBM solidDB Linked Library Access User Guide*

The operation of the solidDB Universal Cache is described in:

- *IBM solidDB Universal Cache User Guide*

High Availability (HotStandby) is described in:

- *IBM solidDB High Availability User Guide*

In-memory database features are described in:

- *IBM solidDB In-Memory Database User Guide*

The operation of Advanced Replication is described in:

- *IBM solidDB Advanced Replication User Guide*

The operation of CDC Replication is described in:

- *IBM solidDB InfoSphere CDC Replication User Guide*

InfoSphere CDC documentation

InfoSphere CDC for solidDB documentation is included in the *IBM solidDB 6.3 Documentation* package. Documentation for InfoSphere CDC Management Console, InfoSphere CDC Access Server, and InfoSphere CDC engine for the other data server is part of the *InfoSphere Change Data Capture v6.3 Documentation* package.

Delivery and location of documentation for InfoSphere CDC components

Table 5. Documentation for InfoSphere CDC components

Component	Document	Documentation delivery
InfoSphere CDC for solidDB	<i>IBM solidDB Universal Cache User Guide</i> <i>IBM solidDB InfoSphere CDC Replication User Guide</i>	Included in the <i>IBM solidDB 6.3 Documentation</i> package
InfoSphere CDC Management Console	<i>InfoSphere Change Data Capture Access Server and Management Console, Installation Guide</i> <i>InfoSphere Change Data Capture Management Console, Administration Guide</i> <i>InfoSphere Change Data Capture API and Management Console Commands Reference</i>	Included in the <i>InfoSphere Change Data Capture v6.3 Documentation</i> package
InfoSphere CDC Access Server	No separate documentation, see the Management Console documentation	No separate documentation, see the Management Console documentation
InfoSphere CDC for the other data server	<i>InfoSphere Change Data Capture, End-User Documentation</i>	Included in the <i>InfoSphere Change Data Capture v6.3 Documentation</i> package

Structure of InfoSphere CDC documentation

- The core functionality of the InfoSphere CDC management tools, that is, Management Console and Access Server, is described in the *InfoSphere Change Data Capture Management Console, Administration Guide*.
- The installation and configuration instructions specific to each InfoSphere CDC for other data server are provided in the data server specific documents called *InfoSphere Change Data Capture, End-User Documentation*.

Similar information for the InfoSphere CDC for solidDB is included in the last chapter of the *IBM solidDB Universal Cache User Guide* and *IBM solidDB InfoSphere CDC Replication User Guide*.

3 Pre-requisites

System requirements

solidDB requires:

- about 48 MB of disk space, including the space for separately installed documentation (The number varies considerably, depending upon the platform.)
- at least 40 MB of RAM in the default configuration
- if you create main-memory tables, then you will need additional memory to store those tables
- adequate disk space for the database(s). An empty database usually requires about 16 MB of disk space.
- enough disk space to accommodate transaction log files preserved for replication recovery (catchup) if InfoSphere CDC technology is used (or, the solidDB Log Reader is enabled). By default, the log retention space required is 10 GB. The size of the log space may be set with a configuration parameter (see the section *Configuring solidDB* in *IBM solidDB Universal Cache User Guide* or *IBM solidDB InfoSphere CDC Replication User Guide*).

InfoSphere CDC for solidDB requires:

- about 220 MB of disk space
- at least 500 MB of RAM

For system requirements of the other InfoSphere CDC components, see the solidDB Web pages at <http://www-01.ibm.com/software/data/soliddb/> or the documentation provided with each InfoSphere CDC component.

Platform support

solidDB is supported on more than 30 different platforms, each understood as a combination of hardware type and operating system. Typically all the current commonly used platforms are supported, as well as some legacy platforms. Support for certain rare platforms may be obtained on request.

For the list of current platforms for the solidDB product family, see the solidDB Web pages at <http://www-01.ibm.com/software/data/soliddb/>.

Available interfaces for client application development

solidDB provides ODBC and JDBC interfaces for clients, which are briefly described in the sections below. For more details, see the *IBM solidDB Programmer Guide*. For information on importing and exporting large amounts of data, see the *IBM solidDB Administrator Guide*, and see the samples in the 'samples/importexport' subdirectory.

JDBC Driver 2.0

Standard Compliance

IBM solidDB JDBC 2.0 Driver supports the JDBC 2.0 specification. Additionally, Connection Pooling, JNDI Data Sources and Rowsets of the JDBC 2.0 Optional Package (known before as Standard Extension) are supported, too.

The JDBC Driver has been successfully tested with JDK versions 1.2.2, 1.3, and 1.4. and certified with the JDBC API Test Suite 1.3.1.

Non-standard features include support for IBM WebSphere® and timeout control extensions. They are discussed below. For information about solidDB JDBC Driver installation, see *JDBC Readme* in the 'doc_html' or 'doc_txt' directory of the solidDB package.

Full documentation is included in the *IBM solidDB Programmer Guide*.

The following features of the Optional Package are currently supported by the solidDB JDBC 2.0 driver :

- Connection pooling (class `solid.jdbc.ConnectionPoolDataSource`)
- Connected RowSet (class `solid.jdbc.rowset.SolidJDBCRowSet`)
- Implemented JDBC data sources:
 - `solid.jdbc.DataSource` (implements `javax.sql.DataSource`)
 - `solid.jdbc.SolidConnectionPoolDataSource` (implements `javax.sql.ConnectionPoolDataSource`)

solidDB JDBC Driver extensions

The non-standard extensions listed below are supported. For more information, see the *IBM solidDB Programmer Guide*.

- **JDBC URL format:** allows to set the connection property values in the URL string
- **Connection timeout:** Connection timeout refers to the response timeout of any JDBC call invoking data transmission over a connection socket. If the response message is not received within the time specified, an I/O exception is thrown. The JDBC standard (2.0/3.0) does not support setting of the connection timeout. The solidDB product has two ways for doing that: one using a non-standard driver manager extension method and the other using the property mechanisms. The time unit in either case is one millisecond.
- **Login timeout:** The timeout fires at the connect time. The setting is implemented with a connection property. The property overrides the login timeouts for JDBC specified by other means (like login timeout parameter in `DriverManager`).
- **Connection idle timeout:** The server closes a connection if it is inactive, for a given time. This is implemented as a connection property and the value overrides the server parameter setting, for this session.
- **Statement Cache:** solidDB JDBC driver enables the user to set the size of a given Connection's statement cache as a property during the connection creation.
- **Transparent Connectivity Support:** solidDB JDBC driver fully supports solidDB Transparent Connectivity (TC) including transparent failover and load balancing. See the *IBM solidDB High Availability User Guide* for more information about usage of Transparent Connectivity.
- **WebSphere Support:** A data sourcer adapter called 'SolidDataStoreHelper' is provided in a separate file 'SolidDataStoreHelper.jar', in the 'jdbc' directory of the solidDB package..

ODBC Driver 3.5.x

solidDB provides two ODBC drivers, one for Unicode and one for ASCII character sets. For more information about these drivers, see the *IBM solidDB Programmer Guide*.

ODBC features not supported in this release

The following functions are not supported in this release:

- SQLBrowseConnect
- SQLSetScrollOptions
- SQLParamOptions
- SQLNativeSql
- SQLMoreResults

ODBC extensions

solidDB ODBC driver incorporates several extensions having to do with timeout control, statement cache behavior and support for Transparent Connectivity. For more information, see the *IBM solidDB Programmer Guide*.

solidDB Java Accelerator

Java Accelerator is a solution allowing to run the solidDB server in the same Java Virtual Machine (VM) process as the Java application. solidDB Java Accelerator is available on the following platforms.

- Linux®
- HP-UX 11
- Microsoft Windows 32-bit
- Solaris 8, 9 and 10

JDK Versions:

- Java environment: JDK 1.4.2
solidDB Java Accelerator has been tested only with Sun J2SE (JDK 1.4.2).

4 Installation

The following sections provide instructions for installing solidDB server and the solidDB Documentation package.

If you are deploying solidDB Universal Cache or InfoSphere CDC replication, see the *IBM solidDB Universal Cache User Guide* and the *IBM solidDB InfoSphere CDC Replication User Guide* for instructions on how to install all the needed InfoSphere CDC components.

Installing solidDB

solidDB can be installed using the interactive installer for graphical user interface or command line, or using a silent installation method.

Graphical interface installation

The GUI installation provides a graphical method for installing solidDB.

Before you begin

In Windows Vista and Windows 2008 Server, you must have Administrator rights to install solidDB.

Procedure

1. Install Java Runtime Environment (JRE) or Java Development Kit (JDK), version 1.4.2 or newer, if not already installed.

JRE or JDK 1.4.2 or newer is needed to run the solidDB installer.

Note: On Linux systems, GNU Compiler for Java (GCJ) is not supported.

2. On the downloaded installation image or the installation DVD, locate the installation program file for your operating system:
 - solidDB-6.3-<platform>.exe (Windows)
 - solidDB-6.3-<platform>.bin (Linux and UNIX)
3. Double-click the installation program file. The solidDB installation wizard starts.
4. Follow the wizard's instructions to complete the installation.

Note: In Linux and UNIX operating systems, you must be able write to the directory that you are using for the installation. If the installation program cannot create the directory, you are prompted to specify a different directory.

Results

You have installed solidDB. To see the solidDB package documentation, open `welcome.html` in the installation directory. Alternatively, you can access the ASCII-formatted files in the `'doc_text'` directory.

What to do next

For information on errors encountered during installation, review the installation log file (IBM_solidDB_6.3_InstallLog.log) located in the installation root directory.

- For information on errors encountered during installation, review the installation log file (IBM_solidDB_6.3_InstallLog.log) located in the installation root directory.
- See also section “Post-installation tasks” on page 22.

Console installation

Use the console installation method to install solidDB from a command-line interface.

Before you begin

In Windows Vista and Windows 2008 Server, you must have Administrator rights to install solidDB.

Procedure

1. Install Java Runtime Environment (JRE) or Java Development Kit (JDK), version 1.4.2 or newer, if not already installed.
JRE or JDK 1.4.2 or newer is needed to run the solidDB installer.

Note: On Linux systems, GNU Compiler for Java (GCJ) is not supported.

2. On the downloaded installation image or the installation DVD, locate the installation program file for your operating system:

- solidDB-6.3-<platform>.exe (Windows)
- solidDB-6.3-<platform>.bin (Linux and UNIX)

3. Start the installation program from the command line. Use the following command:

- Windows

```
<installation_program> -i console
```

For example, in Windows 32-bit operating system:

```
solidDB-6.3-w32.exe -i console
```

- Linux and UNIX

```
sh <installation_program> -i console
```

For example, in Linux 64-bit operating system:

```
sh solidDB-6.3-linux-x86_64.bin -i console
```

4. Follow the displayed instructions to complete the installation.

Note: In Linux and UNIX operating systems, you must be able write to the directory that you are using for the installation. If the installation program cannot create the directory, you are prompted to specify a different directory.

Results

You have installed solidDB. To see the package documentation, open `welcome.html` in the installation directory. Alternatively, you can access the ASCII-formatted files in the 'doc_text' directory.

What to do next

- For information on errors encountered during installation, review the installation log file (IBM_soliddb_6.3_InstallLog.log) located in the installation root directory.
- See also section “Post-installation tasks” on page 22.

Silent installation

Use the silent installation method to install solidDB without any user interaction. This method can be used, for example, for large-scale deployments of solidDB where the silent installation command is embedded in a script.

Procedure

1. Install Java Runtime Environment (JRE) or Java Development Kit (JDK), version 1.4.2 or newer, if not already installed, on all the systems you will be installing solidDB

JRE or JDK 1.4.2 or newer is needed to run the solidDB installer.

Note: On Linux systems, GNU Compiler for Java (GCJ) is not supported.

2. Copy the installation program file from the downloaded installation image or the installation DVD to a temporary directory on your system. Use the installation program for your operating system:

- solidDB-6.3-<platform>.exe (Windows)
- solidDB-6.3-<platform>.bin (Linux and UNIX)

3. Start the installation program from the command line and generate a response file with the following command:

- Windows

```
<installation_program> -r <response-file>
```

For example, in Windows 32-bit operating system:

```
soliddb-6.3-w32.exe -r response.txt
```

- Linux and UNIX

```
sh <installation_program> -r <response-file>
```

For example, in Linux 64-bit operating system:

```
sh solidDB-6.3-linux-x86_64.bin -r response.txt
```

The solidDB installation wizard starts.

4. Follow the displayed instructions to complete the installation.

Note: In Linux and UNIX operating systems, you must be able write to the directory that you are using for the installation. If the installation program cannot create the directory, you are prompted to specify a different directory.

5. After the installation has finished, check that the response file you generated contains the options you need. The response file is located in the installation root directory.

Tip: A sample response file called `installer-response-file.txt` is available in the solidDB installation directory.

6. On another system, perform the silent installation by running the following command:

- Windows

```
<installation_program> -i silent -f <response-file>
```

For example, in Windows 32-bit operating system:

```
soliddb-6.3-w32.exe -i silent -f response.txt
```

- Linux and UNIX

```
sh <installation_program> -i silent -f <response-file>
```

For example, in Linux 64-bit operating system:

```
sh solidDB-6.3-linux-x86_64.bin -i silent -f response.txt
```

Post-installation tasks

After you have installed solidDB, apply the full product license key, check the solidDB JDBC Driver settings, and enable the use of samples.

Procedure

1. If you purchased solidDB or solidDB Universal Cache, copy the license file (solid.lic or soliduc.lic) from the License Certificate image to the solidDB working directory.

Tip: The default installation of solidDB includes an evaluation license (solideval.lic) in the solidDB installation directory. The evaluation license enables you to evaluate solidDB for 90 days.

2. Set the CLASSPATH environment variable for your environment to include the solidDB JDBC Driver .jar file installation path. The solidDB JDBC Driver (SolidDriver2.0.jar) is located in the 'jdbc' directory under the solidDB installation directory.

- **Windows**

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

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

Control Panel > System > Advanced > Environment Variables

- **Linux and UNIX**

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

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

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

If you are using an UNIX shell other than C shell, modify this command to make it appropriate for your shell.

3. Enable the use of samples. In the solidDB installation directory, execute the script called copy_licenses (copy_licenses.bat in Windows). This script copies the evaluation license file from the installation root directory to all appropriate sample directories.

Uninstalling solidDB

In Windows environments, uninstall solidDB through **Control Panel > Add or Remove Programs**. In Linux and UNIX environments, remove the installation directory.

Procedure

- **Windows:**

1. Go to **Control Panel > Add or Remove Programs > Change or Remove Programs**.
2. Select IBM solidDB, and click the **Change/Remove** button.

- **Linux and UNIX:**
 1. Remove the installation directory.

Installing solidDB Documentation package

If you want to access the English version solidDB documentation in PDF format on the same node where your solidDB server is installed, you can download and unzip the solidDB Documentation package into the 'manuals' directory in the solidDB installation directory.

Procedure

1. Locate the *IBM solidDB 6.3 Documentation* package.
 - In physical media deliveries, the documentation package is included in the Quick Start DVD.
 - If you download your software from IBM Passport Advantage, locate the documentation package.
 - If you download your solidDB from the solidDB website at <ftp://ftp.software.ibm.com/software/data/soliddb/info/6.3/man/>, locate the .zip file for the English version manuals.
2. Download and unzip the documentation package into the 'manuals' directory in your solidDB installation directory.

Installing fix packs

Fix packs are distributed as single installation program files.

- “Installing fix packs for solidDB”
- “Installing fix packs for InfoSphere CDC components” on page 24

Installing fix packs for solidDB

Table 6. solidDB point release versions for fix packs

Fix pack	solidDB server point release version
Fix Pack 8	06.30.0049
Fix Pack 7	06.30.0047
Fix Pack 6	06.30.0044
Fix Pack 5	06.30.0040
Fix Pack 4	06.30.0039
Fix Pack 3	06.30.0037
Fix Pack 2	06.30.0033
Fix Pack 1	06.30.0029

1. **Make a backup copy of your database files, log files and the solid.ini configuration file.**
For instructions, see section *Performing backup and recovery* in the *IBM solidDB Administrator Guide*.
2. **Shut down solidDB.**
For instructions, see *Shutting down solidDB*.
3. **Run the solidDB fix pack installer according to the instructions in Installing solidDB.**

solidDB fix packs are delivered as single installation files. To replace your existing solidDB installation, use the same installation directory where your existing solidDB is installed.

If you do not want to run the installer on top of your existing solidDB installation (for example, on your production environment node), use a separate directory and copy the executables, libraries, and drivers manually, as applicable for your setup.

Installing fix packs for InfoSphere CDC components

Before you begin

To ensure that configuration information for your instances, datastores, and subscriptions is retained during the installation, complete the following steps before you start installing any InfoSphere CDC components:

1. End replication on all subscriptions.
2. Disconnect from Access Server.
3. Exit Management Console.
4. Stop all InfoSphere CDC instances in the InfoSphere CDC for solidDB Configuration Tool and InfoSphere CDC for back-end data server Configuration Tool.

For instructions, see

- *IBM InfoSphere Change Data Capture Management Console, Administration Guide*, included in the *InfoSphere Change Data Capture v6.3 Documentation* package
- Stopping InfoSphere CDC

Installing fix packs for Access Server

Important: You must update Management Console and Access Server to the same fix pack level.

1. Uninstall the Access Server.
2. Run the Access Server fix pack installer according to the instructions in the *IBM InfoSphere Access Server and Management Console Installation Guide*, included in the *InfoSphere Change Data Capture v6.3 Documentation* package.

Note:

- Use the same installation directory where the previous version of the Access Server was installed.
- You cannot run two versions of the Access Server on the same node.

Installing fix packs for Management Console

Important: You must update Management Console and Access Server to the same fix pack level.

1. Uninstall the Management Console.
2. Run the Management Console fix pack installer according to the instructions in the *IBM InfoSphere Access Server and Management Console Installation Guide*, included in the *InfoSphere Change Data Capture v6.3 Documentation* package.

Note: Use the same installation directory where the previous version of the Management Console was installed.

Installing fix packs for InfoSphere CDC for solidDB

Important: To install a fix pack for InfoSphere CDC for solidDB, you must already have an installation of InfoSphere CDC for solidDB.

1. Run the InfoSphere CDC for solidDB fix pack installer according to the instructions in Installing InfoSphere CDC.

Note: Use the same installation directory where the previous version of the InfoSphere CDC for solidDB is installed; the installer will prompt you to upgrade the installation.

Installing fix packs for InfoSphere CDC for *back-end data server*

Important: To install a fix pack for InfoSphere CDC for *back-end data server*, you must already have an installation of InfoSphere CDC for *back-end data server*.

1. Run the InfoSphere CDC for *back-end data server* fix pack installer according to the instructions in the *IBM InfoSphere Change Data Capture, End-User Documentation* for the back-end data server, included in the *IBM InfoSphere Change Data Capture v6.3 Documentation* package.

Note: Use the same installation directory where the previous version of the InfoSphere CDC for *back-end data server* is installed; the installer will prompt you to upgrade the installation.

Restarting replication on subscriptions after installation

After you have installed all the fix pack components:

1. Ensure that solidDB and your back-end data server are running.
2. Start the InfoSphere CDC instances in the InfoSphere CDC for solidDB Configuration Tool and InfoSphere CDC for back-end data server Configuration Tool.
3. Log into Management Console.
4. Start replication on the subscriptions.

For instructions, see

- Starting InfoSphere CDC
- *IBM InfoSphere Change Data Capture Management Console, Administration Guide*, included in the *InfoSphere Change Data Capture v6.3 Documentation* package

5 Starting IBM solidDB and creating your first database

solidDB configuration file

When you start solidDB, it reads configuration parameters from `solid.ini` configuration file. You can also use solidDB without the configuration file, in which case the default settings are used.

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

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

Note: Note for evaluators:

The solidDB configuration file will serve most evaluation needs as far as 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 configuration files 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 solidDB components.

Starting solidDB

Procedure

To start solidDB:

- On Windows operating systems, click the icon labeled **Start IBM solidDB server** through the **Start > Programs** menu path.
- On Linux or UNIX operating systems, enter the 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.

Results

When you start solidDB, it checks if a database already exists. If no database is found, the engine automatically creates a new database using the settings in the `solid.ini` configuration file. However, if you want to create a database manually to a specific location, refer to “Creating a database manually (example).”

Creating a database manually (example)

This section provides instructions for creating a new database for evaluation purposes in the '`eval_kit\standalone`' directory in Windows environment.

Procedure

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

```
bin\solid -c eval_kit\standalone
```

A window with the title Creating a new database will open.

Note: If you have already created a database in the '`\eval_kit\standalone`' directory, the command will start the solidDB server without opening the Creating a new database window.

2. 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.

You can use, for example, "dba" for the username and password of evaluation databases, since it is simple and easy to remember. To use it:

- a. Type any name you prefer as the catalog name.
- b. Type "dba" as username and password.
- c. Retype "dba" as the password
- d. Press **OK**.

A new database is created and it is running.

Results

If you successfully created the database, your solidDB process is now running. The process name is solid.exe in Windows operating systems.

Creating a database automatically (example)

When you start solidDB through the **Start > Programs** menu on Windows operating systems and no database is found, the engine automatically creates a new database.

About this task

This section provides instructions for creating a new database for evaluation purposes in the '`eval_kit\standalone`' directory in Windows environment.

Procedure

Start solidDB through the **Start > Programs** menu path. If no database is found, the engine automatically creates a new database using the settings in the `solid.ini` configuration file.

A new database is created and it is running.

Results

If you successfully created the database, your solidDB process is now running. The process name is solid.exe in Windows operating systems.

6 Connecting to solidDB for the first time

If you successfully created the database, your solidDB process is now running, and you can connect to it from any client application by using the communication ports that the solidDB process is listening to. You can verify your setup by trying to connect to the evaluation database. This can be done, for example, by using the command-line oriented solidDB SQL Editor (solsql).

If you have problems in connecting to the database or if any errors occur, see the *IBM solidDB Administrator Guide*.

Using solidDB SQL Editor (solsql)

With solidDB SQL Editor (solsql), SQL statements (including administrative commands prefixed with ADMIN) can be issued at the command line, command prompt, or by executing a script file that contains the SQL statements.

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

```
ADMIN COMMAND 'help';
```

Starting and shutting down the solidDB SQL Editor Procedure

1. To start the solidDB SQL Editor:
 - On Windows operating systems, start the solidDB SQL Editor from the icon in the **Start** menu.
 - In all other environments, enter the SQL editor 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 on the same computer, as where you are executing the command:

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

2. To shut down the solidDB SQL Editor, enter the command:
exit;

Connecting to a database in the solidDB SQL Editor

If you entered a valid user name and password when starting the solidDB SQL Editor, you are connected to the database. If you did not give your user name and password when starting the solidDB SQL Editor, you will be prompted for them.

Procedure

Enter your username and password (for example "dba", "dba") in the solidDB SQL Editor.

Results

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



Figure 3. solidDB SQL Editor Connected

Entering queries in the solidDB SQL Editor

You can execute SQL queries in the solidDB SQL Editor.

Procedure

1. Enter an SQL query in solsql, for example:
`select table_name from tables;`
2. Press **Enter** to execute the query.

Checking the database status

The general server status can be retrieved in the solidDB SQL Editor (solsql).

Procedure

1. Issue the following command in the solidDB SQL Editor:
`ADMIN COMMAND 'status';`

An example output from the solidDB SQL Editor is shown below:

```

c:\ solsql SQL editor
ADMIN command 'status';
RC TEXT
-----
0 IBM solidDB started at 2008-10-29 16:57:24
0 Current directory is C:\Program Files\IBM\solidDB\solidDB6.3\eval_kit\
standalone
0 Using configuration file C:\Program Files\IBM\solidDB\solidDB6.3\eval_
kit\standalone\solid.ini
0 Memory statistics:
0 39158 kilobytes
0 Process size statistics:
0 Resident set size: 16316 kilobytes
0 Virtual size: 42592 kilobytes
0 Transaction count statistics:
0 Commit Abort Rollback Total Read-only Trxbuf Active Validate
0 108 0 1 109 237 0 1 0

0 Cache count statistics:
0 Hit rate Find Read Write
0 100.0 28789 1 31
0 Database statistics:
0 Index writes 3623 After last merge 0
0 Log writes 2268 After last cp 0
0 Active searches 0 Average 1
0 Database size 16128 kilobytes
0 Log size 16 kilobytes
0 User count statistics:
0 Current Maximum Total
0 1 1 1
23 rows fetched.

```

Figure 4. Example output from solidDB SQL Editor (solsql)

2. View all parameter settings.

ADMIN COMMAND 'par';

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

ADMIN COMMAND 'par section_name';

7 Stopping and restarting the database

Stopping the database with solidDB SQL Editor

To stop the database, you must prevent users from connecting to the database, disconnect all connected users, and then shut down solidDB.

About this task

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

Procedure

1. Prevent additional users from connecting to the database:
`admin command 'close';`
2. Throw out all connected users (except the one who issued the command):
`admin command 'throwout all';`
3. Shut down solidDB:
`admin command 'shutdown';`
4. Exit from the solsql tool:
`exit;`

Note: The first three commands can be replaced with this one: `admin command 'shutdown force';`

Results

The solidDB database is shutdown with all previously connected users disconnected.

Note:

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.
```

Restarting the database (example)

You can restart solidDB by issuing commands or by using the **Start > Programs** menu on Windows operating systems.

About this task

This section provides instructions for restarting an evaluation database in the `eval_kit\standalone` directory.

Procedure

Choose one of the following methods to restart the database:

- In Windows, use the **Start > Programs** menu.
- In Linux, UNIX or Windows command prompt, move to the solidDB installation root directory and enter the following command:
`bin\solid -c .\eval_kit\standalone`
- If you have the solidDB bin directory in your path, go to the 'eval_kit\standalone' directory to make it your current working directory, and start solidDB by executing the command `solid`.

8 Running samples

You will find sample programs and scripts in the 'samples' directory in the solidDB installation directory. Each sample directory includes also a *readme.txt* file that provides instructions for how to use the samples.

About this task

In order to be able to run samples, you must copy a license file into each of the sample run directories. You can use the evaluation license (*solideval.lic*) or the full product license (*solid.lic* or *soliduc.lic*).

Tip: The evaluation license enables you to evaluate solidDB for 90 days. For acquiring a permanent license, please contact IBM Corporation.

Procedure

1. Copy the license file to the appropriate sample directories.
 - Execute the script called *copy_licenses* (*copy_licenses.bat* in Windows). This script copies the evaluation license file from the installation root directory to all appropriate sample directories.
 - or
 - Copy the license file to the appropriate sample directories manually.
2. Run the sample according to instructions provided in the *readme.txt* file, located in the sample directory.

Note: Some sample scripts may remove and recreate database files residing in sample subdirectories. Only the database in the 'standalone' directory is always left intact.

9 Using solidDB with graphical SQL clients

solidDB does not include a graphical user interface (GUI). However, graphical SQL clients such as IBM Data Studio Administrator, Eclipse SQL Explorer, or Squirrel SQL Client can be used with solidDB to browse database objects, view and modify table data, and issue SQL commands.

- IBM Data Studio Administrator and Eclipse SQL Explorer clients are ideal if your development environment already uses the Data Studio or Eclipse frameworks.
- Squirrel SQL Client is a standalone program that requires no other components than the client and solidDB to work.

All the above mentioned clients use the JDBC interface to connect to solidDB.

IBM Data Studio Administrator

IBM Data Studio Administrator is a powerful and flexible tool that helps you manage your database objects and simplifies the process of identifying, analyzing, and implementing database schema changes.

You can connect Data Studio Administrator to a solidDB database using the Data Studio Administrator's Generic JDBC database manager options.

For more details on setting up Data Studio Administrator with solidDB, see the following articles on ibm.com[®]:

- *Configuring Data Studio Administrator V2.1 for use with solidDB:*
<http://www-01.ibm.com/support/docview.wss?rs=3457&uid=swg21411299>
- *Configuring Data Studio Administrator V1.2 for use with solidDB:*
<http://www-01.ibm.com/support/docview.wss?rs=3457&uid=swg21411276>

Eclipse SQL Explorer

Eclipse SQL Explorer is a thin SQL client that allows you to query and browse any JDBC compliant database, including solidDB.

The Eclipse SQL Explorer can be added as a plugin to your existing Eclipse environment, or you can run it as a standalone client.

For more details on setting up SQL Explorer with solidDB, see the following article on ibm.com:

- *Configuring Eclipse SQL Explorer for use with solidDB:* <http://www-01.ibm.com/support/docview.wss?rs=3457&uid=swg21411273>

Squirrel SQL Client

Squirrel SQL Client is an open-source Java SQL Client program for JDBC compliant databases. It is a standalone client that can be configured to connect to your solidDB database via the solidDB JDBC driver.

For more details on setting up Squirrel SQL Client with solidDB, see the following article on ibm.com:

- *Configuring Squirrel SQL for use with solidDB*: <http://www-01.ibm.com/support/docview.wss?rs=3457&uid=swg21411262>

10 Upgrading solidDB to a new release level

solidDB supports upgrades from at least the two previous release levels. Upgrades to solidDB 6.3 are supported from the 6.1 and 6.0 release levels; conversion of the database files is however needed. If you need to upgrade from older versions, contact solidDB Technical Support for assistance.

Before you begin

Familiarize yourself with the new and changed features described in the solidDB 6.3 *Release Notes*, available in the **IBM solidDB 6.3 and IBM solidDB Universal Cache 6.3 Information Center** at <http://publib.boulder.ibm.com/infocenter/soliddb/v6r3/index.jsp>.

About this task

The steps below describe the procedure for upgrading a single solidDB server to a new release level. If you are using solidDB in High Availability mode, see section *Upgrading your server by using HotStandby* in the *IBM solidDB High Availability User Guide*.

The 6.1 and 6.0 release level solidDB JDBC and ODBC drivers are compatible with the 6.3 release level solidDB server. The drivers need to be upgraded only if you plan to take into use the new features introduced at release level 6.3.

Procedure

1. Make a backup copy of your database files and solid.ini configuration files.

For instructions, see section *Performing backup and recovery* in the *IBM solidDB Administrator Guide*.

2. Shut down solidDB.

3. Install the new version of solidDB.

solidDB is delivered as a single installation file. If you do not wish to run the installer on your production environment node, install solidDB on a separate node and copy the executables, libraries, and drivers manually to your production node, as applicable for your setup.

For a description of the solidDB server package contents, see “solidDB server package” on page 9.

For instructions on how to run the solidDB installer, see “Installing solidDB” on page 19.

4. Update the license file.

Copy the new license file (solid.lic or soliduc.lic) from the License Certificate image to the solidDB working directory.

5. Check the changes in parameter factory values and update the solid.ini configuration files if necessary.

Changes in the parameter values are described in the Release Notes.

6. Check and update the solidDB JDBC and ODBC driver related settings as applicable to your environment.

- solidDB JDBC Driver

If you installed the new solidDB JDBC Driver (SolidDriver2.0.jar) into a different location than the previous one, add the new driver location to the CLASSPATH System environmental variable. By default, the solidDB JDBC driver is installed to the 'jdbc' directory in the solidDB installation directory.

- solidDB ODBC Driver
 - In Windows and Linux operating systems, define the new ODBC data sources. In order to use the new driver, you must either modify the existing data sources or create new ones.
 - If the application links to the solidDB ODBC driver directly, you must recompile your application.

7. Convert the existing database files by starting solidDB with the -x autoconvert or -x convert command line option.

In the solidDB working directory, execute the following command:

```
solid -x autoconvert
```

or

```
solid -x convert
```

The -x autoconvert option converts the database and starts solidDB.

The -x convert option converts the database and exits solidDB.

Index

Special characters

- xautoconvert command line option
 - upgrading 39
- xconvert command line option
 - upgrading 39

A

- automatic creation of a database 28

C

- Client-Server architecture 1
 - client 1
 - server 1
- connecting
 - solidDB databases 29
- creating
 - databases
 - automatically 28
 - manually 27

J

- JDBC 1

L

- linked library access 1

O

- ODBC 1

R

- restarting
 - solidDB databases 33

S

- solid.ini 27
- solidDB configuration file
 - overview 27
- solidDB directory structure 9
- solidDB SDK 9
- solidDB server package 9
- solidDB SQL Editor (solsql)
 - entering queries 30
 - starting 29
 - stopping 29
- SQL queries
 - executing in solidDB SQL Editor 30
- stopping
 - solidDB databases 33

T

- TCP/IP 1

U

- uninstall 22
- upgrading
 - xautoconvert command line option 39
 - xconvert command line option 39

Notices

© Copyright Oy International Business Machines Ab 1993, 2011.

All rights reserved.

No portion of this product may be used in any way except as expressly authorized in writing by Oy International Business Machines Ab.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the

names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

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

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

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

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

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

Trademarks

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

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

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

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

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

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



Printed in USA

GI11-9212-03

