IBM solidDB
IBM solidDB Universal Cache
Version 6.5

Getting Started Guide

GI11-9229-06

Note
Before using this information and the product it supports, read the information in "Notices" on page 51.

First edition, sixth revision

This edition applies to version 6, release 5, Fix Pack 14 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.

Note: Information in this document might have been updated since this document was made available. For the latest available information, see the *IBM solidDB and IBM solidDB Universal Cache V6.5 Information Center* at http://pic.dhe.ibm.com/infocenter/soliddb/v6r5/index.jsp.

Contents

Figures

Tables

Summary of changes

Changes for revision 06

- Editorial corrections.

Changes for revision 05

- Section Installing solidDB® ODBC Driver updated: you can use the installation program to install the ODBC driver only on Windows environments. On Linux and UNIX environments, to install the ODBC driver to a different machine than where solidDB is installed, you must copy the ODBC driver library files manually.

Changes for revision 04

- Editorial corrections.

Changes for revision 03

- New sections added:
  - solidDB supported platforms
  - solidDB Universal Cache supported platforms
  - Supported back-end data servers for solidDB Universal Cache

Changes for revision 02

- Editorial corrections

Changes for revision 01

- Section Library file names updated with library file naming conventions.
- OS user limit requirements (Linux and UNIX) added.
- Section Post-installation tasks updated: it is no longer necessary to execute the `copy_licenses` script manually after installation. The solidDB installation program now places a copy of the evaluation license to each of the samples directories.
- New sections added:
  - IBM® Data Server Driver for ODBC and CLI package
  - Using solidDB with graphical SQL clients

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 about 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.
argument	Words emphasized like this indicate information that the user or the application must provide.
Administrator Guide	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` *`table_name`*	Syntax descriptions are on this font. Replaceable sections are on *this* 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 Overview of 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 6.5 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.

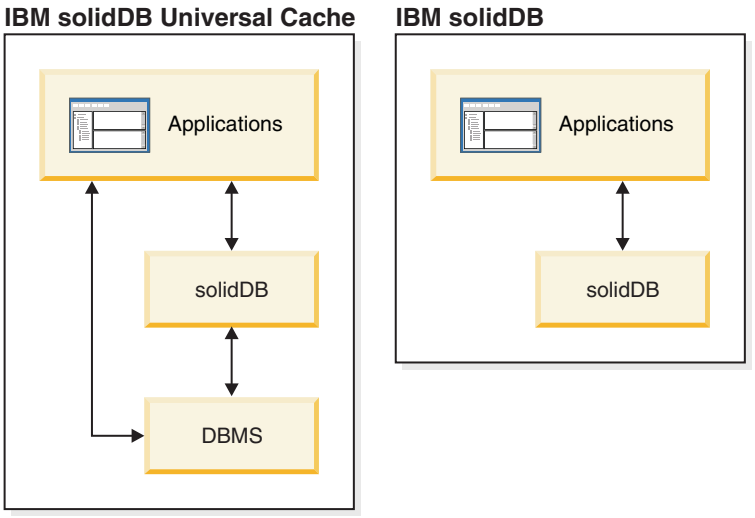

Figure 1. solidDB 6.5 products

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 *shared memory access* (SMA) and *linked library access* (LLA). In this case, the 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. By replacing the network connection with local function calls, performance is improved significantly. For more information, see the *IBM solidDB Shared Memory Access and 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.

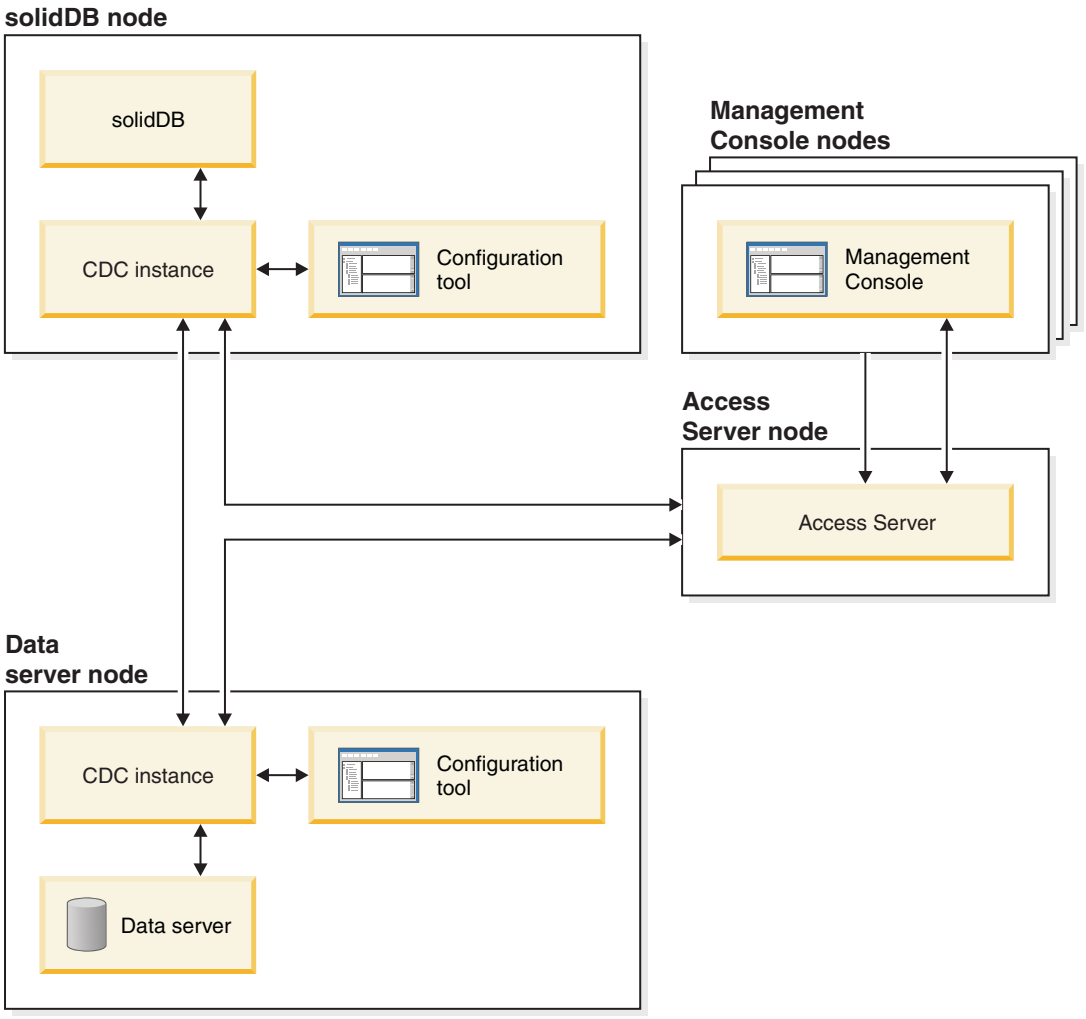

Figure 2. solidDB Universal Cache architecture

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 runtime 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 can 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 data server, 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 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 can 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 can 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.

Shared memory access (SMA) and Linked library access (LLA)

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

For more information, see the *IBM solidDB Shared Memory Access and 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 can be 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 notebooks, 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 Replication with InfoSphere CDC User Guide*.

- **In solidDB Universal Cache**, the InfoSphere CDC technology can be used to set up 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 following table 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	X	X
InfoSphere Change Data Capture solidDB	X[1]	X
InfoSphere Change Data Capture Management Console	X[1]	X
InfoSphere Change Data Capture Access Server	X[1]	X
InfoSphere Change Data Capture *back-end data server* One of the following: • InfoSphere Change Data Capture DB2 Linux, UNIX, and Windows • InfoSphere Change Data Capture Informix® • InfoSphere Change Data Capture Microsoft SQL Server • InfoSphere Change Data Capture Oracle Trigger • InfoSphere Change Data Capture Oracle Redo • InfoSphere Change Data Capture Sybase • InfoSphere Change Data Capture DB2 z/OS® • InfoSphere Change Data Capture DB2 iSeries®		X
IBM Data Server Driver for ODBC and CLI V9.7		X[2]

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

[2] Needed only in those solidDB Universal Cache configurations deploying SQL passthrough in which the back-end data server is an IBM data server.

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, *solideval.lic*. The evaluation license enables you to evaluate solidDB for 90 days. For acquiring a permanent license, contact IBM Corporation.

Directory structure

The default installation of solidDB 6.5 creates a directory called `solidDB6.5`.

The files and subdirectories in the `solidDB6.5` installation directory are explained in the following table.

Table 4. `solidDB6.5` directory structure

Location	Explanation
Root directory	The root directory contains, for example: • a script used to facilitate running samples in the database evaluation phase • the evaluation license file • the `welcome.html` file for accessing the package documentation
`bin`	solidDB binary files and dynamic library files
`doc_html`, `doc_txt`	Package documentation in HTML and text format
`eval_kit/standalone`	Working directory for an evaluation version of the solidDB server. This directory contains a sample `solid.ini` configuration file and an evaluation license file (`solideval.lic`).
`eval_kit/cdc`	Working directory for an evaluation version of the solidDB server for use with Universal Cache or InfoSphere CDC replication. This directory contains a sample `solid.ini` configuration file and an evaluation license file (`solideval.lic`).
`include`	C program headers
`jdbc`	JDBC Driver for solidDB Data store helper archive for use with WebSphere® (`SolidDataStoreHelper.jar`) solidDB dialect for Hibernate (`SolidSQLDialect.jar`)
`lib`	Static linkable library files
`lib32`	32-bit static linkable library files – 64-bit AIX® and Solaris packages only
`licence`	License and notices files
`manuals`	The English versions of the manuals in PDF format can be downloaded to this folder and then accessed through the **Manuals** link on the **Welcome**-page
`procedures`	SQL scripts for creating and running stored procedures for data aging and refresh

Table 4. solidDB6.5 directory structure (continued)

Location	Explanation
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 the library files fall into one of the following categories:

- ODBC drivers
- solidDB shared memory access and 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 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. In Windows environments, the lib directory also contains the import libraries.

The exact library file names depend on the platform. See the following tables for examples in Windows and Linux environments:

Table 5. Example: solidDB library files in Windows 32-bit package

File name	Description
bin\	
sacw3265.dll	ODBC library - ASCII
snpw3265.dll	NamedPipes communication protocol link library
socw3265.dll	ODBC library - Unicode
sosw3265.dll	ODBC Driver Manager setup library
ssaw3265.dll	solidDB SA API library
ssolidac65.dll	Linked library access (LLA) dynamic library
stcw3265.dll	TCP/IP communication protocol link library
lib\	

Table 5. Example: solidDB library files in Windows 32-bit package (continued)

File name	Description
solidctrlstub.lib	solidDB Control API (SSC) stub library.
	This static library is used if you want to write code that can be run either locally with the linked library access library, or remotely without the linked library access.
solidlight.lib	solidDB Light Client library
solidimpac.lib	Linked library access (LLA) import library
solidimpodbca.lib	ODBC import library - ASCII
solidimpodbcu.lib	ODBC import library - Unicode
solidimpsa.lib	solidDB SA API import library

Table 6. Example: solidDB library files in Linux 32-bit package

File name	Description
bin\	
sacl2x65.so	ODBC shared library - ASCII
socl2x65.so	ODBC shared library - Unicode
ssal2x65.so	solidDB SA API library
ssolidac65.so	Linked library access (LLA) shared library
ssolidsma65.so	Shared memory access (SMA) shared library
lib\	
solidctrlstub.a	solidDB Control API (SSC) stub library.
	This static library is used if you want to write code that can be run either locally with the linked library access library, or remotely without the linked library access.
solidlight.a	solidDB Light Client library
solidac.a	Linked library access (LLA) static library
solidodbca.a	ODBC static library - ASCII
solidodbcu.a	ODBC static library - Unicode
solidsa.a	solidDB SA API static library
libssolidac65.so	Symbolic link for shared LLA library
libssolidsma65.so	Symbolic link for shared SMA library
libsacl2x65.so	Symbolic link for shared ODBC library - ASCII
libsocl2x65.so	Symbolic link for shared ODBC library - Unicode
libssal2x65.so	Symbolic link for shared solidDB SA API library
libsolidodbca.a	Symbolic link for static ODBC library - ASCII
libsolidodbcu.a	Symbolic link for static ODBC library - Unicode
libsolidsa.a	Symbolic link for static solidDB SA API library
libsolidac.a	Symbolic link for static LLA library

For a list of the library file names in your installation of solidDB, see the SDK Notes in the solidDB package, accessible through the **Welcome**-page in your solidDB installation directory.

Dynamic library file naming conventions

Dynamic library files use the following naming convention:

```
sLLpppVV.eee
```

where

- `LL` = library's purpose
  - ac: ODBC library - ASCII
  - np: NamedPipes communication protocol link library
  - oc: ODBC library - Unicode
  - os: ODBC Driver Manager setup (for Windows only)
  - sa: solidDB SA API library
  - solidac: Linked library access (LLA) dynamic library
  - solidsma: Shared memory access (SMA) dynamic library
  - tc: TCP/IP communication protocol link library
- `ppp` = platform
  - a5x64: AIX, 64-bit
  - hia64: HP-UX 11 64-bit (IA64)
  - l2x: Linux for x86
  - l2x64: Linux for x86, 64-bit
  - s0x64: Solaris 10 (SPARC, 64-bit)
  - s0xi64: Solaris 10 (ix86, 64-bit)
  - w32: Windows 32-bit (x86)
  - w64: Windows 64-bit (x86)
- `VV` = first two digits of the solidDB version, for example 65 for version 6.5, 63 for version 6.3
- `eee` = platform-specific filename extension:
  - `*.dll` Dynamic Link Library for Windows
  - `* .so` (Shared Object) for AIX, Solaris and Linux
  - `* .sl` (Shared Library) for HP-UX

ODBC, JDBC, and proprietary programming interfaces

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*.

solidDB JDBC Driver 2.0

Table 7. solidDB JDBC Driver 2.0 key information

Compatibility	JDBC 2.0, with selected features of JDBC 2.0 Optional Package
Driver location	`<solidDB installation directory>/jdbc/SolidDriver2.0.jar`

JDBC URL format	`jdbc:solid://<hostname>:<port>/<username>/<password>[?<property-name>=<value>]...`
	
	For example:
	`"jdbc:solid://localhost:1964/dba/dba"`

Driver class name	solid.jdbc.SolidDriver

Standard Compliance

The 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 solidDB 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.

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.sqlDataSource)
  - solid.jdbc.SolidConnectionPoolDataSource (implements javax.sql.ConnectionPoolDataSource)
- JTA (Java Transaction API) – XA interface for Java (implements javax.transaction.xa.XAResource and javax.transaction.xa.Xid)

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

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 1 ms.
- **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 Driver Manager).
- **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.
- **Shared memory access (SMA) connection:** allows to connect to a SMA server
- **SQL passthrough:** allows to set the SQL passthrough mode

- **Catalog and schema name properties:** allows to set the catalog name and schema name for the solidDB database
- **WebSphere Support:** A data source adapter called SolidDataStoreHelper is provided in a separate file `SolidDataStoreHelper.jar`, in the `'jdbc'` directory of thesolidDB package.

solidDB 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*.

The following functions are not supported:

- 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*.

Proprietary interfaces

The solidDB Application Programming Interface (SA API) and solidDB Server Control API (SSC API) allow, for example, C programs to directly call functions inside the database server. These proprietary interfaces are provided with the solidDB shared memory access (SMA) and linked library access (LLA) libraries.

System tools and utilities

solidDB includes two console tools, solidDB Remote Control (**`solcon`**) and solidDB SQL Editor (**`solsql`**), for data management and administration, as well as command-line utilities for exporting and loading data to solidDB databases.

The solidDB data management tools are included in the solidDB package, available in the `'bin'` directory in the solidDB installation directory.

Console tools: solidDB SQL Editor (solsql) and solidDB Remote Control (solcon)

- solidDB SQL Editor (**`solsql`**) is a console tool used to issue SQL statements and solidDB ADMIN COMMANDs at the command prompt, or by executing a script file that contains the SQL statements.
- solidDB Remote Control (**`solcon`**) is a console tool for administration; users with administrator rights can issue ADMIN COMMANDs at the command prompt or by executing a script file that contains the commands. With **`solcon`**, the ADMIN COMMANDs can be issued as part of the **`solcon`** startup command line.

Tools for exporting and loading data

solidDB provides the following tools for exporting and loading data:

- solidDB Speed Loader (**solloado** or **solload**) loads data from external files into a solidDB database.
- solidDB Export (**solexp**) exports data from a solidDB database to files. It also creates control files used by solidDB Speed Loader (**solloado** or **solload**) to perform data load operations.
- solidDB Data Dictionary (**soldd**) exports the data dictionary of a database. It produces an SQL script that contains data definition statements that describe the structure of the database.

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 packages

The InfoSphere CDC components are delivered as separately deployable packages.

Table 8. InfoSphere CDC installation images

InfoSphere CDC component name	Installation package	Contents
InfoSphere Change Data Capture solidDB (InfoSphere CDC for solidDB)	**Linux and UNIX:** setup-<platform>-solid.bin For example: setup-linux-x86-solid.bin **Windows:** setup-x86-solid.exe	• Software for the configuration tool and the InfoSphere CDC instance for solidDB • solidDB JDBC Driver (`SolidDriver2.0.jar` in the `/lib` directory) • Tools, utilities and samples (`/samples` directory) – Automation tools, utilities, and samples for scripting most common InfoSphere CDC tasks (`ucutils`, `ucpassthrough`, and `uchsbmonitor` directories) – Generic InfoSphere CDC samples for Java user exits and SQL scripts • InfoSphere CDC API documentation (`/docs` directory)
InfoSphere Change Data Capture *for a back-end data server* **Note:** Included in solidDB Universal Cache only	**Linux and UNIX:** setup-<platform>-<back-end_dataserver>.bin For example: setup-aix-power-udb.bin **Windows:** setup-x86-<back-end_dataserver>.exe	• Software for the configuration tool and the InfoSphere CDC instance for *back-end data server* • PDF format *InfoSphere Change Data Capture, End-User Documentation* (`/docs` directory) • Sample Java user exits and SQL scripts (`/samples` directory) • InfoSphere CDC API documentation (`/docs` directory)

Table 8. InfoSphere CDC installation images (continued)

InfoSphere CDC component name	Installation package	Contents
InfoSphere Change Data Capture Access Server	**Linux and UNIX:** setup-<platform>-accessserver.bin For example: setup-solaris-sparc-accessserver.bin **Windows:** setup-win-x86-accessserver.exe	• Software for controlling access to your replication environment
InfoSphere Change Data Capture Management Console	**Linux and UNIX:** N/A - available only for Windows **Windows:** setup-x86-dmclient.exe	• Software for configuring and monitoring InfoSphere CDC user access and replication subscriptions in the solidDB Universal Cache and InfoSphere CDC Replication configuration setups • PDF format *InfoSphere Change Data Capture Management Console, Administration Guide* (`/documentation` directory) • Online help (accessible through Help menu in the Management Console user interface) • IBM Java SDK and Runtime Environment Guides (`/docs` directory)

IBM Data Server Driver for ODBC and CLI package

The IBM Data Server Driver for ODBC and CLI is delivered as a compressed file. It is used with the SQL passthrough feature in solidDB Universal Cache if the back-end data server is an IBM data server.

- Windows operating systems:
  `ibm_data_server_driver_for_odbc_cli_<platform>.zip`
- Linux and UNIX operating systems:
  `ibm_data_server_driver_for_odbc_cli_<platform>.tar.Z`

There is no installation program for the IBM Data Server Driver for ODBC and CLI. Instead, you must install the driver manually by uncompressing the file.

Documentation packages

Documentation for solidDB is composed of an *IBM solidDB Documentation* package and an *InfoSphere Change Data Capture 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 components for the back-end data servers.

solidDB documentation

solidDB documentation is available online in the solidDB 6.5 and solidDB Universal Cache 6.5 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 6.5 and solidDB Universal Cache 6.5 Information Center

The most up-to-date solidDB documentation is available in the information center format at http://publib.boulder.ibm.com/infocenter/soliddb/v6r5/.

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/6.5/man/.
- IBM Publications Center at http://www.elink.ibmlink.ibm.com/publications/servlet/pbi.wss

In addition, the PDF format manuals are available as the *IBM solidDB 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 31.

Structure of solidDB documentation

Common functionality of solidDB as a standalone 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 Shared Memory Access and 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 InfoSphere CDC Replication is described in:

- *IBM solidDB Replication with InfoSphere CDC User Guide*

InfoSphere CDC documentation

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

Delivery and location of documentation for InfoSphere CDC components

The *InfoSphere Change Data Capture Documentation* package is available in information center and PDF format:

- IBM InfoSphere Change Data Capture version 6.3.1 Fix Pack 3 Information Center
- InfoSphere Change Data Capture 6.3.1 Fix Pack 3 End-User Documentation in PDF format - IBM Software Support Portal
- Embedded Help accessible through the Management Console **Help** menu
- *InfoSphere Change Data Capture Documentation* installation package (PDF format), available at Passport Advantage

3 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 Replication with InfoSphere CDC User Guide* for instructions on how to install all the needed InfoSphere CDC components.

System requirements

The solidDB product family supports 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.

IBM solidDB supported platforms

The following table shows the supported platforms for the components included in the IBM solidDB 6.5 product offering.

Important:

- solidDB server is supported on all the operating system and software versions or editions shown in the table.
- The InfoSphere CDC components are supported on selected platforms or versions and editions. For example, InfoSphere CDC Access Server is not supported on AIX 7.1 platforms.

  To confirm that all the needed components are supported on your planned environment, contact your IBM representative.

Table 9. Commonly used solidDB platforms

Operating system		Hardware	solidDB	InfoSphere CDC for solidDB	MC	AS
AIX	AIX 7.1 AIX 6.1AIX 5L™ 5.3 patch level 8	64-bit systems with POWER5, POWER6®, or POWER7	X	X		X
HP-UX	HP-UX 11i v3 HP-UX 11i v2	Itanium-based HP Integrity Series systems	X	X		X
Linux	Red Hat Enterprise Linux (RHEL) 6 Red Hat Enterprise Linux (RHEL) 5 Red Hat Enterprise Linux (RHEL) 4 SUSE Linux Enterprise Server (SLES) 11 SUSE Linux Enterprise Server (SLES) 10 SUSE Linux Enterprise Server (SLES) 9	32-bit and 64-bit systems based on Intel or AMD processors that are capable of running the supported Linux operating systems (x86 and x64 systems)	X	X		X
Solaris	Solaris 10	64-bit Systems with UltraSPARC or x86 processors	X	X		X

Table 9. Commonly used solidDB platforms (continued)

Operating system		Hardware	solidDB	InfoSphere CDC for solidDB	MC	AS
Windows	Windows Server 2008 (Standard Server, Enterprise Server, and Datacenter Editions) Windows Server 2003 (Standard Server, Enterprise Server, and Datacenter Editions) Windows Server 2000 (Standard Server, Enterprise Server, and Datacenter Editions) Windows Vista (Business, Enterprise, and Ultimate editions) Windows XP (Professional Edition)	32-bit and 64-bit systems based on Intel or AMD processors that are capable of running the supported Windows operating systems (x86 and x64 systems)	X	X	X	X

Related concepts:

"Additional solidDB server installation requirements" on page 24

IBM solidDB Universal Cache supported platforms

The following table shows the supported platforms for the components included in the IBM solidDB Universal Cache 6.5 product offering.

Important:

- solidDB server is supported on all the operating system and software versions or editions shown in the table, except iSeries and z/OS.
- The InfoSphere CDC components might not be supported on every operating system version or edition mentioned. For example, InfoSphere CDC Access Server is not supported on AIX 7.1 platforms. Check with your IBM representative about your specific configuration.

Table 10. solidDB Universal Cache supported platforms

Operating system		Hardware	solidDB server	InfoSphere CDC											
				solidDB	MC	AS	DB2	DB2 z/OS	DB2 iSeries	IDS	OR	OT	MS SQL	Sybase	ODBC
AIX	AIX 7.1 AIX 6.1 AIX 5L 5.3 patch level 8	64-bit systems with POWER5, POWER6, or POWER7	X	X		X	X			X	X				X
HP-UX	HP-UX 11i v3 HP-UX 11i v2	Itanium-based HP Integrity Series systems	X	X		X				X	X			X	X
Linux	Red Hat Enterprise Linux (RHEL) 6, 5, 4 SUSE Linux Enterprise Server (SLES) 11, 10, 9	32-bit and 64-bit systems based on Intel or AMD processors that are capable of running the supported Linux operating systems (x86 and x64 systems)	X	X		X	X			X	X			X	X
	Red Hat Enterprise Linux (RHEL) 5, 4 System z SUSE Linux Enterprise Server (SLES) 10, 9 System z	System z										X			
Solaris	Solaris 10	64-bit Systems with UltraSPARC or x86 processors	X	X		X	X				X			X	X

Table 10. solidDB Universal Cache supported platforms (continued)

Operating system		Hardware	solidDB server	InfoSphere CDC											
Windows	Windows Server 2008, 2003, 2000 (Standard Server, Enterprise Server, and Datacenter Editions) Windows Vista (Business, Enterprise, and Ultimate editions) Windows XP (Professional Edition)	32-bit and 64-bit systems based on Intel or AMD processors that are capable of running the supported Windows operating systems (x86 and x64 systems)	X	X	X	X	X			X			X	X	X
iSeries	OS/400 5.2	POWER System							X						
z/OS	z/OS 1.4	System z						X							

MC = InfoSphere Change Data Capture Management Console

AS = InfoSphere Change Data Capture Access Server

DB2 = InfoSphere Change Data Capture DB2 Linux, UNIX, and Windows

DB2 z/OS = InfoSphere Change Data Capture DB2 z/OS

DB2 iSeries = InfoSphere Change Data Capture DB2 iSeries

IDS = InfoSphere Change Data Capture Informix

OR = InfoSphere Change Data Capture Oracle Redo

OT = InfoSphere Change Data Capture Oracle Trigger

MS SQL = InfoSphere Change Data Capture Microsoft SQL Server

Sybase = InfoSphere Change Data Capture Sybase

ODBC = IBM Data Server Driver for ODBC and CLI

Related concepts:

"Additional solidDB server installation requirements" on page 24

Supported back-end data servers for solidDB Universal Cache

solidDB Universal Cache supports a number of IBM and other data servers.

IBM DB2 for Linux, UNIX, and Windows

- DB2 Enterprise Server Edition V9.7
- DB2 Enterprise Server Edition V9.5
- DB2 Enterprise Server Edition V9.1

IBM DB2 for iSeries

- DB2 for i/OS V6R1
- DB2 for i/OS V5R4

IBM DB2 for z/OS

- DB2 for z/OS V9
- DB2 for z/OS V8
- DB2 for z/OS V7

IBM Informix Dynamic Server (IDS)

- IDS V11.50.3 Enterprise Edition

Oracle Database

- Oracle Database 11g

- Oracle Database 10g
- Oracle Database 9g

Microsoft SQL Server

- Microsoft SQL Server 2008
- Microsoft SQL Server 2005
- Microsoft SQL Server 2000

Sybase Adaptive Server Enterprise (ASE)

- Sybase ASE V15
- Sybase ASE V12.5.4

Additional solidDB server installation requirements

To install solidDB, you need:

- About 48 MB of disk space, including the space for separately installed documentation – the number varies considerably, depending on the platform
- At least 40 MB of RAM in the default configuration
- Adequate disk space for your database – an empty database typically requires about 16 MB of disk space
- If you use in-memory tables, additional memory to store those tables
- If you use InfoSphere CDC technology (or, the solidDB log reader is enabled), enough disk space to accommodate transaction log files preserved for replication recovery (catchup) – by default, the required log retention space is 10 GB
- Java Runtime Environment (JRE) or Java Development Kit (JDK), version 1.4.2 or newer, is required for
  - solidDB installation program

    Note: On Linux systems, the installation program does not support GNU Compiler for Java (GCJ).
  - Shared memory access (SMA) and linked library access (LLA)

User process resource limits (ulimits) considerations in Linux and UNIX environments

In Linux and UNIX environments, you might need to modify the settings for the user process resource limits (ulimits) of your system. For details, see *OS user limit requirements (Linux and UNIX)*.

Security-enhanced Linux considerations

On Red Hat Enterprise Linux (RHEL) operating systems, if Security-enhanced Linux (SELinux) is enabled and in enforcing mode, the installer might fail due to SELinux restrictions.

To determine if SELinux is installed and in enforcing mode, you can do one of the following:

- Check the `/etc/sysconfig/selinux` file.
- Run the **`sestatus`** command.
- Check the `/var/log/messages` file for SELinux notices.

To disable SELinux, you can do one of the following:

- Set SELinux in permissive mode and run the **setenforce 0** command as a superuser.
- Modify `/etc/sysconfig/selinux` and reboot the machine.

If solidDB installs successfully on a RHEL system, all solidDB processes will run in the unconfined domain. To assign the processes to their own domains, so that also confined users can run them, you need to modify the policy modules. For more information, see *Using solidDB with SELinux* in the *IBM solidDB Administrator Guide*.

OS user limit requirements (Linux and UNIX)

In Linux and UNIX environments, you might need to modify the settings for the user process resource limits (ulimits) of your system.

If required ulimit values are not met, solidDB can encounter unexpected operating system resource shortage errors, such as `SOLID Communication Error 21309: Failed to accept a new client connection, out of TCP/IP resources`.

To set the resource limits permanently on your system, modify the hard operating system ulimits for the **data**, **nofiles**, and **fsize** resources (root user or a System Administrator rights required).

Table 11. Recommended ulimit values

Hard ulimit resource	Description	Recommended value
data	Maximum private memory allowed for a process	Unlimited
nofiles	Maximum number of open files/file descriptors allowed for a process (related to sockets available to the operating system and applications)	65536 or unlimited
fsize	Maximum file size allowed	Unlimited

For instructions on how to query and modify the ulimit values, see your operating system documentation.

Example

To query the hard ulimit values:

```
ulimit -Ha
-t: cpu time (seconds) unlimited
-f: file size (blocks) unlimited
-d: data seg size (kbytes) unlimited
-s: stack size (kbytes) 100000
-c: core file size (blocks) unlimited
-n: file descriptors 2048
-v: virtual memory size (kb) unlimited
```

To set the hard **nofiles** value to 65536:

```
ulimit -Hn 65536
```

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-6.5-<platform>.exe` (Windows)
   - `solidDB-6.5-<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 about errors encountered during installation, review the installation log file (`IBM_solidDB_6.5_InstallLog.log`) located in the installation root directory.
- See also section “Post-installation tasks” on page 28.

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-6.5-<platform>.exe` (Windows)
   - `solidDB-6.5-<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-6.5-w32.exe -i console
     ```

   - Linux and UNIX

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

     For example, in Linux 64-bit operating system:

     ```
     sh solidDB-6.5-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 about errors encountered during installation, review the installation log file (`IBM_solidDB_6.5_InstallLog.log`) located in the installation root directory.
- See also section “Post-installation tasks” on page 28.

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-6.5-<platform>.exe` (Windows)
   - `solidDB-6.5-<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-6.5-w32.exe -r response.txt
     ```

   - Linux and UNIX

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

     For example, in Linux 64-bit operating system:

     ```
     sh solidDB-6.5-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-6.5-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-6.5-linux-x86_64.bin -i silent -f response.txt
     ```

Post-installation tasks

After you have installed solidDB, you may want to 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. If you are planning to use the solidDB JDBC Driver, 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.

Installing JDBC and ODBC drivers

The solidDB JDBC Driver and solidDB ODBC Driver are included in the solidDB server installation package. Both drivers are installed by default during solidDB installation.

Installing solidDB JDBC Driver

The solidDB JDBC Driver (`SolidDriver2.0.jar`) is installed during solidDB server installation. Depending on your environment, you might need to set various configuration settings before using the solidDB JDBC Driver.

Default installation directory

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

If your application is located on a different computer than the solidDB server, you need to copy the JDBC driver file to the computer where the application is located.

Tip:

- The `jdbc` directory contains also the solidDB Data Store Helper Class (`SolidDataStoreHelper.jar`) for use with WebSphere.
- The `samples/jdbc` directory in the solidDB installation directory contains Java code samples that use the solidDB JDBC Driver. Instructions for running the sample are available in the `readme.txt` file, which located in the same directory.

Requirements for Java environment

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

Setting the CLASSPATH environmental variable

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

- **Windows**

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

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

  Control Panel > **System** > **Advanced** > **Environment Variables**
- **Linux and UNIX**

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

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

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

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

Installing solidDB ODBC Driver

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

Windows

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

- Windows 32-bit operating systems:
  - IBM solidDB 6.5 32-bit – ANSI
  - IBM solidDB 6.5 32-bit – Unicode
- Windows 64-bit operating systems:
  - IBM solidDB 6.5 64-bit – ANSI
  - IBM solidDB 6.5 64-bit – Unicode

Linux and UNIX

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

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

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

Installing ODBC drivers without solidDB installation (Windows)

To install the ODBC drivers without installing solidDB in Windows environments:

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

Installing ODBC drivers without solidDB installation (Linux and UNIX)

To install the ODBC drivers without installing solidDB in Linux and UNIX environments:

1. Install solidDB using the installation program.
2. Copy the ODBC driver library file to your client node.

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 extract the solidDB Documentation package into the `'manuals'` directory in the solidDB installation directory.

Procedure

1. Locate the *IBM solidDB 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 FTP site at ftp://ftp.software.ibm.com/software/data/soliddb/info/6.5/man/, locate the .zip file for the English version manuals.
2. Download and extract 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 32

Installing fix packs for solidDB

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 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 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 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 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 Documentation* package

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.

4 Verifying your solidDB installation

After you have installed solidDB server, you can start up solidDB and create your first database to verify that the installation was successful. You can also use the samples provided in the package to explore the solidDB features and functionality.

After a successful installation, you can find the following directories below your solidDB installation directory:

```
<installation directory>
    bin\
    ..
    eval_kit\
        standalone\
        cdc\
    ..
    samples
    ..
```

The `standalone` directory will work as the *working directory* for starting solidDB and creating your first database in the way shown in the sections below. The `standalone` directory contains a sample `solid.ini` configuration file and an evaluation license file (`solideval.lic`). It will also hold your first database (`solid.db`) as well as the related transaction log, message, and trace files.

The `cdc` directory contains a sample `solid.ini` configuration file for use with solidDB Universal Cache. The `cdc` directory can be used as the working directory when setting up solidDB Universal Cache.

Both the `standalone` directory and the `cdc` directory contain a license file that enables you to evaluate solidDB for a limited time. All solidDB and solidDB Universal Cache product options and features are available for evaluation.

Working directory and `solid.ini` configuration file

A working directory is the directory that contains the files related to running a particular solidDB instance, including the `solid.ini` configuration file. The `solid.ini` file specifies parameters that help customize and optimize solidDB.

Working directory

Typically the working directory contains the following files:

- license file
- `solid.ini` configuration file
- database files
- transaction log files
- message and trace files

In typical setups, the working directory contains the valid license file and the `solid.ini` configuration file. You can also specify a different directory for the license file and the `solid.ini` file by creating a SOLIDDIR environment variable that specifies the location.

When you start solidDB, it reads configuration parameters from `solid.ini`. You can also use solidDB without the configuration file, in which case the default settings (factory values) are used. Also, if `solid.ini` exists, but a value for a particular parameter is not set, solidDB will use a factory value for that parameter. The factory values might depend on the operating system you are using.

The database, transaction log, message, and trace files can be located in other directories of your choice; in that case, the directory paths and file names must be defined in the `solid.ini` configuration file. For example, the **`FileSpec`** parameter in the `[IndexFile]` section of the `solid.ini` file specifies the directory and files name for the database file.

`solid.ini` configuration file for evaluation setups

The parameter settings in the sample `solid.ini` configuration file in the `eval_kit\standalone` directory serve most evaluation needs for the solidDB server.

If you want to use other components such as the solidDB Universal Cache or the High Availability feature, the configuration file must be modified accordingly.

The samples also contain `solid.ini` configuration files which contain configuration settings that are needed to run the samples. When evaluating solidDB with the different features and functionality, the `solid.ini` configuration files provided with the corresponding samples can be used as a starting point.

Related topics

For more information about the `solid.ini` configuration file and how to set parameters through it to configure solidDB, refer to the *IBM solidDB Administrator Guide*.

Starting solidDB and creating your first database

This section describes how to start an evaluation version of solidDB, using the license file and the `solid.ini` configuration file in the `eval_kit\standalone` directory.

To start solidDB, a valid license file must be located in your working directory.

Starting solidDB in Linux and UNIX environments

1. If you have several solidDB installations on the same machine, verify that the version 6.5 installation is your default solidDB installation.
   a. To check which **solid** executable is defined in the PATH environment variable, enter the following command at the command prompt:
      ```
      which solid
      ```
   b. To verify the version of the **solid** executable, enter the following command at the command prompt:
      ```
      solid -x version
      ```
2. In the `eval_kit/standalone` directory, enter the command `solid -f` at the command prompt.
   The option `-f` forces the server to run in the foreground.
   When you start solidDB, it checks if a database already exists. If no database is found, solidDB prompts you to create a new database.
3. Type `y` to create a new database.

4. Provide the username, password, and a name for the default database catalog.

   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.

   For example, you can use `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.

Result

solidDB creates a new database called `solid.db` in the `eval_kit/standalone` directory.

The solidDB process is left running in the foreground. To make solidDB run in the background:

1. Press **CTRL+Z** to suspend the process.
2. Enter `bg` to force the process to the background.

Tip: Alternatively, you can start solidDB and create a new database by specifying the database and login data as command-line options.

For example:

```
solid -Udba -Pdba -Cdba
```

This starts solidDB in the background and creates a new database with the following login data:

- Username: dba
- Password: dba
- Catalog name: dba

Starting solidDB in Windows environments

Click the icon labeled **Start IBM solidDB server** through the **Start** > **Programs** > **IBM solidDB** menu path.

Result

When you start solidDB, it checks if a database already exists. If no database is found, solidDB creates a new database called `solid.db` in the `eval_kit\standalone` directory automatically, with the following login data:

- Username: dba
- Password: dba
- Catalog name: dba

Tip: If you would like the **Start IBM solidDB server** icon to start solidDB and create a database in a different directory or with different login data and catalog name, you can modify the installation default settings by right-clicking the **Start IBM solidDB server** icon and editing the shortcut **Properties**. For more details, see section *Modifying Windows shortcuts* in the *IBM solidDB Administrator Guide*.

Connecting to solidDB for the first time

After you have created your first database and your solidDB process is running, you can connect to it from any client application by using the communication ports that the solidDB process is listening to. This can be done, for example, by using the command-line tool solidDB SQL Editor (solsql).

The communication port is defined with the **Listen** parameter in the [Com] section of the solid.ini file.

For the evaluation version of solidDB in the eval_kit\standalone directory, the following communication ports are defined:

```
[Com]
Listen=tcpip 2315, tcpip 1315, tcpip 1964
```

Connecting to a database with the solidDB SQL Editor

To connect to a database with solidDB SQL Editor (solsql), enter the solsql start command at your operating system command prompt, specifying the network name and credentials for the solidDB server you want to connect to. On Windows operating systems, you can also use the **Start** > **Programs** > **IBM solidDB** > **solidDB SQL Editor** menu path to start the editor.

Procedure

To start the solidDB SQL Editor:

- On Windows operating systems, click the icon labeled **solsql SQL Editor** through the **Start** > **Programs** > **IBM solidDB** menu path.

  This starts the solidDB SQL Editor, connecting to solidDB server using the communication port "tcp 2315" and username "dba" and password "dba".

  Tip: If you would like the **solsql SQL Editor** icon to start solsql with different login data or communication port, you can modify the installation default settings by right-clicking the **solsql SQL Editor** icon and editing the shortcut **Properties**. For more details, see section *Modifying Windows shortcuts* in the *IBM solidDB Administrator Guide*.
- 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 2315" dba dba
  ```

  Tip: If you do not give your user name and password when starting the solidDB SQL Editor, you are prompted for them.

Results

If you entered a valid user name and password when starting the solidDB SQL Editor, you are connected to the database. The figure below shows the solidDB SQL Editor after a successful connection to the database.

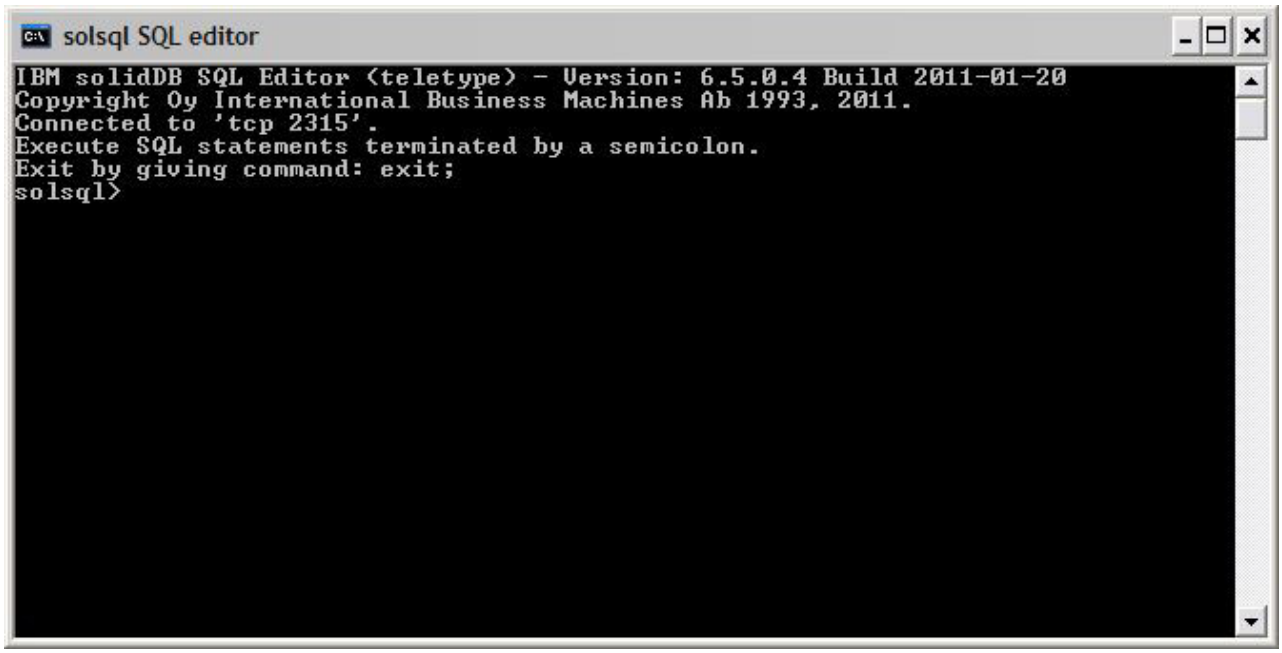

Figure 3. solidDB SQL Editor connected

Viewing database and configuration status

You can use the solidDB SQL Editor (**solsql**), for example, to view the status of the database, check the parameter settings, or display the available administration commands (ADMIN COMMANDs).

About this task

When you issue ADMIN COMMANDs commands in **solsql**:

- use a semicolon at the end of the command, and
- use single quotation marks, not double quotation marks when specifying the command name.

Procedure

1. To view the status of your database, issue the following command in the solidDB SQL Editor:

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

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

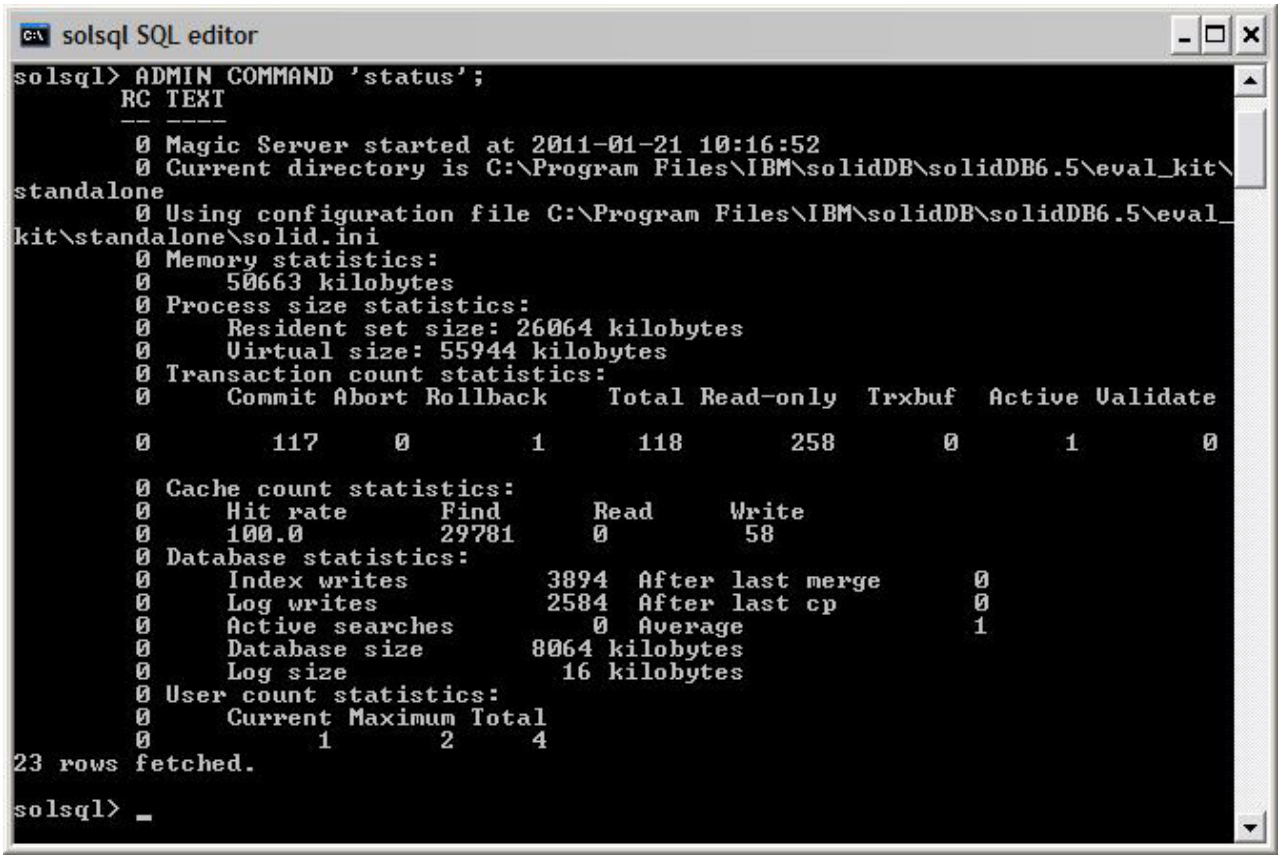

*Figure 4. Example output from solidDB SQL Editor (**solsql**)*

2. To view all parameter settings, issue the following command:

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

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

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

   For example:

   ```
   ADMIN COMMAND 'par IndexFile;
   ```

3. To view the available ADMIN COMMANDs, issue the following command:

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

Executing SQL statements with solidDB SQL Editor

You can execute SQL statements in the solidDB SQL Editor. This section shows an example of creating, viewing, and deleting a test table.

Procedure

1. Create a table by issuing the following commands in solsql:

   ```
   create table testtable (value integer, name varchar);
   commit work;

   insert into testtable (value, name) values (31, 'Duffy Duck');
   commit work;

   select value, name from testtable;
   commit work;
   ```

2. View the table with the following command:

```
select value, name from testtable;

    VALUE NAME
    ----- ----
       31 Duffy Duck
1 rows fetched.

commit work;
```

3. Delete the table with the following command:

```
drop table testtable;
commit work;
```

Tip: To ensure that every database operation gets committed, remember to execute the COMMIT WORK statement.

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 quotation marks, not double quotation marks.

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 shut down with all previously connected users disconnected.

Note:

When you shut down the server, it breaks the connection to **solsql**, and your **solsql** might 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, click the icon labeled **Start IBM solidDB server** through the **Start** > **Programs** > **IBM solidDB** menu path.
- 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
  ```

  The command-line option `-c <directory_path>` defines the working directory.
- 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`.

Stopping and restarting the database

5 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.ibm.com/support/docview.wss?rs=3457&uid=swg21411299
- *Configuring Data Studio Administrator V1.2 for use with solidDB*: http://www.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.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.ibm.com/support/docview.wss?rs=3457&uid=swg21411262

6 Running samples

You can 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.

Before you begin

To run samples:

- The sample directory must contain a valid license file. By default, each sample directory contains a copy of the evaluation license (solideval.lic).

  If you want to use the full product license (solid.lic or soliduc.lic) with the samples:

  1. Copy the license file to the root of the solidDB installation directory.
  2. Execute the script called copy_licenses (copy_licenses.bat in Windows). This script copies the license file from the installation root directory to all appropriate sample directories.

Procedure

Run the sample according to instructions provided in the *readme.txt* file, located in the sample directory.

Note: Some sample scripts remove and re-create database files residing in sample subdirectories. Only the databases in the eval_kit/standalone and eval_kit/cdc directories are always left intact.

7 Upgrading solidDB to a new release level

solidDB supports upgrades from at least the two previous release levels. Upgrades to solidDB 6.5 are supported from the 6.3 and 6.1 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 6.5 and IBM solidDB Universal Cache 6.5 Information Center** at http://publib.boulder.ibm.com/infocenter/soliddb/v6r5/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 HotStandby servers* in the *IBM solidDB High Availability User Guide*.

The previous release level solidDB JDBC and ODBC drivers are compatible with the 6.5 release level solidDB server. However, the 6.5 versions of the drivers cannot be used with older servers. Thus, the servers must be upgraded before the drivers.

The drivers need to be upgraded only if you plan to take into use the new features introduced at release level 6.5.

Note: If you want to convert your 6.3 or 6.1 database to support the full Unicode mode introduced in version 6.5, see section *Converting partial-Unicode databases to Unicode* in the *IBM solidDB Programmer Guide*.

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 want 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 25.
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

C

D

E

J

L

R

S

T

U

W

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

This information 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

GI11-9229-06