

Note

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 93.

Edition Notice

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corporation 1993, 2010.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication v

Chapter 1. IBM data server clients 1

Introduction to IBM data server clients 1
IBM data server clients and drivers overview . . . 1
IBM data server client types 2
Installation methods for IBM data server clients . 5
Options for connecting to DB2 databases 6

Chapter 2. Installing IBM data server clients 11

IBM data server client installation requirements . . 11
Disk and memory requirements 11
Installation requirements for DB2 servers and IBM data server clients (AIX) 11
Installation requirements for DB2 servers and IBM data server clients (HP-UX) 13
Installation requirements for DB2 servers and IBM data server clients (Linux) 17
Installation requirements for DB2 servers and IBM data server clients (Solaris) 20
Installation requirements for DB2 servers and IBM data server clients (Windows) 22
DB2 Connect product installation requirements for host and midrange systems 24
Installing IBM data server clients 24
Installing IBM data server clients (Windows) . . 24
Installing IBM data server clients (Linux and UNIX) 27
Non-root installation overview (Linux and UNIX) 28
IBM Data Server Drivers 36
IBM Data Server Driver restrictions 36
db2dsdriver configuration file 37
db2dsdcfgfill - Create configuration file
db2dsdriver.cfg 39
Copying existing database directory information into the db2dsdriver configuration file 41
Installing IBM Data Server Driver Package (Windows). 41
Installing IBM Data Server Driver Package (Linux and UNIX) 46

Chapter 3. Database connections for IBM data server clients. 49

Client-to-server communications configuration overview 49
Supported combinations of client and server versions 51
Communication protocols supported 53
Adding database connections using the Configuration Assistant 53
Configuring client-to-server connections using the command line processor 58

Chapter 4. IBM data server client deployment in a thin client topology (Windows) 67

Thin client topology overview (Windows) 67
Thin client setup overview (Windows) 68
Installing IBM Data Server Client or DB2 Connect Personal Edition on the code server (Windows). 69
Making the code directory available to all thin client workstations (Windows) 69
Creating a thin client response file (Windows). . 70
Mapping a network drive from each thin client to the code server (Windows) 70
Setting up thin clients using the thnsetup command (Windows) 71

Chapter 5. Merge Modules 73

Types of merge modules 73
Non-DB2 instance merge modules (Windows). . 73
DB2 instance merge modules (Windows) . . . 74

Chapter 6. Additional installation options 77

Installation command-line options 77
IBM Data Server Runtime Client installation command-line options 77
Command line options to install IBM Data Server Driver Package (Windows) 78

Chapter 7. Uninstalling 81

Uninstalling an IBM data server client 81

Appendix A. Overview of the DB2 technical information 83

DB2 technical library in hardcopy or PDF format. . 83
Ordering printed DB2 books. 86
Displaying SQL state help from the command line processor 87
Accessing different versions of the DB2 Information Center 87
Displaying topics in your preferred language in the DB2 Information Center 87
Updating the DB2 Information Center installed on your computer or intranet server 88
Manually updating the DB2 Information Center installed on your computer or intranet server . . . 89
DB2 tutorials 91
DB2 troubleshooting information 91
Terms and Conditions 92

Appendix B. Notices 93

Index 97

About this publication

If you are interested in installing and configuring an IBM® data server clients or driver, setting up a thin client, or DB2 Connect™ thin client environment, you should read this book.

Chapter 1. IBM data server clients

Introduction to IBM data server clients

IBM data server clients and drivers overview

This topic outlines information available about clients and drivers, and provides links to further details. This topic will help you to perform the following tasks:

1. Choose the appropriate IBM data server client or driver to enable connections between your system and remote databases.
2. Choose the most suitable method for installing your client or driver.
3. Complete the steps and address the considerations needed to set up a client or driver.

Connection options

Options for connecting a system to a remote database include various IBM data server clients and drivers. The options available depend on whether the system connecting to the remote database is:

- An application located on a business user's machine or an application server
- An application development workstation
- A database administrator workstation

There are additional options to consider if you need to also connect to midrange or mainframe databases.

IBM data server client and driver types

The following lists the IBM data server clients and drivers:

- IBM Data Server Driver Package
- IBM Data Server Client
- IBM Data Server Runtime Client
- IBM Data Server Driver for JDBC and SQLJ.
- IBM Data Server Driver for ODBC and CLI

The IBM Data Server Driver is a lightweight solution and the recommended best practice package for end user code deployment. It provides robust runtime support for applications using ODBC, CLI, .NET, OLE DB, PHP, Ruby, JDBC, or SQLJ without the need of installing Data Server Runtime Client or Data Server Client. The IBM Data Server Driver for ODBC and CLI solution is designed mainly for independent software vendor (ISV) deployments.

It is also recommended that the IBM Data Server Driver Package be installed first. It can then be configured to work in conjunction with DB2 Connect.

In addition, a separate product, DB2 Connect Personal Edition, includes all the functionality of IBM Data Server Client plus the capability to connect to midrange and mainframe databases.

DB2 Connect capability can be added to any client or driver.

Installation methods

The common method for installing a client or driver is to run the installation program provided on a product DVD.

The IBM Data Server Driver Package on Windows® can be installed by following the prompts. There is no installation program for IBM Data Server Driver for ODBC and CLI or for IBM Data Server Driver Package on Linux® and UNIX®. You must install the driver manually.

Other installation methods are also available. Some methods are designed to automate the deployment of large numbers of clients. Other methods use various Windows operating system capabilities. For example, on Windows operating systems, you can use merge modules to embed the functionality of Data Server Runtime Client or IBM Data Server Driver Package in your application.

Setting up a client or driver

After you decide which client to use, set up the client by performing the following steps:

1. Ensure that system prerequisites are satisfied.
2. Perform the installation.

For systems where a Version 8 client or a DB2® Version 9 client already exists, consider whether to upgrade the existing client to a Version 9.7 Data Server Client, or, keep the pre-Version 9.7 client and install the Version 9.7 Data Server Client as an additional client. Upgrading to a Version 9.7 client is strongly recommended. Running multiple copies of the client packages is for advanced users only.

Note: The option to upgrade and replace the existing client applies to Data Server Client only.

IBM data server client types

There are several types of IBM data server clients and drivers available. Each provides a particular type of support.

The following lists the IBM data server client and driver types available to you:

- IBM Data Server Driver for JDBC and SQLJ
- IBM Data Server Driver for ODBC and CLI
- IBM Data Server Driver Package
- IBM Data Server Runtime Client
- IBM Data Server Client

Each IBM data server client and driver provides a particular type of support:

- For Java™ applications only, use IBM Data Server Driver for JDBC and SQLJ.
- For applications using ODBC or CLI only, use IBM Data Server Driver for ODBC and CLI. (Also referred to as cli driver.)
- For applications using ODBC, CLI, .NET, OLE DB, PHP, Ruby, JDBC, or SQLJ, use IBM Data Server Driver Package.
- For applications using DB2CI, use IBM Data Server Client.
- If you need DB2 Command Line Processor Plus (CLPPlus) support, use IBM Data Server Driver Package.

- If you need DB2 command line processor (CLP) support and basic client support for running and deploying applications, use IBM Data Server Runtime Client.
- If you need support for database administration, and application development using an application programming interface (API), such as ODBC, CLI, .NET, or JDBC, use IBM Data Server Client.

IBM Data Server Driver for JDBC and SQLJ

IBM Data Server Driver for JDBC and SQLJ is the default driver for Java stored procedures and user-defined functions. This driver provides support for client applications and applets that are written in Java using JDBC to access local or remote servers, and SQLJ for embedded static SQL in Java applications. This driver is a prerequisite for IBM Optim™ pureQuery Runtime, which provides static support for Java, enables optimized data access using the pureQuery API, and is supported by a full integrated development environment (IDE) for Java database application development using Optim Development Studio. (Both Optim products are available separately.)

IBM Data Server Driver for ODBC and CLI

Data Server Driver for ODBC and CLI is a lightweight deployment solution designed for independent software vendors (ISV) deployments. This driver, also referred to as cli driver, provides runtime support for applications using ODBC API, or CLI API without need of installing the Data Server Client or the Data Server Runtime Client. This driver is available only as a tar file, not as an installable image. Messages are reported only in English.

The IBM Data Server Driver for ODBC and CLI provides:

- runtime support for the CLI API;
- runtime support for the ODBC API;
- runtime support for the XA API;
- database connectivity;
- support for DB2 Interactive Call Level Interface (db2cli);
- LDAP Database Directory support; and
- tracing, logging, and diagnostic support.

Register Data Server Driver for ODBC and CLI with the Microsoft® ODBC driver manager using the db2oreg1.exe utility.

IBM Data Server Driver Package,

IBM Data Server Driver Package is a lightweight deployment solution that provides runtime support for applications using ODBC, CLI, .NET, OLE DB, PHP, Ruby, JDBC, or SQLJ without the need of installing Data Server Runtime Client or Data Server Client. This driver has a small footprint and is designed to be redistributed by independent software vendors (ISVs), and to be used for application distribution in mass deployment scenarios typical of large enterprises.

The IBM Data Server Driver Package capabilities include:

- The DB2 Command Line Processor Plus (CLPPlus) for dynamically creating, editing, and running SQL statements and scripts.
- Support for applications that use ODBC, CLI, PHP, or Ruby to access databases.
- Support for client applications and applets that are written in Java using JDBC, and for embedded SQL for Java (SQLJ).

- IBM Informix[®] Dynamic Server support for .NET, PHP, and Ruby.
- Support for running embedded SQL applications. No precompiler or bind capabilities are provided.
- Application header files to rebuild the PHP, Ruby, Python, and Perl drivers. The Python and Perl drivers are not available in IBM Data Server Driver Package; however, you can download and build these drivers using the header files.
- Support for DB2 Interactive Call Level Interface (db2cli).
- Support for the for DRDA[®] traces (db2drdat).
- On Windows operating systems, IBM Data Server Driver Package also provides support for applications that use .NET or OLE DB to access databases. In addition, this is available as an installable image, and merge modules are available to allow you to easily embed the driver in a Windows Installer-based installation.

IBM Data Server Runtime Client

The IBM Data Server Runtime Client provides a way to run applications on remote databases. GUI tools are not shipped with the IBM Data Server Runtime Client.

Capabilities include:

- The DB2 command line processor (CLP) for issuing commands. The CLP also provides a basic way to perform remote administration of servers.
- Base client support to handle database connections, SQL statements, XQuery statements, and commands.
- Support for common database access interfaces: JDBC, ADO.NET, OLE DB, ODBC, Command Line Interface (CLI), PHP, and Ruby. This support includes drivers and capabilities to define data sources. For example, for ODBC, installing an IBM data server client installs the ODBC driver and registers the driver. Application developers and other users can use the Windows ODBC Data Source Administrator tool to define data sources.
- Lightweight Directory Access Protocol (LDAP) exploitation.
- Support for common network communication protocols: TCP/IP, and Named Pipe.
- Support for installing multiple copies of a client on the same computer. These copies can be the same or different versions.
- License terms that allow free redistribution of IBM Data Server Runtime Client with your application.
- smaller deployment footprint compared to that of the full IBM Data Server Client in terms of installation image size and disk space required.
- A catalog that stores information for connecting to databases and servers.
- Packing advantages on Windows operating systems: You can package the client with your application to provide connectivity for that application. Also, the client is available as Windows Installer merge modules that enable you to include the RTCL DLL files in your application installation package. This approach also enables you to include only the parts of the client that you need with your application.
- IBM Informix Dynamic Server support for PHP, Ruby, .NET, and JDBC

IBM Data Server Client

IBM Data Server Client includes all the functionality of IBM Data Server Runtime Client, plus functionality for database administration, application development, and client/server configuration.

Capabilities include the following:

- A larger deployment footprint compared to that of IBM Data Server Runtime Client in terms of installation image size and disk space required. However, on Windows operating systems, you can prune the IBM Data Server Client image to reduce the installation image size.
- The configuration Assistant to assist with cataloging databases and configuring the database server.
- Control Center and other graphical tools for database implementation and for database administration. These tools are available for versions of Windows on x86 (32-bit only), Windows on x64 (AMD64/EM64T), Linux on x86, and Linux on AMD64/EM64T (x64).
- First Steps documentation for new users.
- Visual Studio tools
- Application header files
- Precompilers for various programming languages
- Bind support
- Samples and tutorials
- IBM Informix Dynamic Server support for PHP, Ruby, .NET, JCC, and JDBC

Installation methods for IBM data server clients

Common and alternative methods for installing an IBM data server client or driver.

Clients are commonly installed on machines where there is no server present. You do not need to install a client if you already installed a server product is already installed because the server includes all the functionality present in an IBM data server client.

Common situations

The common method for installing an IBM data server client or driver is to run the installation program provided on a product DVD (setup command on Windows operating systems and db2setup command on Linux and UNIX operating systems). The IBM Data Server Client installation image is included in the database server installation image.

Automation of large-scale deployments

One group of methods automates the deployment of large numbers of clients:

- **Response file.** You can automate client installation by using the response file installation method. A response file installation lets you install database products without any user interaction.
- **Third-party deployment tools.** You can install clients using deployment tools or methods such as Windows Active Directory, Windows Systems Management Server (SMS), or Tivoli® products.

Use of Windows operating system capabilities

Another group of options uses Windows operating system capabilities:

- **Windows thin client topology.** This option is supported for the IBM Data Server Client and DB2 Connect Personal Edition. A thin client topology is where client code is installed in a shared Windows directory on a single code server rather than on the local hard disk of each client workstation. Individual client workstations connect to the shared Windows directory on the code server to run the Data Server Client code.
- **A Windows non-administrator ID.** The common installation method uses a Windows administrator user ID: that is, a user ID in the Administrators group. However, you can also install an IBM data server client using a user ID that is part of the Windows Power Users group or Users group. This method is suitable when the user ID performing the installation does not have administrator privileges. The DB2 product also supports the Windows Elevated Privileges mechanism. It is recommended to use Windows Elevated Privileges to allow a non-administrator to install an IBM data server client.

Linux and UNIX alternatives

On Linux and UNIX operating systems, an alternative installation method provided for database servers is also applicable to clients: the `db2_install` script.

Also, IBM Data Server Driver for ODBC and CLI is available as a tar file.

Separate client instances

If a database server product is installed, you can use a separate client instance instead of using a server instance that also serves as a client instance.

To create a separate client instance, use the `db2icrt` command with the `-s` option, as shown in the following example:

```
db2icrt -s client instname
```

Options for connecting to DB2 databases

This section outlines the options for what you can install on a machine (the local system) to enable it to connect to a database on a different machine (the remote system). To select a suitable option, you first need to consider whether the local system is:

- a system that runs business applications on a business user's system or on an application server.
- an application development workstation.
- a database administrator workstation.

You also need to determine where the databases reside that you want to connect to. The databases could be located:

- on same machine, that is, on the local system. This includes databases located in a single DB2 instance or in various DB2 instances.
- on different machines, that is, on remote systems.
- on different machines that are midrange or mainframe servers.

Options for business user systems or application servers

Typically, when a business application connects to a database, one of the following topologies is involved:

- An application connects to a DB2 database where both are located on the same machine. An example is an application used by a single user on their personal workstation.
- An application connects to a DB2 database on another machine.
- An application client connects to an application server which, in turn, connects to one or more DB2 databases located:
 - on the same machine only.
 - on one or more other machines only.
 - a combination of the above.

If a DB2 server is installed on the same machine as the application, there is no need to install a separate client. The DB2 server product includes functionality to enable applications to connect to local databases and to connect to databases on remote machines, as well.

If the machine with the application does not also have a DB2 server, you have the following options to enable applications to connect to remote DB2 databases:

- **DB2 application driver.** With a DB2 application driver, the information needed to connect to a database is included in the application or the application prompts the user to provide it. This approach differs from an IBM data server client which maintains this information in its catalog. The application driver is deployed as a file in the application directory so no separate DB2-specific installation or setup is required. Typically, an application driver is packaged with an application in a manner that provides connectivity only for that application. A DB2 application driver can coexist on the same machine with other DB2 application drivers or with an IBM data server client. DB2 products provide drivers for Java (JDBC and SQLJ) and for ODBC, CLI, .NET, OLE DB, or open source applications. Drivers can be obtained by copying driver files from an IBM Data Server Driver Package installation image or by downloading the driver files from developerWorks®.
- **IBM data server client.** This option involves installing and configuring one of the clients included with the DB2 product. The IBM data server client is installed on any machine that connects directly to the DB2 database. Depending on the application topology, the client is installed on each business user workstation or on an application server. A single IBM data server client can enable all applications on the machine to connect to one or more DB2 databases on other machines.
- **DB2 instance merge modules.** These merge modules create a DB2 instance environment. This approach provides a way to deploy the IBM Data Server Runtime Client by including the files in the corresponding modules. This approach is targeted for use with Windows Installer and other install tools that support Windows Installer merge modules. With this approach, a single installation program installs both the application and the Data Server Runtime Client. If you do not require an instance environment or a Command Line Processor (CLP) you should use the non-DB2 instance merge modules to avoid instance management.
- **Non-DB2 instance merge modules.** These merge modules create a non-DB2 instance environment. This approach provides a way to deploy the IBM Data Server Driver Package by including the client DLL files in the application deployment package. This approach is targeted for use with Windows Installer and other install tools that support Windows Installer merge modules. With this approach, a single installation program installs both the application and the IBM Data Server Driver Package.

Options for application development workstations

The IBM Data Server Client provides all the functionality of the IBM Data Server Runtime Client plus tools used for client-server configuration, database administration and application development. The points below describe the role and setup of the Data Server Client in light of the other tools and products used by application developers.

There are several tools and products typically used by application developers who write code to access a DB2 database. Each developer workstation typically includes the following components:

- An integrated development environment (IDE) such as IBM Data Studio, IBM Optim Development Studio, Rational® Application Developer or Microsoft Visual Studio.
- A DB2-specific development tool related to the IDE such as IBM Database Add-Ins for Microsoft Visual Studio.
- Access to a database server to host the database they are developing. This database server can reside in one or both of the following locations:
 - On each developer's workstation, so each developer has their own local copy of the database.
 - On a workgroup server, so multiple developers work on the same copy of the database.

With the foregoing as context, the value of the Data Server Client is that it provides headers and libraries required to compile applications and provides tools for database administration. However, it is not always necessary to install the Data Server Client to obtain these tools. Any time a DB2 server is installed on a machine, there is no need to install a separate IBM data server client. The DB2 server product includes all functionality available in a standalone Data Server Client.

Options for administrator workstations

A database administrator can perform administration tasks on remote databases in one of two major ways. They use a tool, such as telnet, to connect to an individual database server machine and run DB2 administrator commands locally. Alternatively, they can run tools and commands from their workstation that make connections to the remote databases. This section focuses on the second case which involves choices about which IBM data server client to use and where to install it.

DB2 provides various means to perform administration activities on one or more DB2 servers remotely from an administrator workstation. The options and considerations are the following:

- **Install the IBM Data Server Runtime Client.** This option is suitable if you only use the Command Line Processor (CLP) for administration.
- **Install the IBM Data Server Client.** This client includes all the functionality of the Data Server Runtime Client plus tools used for client-server configuration, database administration and application development. Functionality includes the graphical administration tools such as the Configuration Assistant and Control Center (available on supported platforms). These tools use the DB2 Administration Server (DAS) on the DB2 server, which is setup by default during DB2 server installation.

- **Install a DB2 server product.** Any time a DB2 server is installed on a machine, there is no need to install a separate IBM data server client. The DB2 server product includes all the functionality available in any standalone IBM data server client.

Options for connecting to midrange and mainframe databases

With DB2 Connect products, you can connect to DB2 databases on mainframe and midrange platforms, namely OS/390® and z/OS®, System i®, VSE, and VM. You can also connect to non-IBM databases that comply with the Distributed Relational Database Architecture™ (DRDA). With DB2 Connect, you can connect from a user's workstation or from a DB2 for Linux, UNIX, or Windows server.

The options and considerations are the following:

- **IBM Data Server Driver Package.** This lightweight deployment solution that provides runtime support for applications using ODBC, CLI, .NET, OLE DB, PHP, Ruby, JDBC, or SQLJ without the need of installing Data Server Runtime Client or Data Server Client. To connect to a z/OS server or a System i server, you must register a DB2 Connect license key. (Retrieve the license file from your Passport Advantage® distribution, for example db2conpe.lic, then copy the license file to the license directory under the directory where the driver was installed.)
- **DB2 Connect Personal Edition.** This product is installed on a workstation and provides connectivity from that workstation. This option is targeted for situations where there is a need to connect from a workstation directly to one or more hosts. This product is available for Linux, Solaris and Windows operating systems.
- **DB2 Connect Server Editions.** A server edition of the DB2 Connect product is usually installed on a connectivity server and serves as a gateway to supported mainframe or midrange database servers. A server edition of the DB2 Connect product enables workstations with IBM data server clients to setup access to host systems via the DB2 Connect gateway as if they were databases on DB2 on Windows, Linux and UNIX.

Both options can be used concurrently. For example, a workstation can connect to a midrange or mainframe database in the following ways:

- Install DB2 Connect Personal Edition locally and use it to connect directly to a host.
- Connect to the same or different host through an intermediate DB2 Connect server gateway.

Chapter 2. Installing IBM data server clients

IBM data server client installation requirements

Disk and memory requirements

Disk requirements

The disk space required for your product depends on the type of installation you choose and the type of file system you have. The DB2 Setup wizard provides dynamic size estimates based on the components selected during a typical, compact, or custom installation.

Remember to include disk space for required databases, software and communication products.

On Linux and UNIX operating systems, 2 GB of free space in the /tmp directory is recommended.

Memory requirements

At a minimum, a DB2 database system requires 256 MB of RAM. For a system running just a DB2 product and the DB2 GUI tools, a minimum of 512 MB of RAM is required. However, 1 GB of RAM is recommended for improved performance. These requirements do not include any additional memory requirements for other software that is running on your system.

When determining memory requirements, be aware of the following:

- DB2 products that run on HP-UX Version 11i for Itanium-based systems require 512 MB of RAM at a minimum.
- For IBM data server client support, these memory requirements are for a base of five concurrent client connections. You will need an additional 16 MB of RAM per five client connections.
- Memory requirements are affected by the size and complexity of your database system, and the extent of database activity and the number of clients accessing your system.

For DB2 server products, the self-tuning memory feature simplifies the task of memory configuration by automatically setting values for several memory configuration parameters. When enabled, the memory tuner dynamically distributes available memory resources among several memory consumers including sort, the package cache, the lock list and buffer pools.

- On Linux and UNIX operating systems, SWAP space at least twice as large as RAM is recommended.

Installation requirements for DB2 servers and IBM data server clients (AIX)

Before you install DB2 database products on AIX[®] operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems. See com.ibm.db2.luw.qb.server.doc/doc/r0023732.dita for a list of DB2 database products that are available on each operating system.

To install a DB2 database product, the following requirements must be met:

Table 1. AIX installation requirements

Operating System	Hardware
<p>AIX Version 5.3</p> <ul style="list-style-type: none"> • 64-bit AIX kernel is required • AIX 5.3 Technology Level (TL) 9 and Service Pack (SP) 2 • Minimum C++ runtime level requires the xlC.rte 9.0.0.8 and xlC.aix50.rte 9.0.0.8 (or later) filesets. These filesets are included in the June 2008 IBM C++ Runtime Environment Components for AIX package. 	<p>64-bit Common Hardware Reference Platform (CHRP) architecture, excluding POWER3™ processor-based systems.¹</p> <p>All processors that are capable of running the supported AIX operating systems.</p>
<p>AIX Version 6.1²</p> <ul style="list-style-type: none"> • 64-bit AIX kernel is required • AIX 6.1 TL 2 • Note: POWER7™ systems³ require AIX 6.1 TL4 SP2. • Minimum C++ runtime level requires the xlC.rte 9.0.0.8 and xlC.aix61.rte 9.0.0.8 (or later) filesets. These filesets are included in the June 2008 IBM C++ Runtime Environment Components for AIX package. 	
<p>AIX Version 7.1</p> <ul style="list-style-type: none"> • 64-bit AIX kernel is required • AIX 7.1.0 General Availability (GA) • Minimum C++ runtime level requires the xlC.rte 11.1.0.0 and xlC.aix61.rte 11.1.0.0 (or later) filesets. These filesets are included in the April 2010 IBM C++ Runtime Environment Components for AIXV11.1 package. 	

- ¹To verify that it is a CHRP architecture system, issue the command `lscfg` and look for the following output: `Model Architecture: chrp`. For POWER3 processor-based systems, first upgrade to POWER4™ processor-based systems before installing DB2 Version 9.7. POWER3 processor-based systems are not supported in DB2 Version 9.7.
- ²In AIX 6.1 there are two types of Workload Partitions (WPARs): system WPARs and application WPARs. DB2 installation is supported only on a system WPAR. AIX 6.1 also supports the ability to encrypt a JFS2 file system or set of files.
- ³Additional considerations for POWER7 systems:
 - If you have IBM Tivoli System Automation for Multiplatforms (SA MP) 2.2..x or SA MP 3.1.x Reliable Scalable Cluster Technology (RSCT) 2.5.4.0 installed, you require the SA MP 3.1.0.6 special package with RSCT 2.5.4.2. Version 9.7

Fix Pack 2 and later fix packs include an updated version of SA MP which includes an RSCT level that supports POWER7 systems.

- DB2 products running on POWER7 systems support Active Memory™ Expansion.

Note: If the minimum operating system requirements were applied using an operating system upgrade rather than using a new operating system installation, you must install I/O completion ports (IOCP) separately. You can obtain the IOCP libraries from the AIX installation CD. Then, configure IOCP on the DB2 server or data server client. For more information, see “configuring IOCP on AIX” in *Troubleshooting and Tuning Database Performance*.

Software considerations

- For application development and runtime considerations, see the topics in Supported programming languages and compilers for database application development.
- You can download the latest IBM C++ Runtime Environment Components for AIX at the IBM AIX XL C and C++ support Web site.
- (Clients only) If you plan to use Kerberos Authentication, you require IBM Network Authentication Service client v1.4 or later.
- Use the bosboot command to switch to the 64-bit kernel.

To switch to a 64-bit kernel, you require root authority and enter the following commands:

```
ln -sf /usr/lib/boot/unix_64 /unix
ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix
bosboot -a
shutdown -Fr
```

- One of the following browsers is required to view online help and to run First Steps (db2fs):
 - Firefox 2.0 and later
 - Mozilla-based browsers 1.7 and later
- An X Window System software capable of rendering a graphical user interface is required if:
 - you want to use the DB2 Setup wizard to install a DB2 product on Linux or UNIX operating systems
- For details regarding known AIX issues, see www.ibm.com/support/docview.wss?&uid=swg21165448

Mounting considerations

For root installations of DB2 products on UNIX operating systems, do not mount your file system with the *noetuid* option.

Installation requirements for DB2 servers and IBM data server clients (HP-UX)

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems. See com.ibm.db2.luw.qb.server.doc/doc/r0023732.dita for a list of DB2 database products that are available on each operating system.

To install a DB2 product on an HP-UX operating system, the following operating system, hardware, and communications requirements must be met:

Table 2. HP-UX installation requirements

Operating System	Hardware
<p>DB2 products are supported on:</p> <ul style="list-style-type: none"> • HP-UX 11iv2 (11.23.0505) with: <ul style="list-style-type: none"> – May 2005 Base Quality (QPKBASE) bundle – May 2005 Applications Quality (QPKAPPS) bundle – PHCO_38637 - libc cumulative patch • HP-UX 11iv3 (11.31) with: <ul style="list-style-type: none"> – PHCO_38658 - libc cumulative patch <p>To support SDK6, the following patches are also required:</p> <ul style="list-style-type: none"> • HP-UX 11iv2: PHSS_37201 • HP-UX 11iv3: PHSS_37202 	<p>Itanium® based HP Integrity Series Systems</p> <p>Any hardware that is explicitly compatible and fully capable of running the specified operating system, all the corresponding supporting software, and any associated applications unmodified.</p>

Kernel configuration considerations

A system restart is required if you update the kernel configuration parameters. The kernel configuration parameters are set in `/etc/system`. Depending on the values of your kernel configuration parameters, modify some of them before you install the Version 9 client or DB2 server products. If the kernel parameter being modified is not listed as dynamic, a system reboot is required to make the changes to `/etc/system` take effect.

Software considerations

- One of the following browsers is required to view online help and to run First Steps (db2fs):
 - Firefox 2.0 and later
 - Mozilla-based browsers 1.7 and later
- An X Window System software capable of rendering a graphical user interface is required if you want to use the DB2 Setup wizard to install a DB2 product.
- For details regarding known HP-UX issues, see www.ibm.com/support/docview.wss?&uid=swg21257602

Note: Starting in DB2 V9.7 Fix Pack 1 and later, DB2 products installed on the HP-UX operating system support long host names. The length has been extended to 255 bytes, in any combination of characters or digits. The minimum operating system requirement is HPUX 11iv2 (and above).

To enable long host name support, complete the following tasks:

1. Turn on the kernel tunable parameter `expanded_node_host_name`.


```
Kctune expanded_node_host_name=1
```
2. Compile applications requiring long host name support with the `-D_HPUX_API_LEVEL=20040821` option.

Mounting considerations

For root installations of DB2 products on UNIX operating systems, do not mount your file system with the `nosetuid` option.

Recommended kernel configuration parameters (HP-UX)

For HP-UX systems running a DB2 64-bit database system, run the `db2osconf` command to suggest appropriate kernel configuration parameter values for your system. The `db2osconf` utility can only be run from `$DB2DIR/bin`, where `$DB2DIR` is the directory where you installed your DB2 product.

Modifying kernel parameters (HP-UX)

For your DB2 product to perform properly on HP-UX, you might need to update your system's kernel configuration parameters. If you update your kernel configuration parameter values, you must restart your computer.

You must have root authority to modify kernel parameters.

To modify kernel parameters:

1. Enter the **sam** command to start the System Administration Manager (SAM) program.
2. Double-click the **Kernel Configuration** icon.
3. Double-click the **Configurable Parameters** icon.
4. Double-click the parameter that you want to change and type the new value in the **Formula/Value** field.
5. Click **OK**.
6. Repeat these steps for all of the kernel configuration parameters that you want to change.
7. When you are finished setting all of the kernel configuration parameters, select **Action --> Process New Kernel** from the action menu bar.

The HP-UX operating system automatically restarts after you change the values for the kernel configuration parameters.

Installation requirements for DB2 servers and IBM data server clients (Linux)

Before you install DB2 database products on Linux operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems.

DB2 database products are supported on the following hardware:

- x86 (Intel® Pentium®, Intel Xeon®, and AMD) 32-bit Intel and AMD processors
- x64 (64-bit AMD64 and Intel EM64T processors)
- POWER® (IBM eServer™ OpenPower®, iSeries®, pSeries®, System i, System p®, and POWER Systems that support Linux)
- eServer System z®, System z9® or System z10

The minimum supported operating systems for Linux include:

- Red Hat Enterprise Linux (RHEL) 5 Update 2
- SUSE Linux Enterprise Server (SLES) 10 Service Pack 2
- SUSE Linux Enterprise Server (SLES) 11
- Ubuntu 8.0.4.1

For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Note:

- Version 9.7 Fix Pack 2 and later fix packs include an updated version of IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component that you can use in environments with SLES 11 or POWER7 systems. For more information, see “Installing IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component” or “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component”.

Multithreaded architecture limitations

If you are installing a DB2 32-bit database product on a Linux operating system, consider upgrading to a 64-bit operating system and installing the DB2 64-bit database product instead. The multithreaded architecture generally simplifies memory configuration. However, this could affect the memory configuration of 32-bit DB2 database servers. For example:

- Private memory for agent threads is allocated within a single process. The aggregate of all private memory allocations for database agents might not fit in a single process memory space.
- Support for multiple databases is limited because all database shared memory segments for all databases are allocated in a single process. You might need to reduce the memory usage for some databases in order to activate all databases successfully at the same time. However, the database manager performance might be impacted. Another alternative is to create multiple instances and catalog the databases across the instances. However, sufficient system resources is required to support this configuration.

Distribution Requirements

You should update your kernel configuration parameters before installing a DB2 database product. The default values for particular kernel parameters might not be sufficient when running a DB2 database system.

You might also have other products or applications that require Linux system resources. You should modify the kernel configuration parameters based on the needs of your Linux system working environment.

The kernel configuration parameters are set in `/etc/sysctl.conf`.

Refer to your operating system manual for information about setting and activating these parameters using the `sysctl` command.

Package requirements

The following tables list the package requirements for SLES and RHEL distributions:

- `libaio.so.1` is required for DB2 database servers using asynchronous I/O.
- `libstdc++.so.5` is required for DB2 Net Search Extender
- `libstdc++.so.6` is required for DB2 database servers and clients. However, if IBM Tivoli System Automation for Multiplatforms, or, IBM Tivoli Monitoring for Databases: DB2 Agent is used, `libstdc++.so.5` is needed.

Package requirements for SLES and RHEL

Package name	Description
libaio	Contains the asynchronous library required for DB2 database servers.
compat-libstdc++	Contains libstdc++.so.6 (not required for Linux on POWER)

The following tables list the package requirements for SUSE Linux and Red Hat distributions for DB2 partitioned database servers.

- The ksh93 Korn shell is required for SUSE10 and RHEL5 systems. The pdksh Korn Shell package is required for all other DB2 database systems.
- A remote shell utility is required for partitioned database systems. DB2 database systems support the following remote shell utilities:
 - rsh
 - ssh

By default, DB2 database systems use rsh when executing commands on remote DB2 nodes, for example, when starting a remote DB2 database partition. To use the DB2 database system default, the rsh-server package must be installed (see table below). More information about rsh and ssh is available in the DB2 Information Center.

If you choose to use the rsh remote shell utility, inetd (or xinetd) must be installed and running as well. If you choose to use the ssh remote shell utility, you need to set the **DB2RSHCMD** communication variable immediately after the DB2 installation is complete. If this registry variable is not set, rsh is used.

- The nfs-utils Network File System support package is required for partitioned database systems.

All required packages should be installed and configured before continuing with the DB2 database system setup. For general Linux information, see your Linux distribution documentation.

Package requirements for SUSE Linux

Package name	Description
pdsh or ksh93	Korn Shell. This package is required for partitioned database environments.
openssh	This package contains a set of server programs which allow users to run commands on (and from) remote computers via a secure shell. This package is not required if you use the default configuration of DB2 database systems with rsh.
rsh-server	This package contains a set of server programs which allow users to run commands on remote computers, login in to other computers, and copy files between computers (rsh, rexec, rlogin, and rcp). This package is not required if you configure DB2 database systems to use ssh.
nfs-utils	Network File System support package. It allows access to local files from remote computers.

Package requirements for Red Hat

Directory	Package name	Description
/System Environment/Shell	pdsh or ksh93	Korn Shell. This package is required for partitioned database environments.

Package requirements for Red Hat

Directory	Package name	Description
/Applications/Internet	openssh	This package contains a set of client programs which allow users to run commands on a remote computer via a secure shell. This package is not required if you use the default configuration of DB2 database systems with rsh.
/System Environment/ Daemons	openssh-server	This package contains a set of server programs which allow users to run commands from a remote computer via a secure shell. This package is not required if you use the default configuration of DB2 database systems with rsh.
/System Environment/ Daemons	rsh-server	This package contains a set of programs which allow users to run commands on a remote computer. Required for partitioned database environments. This package is not required if you configure DB2 database systems to use ssh.
/System Environment/ Daemons	nfs-utils	Network File System support package. It allows access to local files from remote computers.

Software considerations

- (Clients only) If you plan to use Kerberos Authentication, you require IBM Network Authentication Service client v1.4 or later.
- One of the following browsers is required to view online help and to run First Steps (db2fs):
 - Firefox 2.0 and later
 - Mozilla-based browsers 1.7 and later
- An X Window System software capable of rendering a graphical user interface is required if:
 - you want to use the DB2 Setup wizard to install a DB2 database product on Linux or UNIX operating systems, or
 - you want to use any DB2 graphical tools on Linux for x86 and Linux on AMD 64/EM64T.
- Micro Focus does not offer support for any of its COBOL compiler products on SLES 11.

Security-enhanced Linux considerations

On RHEL 5 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 it in permissive mode and run the **setenforce 0** command as a superuser
- modify `/etc/sysconfig/selinux` and reboot the machine.

If your DB2 database product installs successfully on a RHEL 5 system, DB2 processes will run in the unconfined domain. To assign DB2 processes to their own domains, modify the policy. A sample SELinux policy is provided in the `sqllib/samples` directory.

Modifying kernel parameters (Linux)

Starting with Version 9.7 Fix Pack 2, you no longer need to update Linux kernel parameters related to interprocess communication (IPC). For Version 9.7 Fix Pack 1 or earlier, the *enforced minimum settings* for particular kernel parameters might not be sufficient to run a DB2 database system and you might need to update them before installing a DB2 database product.

You must have root authority to modify kernel parameters.

To update kernel parameters on Red Hat and SUSE Linux:

1. Run the `ipcs -l` command to list the current kernel parameter settings.
2. Analyze the command output to determine whether you have to change kernel settings or not by comparing the current values with the *enforced minimum settings* for Version 9.7 Fix Pack 2 or later fix packs at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.qb.server.doc/doc/c0057140.html>. The following text is an example of the `ipcs` command output with comments added after `//` to show what the parameter names are:

```
# ipcs -l

----- Shared Memory Limits -----
max number of segments = 4096           // SHMMNI
max seg size (kbytes) = 32768           // SHMMAX
max total shared memory (kbytes) = 8388608 // SHMALL
min seg size (bytes) = 1

----- Semaphore Limits -----
max number of arrays = 1024             // SEMMNI
max semaphores per array = 250          // SEMMSL
max semaphores system wide = 256000     // SEMMNS
max ops per semop call = 32             // SEMOPM
semaphore max value = 32767

----- Messages: Limits -----
max queues system wide = 1024           // MSGMNI
max size of message (bytes) = 65536     // MSGMAX
default max size of queue (bytes) = 65536 // MSGMNB
```

3. Modify the kernel parameters that you have to adjust by editing the `/etc/sysctl.conf` file. If this file does not exist, create it. The following lines are examples of what should be placed into the file:

```
#Example for a computer with 16GB of RAM:
kernel.shmmni=4096
kernel.shmmax=17179869184
kernel.shmall=8388608
#kernel.sem=<SEMMSL> <SEMMNS> <SEMOPM> <SEMMNI>
kernel.sem=250 256000 32 4096
kernel.msgmni=16384
kernel.msgmax=65536
kernel.msgmnb=65536
```

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


```
sysctl -p
```
5. To make the changes effective after every reboot:
 - (SUSE Linux) Make `boot.sysctl` active
 - (Red Hat) The `rc.sysinit` initialization script will read the `/etc/sysctl.conf` file automatically

Installation requirements for DB2 servers and IBM data server clients (Solaris)

Before you install a DB2 database product on Solaris, ensure that your system meets the operating system, hardware and software requirements.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems. See `com.ibm.db2.luw.qb.server.doc/doc/r0023732.dita` for a list of DB2 database products that are available on each operating system.

Table 3. Solaris installation requirements

Operating System	Hardware
Solaris 9 <ul style="list-style-type: none"> • 64- bit kernel • Patches 111711-12 and 111712-12 • If raw devices are used, patch 122300-11 • 64-bit Fujitsu PRIMEPOWER and Solaris 9 Kernel Update Patch 112233-01 or later to get the fix for patch 912041-01 Solaris 10 <ul style="list-style-type: none"> • 64- bit kernel • Patch 118822-25 • If raw devices are used, patch 125100-07 	UltraSPARC or SPARC64 processors
Solaris 10 <ul style="list-style-type: none"> • 64- bit kernel • Patch 127128-11 	Solaris x64 (Intel 64 or AMD64)

- Support is only for DB2 to be installed on local zones. Installation on the global zone is not supported by DB2 at this time.

Kernel configuration considerations

The kernel configuration parameters are set in `/etc/system`. If the kernel parameter being modified is not listed as dynamic, a system reboot is required to make the changes to `/etc/system` take effect. These parameters must be set before you install an IBM data server client.

Software considerations

- (Clients only) If you plan to use Kerberos Authentication, you require Solaris 9 or higher with IBM Network Authentication Service (NAS) client v1.4 or later.
- One of the following browsers is required to view online help and to run First Steps (db2fs):
 - Firefox 2.0 and later

- Mozilla-based browsers 1.7 and later
- An X Window System software capable of rendering a graphical user interface is required if you want to use the DB2 Setup wizard to install a DB2 database product
- For details regarding known Solaris issues, see www.ibm.com/support/docview.wss?&uid=swg21257606

Security patches can be obtained from the <http://sunsolve.sun.com> Web site. From the SunSolve Online Web site, click **Patches** in the left panel.

The Java2 Standard Edition (J2SE) Solaris Operating System Patch Clusters and the SUNWlibC software are also required and can be obtained from the <http://sunsolve.sun.com> Web site.

For DB2 database products on 64-bit Fujitsu PRIMEPOWER systems, you require the following:

- Solaris 9 Kernel Update Patch 112233-01 or later to get the fix for patch 912041-01.

The Fujitsu PRIMEPOWER patches for the Solaris Operating System can be downloaded from FTSI at: <http://download.ftsi.fujitsu.com/>.

DB2 database products support the following Solaris concepts:

- Solaris Logical Domains (LDoms)
- Solaris Zones
- ZFS filesystems

For details about virtualization technology supported by DB2 products, see <http://www.ibm.com/developerworks/wikis/display/im/DB2+Virtualization+Support>.

Mounting considerations

For root installations of DB2 products on UNIX operating systems, do not mount your file system with the *noatime* option.

Modifying kernel parameters (Solaris Operating System)

For the DB2 database system to operate properly, it is recommended that you update your system's kernel configuration parameters. You can use the `db2osconf` utility to suggest recommended kernel parameters. If you want to take advantage of project resource controls (`/etc/project`), consult your Solaris documentation.

You must have root authority to modify kernel parameters.

To use the `db2osconf` command, you must first install the DB2 database system. The `db2osconf` utility can only be run from `$DB2DIR/bin`, where `$DB2DIR` is the directory where you installed your DB2 product.

You must restart your system after modifying kernel parameters.

To set a kernel parameter, add a line at the end of the `/etc/system` file as follows:

```
set parameter_name = value
```

For example, to set the value of the `msgsys:msginfo_msgmax` parameter, add the following line to the end of the `/etc/system` file:

```
set msgsys:msginfo_msgmax = 65535
```

After updating the /etc/system file, restart the system.

Installation requirements for DB2 servers and IBM data server clients (Windows)

Before you install a DB2 database product on Windows operating systems, ensure that the system you choose meets the necessary operating system, hardware, and software requirements.

Table 4. Windows workstation platforms

Operating System	Prerequisites	Hardware
Windows XP Professional (32-bit and 64-bit)	Windows XP Service Pack 2 or later	All Intel and AMD processors capable of running the supported Windows operating systems (32-bit and 64-bit based systems)
Windows Vista Business (32-bit and 64-bit)	IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 2.0 framework runtime.	
Windows Vista Enterprise (32-bit and 64-bit)		
Windows Vista Ultimate (32-bit and 64-bit)		
Windows 7 Professional (32-bit and 64-bit)	64-bit IBM data server provider for .NET applications are supported.	
Windows 7 Enterprise (32-bit and 64-bit)	All Windows Vista service packs are supported.	
Windows 7 Ultimate (32-bit and 64-bit)		

Table 5. Windows server platforms

Operating System	Prerequisites	Hardware
Windows 2003 Datacenter Edition (32-bit and 64-bit)	Service Pack 2 or later. R2 is also supported	All Intel and AMD processors capable of running the supported Windows operating systems (32-bit and 64-bit based systems).
Windows 2003 Enterprise Edition (32-bit and 64-bit)	IBM data server provider for .NET client applications and CLR server-side procedures require .NET 2.0 framework runtime.	
Windows 2003 Standard Edition (32-bit and 64-bit)	64-bit IBM data server provider for .NET applications are supported.	
Windows Server 2008 Datacenter Edition (32-bit and 64-bit) and Windows Server 2008 R2 (64-bit)	IBM data server provider for .NET client applications and CLR server-side procedures require .NET 2.0 framework runtime.	
Windows Server 2008 Enterprise Edition (32-bit and 64-bit) and Windows Server 2008 R2 (64-bit)	64-bit IBM data server provider for .NET applications are supported.	
Windows Server 2008 Standard Edition (32-bit and 64-bit) and Windows Server 2008 R2 (64-bit)	All Windows Server 2008 service packs are supported.	

Note: DB2 database products support the hardware-enforced Data Execution Prevention (DEP) feature that is built into some Windows operating systems.

Additional software considerations

- Windows Installer 3.0 is required. It will be installed by the installer if it is not detected.
- IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 2.0 framework runtime. In an x64 environment, 32-bit IBM data server provider for .NET applications will run in the WOW64 emulation mode.
- MDAC 2.8 is required. The DB2 Setup wizard will install MDAC 2.8 if it is not already installed.

Note: If a previous version of MDAC (for example, 2.7) is already installed, DB2 install will upgrade MDAC to 2.8. For a typical install, MDAC 2.8 is installed. For a custom install, MDAC 2.8 is installed but only if you have not deselected the default which is to install it. If you deselect MDAC as part of a custom install, it will not be installed.

- If you plan to use LDAP (Lightweight Directory Access Protocol), you should use either a Microsoft LDAP client or the IBM Tivoli Directory Server v6 client (also known as the IBM LDAP client which is included with DB2 database products). Before installation of the Microsoft Active Directory, you will need to extend your directory schema using the db2schex utility, which can be found on the installation media under the db2\Windows\utilities directory.

The Microsoft LDAP client is included with Windows operating systems.

- One of the following browsers is required to view online help, run the DB2 install launchpad (setup.exe), and to run First Steps (db2fs):

- Firefox 2.0 and later
- Internet Explorer 6.0 and later
- Mozilla-based browsers 1.7 and later
- Safari 3.0 and later

DB2 Connect product installation requirements for host and midrange systems

DB2 Connect products enable workstations to connect to databases on supported host and midrange platforms (for example, DB2 on z/OS). In some cases, DB2 Connect customers might need to apply patches to the host or midrange database product to enable this functionality.

See “Installing IBM data server clients (Linux and UNIX)” and “Installing IBM data server clients (Windows)” in *Installing IBM Data Server Clients*.

Installing IBM data server clients

Installing IBM data server clients (Windows)

Instructions to install any IBM data server client type, namely the IBM Data Server Client, the IBM Data Server Runtime Client, and the IBM Data Server Driver Package. The main procedure covers a simple, but common, case where there is no DB2 database product already installed.

If the machine already has a prior version of a client installed, you should first review topics that cover upgrading.

If the machine already has a DB2 database server product installed, it is not necessary to install a client because the DB2 database server provides all the capability found in an IBM data server client.

Before installing IBM data server clients or client package:

- You have determined which client best suits your need.
- You have located a DVD or other install image that you need. You can download an image from the IBM Support Fix Central Web site: <http://www.ibm.com/support/fixcentral/>. Data Server client and driver packages are found under the **Information Management** product group and **IBM Data Server Client Packages** product selection. Select the appropriate installed version and platform and click **Continue**. Click **Continue** again on the next screen and you will be presented with a list of all client and driver packages available for Windows. Ensure you have the appropriate 32-bit or 64-bit version, depending on your machine.
- You have a Windows user account that is part of the Administrators group.

Note: If a non-Administrator user account is going to do the product installation, then the VS2005 runtime library must be installed before attempting to install a DB2 database product. The VS2005 runtime library is needed on the operating system before the DB2 database product can be installed. The VS2005 runtime library is available from the Microsoft runtime library download web site. There are two choices: choose `vcredist_x86.exe` for 32-bit systems or `vcredist_x64.exe` for 64-bit systems.

- Your system meets all memory, disk space, and installation requirements. The installation program will check disk space and basis system requirements, and notify you if there is a problem.

Restrictions

- No other DB2 database product can be installed in the same path if one of the following products is already installed:
 - IBM Data Server Runtime Client
 - IBM Data Server Driver Package
 - *DB2 Information Center*
- The DB2 Setup wizard fields do not accept non-English characters.

This procedure covers the simple case. Information for other cases is covered elsewhere in this topic.

To install any IBM data server client on Windows:

1. Log on to the system with the user account that you want to use to perform the installation.
2. Optional: Shut down any other programs.
3. Insert the DVD into the drive. The autorun feature starts the DB2 Setup wizard which determines the system language and starts the setup program for that language.

For the Data Server Client, you can run the DB2 Setup wizard in a language other than the default system language by manually invoking the DB2 Setup wizard and specifying a language code. For example, the `setup -i fr` command runs the DB2 Setup wizard in French. For the Data Server Runtime Client or Data Server Driver Package, there are separate install images for each language.

4. To install an IBM data server client or client package:
 - If you are installing a Data Server Client, launch the DB2 Setup wizard, when the DB2 Launchpad opens, choose **Install a Product**. Follow the DB2 Setup wizard's prompts.
 - If you are installing a Data Server Runtime Client, it does not have a launchpad. See the Related Links for **setup** command parameters.

If you are installing a second copy of the Data Server Runtime Client, the command is:

```
setup /v" TRANSFORMS=:InstanceId1.mst MSINewInstance=1"
```

To install each subsequent copy of the Data Server Runtime Client (up to a maximum of 16 copies), modify the command by incrementing InstanceId*n*, for example:

```
setup /v" TRANSFORMS=:InstanceId2.mst MSINewInstance=1"
```

Note: It is strongly recommended that installing multiple copies is for advanced users.

- If you are installing Data Server Driver Package, run the **setup** command from the product DVD, or install the driver from a fix pack image by downloading the driver from <http://www.ibm.com/support/docview.wss?rs=71&uid=swg27007053>. If you are installing a fix pack image, see the Related Links for the installation command options of the **setup** command.

If you are installing a second copy of Data Server Driver Package, the following methods can be used:

- To perform a new copy installation with a generated default copy name:
setup /o
- If the copy name already exists, perform a maintenance (or upgrade) installation on that copy. Otherwise, perform the new installation by using the specified copy name.
setup /n *copyname*

After installing Data Server Driver Package, you can optionally create and populate the configuration file, `db2dsdriver.cfg`, with database directory information.

5. If you are installing a Data Server Client on a machine that already has a DB2 Universal Database™ (UDB) Version 8 copy installed, you will be presented with the option to install a new copy or to upgrade the DB2 UDB Version 8 copy.

Installing a new copy preserves the DB2 UDB Version 8 copy and installs an additional DB2 Version 9 copy. Choosing to upgrade will copy the DB2 UDB Version 8 client instance settings to the DB2 Version 9 copy then remove the DB2 UDB Version 8 copy.

Note: If a machine already has a DB2 UDB Version 8 copy installed, the Version 9 copies cannot be set to default.

If you are installing a Data Server Runtime Client, the installation program always installs a new copy. To upgrade a DB2 UDB Version 8 client instance as a subsequent step, see topics on migration.

After completing this procedure, the product is now installed at the location you specified during the installation. As part of the IBM Data Server Client installation procedure, an instance of the DB2 database manager is created. The instance is called "DB2" if there is no other instance called "DB2". If you already have a copy of DB2 Version 8 or DB2 Version 9.1 installed, the default instance is DB2_01.

The default installation path of Data Server Client and Data Server Runtime Client is `Program Files\IBM\sql1ib`. If a second copy is installed in the same machine, the default directory name is `Program Files\IBM\sql1ib_01`. In general, the default directory name is `sql1ib_nn` where *nn* is the number of copies installed in that machine minus one.

The default installation path of Data Server Driver Package is `Program Files\IBM\IBM DATA SERVER DRIVER`. If a second copy is installed in the same machine, the default directory name is `Program Files\IBM\IBM DATA SERVER DRIVER_02`. In general, the default directory name is `IBM DATA SERVER DRIVER_nn` where *nn* is the generated number to make this directory unique.

If you want to install more than one copy of the Data Server Driver Package, you can have a maximum of 16 copies. Each copy must be installed to different directories.

The default copy name of the Data Server Client or Data Server Runtime Client is `DB2COPY1`

The default copy name of the Data Server Driver Package is `IBMDBCL1`

This installation does not include product documentation.

After installing your IBM data server client, the next step is to configure it to access remote DB2 database servers.

Notes on installing using a user account that is not a member of the Administrators group

Members of the Power Users group can install an IBM data server client. Members of the Users group can also install an IBM data server client after they have been enabled to do so. To enable members of the Users group to install an IBM data server client, a member of the Administrators group must ensure the installing user has **write** permission for the following:

- HKEY_LOCAL_MACHINE\SOFTWARE registry branch.
- the system directory (for example, c:\WINNT).
- the default install path (c:\Program Files) or another install path.

Of related interest, a non-administrator can also install fix packs if a non-administrator performed the original installation. However, a non-administrator cannot install fix packs if the original installation was performed by an Administrator user account.

Installing IBM data server clients (Linux and UNIX)

This task describes how to install an IBM data server client on Linux or UNIX. The instructions apply to the IBM Data Server Client and the IBM Data Server Runtime Client. The main procedure covers a simple, but common, case where no DB2 database product is already installed.

If the machine already has a prior version of a client installed, you should first review topics that cover upgrading.

If the machine already has a DB2 database server product installed, it is not necessary to install a client because the DB2 database server provides all the capability found in IBM Data Server Client.

- You have determined which client best suits your needs: the Data Server Client or the Data Server Runtime Client.
- You have located a DVD or other install image that you need. You can download an image from the IBM Support Fix Central Web site: <http://www.ibm.com/support/fixcentral/>. Data Server clients and drivers are found under the **Information Management** product group and **IBM Data Server Client Packages** product selection. Select the appropriate installed version and platform and click **Continue**. Click **Continue** again on the next screen and you will be presented with a list of all client and driver packages available for your platform.
- Your system meets all memory, disk space, and installation requirements. The installation program will check disk space and basis system requirements, and notify you if there is a problem.
- Installing an IBM data server client on the Solaris operating system or on HP-UX requires that you update your kernel configuration parameters. This is also recommended for Linux.

To install any IBM data server client on Linux or UNIX:

1. Insert and mount the appropriate DVD.
2. Change to the directory where the DVD is mounted.
3. Enter the `./db2setup` command to start the DB2 Setup wizard.

4. Choose **Install a Product** when the DB2 Launchpad opens.
5. Select the client you want to install.
6. Follow the DB2 Setup wizard's prompts. Help is available in the wizard to guide you through the remaining steps.

When installation is complete, the IBM data server client is installed by default in the following directories:

Linux /opt/ibm/db2/V9.7

UNIX /opt/IBM/db2/V9.7

This installation does not include product documentation.

After installing your IBM data server client, the next step is to configure it to access a remote DB2 server.

Notes on national language installations

You can run the DB2 Setup wizard in a language other than the default system language by manually invoking the DB2 Setup wizard and specifying a language code. For example, the `./db2setup -i fr` command runs the DB2 Setup wizard in French. However, the DB2 Setup wizard fields do not accept non-English characters.

Notes on installing on a machine that has an existing DB2 Version 9.5 client

The default directory name for the first copy is `V9.7`. If a copy is already installed, the second installation shows a default directory name of `V9.7_01`. In general, the default directory name is `V9.7_nn` where *nn* refers to the number of copies installed minus one.

Notes on installing on a machine that has an existing pre-DB2 Version 9.5 client

Installing a Data Server Client or Data Server Runtime Client on a system that already has either a DB2 Universal Database (UDB) Version 8 or DB2 Version 9 client preserves previous copy and installs an additional DB2 Version 9.5 (or higher) copy. For information on upgrading client instances to DB2 Version 9.5 (or higher), see the upgrading topics.

Non-root installation overview (Linux and UNIX)

Before Version 9.5, you could install products, apply and roll back fix packs, configure instances, add features, or uninstall products only if you had root privileges. Now, if you are a non-root user, you can perform these tasks on Linux and UNIX operating systems.

The DB2 installer automatically creates and configures a non-root instance during a non-root installation. As a non-root user, you can customize the configuration of the non-root instance during the installation. You can also use and maintain the installed DB2 database product without root privileges.

The non-root installation of a DB2 database product has one DB2 instance with most features enabled by default.

A non-root installation can be attractive for many groups, such as the following ones:

- Enterprises that have thousands of workstations and users who want to install a DB2 database product without consuming a system administrator's time
- Application developers who are not typically system administrators but use DB2 database products to develop applications
- Independent Software Vendors (ISVs) who develop software that does not require root user authority yet embeds a DB2 database product

Although non-root installations have most of the function of root installations, there are some differences and limitations. You can lift some of the limitations by having a root user run the db2rfe command.

Differences between root installations and non-root installations

In addition to a few limitations, the directory structure of a non-root installation is slightly different than the directory structure of a root installation.

During a root installation, subdirectories and files for the DB2 database product are created in a directory of the root user's choosing.

Unlike root users, non-root users cannot choose where DB2 database products are installed. Non-root installations are always placed in the `$HOME/sql1ib` directory, where `$HOME` represents the non-root user's home directory. The layout of the subdirectories within the `sql1ib` directory of a non-root is similar to that of a root installation.

For root installations, multiple instances can be created. Instance ownership is associated with the user ID under which the instance was created.

Non-root installations can have only one DB2 instance. The non-root installation directory contains all of the DB2 database product files and instance files with no soft links.

The following table summarizes the differences between root installations and non-root installations.

Table 6. Differences between root installations and non-root installations

Criteria	Root installations	Non-root installations
User can select installation directory	Yes	No. DB2 database products are installed under the user's home directory.
Number of DB2 instances allowed	Multiple	One
Files deployed during installation	Program files only. Instances must be created after installation.	Program files and instance files. The DB2 database product is ready for use immediately after installation.
Upgrade version and instance	No	No need to uninstall the old version before installing the new version. Install new version, and upgrade the instance, together.

Limitations of non-root installations

In addition to the differences between root installations and non-root installations, there are several limitations on non-root installations. This topic discusses the limitations to help you decide if you want to use a non-root installation.

Product limitations

Some DB2 database products are not supported in non-root installations:

- DB2 Query Patroller
- DB2 Net Search Extender
- Locally installed *DB2 Information Center*

Note: The locally installed *DB2 Information Center* is not supported in non-root installations because it requires root user authority to start the daemon. However, a non-root-installation DB2 instance can be configured to use a locally installed *DB2 Information Center* if it is installed on the same computer.

Features and tools limitations

The following features and tools are not available in non-root installations:

- The DB2 Administration Server (DAS) and its associated commands: `dasprt`, `dasdrop`, `daslist`, `dasmigr`, and `dasupdt`
- The Configuration Assistant
- The Control Center
- The ability for the `db2governor` to increase priority is not supported
- In the Work Load Manager (WLM), attempts to set agent priority in a DB2 service class in a non-root DB2 instance are allowed. However, the agent priority will not be respected, and no `SQLCODE` error is returned.
- Automatic starting of non-root DB2 instances at system reboot is not supported

Health monitor limitations

The following health monitor features are not supported in non-root installations:

- Running script or task actions on alert occurrences
- Sending alert notifications

Partitioned database limitation

Only single-partition databases are supported in non-root installations. You cannot add additional database partitions.

Listing DB2 database products

The output produced by the `db2ls` command, when run as a non-root user, is different than the output produced when run as a root user. For details, refer to the `db2ls` command topic.

DB2 copies

Each non-root user can have only one copy of a DB2 database product installed.

DB2 instance limitation

In non-root installations, one DB2 instance is created during installation. Additional instances cannot be created.

DB2 instance actions can be performed only by the instance owner

Root installations and non-root installations can coexist on the same computer in different installation paths. However, a non-root instance can

be updated, or dropped (using the `db2_deinstall` command), only by the non-root user who owns the non-root instance.

A DB2 instance created by a user with root user authority can be updated or dropped only by a user with root user authority.

DB2 instance commands

The following DB2 instance commands are unavailable in non-root installations:

db2icrt

When installing a DB2 database product as a non-root user, a single instance is automatically created and configured. Further instances cannot be created in non-root installations. However, if the automatically created instance needs to be configured, you can use the non-root install configuration command, `db2nrcfg`.

db2iupdt

The `db2iupdt` command cannot be used for non-root instances. Instead, use the non-root install configuration command (`db2nrupdt`) to update the non-root DB2 instance. However, updating the non-root instance is normally not required because it gets updated automatically when updating your DB2 database product.

db2idrop

The instance that gets automatically created during non-root installations cannot be dropped. The DB2 database product must be uninstalled to drop the DB2 instance.

db2iupgrade

Upgrading is not supported for non-root installations. To upgrade a non-root instance, use the `db2nrupgrade` command.

Upgrading limitation

Root instances cannot be upgraded to a non-root instance.

Post-installation actions can be performed only by the DB2 instance owner

Root installations and non-root installations can coexist on the same computer. However, only the original non-root user who installed the DB2 database product can perform subsequent actions such as:

- Applying fix packs
- Adding features
- Installing add-on products

Adjusting ulimit values

The `ulimit` command on UNIX and Linux operating systems sets or reports user resource limits, such as data and stack limits. For root instances, the database server dynamically updates required `ulimit` settings without changing the permanent settings. However, for non-root instances, the `ulimit` settings can only be checked during installation. A warning message is displayed if the settings are inadequate. Root user authority is required to change the `ulimit` settings.

Limitations that can be overcome by running db2rfe

There are further limitations on non-root installations which can be overcome by running the `db2rfe` command. The following features and abilities are initially unavailable in non-root installations:

- Operating system-based authentication
- High Availability (HA) feature
- The ability to reserve service names in the `/etc/services` file
- The ability to increase user data limits (ulimits). This ability applies only to AIX. On other operating systems, user data limits must be increased manually.

Run the Enable root features for non-root install command (`db2rfe`) to enable these features and abilities. Running the `db2rfe` command is optional, and must be run by a user with root user authority.

Authentication type in non-root installations

Operating system-based authentication is the default authentication type for DB2 database products. Since non-root installations do not support operating system-based authentication, if you choose not to run the `db2rfe` command after installing your DB2 database product as a non-root user, then you must manually set the authentication type. You can do so by updating the following parameters in the database manager configuration (`dbm cfg`) file:

- `clnt_pw_plugin` (Client userid-password plug-in configuration parameter)
- `group_plugin` (Group plug-in configuration parameter)
- `srvcon_pw_plugin` (Userid-password plug-in for incoming connections at the server configuration parameter)

Installing a DB2 product as a non-root user

Most DB2 database products can be installed as a non-root user.

Before you install any DB2 database product as a non-root user, be aware of the differences between root installations and non-root installations, and the limitations of non-root installations. For more information on non-root installation, see “Non-root installation overview (Linux and UNIX)”.

Prerequisites for installing a DB2 database product as a non-root user are:

- You must be able to mount the installation DVD, or have it mounted for you.
- You must have a valid user ID that can be used as the owner of a DB2 instance. User IDs have the following restrictions and requirements:
 - Must have a primary group other than `guests`, `admins`, `users`, and `local`
 - Can include lowercase letters (a–z), numbers (0–9), and the underscore character (`_`)
 - Cannot be longer than eight characters
 - Cannot begin with `IBM`, `SYS`, `SQL`, or a number
 - Cannot be a DB2 reserved word (`USERS`, `ADMINS`, `GUESTS`, `PUBLIC`, or `LOCAL`), or an SQL reserved word
 - Cannot use any User IDs with root privilege for the DB2 instance ID, DAS ID or fenced ID.
 - Cannot include accented characters
 - If existing user IDs are specified instead of creating new user IDs, make sure that the user IDs:
 - Are not locked
 - Do not have expired passwords
- The hardware and software prerequisites that exist for the product you are installing apply to the non-root user just as they do for root users.

- On AIX Version 5.3, Asynchronous I/O (AIO) must be enabled. It is strongly recommended the system has I/O Completion Ports (IOCP) enabled.
- Your home directory must be a valid DB2 path.

DB2 installation paths have the following rules:

- Can include lowercase letters (a–z), uppercase letters (A–Z), and the underscore character (_)
- Cannot exceed 128 characters
- Cannot contain spaces
- Cannot contain non-English characters

Installing DB2 database products as a non-root user is transparent to the non-root user. In other words, there is nothing special a non-root user needs to do to install a DB2 database product, other than being logged in as a non-root user.

To perform a non-root installation:

1. Log in as a non-root user
2. Install your DB2 database product using any of the methods available to you.

Options include:

- The DB2 Setup wizard (GUI install)
- The `db2_install` command
- The `db2setup` command with a response file (silent install)

Note: Since non-root users cannot choose the directory where DB2 database products are installed, any **FILE** keyword in your response file is ignored.

3. After the DB2 database product is installed, you must open a new login session to use the non-root DB2 instance. Alternatively, you can use the same login session if you set up the DB2 instance environment with `$HOME/sqllib/db2profile` (for Bourne shell and Korn shell users) or `$HOME/sqllib/db2chsrc` (for C shell users), where `$HOME` is the non-root user's home directory.

After the DB2 database product is installed, verify your operating system user process resource limits (ulimits). If the minimum ulimit values are not met, the DB2 engine can encounter unexpected operating resource shortage errors. These errors can lead to a DB2 database system outage.

Enabling root-based features in non-root installations with `db2rfe`

There are several features and abilities in non-root installations that are initially unavailable but can be enabled by running the `db2rfe` command.

This task requires root user authority.

To enable the features and abilities that are initially unavailable in non-root installations:

1. Locate the sample configuration files. Two sample configuration files are provided:
 - `$HOME/sqllib/instance/db2rfe.cfg` is pre-configured with default values for the non-root DB2 instance
 - `$HOME/sqllib/cfg/db2rfe.cfg.sample` is not configured
 where `$HOME` is the non-root user's home directory.
2. Copy one of the sample configuration files to a different location so the original file remains unaltered.

3. Update the copied configuration file as needed. This configuration file is input to the db2rfe command.

An example of a configuration file is:

```
INSTANCENAME=db2inst2
SET_ULIMIT=NO
ENABLE_HA=NO
ENABLE_OS_AUTHENTICATION=NO
RESERVE_REMOTE_CONNECTION=NO
  **SVCENAME=db2c_db2inst2
  **SVCEPORT=48000
RESERVE_TEXT_SEARCH_CONNECTION=NO
  **SVCENAME_TEXT_SEARCH=db2j_db2inst2
  **SVCEPORT_TEXT_SEARCH=55000
```

Note:

- The value for the **INSTANCENAME** parameter is filled in automatically by DB2 installer
- The **SET_ULIMIT** parameter is available only on AIX. On other operating systems, a user with root authority needs to set ulimit values manually.
- The default value for the other keywords is NO
- Child parameters (such as **SVCENAME**) are commented out by default. Comments are denoted with **
- If you set a parameter to YES, and if it has any child parameters, it is recommended that you uncomment the child parameters and provide appropriate values. Any port values that are provided are examples. Ensure the port values you assign are free.

An example is provided below to show an edited configuration file that will enable the following features and abilities:

- High availability
- Operating system-based authentication
- DB2 Text Search, with a service name of db2j_db2inst2 and a port value of 55000

To enable these features and abilities, edit the configuration file as follows:

```
INSTANCENAME=db2inst2
SET_ULIMIT=NO
ENABLE_HA=YES
ENABLE_OS_AUTHENTICATION=YES
RESERVE_REMOTE_CONNECTION=NO
  **SVCENAME=db2c_db2inst2
  **SVCEPORT=48000
RESERVE_TEXT_SEARCH_CONNECTION=YES
  SVCENAME_TEXT_SEARCH=db2j_db2inst2
  SVCEPORT_TEXT_SEARCH=55000
```

4. Log in with root user authority.
5. Navigate to the \$HOME/sql1lib/instance directory, where \$HOME represent's the non-root user's home directory.
6. Run the db2rfe command using the following syntax:

```
db2rfe -f config_file
```

where *config_file* is the configuration file created in Step 3.

To keep root-based features enabled on non-root installations, rerun the db2rfe command after applying fix packs or upgrading to a new version.

Applying fix packs to a non-root installation

The task of applying fix packs to a non-root installation is essentially the same as applying fix packs to a root installation, with a few exceptions.

Before applying fix packs to a non-root installation, you must log on with the user ID that was used to install the non-root installation.

If you enabled root features in your non-root installation using the `db2rfe` command, you should locate the configuration file that was used when running the `db2rfe` command. That configuration file will be needed to re-enable the root features after you apply the fix pack.

To apply a fix pack to a non-root installation:

1. Apply your fix pack according to the “Applying fix packs” topic.

Note: The `installFixPack` command's `-b` option is invalid for non-root installations.

2. Optional: Run the `db2rfe` command. If you had root-based features enabled in your non-root installation, and if you want to re-enable those features, the `db2rfe` command must be rerun. Running this command requires root user authority.

Note: If you edited `$HOME/sqllib/instance/db2rfe.cfg` when you first enabled root features, that configuration file will not have been overwritten when applying the fix pack, so you can reuse that file when running the `db2rfe` command. However, you should also check `$HOME/sqllib/cfg/db2rfe.cfg.sample`. If the fix pack introduced any new root features available to non-root installations, `$HOME/sqllib/cfg/db2rfe.cfg.sample` shows the new features.

Removing non-root DB2 database products using `db2_deinstall` (Linux and UNIX)

This task provides steps for removing non-root DB2 database products or components using the `db2_deinstall` command.

You must stop the non-root instance before running the `db2_deinstall` command.

- This task applies to DB2 database products that were installed without root user authority. A separate task exists for uninstalling DB2 database products that were installed with root user authority.
- As with root users, non-root users can use the `db2_deinstall` command to uninstall DB2 database products. The `db2_deinstall` command for non-root installations has the same options as root installations, and has an extra option: `-f sqllib`.
- It is important to note that running `db2_deinstall` as a non-root user uninstalls the DB2 database product *and* drops the non-root instance. This is different than root installations, where running `db2_deinstall` only uninstalls the DB2 database program files.
- You cannot remove DB2 database products using a native operating system utility, such as `rpm` or `SMIT`.

To uninstall a DB2 database product that was installed by a non-root user:

1. Log in with the user ID that was used to install the DB2 database product.
2. Navigate to the `$HOME/sqllib/install` directory, where `$HOME` is your home directory.

3. Run the `db2_deinstall` command.

Note:

- If you run the `db2_deinstall` command with the `-a` option, the DB2 database program files are removed, but any configuration files are left behind in a backup directory called `sqllib_bk`.
- If you run the `db2_deinstall` command with the `-a -f sqllib` option, the entire `sqllib` subdirectory in your home directory will be removed. If you have any files in `sqllib` you want to keep, be sure to copy them elsewhere before running `db2_deinstall -a -f sqllib`.
- As with root installations, running the `db2_deinstall` command with the `-F` option against a non-root installation allows the non-root user to remove specific DB2 features.

IBM Data Server Drivers

IBM Data Server Driver restrictions

The IBM Data Server Driver Package simplifies application deployment. However, certain restrictions apply.

The following restrictions apply to IBM Data Server Driver Package:

- No other database product can be installed in the same path if IBM Data Server Driver Package is already installed.
- On Windows operating systems, you can install a maximum of 16 copies of IBM Data Server Driver Package.
- To connect to a z/OS server or a System i server, you must register a DB2 Connect license key. (Retrieve the license file from your Passport Advantage distribution, for example `db2conpe.lic`, then copy the license file to the license directory under the directory where the driver was installed.)
- XA connections against a z/OS server are supported. However, XA connections against a System i server are not supported.
- If you use the configuration file `db2dsdriver.cfg` to specify aliases, the following entries must contain a value:
 - DSN_Collection entries (alias, name, host, and port)
 - database entries (name, host, port).

These entries must be specified and cannot be empty.

- The CLI/ODBC configuration keyword **DBNAME** is not supported.
- The CLI LOAD utility statement attribute, `sql_attr_use_load_api`, is not supported.

Functionality supported with restrictions

- There is no local database directory. Lightweight Directory Access Protocol (LDAP) is supported, but the LDAP cache is not saved to disk.
- Runtime support for embedded SQL is available with the following restrictions:
 - Support is runtime only; there is no PREP or BIND capability.
 - Support is available for DDL and DML SQL statements that are called from EXEC SQL or cursors.
 - When reading the configuration file `db2dsdriver.cfg`, embedded SQL applications can only access connection information (for example, database name, host name, and port number).

- Sysplex capabilities are not supported.
- The following APIs are not supported:
 - APIs for loading (db2Load and db2LoadQuery), exporting (db2Export), and importing (db2Import) data.
 - API for getting the current user authorities (sqluadau).
 - API for forcing users and applications off the system (sqlefrce).
- Executing an embedded compound statement with sub-statements is not supported.

Functionality not supported

- DB2 Command Line Processor (CLP)
- administrative APIs
- installation program
- The CLIENT type authentication is not supported by the IBM Data Server Driver for ODBC and CLI and by the IBM Data Server Driver Package.

Known limitations

For information on short term limitations, go to <http://www.ibm.com/developerworks/wikis/display/DB2/IBM+Data+Server+Driver+Limitations>

db2dsdriver configuration file

The configuration file `db2dsdriver.cfg` contains database directory information and client configuration parameters in a human-readable format.

The `db2dsdriver.cfg` configuration file is an ASCII file containing various keywords and values that can be used to make a connection to a supported database through ODBC, CLI, .NET, OLE DB, or open source (PHP or Ruby) and the applications using the keywords. The keywords are associated with the database *alias name*, and affect all the applications that access the database. You can also use this configuration file to specify Sysplex-related settings such as Sysplex workload balancing.

This configuration file can be used with the following data server drivers:

- cli driver:
 - IBM Data Server Driver for ODBC and CLI
- ds driver:
 - IBM Data Server Driver Package
- IBM data server clients:
 - IBM Data Server Client
 - IBM Data Server Runtime Client

In Version 9.7 Fix Pack 2 and earlier, IBM data server clients use the `db2dsdriver.cfg` configuration file only to retrieve Sysplex-related settings, such as WLB and ACR. Database connection information and properties should come from a different source, such as the database catalog, a connection string, the `db2cli.ini` initialization file, or .NET object properties.

You do not need to create and populate the `db2dsdriver.cfg` configuration file for these drivers. They can function without this configuration file. However, instead

of specifying database name, host, port, and configuration parameters information in your applications, use the configuration file to specify aliases.

In addition, if you have an existing IBM data server client (IBM Data Server Runtime Client or IBM Data Server Client), you can copy the existing database directory information into the `db2dsdriver.cfg` configuration file using the `db2dsdcfgfill` command. Using this command, the configuration file is populated based on the contents of the local database directory, node directory, and Database Connection Services (DCS) directory of a specific database manager instance.

In Version 9.7 Fix Pack 3 and later, parameters in the `db2dsdriver.cfg` configuration file can be used when an IBM data server client connects to a cataloged database. Use an IBM data server client to connect to a cataloged database in one of the following ways:

- Specify a data source name in the connection string.
- Specify a value for the **DBALIAS** parameter in the `db2cli.ini` initialization file.

When an IBM data server client connects to a cataloged database, the application identifies database details by looking at a series of sources in the following order of precedence:

1. The connection string.
2. The `db2cli.ini` file.
3. The `db2dsdriver.cfg` file.

The various keywords, values and specific database connection settings specified in the configuration file apply to all ODBC, CLI, .NET or open source application connections made to that database or alias.

The `db2dsdriver.cfg` configuration file can be edited. However, the file must be edited manually. In addition, after editing the file, you must restart the application for the changes to take effect. The configuration file is created in one of the following paths:

- for a data server client, the configuration file is created in a folder named `cfg` under the instance path:
 - On AIX, HP-UX, Linux, or Solaris operating systems, `instance_path/cfg`
 - On Windows operating systems, `instance_path\cfg`
- for a data server driver, the configuration file is created in a folder named `cfg` under the installation path:
 - On AIX, HP-UX, Linux, or Solaris operating systems, `installation_path/cfg`
 - On Windows operating systems, `installation_path\cfg`

For example, if the instance name is `DB2INST1`, the `db2dsdriver.cfg` file is created in a location similar to the following:

- On AIX: `/home/DB2INST1/sql1lib/cfg/db2dsdriver.cfg`
- On Windows: `C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2INST1\DB2\cfg\db2dsdriver.cfg`

The following restrictions apply to the `db2dsdriver.cfg` configuration file:

- The configuration file cannot contain multiple identical entries for a database with the following properties: database name, server name and port number. In addition, the configuration file cannot contain multiple identical database alias entries.

- The DSN_Collection entries (alias, name, host, and port) and the database entries (name, host, port) should not be empty and must contain a value.
- If multiple parameters are defined on a single line, they are ignored.

Note: As part of casing of the group names, `</wlb>` is allowed as a closing tag for `<WLB>`

A sample configuration file is shipped to help you get started. It is located in the `sqllib/cfg` directory. The configuration file contains the following sections:

- the data source name section contained within the DSN_Collection tags,
- the database information section contained within the databases tags,
- the global attributes section contained within the parameters tags.

```
<configuration>
<DSN_Collection>
  <dsn alias="alias1" name="name1" host="server1.net1.com" port="50001"/>
  <!-- Long aliases are supported -->
  <dsn alias="longaliasname2" name="name2" host="server2.net1.com" port="55551">
    <parameter name="Authentication" value="Client"/>
  </dsn>
</DSN_Collection>
<databases>
  <database name="name1" host="server1.net1.com" port="50001">
    <parameter name="CurrentSchema" value="OWNER1"/>
    <wlb>
      <parameter name="enableWLB" value="true"/>
      <parameter name="maxTransports" value="50"/>
    </wlb>
    <acr>
      <parameter name="enableACR" value="true"/>
    </acr>
  </database>
  <!-- Local IPC connection -->
  <database name="name3" host="localhost" port="0">
    <parameter name="IPCInstance" value="DB2"/>
    <parameter name="CommProtocol" value="IPC"/>
  </database>
</databases>
<parameters>
  <parameter name="GlobalParam" value="Value"/>
</parameters>
</configuration>
```

db2dsdcfgfill - Create configuration file db2dsdriver.cfg

Creates and populates the configuration file `db2dsdriver.cfg` automatically.

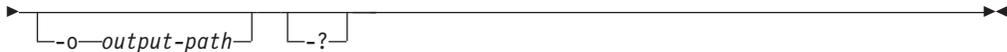
Description

After installing IBM Data Server Driver Package, you can run this command to automatically create and populate the `db2dsdriver.cfg` configuration file in a human-readable format.

The `db2dsdcfgfill` command copies the existing database directory information from either the existing IBM Data Server Client or IBM Data Server Runtime Client, and copies the information to the `db2dsdriver.cfg` configuration file.

Command syntax

```
►► db2dsdcfgfill -i instance-name -p instance-path [-db2cliFile db2cli.ini-path] [-migrateCliIniFor.NET -db2cliFile db2cli.ini-path]
```



Command parameters

-i *instance-name*

The name of the database manager instance whose database directory, node directory, and Database Connection Services (DCS) directory will be used as input by the db2dsdcfgfill command.

Cannot be used in combination with **-p** or **-migrateCliIniFor.NET**.

-p *instance-path*

The full path of the database manager instance directory under which the system database directory, node directory, and DCS directory are located.

Cannot be used in combination with **-i** or **-migrateCliIniFor.NET**.

-migrateCliIniFor.NET

Copies certain entries from db2cli.ini into the db2dsdriver.cfg file. This option is for Microsoft Windows systems only. Only the following keywords will be migrated:

- txnisolation
- connecttimeout
- currentschema

The keywords will be migrated in the following manner:

- entries in the common section of the db2cli.ini will be copied to the global section of the db2dsdriver.cfg file.
- entries that have a database name, host name, and port information will be copied to the database section.
- entries for cataloged databases will be copied to the DSN section.

Cannot be used in combination with **-i** or **-p**.

Note: The IBM Data Server and Runtime Clients do not support the use of the db2dsdriver.cfg file.

-db2cliFile *db2cli.ini-path*

The full path of the db2cli.ini file. This option is for Microsoft Windows systems only.

-o *output-path*

The path where the db2dsdcfgfill command will create the db2dsdriver.cfg configuration file. The configuration file should be created in one of the following paths:

- for IBM Data Server Client or IBM Data Server Runtime Client, the configuration file should be created under the instance path:
 - On AIX, HP-UX, Linux, or Solaris operating systems, *instance_path*
 - On Windows operating systems, *instance_path*
- for IBM Data Server Driver Package, create the configuration file under the installation path:
 - On AIX, HP-UX, Linux, or Solaris operating systems, *installation_path*
 - On Windows operating systems, *installation_path*

-? Displays usage information.

Usage notes

When you run the `db2dsdcfgfill` command, if a `db2dsdriver.cfg` configuration file already exists, the existing `db2dsdriver.cfg` configuration file is overwritten. However, when the `-migrateCliIniFor.NET` option is used with an existing `db2dsdriver.cfg` file, it will merge the information into the existing file instead.

Copying existing database directory information into the `db2dsdriver` configuration file

You can populate the `db2dsdriver.cfg` configuration file with existing database directory information.

You must have an existing Version 9.5 IBM Data Server Client or IBM Data Server Runtime Client installed.

The `db2dsdriver.cfg` configuration file configures the behavior of DB2 CLI, ODBC, open source, or .NET and the applications using them. The keywords are associated with the database alias name, and affect all the applications that access the database.

If you have an existing IBM Data Server Client or IBM Data Server Runtime Client, you can copy the existing database directory information into the `db2dsdriver.cfg` configuration file using the `db2dsdcfgfill` command. Using this command, the configuration file is populated based on the contents of the local database directory, node directory, and Database Connection Services (DCS) directory of a specific database manager instance.

Restrictions

None.

To copy existing database directory information from an IBM Data Server Client or IBM Data Server Runtime Client into the `db2dsdriver` configuration file:

Enter the `db2dsdcfgfill` command. For example, `db2dsdcfgfill -i instance_name -p instance_path -o output_path`. The parameter `-o output_path` indicates the path where the `db2dsdriver.cfg` configuration file is created. The configuration file must be created in one of the following:

- for IBM Data Server Client or IBM Data Server Runtime Client, create the configuration file in a folder named `cfg` under the instance path:
 - On AIX, HP-UX, Linux, or Solaris operating systems, `instance_path/cfg`
 - On Windows operating systems, `instance_path\cfg`
- for IBM Data Server Driver Package, create the configuration file in a folder named `cfg` under the installation path:
 - On AIX, HP-UX, Linux, or Solaris operating systems, `installation_path/cfg`
 - On Windows operating systems, `installation_path\cfg`

Installing IBM Data Server Driver Package (Windows)

On Windows operating systems, IBM Data Server Driver Package simplifies application deployment. This driver, which has a small footprint, is designed to be redistributed by independent software vendors (ISVs) and to be used for application distribution in mass deployment scenarios typical of large enterprises.

IBM Data Server Driver Package is not part of IBM Data Server Client or IBM Data Server Runtime Client. These are different packages and they have their own install images. You need to use the appropriate installer package.

IBM Data Server Driver Package must be installed separately. No other DB2 database product can be installed in the same path if IBM Data Server Driver Package is already installed.

IBM Data Server Driver Package is installed by running the DB2 setup program from the product DVD. There are separate install images for each language.

To install IBM Data Server Driver Package using a response file, or to install it from a fix pack image, you can run the DB2 setup command from the command line. A sample response file is located in the `\samples` subdirectory. For fix pack images, you can download the driver package that contains the setup program from the IBM Support Fix Central Web site: <http://www.ibm.com/support/fixcentral/>. Data Server client and driver packages are found under the **Information Management** product group and **IBM Data Server Client Packages** product selection. Select the appropriate installed version and platform, in this case Windows, and click **Continue**. Click **Continue** again on the next screen and you will be presented with a list of all client and driver packages available for Windows.

The default installation path of IBM Data Server Driver Package is Program Files\IBM\IBM DATA SERVER DRIVER. If a second copy is installed in the same machine, the default directory name is Program Files\IBM\IBM DATA SERVER DRIVER_02. In general, the default directory name is IBM DATA SERVER DRIVER_ *nn* where *nn* is the generated number to make this directory unique.

If you want to install more than one copy of IBM Data Server Driver Package, you can have a maximum of 16 copies. Each copy must be installed to different directories. The default copy name of IBM Data Server Driver Package is IBMDBCL1. If there are further drivers installed on your machine, they will receive default names: IBMDBCL2, IBMDBCL3, and so on.

Note: Installation of multiple copies is an advanced installation method that is not recommended for most users.

The `db2diag.log` path for IBM Data Server Driver Package on Windows is `%SYSTEM APP DATA PATH%\IBM\DB2\DB2COPYName`, where `%SYSTEM APP DATA PATH%` is `Drive:\Documents and Settings\All Users\Application Data\` for Windows 2003, Windows XP and `Drive:\ProgramData\` for Windows 2008 and Windows Vista. The `DB2COPYName` is the name of your IBM Data Server Driver Package copyname provided during the installation.

After installing IBM Data Server Driver Package, you can optionally create and populate the configuration file, `db2dsdriver.cfg`, with database directory information.

Network installation of IBM Data Server Driver Package (Windows)

You can minimize the effort and disk space required to install IBM Data Server Driver Package on client workstations by installing the code on a network share and registering remote client workstations to use the driver as if it is installed locally. This installation method is available on Windows operating systems only.

To set up a network installation of IBM Data Server Driver Package, you install the code on a code server, rather than on each client workstation, and make the code available to remote client workstations through a network share. You run the `db2dsdpreg` utility to set up each remote client workstation and make the required links to the code server. When a registered remote client initiates a database connection, the driver code is dynamically loaded from the code server as required. The remote client then connects to the database in the typical way. When you update the installed IBM Data Server Driver Package, the updated code is automatically available to the remote client workstations.

The following figure shows a network installation of IBM Data Server Driver Package.

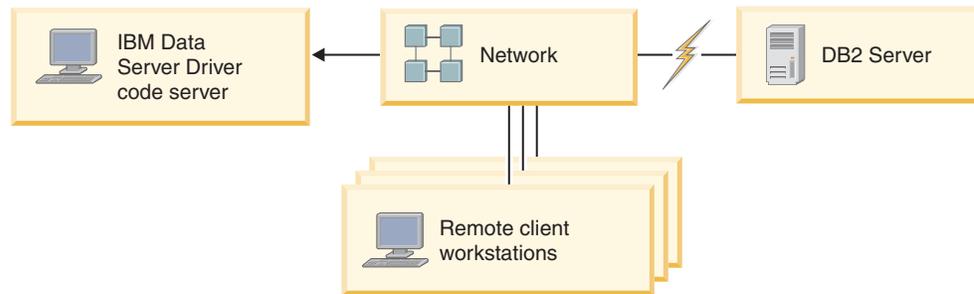


Figure 1. A typical network installation of IBM Data Server Driver Package

Note:

- Remote client workstations that access data on System z or Power Systems™ data servers must have a DB2 Connect license.
- Remote clients on a 32-bit workstation cannot use a 64-bit code server.
- This installation method requires code to be loaded across a LAN connection. The extent of performance loss at program initialization time depends on variables such as the load on and speed of both the network and the code server.

Setting up a network installation of IBM Data Server Driver Package (Windows):

To set up a network installation of IBM Data Server Driver Package, you install the driver on the code server, make the code directory available to client workstations, map a network drive from each client workstation to the code server, and register each client workstation.

Perform the following steps to set up a network installation of IBM Data Server Driver Package on a code server and make the code accessible from a remote client workstation.

To set up a network installation:

1. Install IBM Data Server Driver Package on the code server.
2. Make the code directory on the code server available to all remote client workstations. Refer to “Making the code directory available to remote client workstations (Windows)” on page 44.
3. Map a network drive from each remote client workstation to the code server. Refer to “Mapping a network drive from each remote client workstation to the code server (Windows)” on page 44.

4. Register each remote client workstation by running the db2dsdprep utility. Refer to “Registering remote client workstations by running the db2dsdprep utility (Windows)” on page 45.

Making the code directory available to remote client workstations (Windows):

To load the required code from the code server, each of the remote client workstations must be able to read the directory where IBM Data Server Driver Package is installed.

You must install IBM Data Server Driver Package on the code server.

The following procedure uses Windows XP as an example to show how to make the code directory available to all remote client workstations (in read mode).

To make the code directory available to all remote client workstations:

1. On the code server, start Windows Explorer.
2. Select the directory on the code server that will be used to serve client workstations. For this example, select the C:\Program Files\IBM\IBM DATA SERVER DRIVER directory to set up the share.
3. Select **File** → **Properties** from the menu bar.
4. Click the **Sharing** tab.
5. Click the **Shared This Folder** radio button.
6. In the **Share Name** field, enter a share name that is eight characters or fewer. For example, enter DSDRVRSV.
7. Provide read access to the code directory to all client users:
 - a. Click **Permissions**. The Share Permissions window opens.
 - b. In the **Group or Users Name** list, highlight the Everyone group.

Note: You can give access to the Everyone group, to a group that you have specifically defined for remote client users, or to individual remote client users.

- c. Select **Read**.
- d. Click **OK** until all windows are closed.

Next, map a network drive from each remote client workstation to the network share on the code server.

Mapping a network drive from each remote client workstation to the code server (Windows):

Each remote client workstation must have a network drive that is mapped to the network share on the code server.

- Make the code directory on the code server available to all remote client workstations. Refer to “Making the code directory available to remote client workstations (Windows).”
- Log on to the remote client workstation as a valid user with shared directory access to the code server.

The following procedure uses Windows XP as an example to show how to map a network drive from the remote client workstation to the network share on the code server.

To map a network drive from the remote client workstation:

1. On the remote client workstation, start Windows Explorer.
2. On the **Tools** menu, click **Map Network Drive**.
3. In the **Drive** list, select the drive to which you want to map the location of the code server.
4. In the **Folder** field, specify the location of the share as follows:

```
\\computer_name\share_name
```

where:

```
computer_name
```

represents the computer name of the code server

```
share_name
```

represents the share name of the shared directory on the code server

5. Select the **Reconnect at Logon** check box to make the share persistent.
6. Click **Finish**.

Next, register each remote client workstation to allow remote clients to use the network installation of IBM Data Server Driver Package as if the code is installed locally.

Registering remote client workstations by running the db2dsdpreg utility (Windows):

To access a network installation of IBM Data Server Driver Package, remote client workstations must be registered. Run the db2dsdpreg utility to register each remote client workstation and create the required links to the code server.

- Map a network drive from the remote client workstation to the code server. Refer to “Mapping a network drive from each remote client workstation to the code server (Windows)” on page 44.
- Ensure that either Microsoft Visual C++ 2005 or the appropriate runtime components of the Visual C++ libraries are installed. The runtime libraries are available from the Microsoft runtime library download Web site. For 32-bit systems, install vcredist_x86.exe. For 64-bit systems, install vcredist_x86.exe and vcredist_x64.exe.
- Log on to the remote client workstation as a valid user with shared directory access to the code server. You must have administrator level privileges on the remote client workstation.

Perform the following steps on each remote client workstation that needs to use the network installation of IBM Data Server Driver Package.

Restriction: Remote clients on a 32-bit workstation cannot use a 64-bit code server.

To register a remote client workstation:

1. From a Windows command prompt, issue the following command, where z is the network drive that you mapped to the code server:

```
cd z:\bin
```

2. Issue the following command to run the **db2dsdpreg** utility and write logging information to a file:

```
db2dsdpreg /l c:\db2dsdpreg.log
```

where `c:\db2dsdpreg.log` is the path where the utility will create the log file. If you do not specify a path for the log file, it is created under `My Documents\DB2LOG` for the current user.

Tip: You can view additional registration options, including options to unregister and reregister, by issuing the following command: `db2dsdpreg /h`.

3. When the `db2dsdpreg` utility is done, check the messages in the log file (for example, `c:\db2dsdpreg.log`).

The log file contains error messages for any errors that are encountered when you run the utility.

Installing IBM Data Server Driver Package (Linux and UNIX)

On Linux and UNIX operating systems, IBM Data Server Driver Package simplifies application deployment. This driver, which has a small footprint, is designed to be redistributed by independent software vendors (ISVs) and to be used for application distribution in mass deployment scenarios typical of large enterprises.

To install IBM Data Server Driver Package, you need to obtain the compressed file that contains this driver package. Download this file from the IBM Support Fix Central Web site: <http://www.ibm.com/support/fixcentral/>. Data Server clients and drivers are found under the **Information Management** product group and **IBM Data Server Client Packages** product selection. Select the appropriate installed version and platform and click **Continue**. Click **Continue** again on the next screen and you will be presented with a list of all client and driver packages available for your platform.

On Linux and UNIX operating systems, IBM Data Server Driver Package is installed by running the `installDSDriver` command. This driver package includes database drivers for Java, ODBC/CLI, PHP, and Ruby on Rails, each of which is stored in its own subdirectory. The Java and ODBC/CLI drivers are compressed.

The basic extraction steps for each driver are the same:

1. Uncompress the Data Server Driver Package archive.
2. Copy the files onto the target machine.
3. For the Java and ODBC/CLI drivers, uncompress the driver file into your chosen install directory on the target machine.
4. Optional: remove the compressed driver file.

Java

Path: `jdbc_sqlj_driver/platform`

Filename: `db2_db2driver_for_jdbc_sqlj.zip`

For installation instructions, see: Installing the IBM Data Server Driver for JDBC and SQLJ.

ODBC/CLI

Path: `odbc_cli_driver/platform`

Filename: `ibm_data_server_driver_for_odbc_cli.tar.Z`

For installation instructions, see: Installing the IBM Data Server Driver for ODBC and CLI.

PHP

Path: `php_driver/platform/php32` or `php_driver/platform/php64`

Files: `ibm_db2_n.n.n.so`, `ibm_db2_n.n.n.so`, `pdo_ibm_n.n.n.so`, where *n* represents the version of the extension.

Prerequisite: The PHP drivers require the ODBC/CLI driver also included in this driver package to be installed.

For installation instructions, see: [Setting up the PHP environment on Linux or UNIX](#).

Ruby on Rails

Path: `ruby_driver/platform`

File: `ibm_db-0.10.0.gem`

Prerequisite: The Ruby on Rails driver requires the ODBC/CLI driver also included in this driver package to be installed.

To install the Ruby on Rails driver, from the location of the gem file, run the following command: `gem install ibm_db-0.10.0.gem`. To validate the installation, see: [Verifying installation with IBM Data Servers on Rails applications](#).

After installing Data Server Driver Package, you can optionally create and populate the configuration file, `db2dsdriver.cfg`, with database directory information.

Chapter 3. Database connections for IBM data server clients

Client-to-server communications configuration overview

This topic provides information for choosing a suitable method for configuring client-to-server communications. This topic applies to configuring IBM data server client and server products rather than to database connectivity drivers.

Understanding client-to-server communications: Components and scenarios

The basic components involved in client-to-server communications are described below:

- **Client.** This refers to the initiator of the communications. This role can be filled by any of the following DB2 products or components:
 - IBM Data Server Driver Package
 - IBM Data Server Client or IBM Data Server Runtime Client.
 - DB2 Connect Personal Edition: This product is a superset of the IBM Data Server Client.
 - a DB2 server product: A DB2 server is a superset of the Data Server Client.
- **Server.** This refers to the receiver of the communications request from the client. This role is normally filled by a DB2 for Linux, UNIX, and Windows server product. When DB2 Connect products are present, the term *server* can also mean a DB2 server on a midrange or mainframe platform.
- **Communications protocol.** This refers to the protocol used to send data between the client and server. The DB2 product supports several protocols:
 - TCP/IP. A further distinction can be made between the version: TCP/IPv4 or TCP/IPv6.
 - Named Pipes. This option is available on Windows only.
 - IPC (interprocess communications). This protocol is used for local connections.

There are also some additional components encountered in some environments:

- **DB2 Connect gateway.** This refers to a DB2 Connect Server product that provides a gateway by which IBM data server client can connect to DB2 servers on midrange and mainframe products.
- **LDAP (Lightweight Directory Access Protocol).** In an LDAP-enabled environment, it is not necessary to configure client-to-server communications. When a client attempts to connect to a database, if the database does not exist in the database directory on the local machine then the LDAP directory is searched for information required to connect to the database.

The scenarios listed below illustrate examples of situations covered by client-to-server communications:

- Data Server Client establishes communications with a DB2 server using TCP/IP.
- Data Server Runtime Client establishes communications with a DB2 server using Named Pipes on a Windows network.
- DB2 server establishes communications with another DB2 server via some communications protocol.

- Data Server Client establishes communications with a mainframe DB2 server via a DB2 Connect server using TCP/IP.

When setting up a server to work with development environments (such as IBM Data Studio), you might encounter error message SQL30081N at the initial DB2 connection. A possible root cause is that the firewall at the remote database server has prevented the connection from being established. In this case, verify the firewall is properly configured to accept connection requests from the client.

Understanding client-to-server communications: Types of connections

Generally speaking, references to setting up client-to-server communications refer to *remote connections*, rather than *local connections*.

A *local connection* is a connection between a database manager instance and a database managed by that instance. In other words, the CONNECT statement is issued from the database manager instance to itself. Local connections are distinctive because no communications setup is required and IPC (interprocess communications) is used.

A *remote connection* is one where the client issuing the CONNECT statement to a database is in a different location from the database server. Commonly, the client and server are on different machines. However, remote connections are possible within the same machine if the client and server are in different instances.

Another less common type of connection is a *loopback connection*. This is a type of remote connection where the connection is configured from a DB2 instance (the client) to the same DB2 instance (the server).

Comparison of methods for configuring client-to-server communications

Several methods are available for configuring client-to-server communications. Choosing a suitable method involves answering two questions. The first is *Which tool will you use: Configuration Assistant or command line tools?*

- The Configuration Assistant is a graphical tool provided with versions of the Data Server Client and DB2 server products on Windows and Linux on Intel™ x86 32-bit platforms and AMD64/EM46T platforms. This tool is not provided with the Data Server Runtime Client.
- The command line tools consist of the Command Line Processor (CLP), and the commands db2cfexp (configuration export), and db2cfimp (configuration import).

The second question is: *What type of configuration task do you want to perform?*

Options are:

- Configure a client by entering information manually.
- Configure a client by searching the network for servers to connect to.
- Make databases on a server accessible to one or more clients.
- Use the connection settings for one client as the basis for configuring additional clients.

With answers to these questions, you can use the table below to identify the appropriate configuration method. Links to each method are provided at the end of this topic. Notes follow the table that provide more details.

Table 7. Tools and methods for configuring a client-to-server connection

Type of configuration task	Configuration Assistant	Command line
Configure a client by entering information manually	Configure a database connection manually with the Configuration Assistant	Configure client-to-server connections using the command line processor
Configure a client by searching the network for servers to connect to	Configure a database connection by searching the network with the Configuration Assistant	Not applicable
Use the connection settings for one client as the basis for configuring additional clients	<ol style="list-style-type: none"> 1. Create a client profile using the Configuration Assistant 2. Configure database connections using a client profile with the Configuration Assistant 	create and use a client profile using commands db2cfexp and db2cfimp

Note: *Profiles* are used in some methods for configuring client-to-server communications. A *client profile* is a file that contains settings for a client. Settings can include:

- Database connection information (including CLI or ODBC settings).
- Client settings (including database manager configuration parameters and DB2 registry variables).
- CLI or ODBC common parameters.

A *server profile* is similar to a client profile but contains settings for a server. Profiles can be created and used with the Configuration Assistant or by using the commands db2cfexp (configuration export), and db2cfimp (configuration import).

Note: Configuring a database connection by searching the network with the Configuration Assistant is not a recommended method for DB2 Connect customers connecting to databases on midrange or mainframe platforms.

Supported combinations of client and server versions

This section describes which versions of a client can connect to which versions of a server. This includes support for earlier versions and support for accessing DB2 databases on midrange and mainframe servers.

DB2 client considerations for the IBM DB2 pureScale™ Feature for Enterprise Server Edition

For your application to make full use of DB2 pureScale features, your DB2 server and DB2 client must be at certain release levels:

Server version	Client version	Features available
Version 9.8, or later	Version 9.7, Fix Pack 1 or later	Transaction-level and connection-level workload balancing Automatic client reroute based on workload Client affinities

Server version	Client version	Features available
Version 9.8, or later	Version 9.1, Version 9.5, and Version 9.7 (prior to Fix Pack 1)	Connection-level workload balancing (transaction-level workload balancing is not available) Automatic client reroute based on workload

Combinations of DB2 Universal Database (UDB) Version 8, DB2 Version 9.1, DB2 Version 9.5, and DB2 Version 9.7

DB2 Universal Database (UDB) Version 8, and DB2 Version 9.1 and Version 9.5 clients can access a remote DB2 Version 9.7 server. Note the following restriction:

- There is a restriction when a client is located on the same system as a DB2 server, and they are different versions. In this case, local client-to-server connections using Interprocess Communication (IPC) are not supported. Instead, a connection can be established by treating the connection as a remote connection (called a loopback connection) using TCP/IP.

The following clients and servers can access DB2 Version 9.5, Version 9.1, and DB2 UDB Version 8 servers:

- IBM Data Server Client V9.7
- IBM Data Server Runtime Client V9.7
- IBM Data Server Driver Package V9.7
- IBM Data Server Driver for ODBC and CLI V9.7

However, when a higher version level client accesses a lower version level server, the functionality of the higher version level is not available to the client. For example, IBM Data Server Driver Package Version 9.7 can access a DB2 Version 9.1 server, however, DB2 Version 9.7 functionality is not available to the client.

Access to DB2 Version 9.7 servers from DB2 UDB Version 7 clients

Access from DB2 UDB Version 7 clients is not supported.

Combinations of DB2 Version 9.7 and DB2 products on midrange and mainframe platforms

DB2 Version 9.7 servers support access from the following clients on midrange and mainframe platforms:

- DB2 for z/OS and OS/390 Version 7 or later
- DB2 for i5/OS[®] Version 5 or later
- DB2 for VM and VSE Version 7

The following clients and drivers can access DB2 Connect Version 9.7:

- IBM Data Server Client V9.7
- IBM Data Server Runtime Client V9.7
- IBM Data Server Driver Package V9.7
- IBM Data Server Driver for ODBC and CLI V9.7
- DB2 V9.5 and V9.1 clients and drivers

Communication protocols supported

This topic identifies the supported protocols for connecting from an IBM data server client to a DB2 server. This includes:

- connecting from IBM data server client to midrange or mainframe hosts using DB2 Connect products.
- connecting from mid range or mainframe platforms to databases on DB2 for Linux, UNIX, and Windows.

The TCP/IP protocol is supported on all platforms on which DB2 for Linux, UNIX, and Windows is available. Both TCP/IPv4 and TCP/IPv6 are supported. IPv4 addresses have a four-part structure, for example, 9.11.22.314. IPv6 addresses have an eight-part name, where each part consists of 4 hex digits delimited by a colon. Two colons (::) represents one or more sets of zeros. For example, 2001:0db8:4545:2::09ff:fe7:62dc.

DB2 database products support the SSL protocol and accept SSL requests from applications that use the IBM Data Server Driver for JDBC and SQLJ (type 4 connectivity), IBM Data Server Driver for ODBC and CLI and IBM Data Server Driver Package. Refer to Configuring Secure Sockets Layer (SSL) support in a DB2 instance.

In addition, the Windows Named Pipes protocol is supported on Windows networks. To administer a DB2 database remotely, you must connect using TCP/IP.

Adding database connections using the Configuration Assistant

Configuring client-to-server connections using the Configuration Assistant (CA)

The Configuration Assistant is a graphical tool that can be used to configure database connections between a client and a remote DB2 database.

Important: The Configuration Assistant has been deprecated in Version 9.7 and might be removed in a future release. For more information, see the “Control Center tools and DB2 administration server (DAS) have been deprecated” topic in the *What's New for DB2 Version 9.7* book.

The Configuration Assistant is provided with the IBM Data Server Client and DB2 database products on Windows and Linux (Intel x86 and x64 platforms).

The Configuration Assistant can configure a connection to a database only if the remote database manager is configured to accept inbound client requests. By default, the DB2 database product installation program detects and configures most protocols for inbound client connections.

You can configure a connection to a database using one of the following methods:

“Configuring a database connection by searching the network using the Configuration Assistant” on page 55

Use this method if you don't have any information about the database you want to connect to. This method will search your network and list all the databases available to you. A DB2 Administration Server (DAS) must be running and enabled on the servers for the discovery feature of the CA to return information about DB2 systems.

“Configuring database connections using a client profile with the Configuration Assistant” on page 57

Use this method if you have been given a file that contains all the necessary information to access the target database. This method can also be used to catalog and connect to multiple databases specified in the access profile file.

“Configuring a database connection manually using the Configuration Assistant”

Use this method if you know all the information necessary to connect to the target database. You must know:

- The communication protocols supported by the server on which the target database resides
- The appropriate communication parameters for the server's protocols
- The name of the database

Configuring a database connection manually using the Configuration Assistant

If you have the information for the database you want to connect to and the server upon which it resides, you can manually enter all of the configuration information. This method is analogous to entering commands using the command line processor, however, the parameters are presented graphically.

Before you configure a connection to a database manually using the Configuration Assistant (CA):

- Ensure that you have a valid DB2 user ID for the database you want to connect to.
- If you are configuring a connection from a system that has a DB2 server or DB2 Connect server product installed, ensure that you have a user ID with SYSADM or SYSCTRL authority for the database manager instance.

To configure a connection to a database manually using the CA:

1. Log on to the system with a valid DB2 user ID.
2. Start the CA. The CA can be started from the Start menu on Windows or using the db2ca command.
3. On the CA menu bar, under **Selected**, choose **Add Database Using Wizard**.
4. Select the **Manually configure a connection to a database** radio button and click **Next**.
5. If you are using Lightweight Directory Access Protocol (LDAP), select the radio button that corresponds to the location where you want DB2 directories to be maintained. Click **Next**.
6. Select the radio button that corresponds to the protocol that you want to use from the **Protocol** list. (Note: While APPC, APPN, or NetBIOS might still display as options, they are no longer supported.) If DB2 Connect is installed on your system and you select TCP/IP, you have the option to select **The database physically resides on a host or OS/400 system**. If you select this check box, you have the option of selecting the type of connection that you want to make to the host or OS/400® database:
 - To make a connection through a DB2 Connect gateway, select the **Connect to the server via the gateway** radio button.
 - To make a direct connection, select the **Connect directly to the server** radio button.

Click **Next**.

7. Type the required communication protocol parameters and click **Next**.
8. Type the database alias name of the remote database that you want to add in the **Database name** field and a local database alias name in the **Database alias** field. If you are adding a host or OS/400 database, type the location name for an OS/390 or z/OS database, the RDB name for an OS/400 database, or the DBNAME for a VSE or VM database in the **Database name** field. Optionally, you can add a comment that describes this database in the **Comment**.
Click **Next**.
9. If you are planning to use ODBC, register this database as an ODBC data source. Ensure that ODBC is installed before performing this operation. Click **Next**.
10. In the **Specify the node options** window, select the operating system, and type the remote instance name for the database system you want to connect to.
11. In the **Specify the system options** window, ensure that system name, host name, and operating system are correct. The information on this panel is used to configure the administration node. You can optionally enter a comment. Click **Next**.
12. In the **Specify the security options** window, specify the security option that will be used for authentication.
13. Click **Finish**. You can now use this database. Select the **Exit** menu action to close the CA.

Configuring a database connection by searching the network using the Configuration Assistant

You can use the Configuration Assistant (CA) to search a network for databases.

Before you configure a database connection by searching the network:

- Ensure that you have a valid DB2 user ID.
- If you are configuring a connection from a system that has a DB2 Server or DB2 Connect server product installed, ensure that you have a user ID with SYSADM or SYSCTRL authority for the instance.

The search method feature might be unable to detect a remote system if:

- It is used in a cluster environment.
- The DB2 Administration Server (DAS) is not running on the remote system.
- The search times out. By default, the search will scan the network for 1 second; this might not be long enough to detect the remote system. You can set the **DB2DISCOVERYTIME** registry variable to specify a longer period of time.
- The network that the search is running on is configured so that the search does not reach the remote system required.

The following points apply to cases where you want to explicitly configure an IPv6 address on a network that supports IPv6:

- The system must be listed under **Known Systems**.
- Only the Configuration Assistant Advanced View supports explicitly configure an IPv6 connection.

To configure a database connection by searching the network:

1. Log on to the system with a valid DB2 user ID.
2. Start the CA. The CA can be started from the Start menu on Windows or using the db2ca command on both Windows and UNIX systems.

3. On the CA menu bar, under **Selected**, choose **Add Database Using Wizard**. The **Add Database** Wizard opens.
4. Select the **Search the network** radio button and click **Next**.
5. Double-click the folder beside **Known Systems** to list all the systems known to your client or double-click the folder beside **Other Systems** to list all the systems on the network. If no systems are listed, you can click **Add System** to specify one. After you add a system, it displays in the **Known Systems** list.
6. Expand the entries for the system you are interested in until you see the database you want to add. Select the database. Click **Next**.
7. Type a local database alias name in the **Database alias** field and optionally type a comment that describes this database in the **Comment** field.
8. If you are planning to use ODBC, register this database as an ODBC data source. ODBC must be installed to perform this operation.
9. Click **Finish**. You can now use the database you added. Click **Close** to exit the CA.

Creating a client profile using the Configuration Assistant

This task involves exporting settings from an existing client into a *client profile* using the Configuration Assistant (CA). This task is part of a larger task of setting up one or more clients using the settings from an existing client.

To create a client profile using the CA:

1. Log on to the system with a valid DB2 user ID.
2. Start the CA. The CA can be started from the Start menu on Windows or using the db2ca command.
3. From the **Configure** menu, select **Export Profile**.
4. Select one of the following options:

All If you want to create a profile that contains all of the databases cataloged on your system, and all of the configuration information for this client. Type a name for your client profile and click **Save**.

Database Connections

If you want to create a profile that contains all of the databases cataloged on your system *without* any of the configuration information for this client. Type a name for your client profile and click **Save**.

Customize

If you want to select a subset of the databases that are cataloged on your system, or a subset of the configuration information for this client. In the **Customize Export Profile** window:

- a. Type a name for your client profile.
- b. Select the **Database connections** checkbox to include database connections in the client profile.
- c. From the **Available database aliases** box, select the databases to be exported and click > to add them to the **Selected database aliases** box. To add all of the available databases to the **Selected database aliases** box, click >>.
- d. Select the check boxes that correspond to the options that you want to set up for the target client. Database manager configuration parameters can be updated and customized for a target machine.
- e. Click **Export** to complete this task.
- f. Check your results displayed in the Results tab.

Once you have completed this task, you can configure other clients using the client profile you have created.

Configuring database connections using a client profile with the Configuration Assistant

This task involves configuring a client by using a *client profile* that you have created or obtained previously. This task is part of a larger task of setting up one or more clients using the settings from an existing client. These steps can be repeated for each client you want to configure.

1. Log on to the system with a valid DB2 user ID.
2. Start the CA. The CA can be started from the Start menu on Windows or using the `db2ca` command.
3. From the **Configure** menu, select **Import Profile**.
4. Select one of the following import options. You can choose to import all or a subset of the information in a client profile.

All Select this option to import everything in a client profile. Open the client profile you want to import.

Customize

Select this option to import a subset of the client profile, such as a specific database. From the **Customize Import Profile** window:

- a. Select the client profile you want to import and click **Load**.
- b. Select the databases to be imported from the **Available database aliases** box and click **>** to add them to the **Selected database aliases** box. Click **>>** to add all of the available databases to the **Selected database aliases** box.
- c. Select the check boxes that correspond to the options that you want to customize.
- d. Click **Import** to complete this task.
- e. Check your results displayed in the Results tab.

Testing a database connection using the Configuration Assistant

After configuration, test your database connection in the Configuration Assistant.

To test a database connection:

1. Start the **Configuration Assistant**.
2. Highlight the database in the details view and select **Test Connection** from the **Selected** menu. The Test Connection window opens.
3. Select one or more types of connection that you want to test (**CLI** is the default). You can test more than one type at the same time. Enter a valid user ID and password for the remote database and click **Test Connection**. If the connection is successful, a message confirming the connection appears on the Results page. If the connection test failed, you will receive a help message. To change any settings you might have incorrectly specified, select the database in the details view and select **Change Database** from the **Selected** menu item.

When setting up a server to work with development environments (such as IBM Data Studio), you might encounter error message SQL30081N at the initial DB2 connection. A possible root cause is that the firewall at the remote database server has prevented the connection from being established. In this case, verify the firewall is properly configured to accept connection requests from the client.

LDAP considerations for the Configuration Assistant

In an LDAP-enabled environment, the directory information about DB2 servers and databases is stored in the LDAP directory. When a new database is created, the database is automatically registered in the LDAP directory. During a database connection, the client accesses the LDAP directory to retrieve the required database and protocol information and uses this information to connect to the database.

However, you can still use the CA in the LDAP environment to:

- Manually catalog a database in the LDAP directory.
- Register a database cataloged in LDAP as an ODBC data source.
- Configure CLI/ODBC information about the LDAP server.
- Remove a database cataloged in the LDAP directory.

Configuring client-to-server connections using the command line processor

Configuring client-to-server connections using the command line processor

This task describes how to configure a connection from an IBM data server client to a remote database server using the command line processor (CLP).

Before you configure a client to server connection, ensure:

- Network communications is set up between the machine with the IBM data server client and the machine with the DB2 server. One way to verify this for the TCP/IP protocol is to use the ping command.
- The DB2 server is configured to work on the network. This is normally done as part of installing and configuring the DB2 server product.

Separate topics are provided to guide you through each of the following steps. Some steps have a version for each supported protocol:

1. Identify the communication parameter values for the remote database server. Worksheets are provided:
 - TCP/IP worksheet
 - Named Pipes worksheet
2. If you are using TCP/IP, you have the option to update the client's hosts file and services file with communication parameter values for the remote database server. This step does not apply to Named Pipes.
3. Catalog the server node from the client. Instructions are provided for each communications protocol:
4. Catalog the database that you want to connect to on the client. Refer to .
5. Test the client-to-server connection. Refer to .

Named pipe connections

Named Pipes worksheet for configuring Named Pipes on the client:

Use the following worksheet to help identify the required parameter values for configuring Named Pipes communications.

Table 8. Named Pipes parameter values worksheet

Parameter	Description	Sample Value	Your Value
Computer name (<i>computer_name</i>)	The computer name of the server machine. On the server machine, to locate the value for this parameter, click on Start and select Settings, Control Panel . Double-click on the Network folder and select the Identification tab. Record the computer name.	server1	
Instance name (<i>instance_name</i>)	The name of the instance on the server to which you are connecting.	db2	
Node name (<i>node_name</i>)	A local alias, or nickname, that describes the node to which you are trying to connect. You can choose any name you want; however, all node name values within your local node directory must be unique.	db2node	

Cataloging a Named Pipes node from a client using the CLP:

Cataloging a Named Pipes node adds an entry to the client's node directory to describe the remote node. This entry specifies the chosen alias (*node_name*), the remote server's workstation name (*computer_name*), and the instance (*instance_name*) that the client will use to access the remote DB2 server.

To catalog a Named Pipes node on an IBM data server client, type the following command in the command line processor (CLP):

```
db2 => catalog npipe node node_name
db2 => remote computer_name instance instance_name

db2 => terminate
```

To catalog a remote node called db2node that is located on a server called server1 in the db2 instance, use:

```
db2 => db2 catalog npipe node db2node remote server1 instance db2

db2 => terminate
```

TCP/IP connections

TCP/IP worksheet for configuring a client to server connection:

As you proceed through the configuration steps, use the *Your Value* column in the following table to record the required values.

Table 9. TCP/IP parameter values worksheet

Parameter	Description	Sample Value	Your Value
Version of the IP protocol	Options are: <ul style="list-style-type: none"> IPv4: addresses look like this 9.21.15.235 IPv6: addresses look like this: 2001:0db8:4545:2::09ff:fef7:62dc 	IPv4	
Host name <ul style="list-style-type: none"> Hostname (<i>hostname</i>) or IP address (<i>ip_address</i>) 	To resolve the hostname of the remote system, enter the hostname command at the server. To resolve the IP address, enter the ping hostname command.	myserver or 9.21.15.235 or an IPv6 address	
Service Name <ul style="list-style-type: none"> Connection Service name (<i>svcname</i>) or Port number/Protocol (<i>port_number/tcp</i>) 	Values Required in the services file. The Connection Service name is an arbitrary name that represents the connection port number (<i>port_number</i>) on the client. The port number must be the same as the port number that the <i>svcname</i> parameter maps to in the services file on the server system. (The <i>svcname</i> parameter is located in the database manager configuration file on the server instance.) This value must not be in use by any other applications, and must be unique within the services file. On Linux or UNIX platforms, this value generally must be 1024 or higher. Contact your database administrator for the values used to configure the server.	server1 or 3700/tcp	
Node name (<i>node_name</i>)	A local alias, or nickname, that describes the node to which you are trying to connect. You can choose any name you want; however, all node name values within your local node directory must be unique.	db2node	

Updating hosts and services files for TCP/IP connections:

This task explains when and how to update the hosts file and services file on the client with communication parameter values for the remote database server. This task is optional for connections using TCP/IP and does not apply to connections using Named Pipes. This task is part of the larger task of configuring client-to-server connection using the CLP.

You need to update the hosts file if you want to establish a connection to the remote database server using its hostname and your network does not contain a

DNS (domain name server) that can be used to resolve that hostname to an IP address. This step is not required if you want to refer to the remote database server using its IP address.

You need to update the services file if you want to specify a *connection service* name when establishing a connection to the remote database server. A *connection service* is an arbitrary name that represents the connection port number. This step is not required if you want to refer to the remote database server's port number.

- To update the hosts file on the client to resolve the remote server's hostname to its IP address:

1. Use a text editor to add an entry to the hosts file for the server's IP address.

For example:

```
9.26.13.107          myserver # IPv4 address for myserver
2002:91a:519:13:210:83ff:feff:ca71  myserver # IPv6 address for myserver
```

where:

9.26.13.107
represents the *IPv4 ip_address*

2002:91a:519:13:210:83ff:feff:ca71
represents the *IPv6 ip_address*

myserver
represents the *hostname*

represents a comment describing the entry

Note: Note that IPv6 entries are not needed if your host does not belong on an IPv6 network. For hosts in mixed IPv4 and IPv6 networks, an alternate method is to assign different host names for IPv4 and IPv6 addresses. For example:

```
9.26.13.107          myserver # IPv4 address for myserver
9.26.13.107          myserveripv4 # IPv4 address for myserver
2002:91a:519:13:210:83ff:feff:ca71 myserveripv6 # IPv6 address for myserver
```

If the server is not in the same domain as the IBM data server client, you must provide a fully qualified domain name such as `myserver.spifnet.ibm.com`, where `spifnet.ibm.com` represents the domain name.

- To update the services file on the client to resolve a service name to the remote server's port number:

1. Using a text editor, add the Connection Service name and port number to the services file. For example:

```
server1 50000/tcp # DB2 connection service port
```

where:

server1
represents the Connection Service name

50000
represents the connection port number (50000 is the default)

tcp
represents the communication protocol that you are using

represents the beginning of a comment that describes the entry

The following table lists the location of the hosts file and services file referred to in the preceding procedures.

Table 10. Location of the hosts file and services file

Operating System	Directory
Windows 2000 XP/Windows Server 2003	%SystemRoot%\system32\drivers\etc where %SystemRoot% is a system defined environment variable
Linux or UNIX	/etc

Cataloging a TCP/IP node from a client using the CLP:

Cataloging a TCP/IP node adds an entry to the Data Server Client node directory that describes the remote node. This entry specifies the chosen alias (*node_name*), the *hostname* (or *ip_address*), and the *svcname* (or *port_number*) that the client uses to access the remote host.

You must have System Administrative (SYSADM) or System Controller (SYSCTRL) authority, or have the **catalog_noauth** option set to ON. You cannot catalog a node using root authority.

To catalog a TCP/IP node:

1. Log on to the system as a user with System Administrative (SYSADM) or System Controller (SYSCTRL) authority.
2. If you are using a Linux or UNIX client, set up the instance environment. Run the startup script:

For bash, Bourne or Korn shell

```
. INSTHOME/sql1lib/db2profile
```

For C shell

```
source INSTHOME/sql1lib/db2cshrc
```

where *INSTHOME* represents the home directory of the instance.

3. Start the DB2 command line processor. On Windows, issue the `db2cmd` command from a command prompt. On Linux or UNIX, issue the `db2` command from a command prompt.
4. Catalog the node by entering the following commands in the command line processor:

```
db2 => catalog tcpip node node_name remote hostname|ip_address
server service_name|port_number [remote_instance instance_name]
[system system_name] [ostype os_type]
```

```
db2 => terminate
```

where:

- *node_name* represents a local nickname you can set for the computer that has the database you want to catalog.
- *remote_instance* represents the name of the server instance on which the database resides.
- *system_name* represents the DB2 system name that is used to identify the server.
- *ostype_name* represents the operating system type of the server.

Note:

- a. The terminate command is needed to refresh the directory cache.
- b. Although **remote_instance**, **system**, and **ostype** are optional, they are required for users who want to use the DB2 tools.
- c. The *service_name* used on the client does not have to be the same as the one on the server. However, the port numbers that they map to *must* match
- d. While not shown here, the catalog tcpip node command provides the option to explicitly specify the version of IP, namely IPv4 or IPv6.

To catalog a node that you want to call db2node on a remote server myserver.ibm.com that is using port number 50000, you would enter the following from a **db2** prompt:

```
db2 => catalog tcpip node db2node remote myserver server 50000
DB20000I The CATALOG TCPIP NODE command completed successfully.
DB21056W Directory changes may not be effective until the directory cache is
refreshed.
```

```
db2 => terminate
DB20000I The TERMINATE command completed successfully.
```

Cataloging a database from a client using the CLP

This task describes how to catalog a database from a client using the command line processor (CLP).

Before a client application can access a remote database, the database must be cataloged on the client. When you create a database, the database is automatically cataloged on the server with a database alias that is the same as the database name, unless a different database alias was specified.

The information in the database directory, along with the information in the node directory (unless you are cataloging a local database where a node is not needed), is used on the IBM data server client to establish a connection to the remote database.

- You require a valid DB2 user ID. DB2 does not support using root authority to catalog a database.
- You must have System Administrative (SYSADM) or System Controller (SYSCTRL) authority, or have the **catalog_noauth** option set to ON.
- You need the following information when cataloging a *remote* database:
 - Database name
 - Database alias
 - Node name
 - Authentication type (optional)
 - Comment (optional)

Refer to the parameter values worksheet for cataloging a database for more information about these parameters and to record the values that you use.

- The following parameter values are applicable when cataloging a *local* database:
 - Database name
 - Drive
 - Database alias
 - Authentication type (optional)
 - Comment (optional)

Local databases can be uncataloged and recataloged at any time.

To catalog a database on the client:

1. Log on to the system with a valid DB2 user ID.
2. Optional. Update the Your Value column in the "Parameter values worksheet for cataloging a database."
3. If you are using the DB2 database on a Linux or UNIX platform, set up the instance environment. Run the startup script:

For bash, Bourne or Korn shell

```
. INSTHOME/sql1lib/db2profile
```

For C shell

```
source INSTHOME/sql1lib/db2cshrc
```

where: *INSTHOME* represents the home directory of the instance.

4. Start the DB2 command line processor. On Windows operating systems, issue the db2cmd command from a command prompt. On Linux or UNIX, issue the db2 command from a command prompt.
5. Catalog the database by entering the following commands in the command line processor:

```
db2 => catalog database database_name as database_alias at  
node node_name [ authentication auth_value ]
```

where:

- *database_name* represents the name of the database you want to catalog.
- *database_alias* represents a local nickname for the database you want to catalog.
- *node_name* represents a nickname you can set for the computer that has the database you want to catalog.
- *auth_value* specifies the type of authentication that will take place when connecting to the database. This parameter defaults to the authentication type specified on the server. Specifying an authentication type can result in a performance benefit. Examples of valid values include: SERVER, CLIENT, SERVER_ENCRYPT, KERBEROS, DATA_ENCRYPT, GSSPLUGIN and SERVER_ENCRYPT_AES.

To catalog a remote database called SAMPLE so that it has the local database alias MYSAMPLE, on the node DB2NODE using authentication SERVER, enter the following commands:

```
db2 => catalog database sample as mysample at node db2node  
authentication server  
db2 => terminate
```

Parameter values worksheet for cataloging a database:

Use the following worksheet to record the parameter values required to catalog a database.

Table 11. Catalog database parameter values worksheet

Parameter	Description	Sample Value	Your Value
Database name (<i>database_name</i>)	When a database is created, the database alias is set to the database name unless otherwise specified. For example, when the <code>sample</code> database is created on the server, a database alias of <code>sample</code> is also created. The database name represents the remote database alias (on the server).	<code>sample</code>	
Database alias (<i>database_alias</i>)	An arbitrary local nickname that represents the remote database. If you do not provide one, the default is the same as the database name (<i>database_name</i>). Use this name when you connect to the database from a client.	<code>mysample</code>	
Authentication (<i>auth_value</i>)	The type of authentication required in your environment.	<code>Server</code>	
Node name (<i>node_name</i>)	The name of the node directory entry that describes where the database resides. Use the same value for node name (<i>node_name</i>) that you used to catalog the node.	<code>db2node</code>	

Testing the client-to-server connection using the CLP

After cataloging the node and the database, connect to the database to test the connection. Before testing the connection:

- The database node and database must be cataloged.
- The values for *userid* and *password* must be valid for the system on which they are authenticated. The authentication parameter on the client is set to match the value on the server or it can be left unspecified. If an authentication parameter is not specified, the client will default to `SERVER_ENCRYPT`. If the server does not accept `SERVER_ENCRYPT`, then the client retries using the value returned from the server. If the client specifies an authentication parameter value that doesn't match what is configured on the server, you will receive an error.
- The database manager must be started with the correct protocol defined in the `DB2COMM` registry variable. If it is not started, then you can start the database manager by entering the `db2start` command on the database server.

To test the client to server connection:

1. If you are using a Linux or UNIX platform, set up the instance environment. Run the startup script:

For bash, Bourne or Korn shell

```
. INSTHOME/sql1lib/db2profile
```

For C shell

```
source INSTHOME/sql1lib/db2cshrc
```

where: *INSTHOME* represents the home directory of the instance.

2. Start the DB2 command line processor. On Windows, issue the `db2cmd` command from a command prompt. On Linux or UNIX, issue the `db2` command from a command prompt.
3. Type the following command on the client to connect to the remote database:
`db2 => connect to database_alias user userid`

For example, enter the following command:

```
connect to mysample user jtris
```

You will be prompted to enter your password.

If the connection is successful, you receive a message showing the name of the database to which you have connected. A message similar to the following is given:

```
Database Connection Information
Database server = DB2 9.1.0
SQL authorization ID = JTRIS
Local database alias = mysample
```

You can now work with the database. For example, to retrieve a list of all the table names listed in the system catalog table, enter the following SQL statement:

```
select tabname from syscat.tables
```

When you are finished using the database connection, enter the `connect reset` command to end the database connection.

Chapter 4. IBM data server client deployment in a thin client topology (Windows)

Thin client topology overview (Windows)

This section describes an alternative method for installing an IBM data server client that leverages the Windows support for *thin client topology*. Thin client topologies are only supported in 32-bit environments. You can use this method to install IBM Data Server Client or DB2 Connect Personal Edition on Windows operating systems. This method does not apply to IBM Data Server Runtime Client or IBM Data Server Driver Package.

A thin client topology or *thin client topology environment* consists of one thin client *code server* and one or more *thin clients*. The IBM data server client code is installed on the code server, rather than on each client workstation. On each thin client workstation, only a minimal amount of code and configuration is required. When a thin client initiates a database connection, IBM data server client code is dynamically loaded from the code server as required. The thin client then connects to the database in the usual fashion.

The following figures illustrate the thin client topology. In the first case, Data Server Client is installed on the code server that serves the Data Server Client code to the thin client workstations. These client workstations then connect to one or more DB2 servers.

In the second figure, DB2 Connect Personal Edition is used instead of Data Server Client. DB2 Connect Personal Edition provides the additional capability of enabling clients to connect directly to a DB2 product on midrange and mainframe platforms.

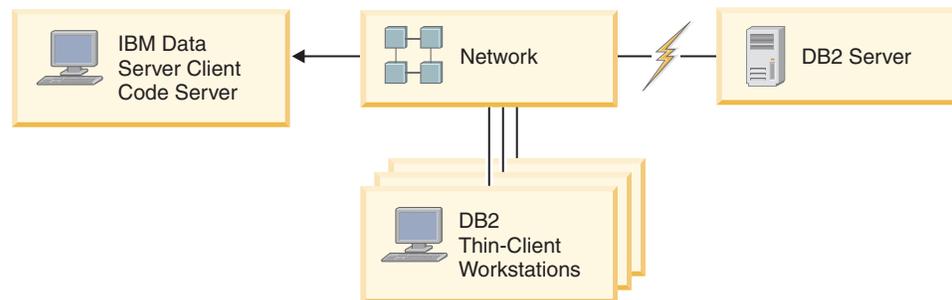


Figure 2. A typical thin client topology using IBM Data Server Client

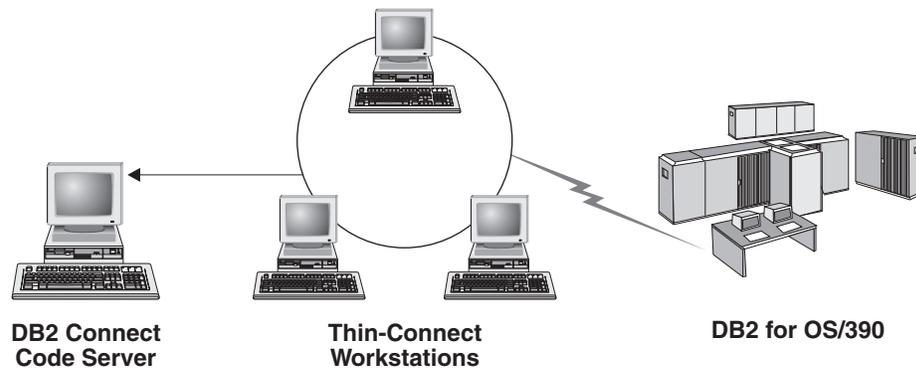


Figure 3. A typical thin client topology using DB2 Connect Personal Edition

Use the thin client method of installing a client when client workstations need only occasional access to a database or when it would be difficult to set up the IBM data server client on each client workstation. If you implement this type of environment, disk space requirements for each workstation are reduced and you can install, update, or upgrade the code on only one machine: that is, the code server.

DB2 programs must be loaded from a code server across a LAN connection. The extent of performance loss at program initialization time depends on variables such as the load on and speed of both the network and the code server.

Note:

- Catalog information must be maintained on each thin client workstation, just as if it were a regular IBM data server client. The catalog files contain all of the information needed for a workstation to connect to a database.
- You can automate the steps to configure database connections for each thin client workstation by using the profile export and import options provided by the Configuration Assistant (CA). After setting up an initial client-to-server connection, export a profile of the configuration settings to all other clients.
- You can avoid the steps to configure database connections for each thin client workstation by using Lightweight Directory Access Protocol (LDAP) in your environment. After you register a database with an LDAP server from a DB2 server, any LDAP-enabled client will retrieve the connection information automatically while connecting.
- The `db2rspgn` command is not supported on the thin client.
- If you are setting up a thin client environment for DB2 Connect Personal Edition, each thin client workstation should have the license for this product.
- If you migrate the DB2 code server from DB2 Version 8 to DB2 Version 9.1 or higher, then you must setup thin clients again after migration.

Thin client setup overview (Windows)

Setting up a thin client environment involves setup on both the code server and each thin client workstation.

To set up a thin client environment:

1. Install an IBM Data Server Client or DB2 Connect Personal Edition on the code server.
2. Make the code directory on the code server available to all thin workstations.
3. Create a thin client response file.

4. Map a network drive from each thin client workstation to the code server.
5. Setup each thin client by running the thnsetup command.

This installation does not include product documentation.

Installing IBM Data Server Client or DB2 Connect Personal Edition on the code server (Windows)

Perform this task to install IBM Data Server Client or DB2 Connect Personal Edition along with product code applicable to the code server. A DB2 thin client workstation can load code only from a DB2 thin client code server, and a DB2 Connect thin client workstation can load code only from a DB2 Connect thin code server. Thin clients are only supported in 32-bit environments.

To install Data Server Client or DB2 Connect Personal Edition on the code server:

1. Locate the appropriate DVD, and start the installation wizard.
2. Select **Custom** installation from the installation wizard.
3. From the Select the Features to Install window, select **Server Support**, and then select **Thin Client Code Server**.
4. Complete the remaining installation wizard steps.

Your next step is to make the code directory on the code server available to all thin workstations.

Making the code directory available to all thin client workstations (Windows)

To load the required code from the code server, each of the target thin client workstations must be able to read the directory where IBM Data Server Client or DB2 Connect Personal Edition source code is installed.

To make the code directory available to all thin client workstations (in read mode) using Windows XP as an example:

1. On the code server, start Windows Explorer.
2. Select the directory on the code server that will be used to serve thin client workstations. For this example, select the d:\sql11ib directory to set up the share.
3. Select **File** → **Properties** from the menu bar.
4. Click the **Sharing** tab.
5. Click the **Shared This Folder** radio button.
6. In the **Share Name** field, enter a share name that is eight characters or fewer. For example, enter NTCODESV.
7. Provide read access to the code directory to all thin client users:
 - a. Click **Permissions**. The Share Permissions window opens.
 - b. In the **Group or Users Name** list, highlight the **Everyone** group.

Note: You can give access to the **Everyone** group, to a group that you have specifically defined for thin client users, or to individual thin client users.

- c. Select **Read**.
- d. Click **OK** until all windows are closed.

Your next step is to create a thin client response file.

Creating a thin client response file (Windows)

Thin clients are only supported in 32-bit environments. A *response file* is used to set up each thin client workstation. A response file is a text file that contains the setup and configuration data to automate an installation. The file consists of a list of keywords and corresponding values. You can create a response file for a thin client installation by editing the sample response file provided with the DB2 database product.

Locate the sample response file `db2thin.rsp` in the `c:\sql11ib\thnsetup` directory, where `c:\sql11ib` represents the location where you installed your thin client code server. In a response file, the asterisk (*) acts like a comment. Any line that is prefixed by an asterisk will be ignored during the installation. To enable a keyword, remove the asterisk. If you do not specify a keyword, or if it is commented out, a default value will be used.

For example, the default entry for the `ODBC_SUPPORT` keyword (used for installing support for ODBC) in the response file is as follows:

```
*COMP =ODBC_SUPPORT
```

To install ODBC, remove the asterisk from the line as shown in this example:

```
COMP =ODBC_SUPPORT
```

For some keywords, you must set values. To enable these keywords, remove the asterisks. However, ensure that you also replace the contents to the right of the equal sign with the value that you want for the keywords.

Following is an example of the entry for `DB2.DIAGLEVEL`:

```
*DB2.DIAGLEVEL = 0 - 4
```

To set this keyword to 4, make the following change:

```
DB2.DIAGLEVEL = 4
```

After you finish editing the response file, save it using a different name to maintain the original sample. For example, call the edited file `test.rsp`, and save it in the same directory in which you set up the shared permissions (for example, `d:\sql11ib`).

You will use this response file in a subsequent step, setting up thin clients using the `thnsetup` command.

Mapping a network drive from each thin client to the code server (Windows)

Each thin client must be mapped to a code server.

You must be logged on to the workstation as a valid user with shared directory access to the code server. You have access to the code server if a locally defined user account was created on the code server.

To map a network drive from the thin client:

1. Launch Windows Explorer.
2. On the **Tools** menu, click **Map Network Drive**.
3. In the **Drive** list, select the drive to which you want to map the location of the code server.

4. In the **Folder** field, specify the location of the share as follows:

`\\computer_name\share_name`

where:

computer_name

represents the computer name of the code server

share_name

represents the share name of the shared directory on the code server

5. Select the **Reconnect at Logon** check box to make the share persistent.

The next step is enabling each thin client.

Setting up thin clients using the thnsetup command (Windows)

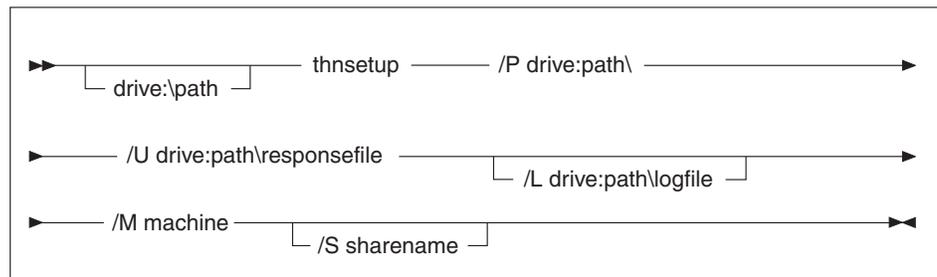
Sets up the thin client workstation and makes the required links to the code server.

Ensure that either Microsoft Visual C++ 2005 or the appropriate runtime components of the Visual C++ libraries are installed before you issue the **thnsetup** command. The runtime library is available from the Microsoft runtime library download Web site. There are two choices: choose `vcredist_x86.exe` for 32-bit systems or `vcredist_x64.exe` for 64-bit systems.

Perform the following steps on each workstation that you want to set up as a thin client.

To set up a thin client:

1. Run the **thnsetup** command. You can specify the following parameters:



where:

- /P** Specifies the path where the DB2 code is installed on the code server. This parameter is required. If you have not already mapped a persistent network drive to the code server. The value of this parameter should be the drive letter used to represent the network drive.
- /U** Specifies the fully qualified response file name. This parameter is required. Normally, the file is located on the code server in the directory `c:\sql\lib\thnsetup`, where `c:\sql\lib\` represents the drive where you installed your thin client code server.
- /L** Specifies the fully qualified log file name, where setup information and any errors occurring during setup will be logged. This parameter is optional. If you do not specify the log file name, the default `db2.log` file name is used. This file will be created in the `db2log` directory, on the drive where your operating system is installed.

- /M** Specifies the name of the code server. This parameter is required.
- /S** Specifies the share name of the code server where you installed the DB2 product. This parameter is necessary only if you did not map a persistent network drive. This parameter is mandatory on Windows XP and Windows Server 2003 operating systems.

2. When the thnsetup command is completed, check the messages in the log file (db2.log in the y:\db2log directory, where y is the drive on which the DB2 code is installed).

Check any error messages. The error messages in the log file depend on the errors that were encountered during the attempted installation. The log file states the reasons for failure.

For example, you might want to create a thin client workstation under the following conditions:

- The shared directory with the share name on a code server is mapped locally to the x drive.
- The response file is called test.rsp.
- The response file is located in the same directory as the code server:

On the thin client workstation, enter the following command from a DOS prompt on the thin workstation:

```
x:\thnsetup\thnsetup /P x: /U x:\thnsetup\test.rsp /M machineName
```

Chapter 5. Merge Modules

Types of merge modules

Non-DB2 instance merge modules (Windows)

Two types of merge modules are available: DB2 instance merge modules and non-DB2 instance merge modules.

It is recommended that you use non-DB2 instance merge modules.

Using non-DB2 instance Windows Installer merge modules, you can easily add IBM Data Server Driver Package functionality to any product that uses the Windows Installer.

When you merge the modules, you will be prompted to supply the copy name. Multiple copies of IBM Data Server Driver Package products can be installed on the same machine; so each copy is known by its unique name. This name will be used when the installation is performed on each target machine. Choose a name that is unlikely to be already used for another IBM data server driver or DB2 copy. Suitable names include the name of your application, for example, myapp_dsdrivercopy_1. If the name is not unique, the installation will fail.

For more information on merge module technology, see the documentation included with your installation authoring product or at <http://msdn.microsoft.com>.

The following merge module is available for your use:

IBM Data Server Driver Package.msm

This module provides support for applications using ODBC, CLI, .NET, OLE DB, PHP, Ruby, JDBC, or SQLJ to access data. It also enables your application to use the IBM Data Server Provider for .NET (DB2 .NET Data Provider and IDS .NET Data Provider). The IBM Data Server Provider .NET is an extension of the ADO.NET interface that enables your .NET applications to quickly and securely access data from DB2 or Informix databases.

DSDRIVER is created by using the merge module described above. Registering IBM Data Server Provider for .NET is based on the version of .NET framework installed on your system. For example, you have to install Microsoft .Net framework 2.0 prior to the installation.

The following merge modules contain language specific messages used by IBM Data Server Driver Package. Depending on the languages of your product, include and install the components in the appropriate merge module.

IBM DSDRIVER Messages - Arabic.msm
IBM DSDRIVER Messages - Bulgarian.msm
IBM DSDRIVER Messages - Chinese(Simplified).msm
IBM DSDRIVER Messages - Chinese(Traditional).msm
IBM DSDRIVER Messages - Croatian.msm
IBM DSDRIVER Messages - Czech.msm

IBM DSDRIVER Messages - Danish.msm
 IBM DSDRIVER Messages - Dutch.msm
 IBM DSDRIVER Messages - English.msm
 IBM DSDRIVER Messages - Finnish.msm
 IBM DSDRIVER Messages - French.msm
 IBM DSDRIVER Messages - German.msm
 IBM DSDRIVER Messages - Greek.msm
 IBM DSDRIVER Messages - Hebrew.msm
 IBM DSDRIVER Messages - Hungarian.msm
 IBM DSDRIVER Messages - Italian.msm
 IBM DSDRIVER Messages - Japanese.msm
 IBM DSDRIVER Messages - Korean.msm
 IBM DSDRIVER Messages - Norwegian.msm
 IBM DSDRIVER Messages - Polish.msm
 IBM DSDRIVER Messages - Portuguese(Brazilian).msm
 IBM DSDRIVER Messages - Portuguese(Standard).msm
 IBM DSDRIVER Messages - Romanian.msm
 IBM DSDRIVER Messages - Russian.msm
 IBM DSDRIVER Messages - Slovak.msm
 IBM DSDRIVER Messages - Slovenian.msm
 IBM DSDRIVER Messages - Spanish.msm
 IBM DSDRIVER Messages - Swedish.msm

DB2 instance merge modules (Windows)

DB2 Version 9.7 supports two types of merge modules: DB2 instance merge modules and non-DB2 instance merge modules. It is recommended that you use non-DB2 instance merge modules.

DB2 instance merge modules require additional overhead and maintenance, but can be used when:

- an application requires a DB2 instance environment, or,
- an application requires function that only exists in a DB2 instance merge module. (The DB2 instance merge modules are listed later in this section.)

Using DB2 instance Windows Installer merge modules, you can easily add IBM Data Server Runtime Client function to any product that uses the Windows Installer.

When you merge the modules, you will be prompted to supply the DB2 copy name. Multiple copies of DB2 products can be installed on the same machine; so each copy is known by its unique name. This name will be used when the installation is performed on each target machine. Choose a name that is unlikely to be already used for another DB2 copy. Suitable names include the name of your application, for example, myapp_db2copy_1. If the name is not unique, the installation will fail.

For more information on merge module technology, see the documentation included with your installation authoring product or at <http://msdn.microsoft.com>.

The following merge modules are available for your use:

DB2 Base Client Merge Module.msm

This module provides the function required for database connection, SQL, and DB2 command function. This module allows you to use the Named

Pipes communication protocol to transfer data in a client-server environment, contains the system bind files which are used for such tasks as creating databases or accessing remote host databases, and, provides various tools used to administer local and remote databases. This module also allows you to create a response file which can be used to configure the DB2 copy within your installation. The configurable option specifies the location of a response file used to configure the DB2 copy. Types of configuration include creating instances, setting database manager configuration parameters or DB2 profile registry variables. You can also use the Command Line Processor (CLP) with this module.

For information about the sequence of various custom actions and the proposed sequences, view the merge module with a tool such as Orca.

DB2 JDBC and SQLJ Support.msm

This module contains JDBC and SQLJ support allowing Java samples to be built and run using the JDBC driver.

DB2 LDAP Exploitation Merge Module.msm

This module allows the DB2 copy to use an LDAP directory to store database directory and configuration information.

DB2 ODBC Support Merge Module.msm

This module provides support for applications that use Open Database Connectivity (ODBC) to access data.

DB2 OLE DB Support Merge Module.msm

This module provides a set of interfaces that allows applications to uniformly access data stored in different data sources.

IBM Data Server Provider for .NET Merge Module.msm

This module enables your application to use the IBM Data Server Provider for .NET. The IBM Data Server Provider for .NET is an extension of the ADO.NET interface that enables your .NET applications to quickly and securely access data from DB2 databases.

The following Microsoft re-distributable merge modules are included with the IBM Data Server Runtime Client merge modules. You must include these Microsoft merge modules when merging Data Server Runtime Client merge modules.

Microsoft NT32:

Microsoft_VC80_CRT_x86.msm

Microsoft_VC80_MFC_x86.msm

policy_8_0_Microsoft_VC80_CRT_x86.msm

policy_8_0_Microsoft_VC80_MFC_x86.msm

Microsoft NT64:

Microsoft_VC80_CRT_x86_x64.msm

Microsoft_VC80_MFC_x86_x64.msm

policy_8_0_Microsoft_VC80_CRT_x86_x64.msm

policy_8_0_Microsoft_VC80_MFC_x86_x64.msm

You can find the Microsoft merge modules on the IBM Data Server Runtime Client DVD under the merge module directory.

The following merge modules contain IBM data server client messages used by the DB2 copy. Depending on the languages of your product, include and install the components in the appropriate merge module.

IBM data server client Messages - Arabic.msm
IBM data server client Messages - Bulgarian.msm
IBM data server client Messages - Chinese(Simplified).msm
IBM data server client Messages - Chinese(Traditional).msm
IBM data server client Messages - Croatian.msm
IBM data server client Messages - Czech.msm
IBM data server client Messages - Danish.msm
IBM data server client Messages - Dutch.msm
IBM data server client Messages - English.msm
IBM data server client Messages - Finnish.msm
IBM data server client Messages - French.msm
IBM data server client Messages - German.msm
IBM data server client Messages - Greek.msm
IBM data server client Messages - Hebrew.msm
IBM data server client Messages - Hungarian.msm
IBM data server client Messages - Italian.msm
IBM data server client Messages - Japanese.msm
IBM data server client Messages - Korean.msm
IBM data server client Messages - Norwegian.msm
IBM data server client Messages - Polish.msm
IBM data server client Messages - Portuguese(Brazilian).msm
IBM data server client Messages - Portuguese(Standard).msm
IBM data server client Messages - Romanian.msm
IBM data server client Messages - Russian.msm
IBM data server client Messages - Slovak.msm
IBM data server client Messages - Slovenian.msm
IBM data server client Messages - Spanish.msm
IBM data server client Messages - Swedish.msm
IBM data server client Messages - Turkish.msm

Chapter 6. Additional installation options

Installation command-line options

IBM Data Server Runtime Client installation command-line options

You can install IBM Data Server Runtime Client using the `db2setup.exe` command on Linux or UNIX operating systems or the `setup.exe` command on Windows operating systems. The parameters of the two commands are different.

The following list describes selected popular standard Windows Installer command line options available when you run `setup.exe` to install IBM Data Server Runtime Client on Windows operating systems. For more information on the available Windows Installer options, see <http://www.msdn.microsoft.com/>.

/w This option forces `setup.exe` to wait until the installation is complete before exiting.

/l*v[*log file name*]

This option allows you to create a log of the installation. You can use the log to troubleshoot any problems that you encountered during the installation. If you specify a directory, this directory must exist or must be created prior installation. Otherwise, the installation returns an error and fails.

/v This option allows you to pass additional command line options and public properties to the Windows Installer. You must specify this option to perform a response file installation.

/qn This option allows you to perform a silent installation with no User Interface (UI), except for a window that the Windows installer displays while it extracts files from the installation package before it starts the actual installation.

/qb! This option displays a basic user interface which shows simple progress and error message handling and hides the **Cancel** button, except in a window that the Windows installer displays while it extracts files from the installation package before it starts the actual installation.

/L This option allow you to change the setup language by specifying the language identifier. For example, to specify French as the setup language, specify the French language identifier, `setup.exe /L 1036` command.

Table 12. Language Identifiers

Language	Identifier
Arabic (Saudi Arabia)	1025
Bulgarian	1026
Chinese (Simplified)	2052
Chinese (Traditional)	1028
Croatian	1050
Czech	1029
Danish	1030

Table 12. Language Identifiers (continued)

Language	Identifier
Dutch (Standard)	1043
English	1033
Finnish	1035
French (Standard)	1036
German	1031
Greek	1032
Hebrew	1037
Hungarian	1038
Italian (Standard)	1040
Japanese	1041
Korean	1042
Norwegian (Bokmal)	1044
Polish	1045
Portuguese (Brazilian)	1046
Portuguese (Standard)	2070
Romanian	1048
Russian	1049
Slovak	1051
Slovenian	1060
Spanish (Traditional Sort)	1034
Swedish	1053
Turkish	1055

Here are the public properties that you can specify to control the installation of Data Server Runtime Client:

- These parameters must be the last parameters in the command line.
- **RSP_FILE_PATH** - This contains the full path to the response file that you use to install Data Server Runtime Client. This is valid only when you specify **/qn**.

To perform a response file installation, the following command line parameter should be used:

```
setup /v"/qn RSP_FILE_PATH=[Full path to the response file]"
```

The example assumes that no copy of the client is already installed. If one or more copies exist, the command is different. To install a second copy using a response file, use the following command:

```
setup /v" TRANSFORMS=:InstanceId1.mst MSINewInstance=1  
/qn RSP_FILE_PATH=[Full path to the response file]"
```

Command line options to install IBM Data Server Driver Package (Windows)

The IBM Data Server Driver Package can be installed by running the DB2 setup command from the command line.

For fix pack images, you can download the driver package that contains the setup program from the IBM Support Fix Central Web site: <http://www.ibm.com/support/fixcentral/>. Data Server client and driver packages are found under the **Information Management** product group and **IBM Data Server Client Packages** product selection. Select the appropriate installed version and platform and click **Continue**. Click **Continue** again on the next screen and you will be presented with a list of all client and driver packages available for Windows.

The following list describes command-line options available for the setup command. For more information about the available Windows Installer options, see <http://www.msdn.microsoft.com/>.

/n [*copy name*]

Specifies the copy name that you want the installation to use. Specifying this option overrides the installation path that is specified in the response file. If the copy exists, a maintenance installation is performed on that copy. Otherwise, a new installation is performed by using the specified copy name.

/o Specifies that a new copy installation with a generated default copy name is to be performed.

/u [*response file*]

Specifies the full path and file name of the response file.

/m Shows the progress dialog during the installation. However, you are not prompted for any input. Use this option with the **/u** option.

/l [*log file*]

Specifies the full path and file name of the log file.

/p [*install-directory*]

Changes the installation path of the product. Specifying this option overrides the installation path that is specified in the response file.

/i *language*

Specifies the two-letter language code of the language in which to perform the installation.

/? Generates usage information.

After installing IBM Data Server Driver Package, you can optionally create and populate the configuration file, `db2dsdriver.cfg`, with database directory information.

The following are some examples of how to use the command-line parameters:

- To install a new copy with a generated default copy name, use the following command:
setup /o
- To install a second copy use the following command:
setup /n "COPY_NAME"
- To perform a response file installation, use the following command:
setup /u "[Full path to the response file]"

A sample response file is located in the `\samples` subdirectory.

Chapter 7. Uninstalling

Uninstalling an IBM data server client

This topic describes how to uninstall an IBM data server client.

Perform one of the following steps to uninstall an IBM data server client:

- To remove an IBM data server client from a Linux or UNIX operating system, run the `db2_deinstall -a` command from the `DB2DIR/install` directory, where `DB2DIR` is the location that you specified when you installed the data server client.
- To remove an IBM data server client from a Windows operating system, use the Add/Remove Programs window, accessible through the Windows Control Panel. Refer to your operating system's help for more information about removing software products from your Windows operating system.

Note: On Windows, the Add/Remove Programs window can be used to uninstall any of the IBM data server clients. If you are removing an IBM Data Server Client, you can run the `db2unins` command. However, this method cannot be used to remove either the IBM Data Server Runtime Client or the IBM Data Server Driver Package. For more information, see the `db2unins` command topic.

Appendix A. Overview of the DB2 technical information

DB2 technical information is available through the following tools and methods:

- DB2 Information Center
 - Topics (Task, concept and reference topics)
 - Help for DB2 tools
 - Sample programs
 - Tutorials
- DB2 books
 - PDF files (downloadable)
 - PDF files (from the DB2 PDF DVD)
 - printed books
- Command line help
 - Command help
 - Message help

Note: The DB2 Information Center topics are updated more frequently than either the PDF or the hardcopy books. To get the most current information, install the documentation updates as they become available, or refer to the DB2 Information Center at ibm.com.

You can access additional DB2 technical information such as technotes, white papers, and IBM Redbooks® publications online at ibm.com. Access the DB2 Information Management software library site at <http://www.ibm.com/software/data/sw-library/>.

Documentation feedback

We value your feedback on the DB2 documentation. If you have suggestions for how to improve the DB2 documentation, send an e-mail to db2docs@ca.ibm.com. The DB2 documentation team reads all of your feedback, but cannot respond to you directly. Provide specific examples wherever possible so that we can better understand your concerns. If you are providing feedback on a specific topic or help file, include the topic title and URL.

Do not use this e-mail address to contact DB2 Customer Support. If you have a DB2 technical issue that the documentation does not resolve, contact your local IBM service center for assistance.

DB2 technical library in hardcopy or PDF format

The following tables describe the DB2 library available from the IBM Publications Center at www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss. English and translated DB2 Version 9.7 manuals in PDF format can be downloaded from www.ibm.com/support/docview.wss?rs=71&uid=swg2700947.

Although the tables identify books available in print, the books might not be available in your country or region.

The form number increases each time a manual is updated. Ensure that you are reading the most recent version of the manuals, as listed below.

Note: The *DB2 Information Center* is updated more frequently than either the PDF or the hard-copy books.

Table 13. DB2 technical information

Name	Form Number	Available in print	Last updated
<i>Administrative API Reference</i>	SC27-2435-02	Yes	September, 2010
<i>Administrative Routines and Views</i>	SC27-2436-02	No	September, 2010
<i>Call Level Interface Guide and Reference, Volume 1</i>	SC27-2437-02	Yes	September, 2010
<i>Call Level Interface Guide and Reference, Volume 2</i>	SC27-2438-02	Yes	September, 2010
<i>Command Reference</i>	SC27-2439-02	Yes	September, 2010
<i>Data Movement Utilities Guide and Reference</i>	SC27-2440-00	Yes	August, 2009
<i>Data Recovery and High Availability Guide and Reference</i>	SC27-2441-02	Yes	September, 2010
<i>Database Administration Concepts and Configuration Reference</i>	SC27-2442-02	Yes	September, 2010
<i>Database Monitoring Guide and Reference</i>	SC27-2458-02	Yes	September, 2010
<i>Database Security Guide</i>	SC27-2443-01	Yes	November, 2009
<i>DB2 Text Search Guide</i>	SC27-2459-02	Yes	September, 2010
<i>Developing ADO.NET and OLE DB Applications</i>	SC27-2444-01	Yes	November, 2009
<i>Developing Embedded SQL Applications</i>	SC27-2445-01	Yes	November, 2009
<i>Developing Java Applications</i>	SC27-2446-02	Yes	September, 2010
<i>Developing Perl, PHP, Python, and Ruby on Rails Applications</i>	SC27-2447-01	No	September, 2010
<i>Developing User-defined Routines (SQL and External)</i>	SC27-2448-01	Yes	November, 2009
<i>Getting Started with Database Application Development</i>	GI11-9410-01	Yes	November, 2009
<i>Getting Started with DB2 Installation and Administration on Linux and Windows</i>	GI11-9411-00	Yes	August, 2009

Table 13. DB2 technical information (continued)

Name	Form Number	Available in print	Last updated
<i>Globalization Guide</i>	SC27-2449-00	Yes	August, 2009
<i>Installing DB2 Servers</i>	GC27-2455-02	Yes	September, 2010
<i>Installing IBM Data Server Clients</i>	GC27-2454-01	No	September, 2010
<i>Message Reference Volume 1</i>	SC27-2450-00	No	August, 2009
<i>Message Reference Volume 2</i>	SC27-2451-00	No	August, 2009
<i>Net Search Extender Administration and User's Guide</i>	SC27-2469-02	No	September, 2010
<i>Partitioning and Clustering Guide</i>	SC27-2453-01	Yes	November, 2009
<i>pureXML Guide</i>	SC27-2465-01	Yes	November, 2009
<i>Query Patroller Administration and User's Guide</i>	SC27-2467-00	No	August, 2009
<i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i>	SC27-2468-01	No	September, 2010
<i>SQL Procedural Languages: Application Enablement and Support</i>	SC27-2470-02	Yes	September, 2010
<i>SQL Reference, Volume 1</i>	SC27-2456-02	Yes	September, 2010
<i>SQL Reference, Volume 2</i>	SC27-2457-02	Yes	September, 2010
<i>Troubleshooting and Tuning Database Performance</i>	SC27-2461-02	Yes	September, 2010
<i>Upgrading to DB2 Version 9.7</i>	SC27-2452-02	Yes	September, 2010
<i>Visual Explain Tutorial</i>	SC27-2462-00	No	August, 2009
<i>What's New for DB2 Version 9.7</i>	SC27-2463-02	Yes	September, 2010
<i>Workload Manager Guide and Reference</i>	SC27-2464-02	Yes	September, 2010
<i>XQuery Reference</i>	SC27-2466-01	No	November, 2009

Table 14. DB2 Connect-specific technical information

Name	Form Number	Available in print	Last updated
<i>Installing and Configuring DB2 Connect Personal Edition</i>	SC27-2432-02	Yes	September, 2010
<i>Installing and Configuring DB2 Connect Servers</i>	SC27-2433-02	Yes	September, 2010

Table 14. DB2 Connect-specific technical information (continued)

Name	Form Number	Available in print	Last updated
DB2 Connect User's Guide	SC27-2434-02	Yes	September, 2010

Table 15. Information Integration technical information

Name	Form Number	Available in print	Last updated
Information Integration: Administration Guide for Federated Systems	SC19-1020-02	Yes	August, 2009
Information Integration: ASNCLP Program Reference for Replication and Event Publishing	SC19-1018-04	Yes	August, 2009
Information Integration: Configuration Guide for Federated Data Sources	SC19-1034-02	No	August, 2009
Information Integration: SQL Replication Guide and Reference	SC19-1030-02	Yes	August, 2009
Information Integration: Introduction to Replication and Event Publishing	GC19-1028-02	Yes	August, 2009

Ordering printed DB2 books

If you require printed DB2 books, you can buy them online in many but not all countries or regions. You can always order printed DB2 books from your local IBM representative. Keep in mind that some softcopy books on the *DB2 PDF Documentation DVD* are unavailable in print. For example, neither volume of the *DB2 Message Reference* is available as a printed book.

Printed versions of many of the DB2 books available on the *DB2 PDF Documentation DVD* can be ordered for a fee from IBM. Depending on where you are placing your order from, you may be able to order books online, from the IBM Publications Center. If online ordering is not available in your country or region, you can always order printed DB2 books from your local IBM representative. Note that not all books on the *DB2 PDF Documentation DVD* are available in print.

Note: The most up-to-date and complete DB2 documentation is maintained in the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7>.

To order printed DB2 books:

- To find out whether you can order printed DB2 books online in your country or region, check the IBM Publications Center at <http://www.ibm.com/shop/publications/order>. You must select a country, region, or language to access publication ordering information and then follow the ordering instructions for your location.
- To order printed DB2 books from your local IBM representative:

1. Locate the contact information for your local representative from one of the following Web sites:
 - The IBM directory of world wide contacts at www.ibm.com/planetwide
 - The IBM Publications Web site at <http://www.ibm.com/shop/publications/order>. You will need to select your country, region, or language to the access appropriate publications home page for your location. From this page, follow the "About this site" link.
2. When you call, specify that you want to order a DB2 publication.
3. Provide your representative with the titles and form numbers of the books that you want to order. For titles and form numbers, see “DB2 technical library in hardcopy or PDF format” on page 83.

Displaying SQL state help from the command line processor

DB2 products return an SQLSTATE value for conditions that can be the result of an SQL statement. SQLSTATE help explains the meanings of SQL states and SQL state class codes.

To start SQL state help, open the command line processor and enter:

```
? sqlstate or ? class code
```

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, ? 08003 displays help for the 08003 SQL state, and ? 08 displays help for the 08 class code.

Accessing different versions of the DB2 Information Center

For DB2 Version 9.8 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r8/>.

For DB2 Version 9.7 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/>.

For DB2 Version 9.5 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/>.

For DB2 Version 9.1 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>.

For DB2 Version 8 topics, go to the *DB2 Information Center* URL at: <http://publib.boulder.ibm.com/infocenter/db2luw/v8/>.

Displaying topics in your preferred language in the DB2 Information Center

The DB2 Information Center attempts to display topics in the language specified in your browser preferences. If a topic has not been translated into your preferred language, the DB2 Information Center displays the topic in English.

- To display topics in your preferred language in the Internet Explorer browser:
 1. In Internet Explorer, click the **Tools** —> **Internet Options** —> **Languages...** button. The Language Preferences window opens.

2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button.

Note: Adding a language does not guarantee that the computer has the fonts required to display the topics in the preferred language.
 - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
 3. Refresh the page to display the DB2 Information Center in your preferred language.
- To display topics in your preferred language in a Firefox or Mozilla browser:
 1. Select the button in the **Languages** section of the **Tools** → **Options** → **Advanced** dialog. The Languages panel is displayed in the Preferences window.
 2. Ensure your preferred language is specified as the first entry in the list of languages.
 - To add a new language to the list, click the **Add...** button to select a language from the Add Languages window.
 - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
 3. Refresh the page to display the DB2 Information Center in your preferred language.

On some browser and operating system combinations, you must also change the regional settings of your operating system to the locale and language of your choice.

Updating the DB2 Information Center installed on your computer or intranet server

A locally installed DB2 Information Center must be updated periodically.

A DB2 Version 9.7 Information Center must already be installed. For details, see the “Installing the DB2 Information Center using the DB2 Setup wizard” topic in *Installing DB2 Servers*. All prerequisites and restrictions that applied to installing the Information Center also apply to updating the Information Center.

An existing DB2 Information Center can be updated automatically or manually:

- Automatic updates - updates existing Information Center features and languages. An additional benefit of automatic updates is that the Information Center is unavailable for a minimal period of time during the update. In addition, automatic updates can be set to run as part of other batch jobs that run periodically.
- Manual updates - should be used when you want to add features or languages during the update process. For example, a local Information Center was originally installed with both English and French languages, and now you want to also install the German language; a manual update will install German, as well as, update the existing Information Center features and languages. However, a manual update requires you to manually stop, update, and restart the Information Center. The Information Center is unavailable during the entire update process.

This topic details the process for automatic updates. For manual update instructions, see the “Manually updating the DB2 Information Center installed on your computer or intranet server” topic.

To automatically update the DB2 Information Center installed on your computer or intranet server:

1. On Linux operating systems,
 - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `/opt/ibm/db2ic/V9.7` directory.
 - b. Navigate from the installation directory to the `doc/bin` directory.
 - c. Run the `ic-update` script:

```
ic-update
```
2. On Windows operating systems,
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `<Program Files>\IBM\DB2 Information Center\Version 9.7` directory, where `<Program Files>` represents the location of the Program Files directory.
 - c. Navigate from the installation directory to the `doc\bin` directory.
 - d. Run the `ic-update.bat` file:

```
ic-update.bat
```

The DB2 Information Center restarts automatically. If updates were available, the Information Center displays the new and updated topics. If Information Center updates were not available, a message is added to the log. The log file is located in `doc\eclipse\configuration` directory. The log file name is a randomly generated number. For example, `1239053440785.log`.

Manually updating the DB2 Information Center installed on your computer or intranet server

If you have installed the DB2 Information Center locally, you can obtain and install documentation updates from IBM.

Updating your locally-installed *DB2 Information Center* manually requires that you:

1. Stop the *DB2 Information Center* on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to apply updates. The Workstation version of the DB2 Information Center always runs in stand-alone mode. .
2. Use the Update feature to see what updates are available. If there are updates that you must install, you can use the Update feature to obtain and install them

Note: If your environment requires installing the *DB2 Information Center* updates on a machine that is not connected to the internet, mirror the update site to a local file system using a machine that is connected to the internet and has the *DB2 Information Center* installed. If many users on your network will be installing the documentation updates, you can reduce the time required for individuals to perform the updates by also mirroring the update site locally and creating a proxy for the update site.

If update packages are available, use the Update feature to get the packages. However, the Update feature is only available in stand-alone mode.

3. Stop the stand-alone Information Center, and restart the *DB2 Information Center* on your computer.

Note: On Windows 2008, Windows Vista (and higher), the commands listed later in this section must be run as an administrator. To open a command prompt or graphical tool with full administrator privileges, right-click the shortcut and then select **Run as administrator**.

To update the *DB2 Information Center* installed on your computer or intranet server:

1. Stop the *DB2 Information Center*.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:
`/etc/init.d/db2icdv97 stop`
2. Start the Information Center in stand-alone mode.
 - On Windows:
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the `Program_Files\IBM\DB2 Information Center\Version 9.7` directory, where `Program_Files` represents the location of the Program Files directory.
 - c. Navigate from the installation directory to the `doc\bin` directory.
 - d. Run the `help_start.bat` file:
`help_start.bat`
 - On Linux:
 - a. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the `/opt/ibm/db2ic/V9.7` directory.
 - b. Navigate from the installation directory to the `doc/bin` directory.
 - c. Run the `help_start` script:
`help_start`

The systems default Web browser opens to display the stand-alone Information Center.

3. Click the **Update** button (🔄). (JavaScript™ must be enabled in your browser.) On the right panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check the selections you want to install, then click **Install Updates**.
5. After the installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center:
 - On Windows, navigate to the installation directory's `doc\bin` directory, and run the `help_end.bat` file:
`help_end.bat`

Note: The `help_end` batch file contains the commands required to safely stop the processes that were started with the `help_start` batch file. Do not use `Ctrl-C` or any other method to stop `help_start.bat`.

- On Linux, navigate to the installation directory's `doc/bin` directory, and run the `help_end` script:
`help_end`

Note: The `help_end` script contains the commands required to safely stop the processes that were started with the `help_start` script. Do not use any other method to stop the `help_start` script.

7. Restart the *DB2 Information Center*.

- On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click **DB2 Information Center** service and select **Start**.
- On Linux, enter the following command:
`/etc/init.d/db2icdv97 start`

The updated *DB2 Information Center* displays the new and updated topics.

DB2 tutorials

The DB2 tutorials help you learn about various aspects of DB2 products. Lessons provide step-by-step instructions.

Before you begin

You can view the XHTML version of the tutorial from the Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>.

Some lessons use sample data or code. See the tutorial for a description of any prerequisites for its specific tasks.

DB2 tutorials

To view the tutorial, click the title.

“pureXML®” in *pureXML Guide*

Set up a DB2 database to store XML data and to perform basic operations with the native XML data store.

“Visual Explain” in *Visual Explain Tutorial*

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2 database products.

DB2 documentation

Troubleshooting information can be found in the *Troubleshooting and Tuning Database Performance* or the Database fundamentals section of the *DB2 Information Center*. There you will find information about how to isolate and identify problems using DB2 diagnostic tools and utilities, solutions to some of the most common problems, and other advice on how to solve problems you might encounter with your DB2 database products.

DB2 Technical Support Web site

Refer to the DB2 Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes,

Authorized Program Analysis Reports (APARs or bug fixes), fix packs, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the DB2 Technical Support Web site at http://www.ibm.com/software/data/db2/support/db2_9/

Terms and Conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Appendix B. Notices

This information was developed for products and services offered in the U.S.A. Information about non-IBM products is based on information available at the time of first publication of this document and is subject to change.

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 character set (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/region 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.

Licenses 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 that has been exchanged, should contact:

IBM Canada Limited
U59/3600
3600 Steeles Avenue East
Markham, Ontario L3R 9Z7
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 may contain 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.

Trademarks

IBM, the IBM logo, and `ibm.com`[®] 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.

The following terms are trademarks or registered trademarks of other companies

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Intel, Intel logo, Intel Inside[®], Intel Inside logo, Intel[®] Centrino[®], Intel Centrino logo, Celeron[®], Intel[®] Xeon[®], Intel SpeedStep[®], Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- Microsoft, Windows, Windows NT[®], and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Index

A

- AIX
 - installing
 - DB2 server products 12
 - IBM data server clients 12

B

- books
 - ordering 86

C

- cataloging
 - database parameter values worksheet 64
 - databases 63
 - host databases 63
 - Named Pipes 59
 - TCP/IP nodes 62
- client profiles
 - configuring using the import function 57
 - creating using the export function 56
- client-to-server communication
 - connections
 - configuring 49
 - testing using CLP 65
 - TCP/IP parameter values worksheet 59
- clients
 - configurations
 - non-supported 51
 - supported 51
 - server connections
 - configuring using CLP 58
 - configuring using Configuration Assistant 53
- code directory
 - thin clients 69
- code servers
 - installing
 - DB2 Connect Personal Edition 69
 - IBM Data Server Client 69
 - mapping network drives
 - remote client 44
 - thin client 70
- command line processor (CLP)
 - cataloging
 - databases 63
 - nodes 62
 - configuring
 - client-to-server connections 58
 - TCP/IP 60
- command-line options
 - IBM Data Server Driver Package installation 79
 - IBM Data Server Runtime Client installation 77
- commands
 - catalog database 63
 - catalog npipe 59
 - catalog tcpip 62
 - db2dsdcfgfill 39
 - db2dsdpreg 45

- commands (*continued*)
 - db2osconf
 - determining kernel configuration parameter values 15
 - db2rfe
 - enabling root-based features 33
 - overcoming limitations on non-root installations 30
 - db2setup
 - installing data server clients 27
 - thnsetup 71
- communication protocols
 - overview 53
- configuration
 - client-to-server connections
 - command line processor (CLP) 58
 - Configuration Assistant (CA) 53
 - TCP/IP worksheet 59
 - TCP/IP
 - client 60
- Configuration Assistant (CA)
 - cataloging databases 49
 - client profiles 56
 - configuring
 - client profiles 57
 - client-to-server communications 49
 - client-to-server connections 53
 - database connections 54
 - Discovery feature 55
 - LDAP 58
 - testing database connections 57
- create configuration file command 39

D

- databases
 - adding manually 54
 - cataloging
 - command line processor (CLP) 63
 - connections
 - configuring 54, 55
 - testing 57
- DB2 Connect
 - installing
 - prerequisites 24
 - Personal Edition
 - installing (Windows) 69
 - thin client
 - code directory 69
 - installing 68
 - mapping network drive to code server 70
 - response files 70
 - setup 67
 - topology overview 67
- DB2 Information Center
 - languages 87
 - updating 88, 89
 - versions 87
- db2dsdcfgfill command
 - copying database directory information 41
 - details 39
- db2dsdpreg command 45
- db2dsdriver.cfg file 41

- db2osconf command
 - determining kernel configuration parameter values 15
- db2rfe command
 - enabling root features 30, 33
- directory structures
 - root installations compared to non-root installations 29
- discovery feature
 - configuring database connection 55
- disk space
 - requirements 11
- documentation
 - overview 83
 - PDF files 83
 - printed 83
 - terms and conditions of use 92
- dropping
 - non-root instances 35

E

- examples
 - connecting to a remote database 65

F

- fix packs
 - non-root installations 35

H

- hardware
 - requirements
 - DB2 server products (AIX) 12
 - DB2 server products (HP-UX) 13
 - DB2 server products (Linux) 15
 - DB2 server products (Solaris Operating Environment) 20
 - IBM data server clients (AIX) 12
 - IBM data server clients (HP-UX) 13
 - IBM data server clients (Linux) 15
 - IBM data server clients (Solaris Operating Environment) 20
 - Windows 22

- help
 - configuring language 87
 - SQL statements 87

- host databases
 - client connections 24

- HP-UX
 - installing
 - DB2 servers 13
 - IBM data server clients 13
 - kernel configuration parameters
 - modifying 15
 - recommended values 15

I

- IBM data server clients
 - cataloging
 - Named Pipes nodes 59
 - TCP/IP nodes 62
 - connecting to 24
 - IBM Data Server Client 1, 2
 - IBM Data Server Driver Package 1

- IBM data server clients (*continued*)
 - IBM Data Server Runtime Client 1, 2
 - installing
 - code servers 69
 - overview 5, 6
 - UNIX 27
 - Windows 24, 37
 - overview 1
 - types 2
 - user accounts 24

- IBM Data Server Driver Package
 - configuration file 41
 - installing
 - command-line options 79
 - Linux 46
 - network 43
 - UNIX 46
 - Windows 42
 - network share 44
 - remote client
 - setting up 43
 - topology overview 43
 - restrictions 36

- IBM data server drivers
 - types 2

- IBM Data Server Runtime Client
 - installing 77

- import function
 - configuring client profiles 57

- installation
 - DB2 products
 - as non-root user 32
 - requirements
 - AIX 12
 - HP-UX 13
 - Linux 15
 - Solaris Operating Environment 20

- installing
 - requirements
 - Windows 22

- instances
 - removing 35

K

- kernel configuration parameters
 - HP-UX
 - db2osconf command 15
 - modifying 15
 - recommended 15
 - modifying kernel parameters on Linux 19
 - Solaris Operating System 21

L

- Lightweight Directory Access Protocol (LDAP)
 - directory support 58

- Linux
 - installing
 - DB2 servers 15
 - IBM data server clients 15
 - IBM Data Server Driver Package 46
 - libraries
 - libaio.so.1 15
 - libstdc so.5 15
 - modifying kernel parameters 19

Linux (*continued*)
 uninstalling DB2
 non-root instances 35

M

mapping
 network drives from thin clients to code server 70
memory
 requirements
 overview 11
merge modules
 DB2 instance 74
 non-DB2 instance 73
modifying kernel parameters
 Linux 19

N

Named Pipes
 parameter values worksheet 58
 supported protocol 53
network drives
 mapping remote clients to code server 44
 mapping thin clients to code server 70
network share
 IBM Data Server Driver Package 44
non-root installations
 differences 29
 directory structure 29
 enabling root-based features 33
 fix packs 35
 installing 32
 limitations 30
 overview 28
 uninstalling 35
non-root instances
 dropping 35
 removing 35
notices 93

O

operating system requirements
 Windows 22
operating systems
 installation requirements
 AIX 12
 HP-UX 13
 Linux 15
 Solaris Operating Environment 20
ordering DB2 books 86

P

parameters
 values worksheets
 cataloging databases 64
 configuring client-to-server connections 59
 Named Pipes 58
problem determination
 information available 91
 tutorials 91

R

remote clients
 IBM Data Server Driver Package 43, 45
 mapping network drive to code server 44
response files
 creating
 thin client 70
root installations
 comparison with non-root installations 29
 directory structure 29
root-based features
 non-root installations 33

S

servers
 client connections 53, 58
software requirements
 AIX 12
 HP-UX 13
 Linux 15
 Solaris Operating Environment 20
 Windows 22
Solaris Operating Environment
 installation requirements
 DB2 servers 20
 IBM data server clients 20
Solaris Operating System
 modifying kernel parameters 21
SQL statements
 help
 displaying 87
SSL
 support 53

T

TCP/IP
 configuring
 clients 60
 platforms supported 53
 TCP/IP6 support 53
terms and conditions
 publications 92
testing
 client-to-server connections 65
 database connections 57
thin clients
 code directory 69
 enabling 71
 installing 68
 mapping network drives from clients to code server 70
 response files 70
 topology 67
 typical setup 67
thnsetup command 71
troubleshooting
 online information 91
 tutorials 91
tutorials
 list 91
 problem determination 91
 troubleshooting 91
 Visual Explain 91

U

uninstalling

- IBM data server clients 81
- non-root installations 35

UNIX

installing

- IBM data server clients 27
- IBM Data Server Driver Package 46

removing

- DB2 non-root instances 35

updates

- DB2 Information Center 88, 89

user accounts

- IBM data server clients 24

W

Windows

installing

- IBM data server clients (procedure) 24, 37
- IBM Data Server Driver Package 42

Windows operating systems

installing

- DB2 servers (requirements) 22
- IBM data server clients (requirements) 22



Printed in USA

GC27-2454-01



Spine information:

IBM DB2 9.7 for Linux, UNIX, and Windows **Version 9 Release 7**

Installing IBM Data Server Clients

