

DB2 Version 9.5
for Linux, UNIX, and Windows



Quick Beginnings for IBM Data Server Clients
Updated March, 2008

DB2 Version 9.5
for Linux, UNIX, and Windows



Quick Beginnings for IBM Data Server Clients
Updated March, 2008

Note

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 95.

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 International Business Machines Corporation 1993, 2008. All rights reserved.

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

Contents

About this publication v

Part 1. IBM data server clients 1

Chapter 1. Introduction to IBM data server clients 3

IBM data server clients setup overview 3
IBM data server client types 4
Installation methods for IBM data server clients . . . 6
Options for connecting to DB2 databases 7

Part 2. Installing IBM data server clients 11

Chapter 2. IBM data server client installation requirements 13

Disk and memory requirements 13
Installation requirements for DB2 servers and IBM data server clients (AIX) 13
Installation requirements for DB2 servers and IBM data server clients (HP-UX) 15
 Recommended kernel configuration parameters (HP-UX) 16
 Modifying kernel parameters (HP-UX) 16
Installation requirements for DB2 servers and IBM data server clients (Linux) 19
 Modifying kernel parameters (Linux) 21
Installation requirements for DB2 servers and IBM data server clients (Solaris Operating Environment) . 22
 Modifying kernel parameters (Solaris Operating System) 24
Installation requirements for DB2 servers and IBM data server clients (Windows) 25
DB2 Connect product installation requirements for host and midrange systems 26

Chapter 3. Installing IBM data server clients 27

Installing IBM data server clients (Windows) . . . 27
Installing IBM data server clients (Linux and UNIX) 30
Non-root installation overview (Linux and UNIX) 31
 Differences between root installations and non-root installations 32
 Limitations of non-root installations 32
 Installing a DB2 product as a non-root user. . . 35
 Enabling root-based features in non-root installations with db2rfe 36
 Applying fix packs to a non-root installation . . 37
 Removing non-root DB2 products using db2_deinstall (Linux and UNIX) 38

Part 3. Database connections for IBM data server clients 39

Chapter 4. Client-to-server communications configuration overview 41

Supported combinations of client and server versions 43
Communication protocols supported 44
Adding database connections using the Configuration Assistant 44
 Configuring client-to-server connections using the Configuration Assistant (CA) 44
 Configuring a database connection manually using the Configuration Assistant 45
 Configuring a database connection by searching the network using the Configuration Assistant. . 46
 Creating a client profile using the Configuration Assistant 47
 Configuring database connections using a client profile with the Configuration Assistant 48
 Testing a database connection using the Configuration Assistant 49
 LDAP considerations for the Configuration Assistant 49
Configuring client-to-server connections using the command line processor 49
 Configuring client-to-server connections using the command line processor. 49
 Named pipe connections 50
 TCP/IP connections 51
 Cataloging a database from a client using the CLP 54
 Testing the client-to-server connection using the CLP 56

Part 4. IBM data server client deployment in a thin client topology (Windows) 59

Chapter 5. Thin client topology overview (Windows) 61

Thin client setup overview (Windows) 62
Installing IBM Data Server Client or DB2 Connect Personal Edition on the code server (Windows) . . 62
Making the code directory available to all thin client workstations (Windows) 63
Creating a thin client response file (Windows). . . 63
Mapping a network drive from each thin client to the code server (Windows) 64
Setting up thin clients using the thnsetup command (Windows). 65

Part 5. Merge Modules	67
Chapter 6. Types of merge modules	69
Non-DB2 instance merge modules (Windows)	69
DB2 instance merge modules (Windows)	70
Part 6. Additional installation options	73
Chapter 7. Installation command-line options	75
IBM Data Server Runtime Client installation command-line options	75
IBM Data Server Driver for ODBC, CLI, and .NET installation command line options (Windows)	76
Part 7. Uninstalling	79
Chapter 8. Uninstalling an IBM data server client	81

Part 8. Appendixes.	83
Appendix A. Overview of the DB2 technical information	85
DB2 technical library in hardcopy or PDF format	85
Ordering printed DB2 books.	88
Displaying SQL state help from the command line processor	88
Accessing different versions of the DB2 Information Center	89
Displaying topics in your preferred language in the DB2 Information Center	89
Updating the DB2 Information Center installed on your computer or intranet server	90
DB2 tutorials	91
DB2 troubleshooting information	92
Terms and Conditions	92
Appendix B. Notices	95
Index	99

About this publication

If you are interested in installing and configuring IBM® Data Server Client, IBM Data Server Runtime Client, or IBM Data Server Driver for ODBC, CLI, and .NET or setting up a thin client or DB2 Connect™ thin client environment.

Part 1. IBM data server clients

Chapter 1. Introduction to IBM data server clients

IBM data server clients setup overview

This topic outlines information available about clients 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 other means to enable connections between your system and remote DB2[®] databases.
2. Choose the most suitable method for installing your client.
3. Complete the steps and address the considerations needed to set up a client.

Connection options

Options for connecting a system to a remote DB2 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 clients types

DB2 products support the following IBM data server clients:

- IBM Data Server Client
- IBM Data Server Runtime Client
- IBM Data Server Driver for ODBC, CLI, and .NET

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

See the related links for details regarding types of IBM data server clients.

Installation methods

The common method for installing Data Server Client or Data Server Runtime Client is to run the installation program provided on a product DVD. The common method for installing Data Server Driver for ODBC, CLI, and .NET is to download the **setup.exe** command from https://www14.software.ibm.com/webapp/iwm/web/pick.do?lang=en_US&source;=swg-datasc and to run the **setup.exe** command.

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 Data Server Driver for ODBC, CLI, and .NET in your application.

Setting up a client

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.
3. Catalog databases and configure connections to remote servers (not required for Data Server Driver for ODBC, CLI, and .NET).

For systems where a DB2 Universal Database™ (UDB) Version 8 client or a DB2 Version 9 client already exists, consider whether to migrate the existing client to a Version 9.5 Data Server Client, or, leave the DB2 UDB Version 8 client and Version 9 client and install the Version 9.5 Data Server Client as an additional client.

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

IBM data server client types

The following types of IBM data server clients are available:

- IBM Data Server Client
- IBM Data Server Runtime Client
- IBM Data Server Driver for ODBC, CLI, and .NET

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

- Use IBM Data Server 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 Runtime Client if you need command line processor (CLP) support and basic client support for running and deploying applications.
- Use IBM Data Server Driver for ODBC, CLI, and .NET if you need runtime support for DB2 CLI API, ODBC API, and .NET API for Windows applications. This client is also a lightweight solution for deploying Windows applications.

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
- IBM Data Studio
- Application header files
- Precompilers for various programming languages
- Bind support
- Samples and tutorials
- IBM Informix[®] Dynamic Server support for PHP, Ruby, .NET, and JDBC

IBM Data Server Runtime Client

The IBM Data Server Runtime Client provides a way to run applications on remote DB2 databases. GUI tools are not shipped with the IBM Data Server Runtime Client.

Capabilities include:

- The command line processor (CLP) for issuing DB2 commands. The CLP also provides a basic way to perform remote administration of DB2 servers.
- Base client support to handle database connections, SQL statements, XQuery statements, and DB2 commands.
- Support for common database access interfaces: JDBC, ADO.NET, OLE DB, ODBC, DB2 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 DB2 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 DB2 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 Driver for ODBC, CLI, and .NET

IBM Data Server Driver for ODBC, CLI, and .NET is a lightweight deployment solution for Windows applications. It provides runtime support for applications using DB2 CLI API, ODBC API, or .NET API without need of installing the Data Server Client or the Data Server Runtime Client.

Capabilities include the following:

- Support for applications that use CLI, ODBC, .NET, PHP, or Ruby to access DB2 databases.
- Packaging advantages on Windows operating systems: The client is available as an installable image. Also, merge modules are available that allow you to easily embed the client in a Windows Installer-based installation
- On Linux and UNIX[®] operating systems, a separate deliverable called IBM Data Server Driver for ODBC and CLI provides a similar lightweight deployment solution for Linux and UNIX applications. The differences are that there is no .NET support and that this driver is available only as a tar file, not as an installable image.
- IBM Informix Dynamic Server support for .NET, PHP, and Ruby

Installation methods for IBM data server clients

This section outlines common and alternative methods for installing IBM Data Server Client, IBM Data Server Runtime Client, and IBM Data Server Driver for ODBC, CLI, and .NET.

Clients are commonly installed on machines where there is no DB2 server present. You do not need to install a client if you already installed a DB2 server product is already installed because the DB2 server includes all the functionality present in an IBM data server client.

Common situations

The common method for installing an IBM data server client 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 DB2 server installation image.

Automation of large-scale deployments

One group of methods is designed to automate the deployment of large numbers of clients:

- **Response file.** You can automate client installation by using the response file installation method. A DB2 response file installation lets you install DB2 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 DB2 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 DB2 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:

- **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 for ODBC, CLI, and .NET 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 for ODBC, CLI, and .NET.
- **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 and CLI applications. Drivers can be obtained by copying driver files from a Data Server Driver for ODBC, CLI, and .NET installation image or by downloading the driver files from developerWorks®.

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 Rational® Application Developer or Microsoft® Visual Studio.
- A DB2-specific development tool related to the IDE such as:
 - IBM Database Developer Add-ins for Visual Studio .NET
 - IBM Data 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:

- **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.

Part 2. Installing IBM data server clients

Chapter 2. 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, as well as by 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 operating system, 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.

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

Table 1. AIX installation requirements

Operating System	Hardware
AIX Version 5.3 <ul style="list-style-type: none"> • 64-bit AIX kernel is required • AIX 5.3 Technology Level (TL) 6 and Service Pack (SP) 2 plus APAR IZ03063 • Minimum C++ runtime level is xLC.rte 9.0.0.1 and xLC.aix50.rte 9.0.0.1 AIX Version 6.1 ² <ul style="list-style-type: none"> • 64-bit AIX kernel is required 	64-bit Common Hardware Reference Platform (CHRP) architecture ¹ All processors that are capable of running the supported AIX operating systems.

- ¹To verify that it is a CHRP architecture system, issue the command `lscfg` and look for the following output: `Model Architecture: chrp`
- 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. This feature is not supported if you are using multiple partition instances.

Software considerations

- (Clients only) If you plan to use Kerberos Authentication, you require IBM Network Authentication Service client v1.4 or later. The NAS client can be downloaded from <https://www6.software.ibm.com/dl/dm/dm-nas-p>.
- Use the `bosboot` command to switch to the 64-bit kernel.
To switch to a 64-bit kernel, you require root authority and should 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):
 - Mozilla 1.4 and up
 - Firefox 1.0 and up
 - Netscape 7.0 and up
- 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

DB2 product installation on NFS (Network File System)

The installation of DB2 products on NFS (Network File System) is not recommended. Running DB2 products on NFS (for example, NFS mounting `/opt/IBM/db2/V9.5` and then running off code that was physically installed on a remote system) requires several manual setup steps. There are also a number of potential issues with setting up NFS for a DB2 server. These include possible problems that involve:

- Performance (impacted by network performance)

- Availability (you are allowing a single point of failure)
- Licensing (there is no checking done across machines)
- Diagnosing NFS errors can be difficult

As mentioned, the setup for NFS will require several manual actions including:

- Ensuring that the mount point preserve the install path
- Permission must be controlled (for example, write permission should not be given to the mounting machine)
- DB2 registries have to be set up manually and maintained across all mounting machines
- The db2ls command, which lists installed DB2 products and features, must be set up and maintained properly if you need to detect DB2 products and features
- More care is required when updating your DB2 product environment
- More steps are required when cleaning up on the exporting machine and the mounting machine

For detailed instructions, see the "Setting up DB2 for UNIX and Linux on NFS mounted file systems" white paper in <http://www.ibm.com/developerworks/db2/library/long/dm-0609lee>.

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

To install a DB2 product, the following operating system, hardware, and communications requirements must be met:

Table 2. HP-UX installation requirements

Operating System	Hardware
DB2 products are supported on: <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 • HP-UX 11iv3 (11.31) 	Itanium® based HP Integrity Series Systems

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, you might need to 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):
 - Mozilla 1.4 and up
 - Firefox 1.0 and up

- Netscape 7.0 and up
- 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 HP-UX issues, see www.ibm.com/support/docview.wss?&uid=swg21257602

DB2 product installation on NFS (Network File System)

The installation of DB2 products on NFS (Network File System) is not recommended. Running DB2 products on NFS (for example, NFS mounting /opt/IBM/db2/V9.5 and then running off code that was physically installed on a remote system) requires several manual setup steps. There are also a number of potential issues with setting up NFS for DB2. These include possible problems that involve:

- Performance (impacted by network performance)
- Availability (you are allowing a single point of failure)
- Licensing (there is no checking done across machines)
- Diagnosing NFS errors can be difficult

As mentioned, the setup for NFS will require several manual actions including:

- Ensuring that the mount point preserve the install path
- Permission must be controlled (for example, write permission should not be given to the mounting machine)
- DB2 registries have to be set up manually and maintained across all mounting machines
- The db2ls command, which lists installed DB2 products and features, must be set up and maintained properly if you need to detect DB2 products and features
- More care is required when updating your DB2 product environment
- More steps are required when cleaning up on the exporting machine and the mounting machine

For detailed instructions, see the “Setting up DB2 for UNIX and Linux on NFS mounted file systems” white paper in <http://www.ibm.com/developerworks/db2/library/long/dm-0609lee>.

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 may need to update your system’s kernel configuration parameters. You must restart your computer if you update your kernel configuration parameter values.

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)

For the latest information on supported Linux distributions, point your browser to <http://www.ibm.com/software/data/db2/linux/validate/>.

To install DB2 products, the distribution requirements, hardware, and communications prerequisites must be met.

DB2 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™, System i or pSeries® systems that support Linux)
- eServer System z™ or System z9™

The supported operating systems for Linux include:

- Red Hat Enterprise Linux (RHEL) 4 Update 4
- Red Hat Enterprise Linux (RHEL) 5
- SUSE Linux Enterprise Server (SLES) 9 Service Pack 3
- SUSE Linux Enterprise Server (SLES) 10 Service Pack 1

Note: POWER requires a minimum of SLES 10 Service Pack 1 or RHEL 5

Multithreaded architecture limitations

If you are installing a DB2 Version 9.5 32-bit database product on a Linux operating system, consider upgrading to a 64-bit operating system and installing the DB2 Version 9.5 64-bit database product instead. The multithreaded architecture generally simplifies memory configuration. However, this could affect the memory configuration of 32-bit DB2 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 in preparation for your Linux distribution. 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 on setting and activating these parameters using the `sysctl` command.

Package requirements

The following tables list the package requirements for SLES and RHEL distributions for DB2 Version 9.5:

- `libaio.so.1` is required for DB2 servers using asynchronous i/o.
- `libstdc++so.5` is required for DB2 servers and clients.

Package requirements for SLES and RHEL

Package name	Description
<code>libaio</code>	contains the asynchronous library required for DB2 servers.
<code>compat-libstdc++</code>	contains <code>libstdc++so.5</code> (not required for Linux on POWER)

The following tables list the package requirements for SUSE Linux and Red Hat distributions for DB2 Version 9.5 partitioned servers.

- The `pdksh` Korn Shell package is required for all DB2 systems.
- A remote shell utility is required for partitioned database systems. DB2 supports the following remote shell utilities:
 - `rsh`
 - `ssh`

By default, DB2 uses `rsh` when executing commands on remote DB2 nodes, for example, when starting a remote DB2 database partition. To use the DB2 default, the `rsh-server` package must be installed (see table below). More information on `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 setup. For general Linux information, see your Linux distribution documentation.

Package requirements for SUSE Linux

Package name	Description
pdksh	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 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 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	pdksh	Korn Shell. This package is required for partitioned database environments.
/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 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 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 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. The NAS client can be downloaded from <https://www6.software.ibm.com/dl/dm/dm-nas-p>.
- One of the following browsers is required to view online help and to run First Steps (db2fs):
 - Mozilla 1.4 and up

- Firefox 1.0 and up
- Netscape 7.0 and up
- 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, or
 - you want to use any DB2 graphical tools on Linux for x86 and Linux on AMD 64/EM64T.

DB2 product installation on NFS (Network File System)

The installation of DB2 products on NFS (Network File System) is not recommended. Running DB2 products on NFS (for example, NFS mounting /opt/IBM/db2/V9.5 and then running off code that was physically installed on a remote system) requires several manual setup steps. There are also a number of potential issues with setting up NFS for DB2. These include possible problems that involve:

- Performance (impacted by network performance)
- Availability (you are allowing a single point of failure)
- Licensing (there is no checking done across machines)
- Diagnosing NFS errors can be difficult

As mentioned, the setup for NFS will require several manual actions including:

- Ensuring that the mount point preserve the install path
- Permission must be controlled (for example, write permission should not be given to the mounting machine)
- DB2 registries have to be set up manually and maintained across all mounting machines
- The db2ls command, which lists installed DB2 products and features, must be set up and maintained properly if you need to detect DB2 products and features
- More care is required when updating your DB2 product environment
- More steps are required when cleaning up on the exporting machine and the mounting machine

For detailed instructions, see the “Setting up DB2 for UNIX and Linux on NFS mounted file systems” white paper in <http://www.ibm.com/developerworks/db2/library/long/dm-0609lee>.

Security-enhanced Linux considerations

On RHEL 4 and 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 (Notice format might differ between RHEL 4 and RHEL 5.)

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 product installs successfully on a RHEL 4 or 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)

Before installing a DB2 database system, you should update your Linux kernel parameters. The default values for particular kernel parameters on Linux are not sufficient when running a DB2 database system.

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
2. Analyze the output to determine if there are any necessary changes required for your system. Comments have been added following the `//` to show what the parameter names are.

```
# ipcs -l

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

- Beginning with the first section on Shared Memory Limits, SHMMAX and SHMALL are the parameters that need to be looked at. SHMMAX is the maximum size of a shared memory segment on a Linux system whereas SHMALL is the maximum allocation of shared memory pages on a system.
 - It is recommended to set the SHMMAX value to be equal to the amount of physical memory on your system. However, the minimum required on x86 systems is 268435456 (256 MB) and for 64-bit systems, it is 1073741824 (1 GB).
 - SHMALL is set to 8 GB by default (8388608 KB = 8 GB). If you have more physical memory than this, and it is to be used for DB2, then this parameter should be increased to approximately 90% of your computer's physical memory. For instance, if you have a computer system with 16 GB of memory to be used primarily for DB2, then SHMALL should be set to 3774873 (90% of 16 GB is 14.4 GB; 14.4 GB is then divided by 4 KB, which is the base page size). The `ipcs` output has converted SHMALL into kilobytes. The kernel requires this value as a number of pages.
- The next section covers the amount of semaphores available to the operating system. The kernel parameter `sem` consists of 4 tokens, SEMMSL, SEMMNS,

SEMOPM and SEMMNI. SEMMNS is the result of SEMMSL multiplied by SEMMNI. The database manager requires that the number of arrays (SEMMNI) be increased as necessary. Typically, SEMMNI should be twice the maximum number of agents expected on the system multiplied by the number of logical partitions on the database server computer plus the number of local application connections on the database server computer.

- The third section covers messages on the system.
 - MSGMNI affects the number of agents that can be started, MSGMAX affects the size of the message that can be sent in a queue, and MSGMNB affects the size of the queue.
 - MSGMAX should be change to 64 KB (that is, 65535 bytes), and MSGMNB should be increased to 65535.
3. To modify these kernel parameters, edit 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:


```
kernel.sem=250 256000 32 1024
#Example shmmax for a 64-bit system
kernel.shmmax=1073741824
#Example shmall for 90 percent of 16 GB memory
kernel.shmall=3774873
kernel.msgmax=65535
kernel.msgmnb=65535
```
 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 Operating Environment)

To install a DB2 product, the following operating system, hardware, and communications requirements must be met:

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 • If raw devices are used, patch 125100-07 	UltraSPARC

Table 3. Solaris installation requirements (continued)

Operating System	Hardware
Solaris 10 <ul style="list-style-type: none"> • 64-bit kernel • Patch 118855-33 • If raw devices are used, patch 125101-07 	Solaris x64 (Intel 64 or AMD64)

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 higher. The NAS client can be downloaded from Web site: <https://www6.software.ibm.com/dl/dm/dm-nas-p>.
- One of the following browsers is required to view online help and to run First Steps (db2fs):
 - Mozilla 1.4 and up
 - Firefox 1.0 and up
 - Netscape 7.0 and up
- 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 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 on the "Patches" menu item 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 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 Environment can be downloaded from FTSI at: <http://download.ftsi.fujitsu.com/>.

DB2 product installation on an NFS-mounted directory

The installation of DB2 products on an NFS (Network File System) mounted directory is not recommended. Running DB2 products on an NFS-mounted directory (for example, mounting an NFS-exported directory on

/opt/IBM/db2/V9.5) requires several manual setup steps. There are also a number of potential issues with setting up NFS for DB2. These include possible problems that involve:

- Performance (impacted by network performance)
- Availability (you are allowing a single point of failure)
- Licensing (there is no checking done across machines)
- Diagnosing NFS errors can be difficult

As mentioned, the setup for NFS will require several manual actions including:

- Ensuring that the mount point preserve the install path
- Permission must be controlled (for example, write permission should not be given to the mounting machine)
- DB2 registries have to be set up manually and maintained across all mounting machines
- The db2ls command, which lists installed DB2 products and features, must be set up and maintained properly if you need to detect DB2 products and features
- More care is required when updating your DB2 product environment
- More steps are required when cleaning up on the exporting machine and the mounting machine

For detailed instructions, see the "Setting up DB2 for UNIX and Linux on NFS mounted file systems" white paper in <http://www.ibm.com/developerworks/db2/library/long/dm-0609lee>.

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), you should 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)

To install a DB2 product, the following operating system, software, and hardware requirements must be met:

Table 4. Windows workstation platforms

Operating System	Prerequisites	Hardware
Windows XP Professional (32-bit and x64)	Windows XP Service Pack 2 or later	All Intel and AMD processors capable of running the supported Windows operating systems (32-bit and x64 based systems)
Windows Vista Ultimate (32-bit and x64)	IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 1.1 SP1 or .NET 2.0 framework runtime	
Windows Vista Business (32-bit and x64)		
Windows Vista Enterprise (32-bit and x64)		
	64-bit IBM data server provider for .NET applications are supported	

Table 5. Windows server platforms

Operating System	Prerequisites	Hardware
Windows 2003 Standard Edition (32-bit and x64)	Service Pack 1 or later.	All Intel and AMD processors capable of running the supported Windows operating systems
Windows 2003 Enterprise Edition (32-bit and x64)	R2 is also supported	
Windows 2003 Datacenter Edition (32-bit and x64)	IBM data server provider for .NET client applications and CLR server-side procedures require .NET 1.1 SP1 or .NET 2.0 framework runtime	
	64-bit IBM data server provider for .NET applications are supported	

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 1.1 SP1 or .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 products). Prior to 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):
 - Internet Explorer 6 and up
 - Mozilla 1.4 and up
 - Firefox 1.0 and up
 - Netscape 7.0 and up

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 the related links for information on supported versions and patches.

Chapter 3. Installing IBM data server clients

Installing IBM data server clients (Windows)

This task describes how to install an IBM data server client on the Windows operating system. The instructions apply to any IBM data server client type, namely the IBM Data Server Client, the IBM Data Server Runtime Client, and the IBM Data Server Driver for ODBC, CLI, and .NET. The main procedure covers a simple, but common, case where there is no DB2 product already installed. Related links are provided for information such as alternate methods for installing an IBM data server client.

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

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

Prerequisites

Before installing IBM data server clients:

- You have determined which client best suits your need.
- You have located a DVD or other install image that you need. 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 product. The VS2005 runtime library is needed on the operating system before the DB2 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 product can be installed in the same path if either of the following products is already installed:
 - IBM Data Server Runtime Client or
 - IBM Data Server Driver for ODBC, CLI, and .NET
- 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 auto-run feature starts the DB2 Setup wizard which determines the system language and starts the setup program for that language.
4. If you are installing a Data Server Client, choose **Install a Product** when the DB2 Launchpad opens. This step does not apply to the Data Server Runtime Client or the Data Server Driver for ODBC, CLI, and .NET because it does not have a launchpad. See the Related Links for command line options.
5. Follow the DB2 Setup wizard's prompts.

After completing this procedure, the product is now installed at the location you specified during the installation. The default installation path of Data Server Client and Data Server Runtime Client is Program Files\IBM\sqllib. The default installation path of Data Server Driver for ODBC, CLI, and .NET is Program Files\IBM\IBM DATA SERVER DRIVER

As part of the 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 have a copy of DB2 Version 8 or DB2 Version 9.1 installed, the default instance is DB2_01.

This installation does not include product documentation. See the related links for options for installing or accessing the DB2 Information Center.

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

Notes on national language installations

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 for ODBC, CLI, and .NET, there are separate install images for each language.

Notes on installing on a machine that has an existing DB2 Version 9 product

When installing Data Server Runtime Client or Data Server Client, the default installation path for the first copy installed of a DB2 product is Program Files\IBM\sqllib. If a second copy is installed in the same machine, the default directory name is Program Files\IBM\sqllib_01. In general, the default directory name is sqllib_*nn* where *nn* is the number of copies installed in that machine minus one.

When installing Data Server Driver for ODBC, CLI, and .NET, the default installation path for the first copy installed 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 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"
```

See the Related Links for additional **setup** command parameters.

If you are installing a second copy of the Data Server Driver for ODBC, CLI, and .NET (up to a maximum of 16 copies), 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*

See the Related Links for additional **setup** command parameters.

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

The default copy name of the Data Server Driver for ODBC, CLI, and .NET is IBMDBCL1

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

Notes on installing on a machine that has an existing DB2 Universal Database (UDB) Version 8 client

When installing a Data Server Client on a machine that already has a DB2 Universal Database (UDB) Version 8 copy installed, users will be presented with the option to install a new copy or to migrate 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 migrate will copy the DB2 UDB Version 8 client instance settings to the DB2 Version 9 copy then remove the DB2 UDB Version 8 copy.

If a machine already has a DB2 Universal Database (UDB) Version 8 copy installed, the Version 9 copies cannot be set to default.

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

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 product is already installed. A summary of the prerequisites is presented below. If you need more detailed information, see the related links at the end of this section.

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

If the machine already has a DB2 server product installed, it is not necessary to install a client because the DB2 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.
- 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/V9.5

UNIX /opt/IBM/db2/V9.5

This installation does not include product documentation.

See the related links for options for installing or accessing the DB2 Information Center.

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` 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 V9.5. If a copy is already installed, the second installation shows a default directory name of V9.5_01. In general, the default directory name is V9.5_*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 copy. For information on migrating client instances to DB2 Version 9.5, see the migration 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 platforms.

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 product without root privileges.

The non-root installation of a DB2 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 product without consuming a system administrator's time
- Application developers who are not typically system administrators but use DB2 products to develop applications
- Independent Software Vendors (ISVs) who develop software that does not require root authority yet embeds a DB2 product

Although non-root installations have most of the functionality 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 product are created in a directory of the root user's choosing.

Unlike root users, non-root users cannot choose where DB2 products are installed. Non-root installations are always placed in the \$HOME/sqllib directory, where \$HOME represents the non-root user's home directory. The layout of the subdirectories within the sqllib 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 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 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 product is ready for use immediately after installation.

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 products are not supported in non-root installations:

- IBM Data Studio
- DB2 Embedded Application Server (DB2 EAS)
- 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 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 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 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 privilege can be updated or dropped only by a user with root privilege.

DB2 instance commands

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

db2icrt

When installing a DB2 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 (`db2nrcfg`) 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 product.

db2idrop

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

db2imigr

Migration is not supported for non-root installations.

Migration limitation

Root instances cannot be migrated 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 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 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 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 platforms, 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 authority.

Authentication type in non-root installations

Operating system-based authentication is the default authentication type for DB2 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 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 products can be installed as a non-root user.

Before you install any DB2 product as a non-root user, you should be aware of the differences between root installations and non-root installations, and the limitations of non-root installations. Refer to the Related Links at the end of this topic for details.

Prerequisites for installing a DB2 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.
- 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 products as a non-root user should be transparent to the non-root user. In other words, there is nothing special a non-root user needs to do to install a DB2 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 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 products are installed, any FILE keyword in your response file is ignored.

Refer to the Related Links at the end of this topic for details.

3. After the DB2 product is installed, you need to open a new login session to use the non-root DB2 instance. Alternatively, you can use the same login session if you source 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 product is installed, you should verify your operating system user process resource limits (ulimits). If the minimum ulimit values are not met, the DB2 engine could encounter unexpected operating resource shortage errors. These errors can lead to a DB2 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 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 as a user with root authority
5. Navigate to the `$HOME/sqllib/instance` directory, where `$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 on page 36.

You must rerun the `db2rfe` command after applying fix packs to keep root-based features enabled on non-root installations.

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 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 products using `db2_deinstall` (Linux and UNIX)

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

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

Note:

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

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

1. Log in with the user ID that was used to install the DB2 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 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. In non-root installations, however, you can also remove specific DB2 features by running the `db2nrupdt` command.

Part 3. Database connections for IBM data server clients

Chapter 4. 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 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.

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

DB2 Universal Database (UDB) Version 8 and DB2 Version 9.1 clients can access a remote DB2 Version 9.5 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.

IBM Data Server Client, IBM Data Server Runtime Client, and IBM Data Server Driver for ODBC, CLI, and .NET Version 9.5 can access DB2 Version 9.1 and DB2

UDB Version 8 servers. However, new DB2 Version 9.5 functionality is not available.

Access to DB2 Version 9.1 or Version 9.5 servers from DB2 UDB Version 7 clients

Access from DB2 UDB Version 7 clients is not supported.

Combinations of DB2 Version 9.1 or Version 9.5, and, DB2 products on midrange and mainframe platforms

DB2 Version 9.5 and Version 9.1 *servers* support access from the following *clients* on midrange and mainframe platforms:

- DB2 for z/OS Version 7 and Version 8.
- DB2 for iSeries™ Version 5.
- DB2 for VM and VSE Version 7.

IBM Data Server Client Version 9.5, IBM Data Server Runtime Client Version 9.5, and DB2 Version 9.1 clients can access DB2 Connect Version 9.5, Version 9.1, and Version 8.

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). Refer to Configuring Secure Socket Layer (SSL) support in a DB2 instance.

In addition, the Windows Named Pipes protocol is supported on Windows networks. To administer a Version 9 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.

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 46

Use this method if you don’t have any information on 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 48

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’ll need to 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 appear 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:

- 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 desired.

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 on the folder beside **Known Systems** to list all the systems known to your client or double-click on 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. Once you add a system, it will appear 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, your database connection should be tested.

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 would like 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 on 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 steps below. 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:
 - Catalog the TCP/IP node from the client.
 - Catalog the Named Pipes node from the client.
4. Catalog the database that you want to connect to on the client.
5. Test the client-to-server connection.

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:fe77:62dc 	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	

Table 9. TCP/IP parameter values worksheet (continued)

Parameter	Description	Sample Value	Your Value
Service Name • 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.

Procedure

- 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:

```
9.21.15.235    myserver    # IP address for myserver
```

where:

9.21.15.235
represents the *ip_address*

myserver
represents the *hostname*

represents a comment describing the entry

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 start-up 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 start-up 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, 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`, and `KERBEROS`.

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 <i>sample</i> database is created on the server, a database alias of <i>sample</i> is also created. The database name represents the remote database alias (on the server).	<i>sample</i>	
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.	<i>mysample</i>	
Authentication (<i>auth_value</i>)	The type of authentication required in your environment.	<i>Server</i>	
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.	<i>db2node</i>	

Testing the client-to-server connection using the CLP

After cataloging the node and the database, you should 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 should be set to match the value on the server or it should 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 isn't 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 start-up script:

For bash, Bourne or Korn shell

```
. INSTHOME/sqllib/db2profile
```

For C shell

```
source INSTHOME/sqllib/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.

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

Chapter 5. 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 for ODBC, CLI, and .NET.

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.

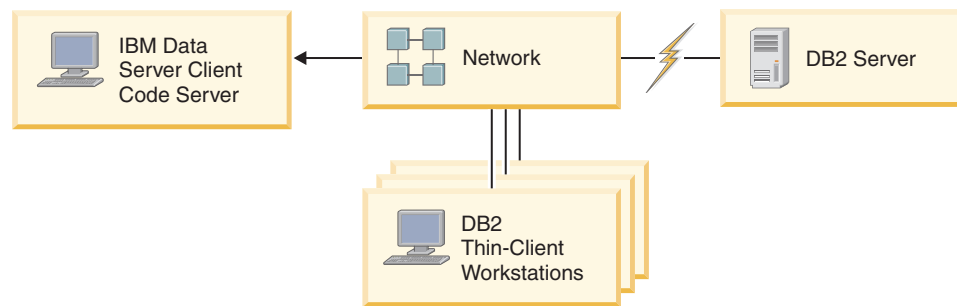


Figure 1. A typical thin client topology using IBM Data Server Client




Figure 2. 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 need to install, update, or migrate 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.

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 launch 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, launch 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:\sqllib` 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 User sName** 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 product.

Locate the sample response file `db2thin.rsp` in the `c:\sqllib\thnsetup` directory, where `c:\sqllib` 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:\sql11b`).

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.

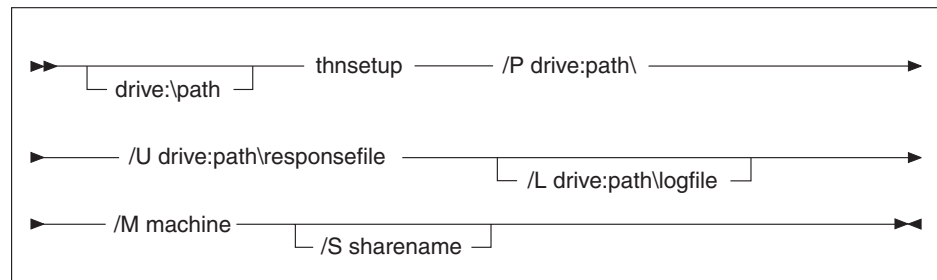
Setting up thin clients using the thnsetup command (Windows)

This command sets up the thin client workstation and makes the required links to the code server.

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

To set up a thin client:

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:\sqllib\thnsetup`, where `c:\sqllib\` 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.

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
```

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.

Part 5. Merge Modules

Chapter 6. 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. See the related links for details on DB2 instance merge modules.

Using non-DB2 instance Windows Installer merge modules, you can easily add IBM Data Server Driver for ODBC, CLI, and .NET 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 for ODBC, CLI, and .NET 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 modules are available for your use:

IBM Data Server Driver for ODBC and CLI Merge Module.msm

This module provides support for applications that use Open Database Connectivity (ODBC) to access data. This module is a prerequisite of other merge modules.

IBM Data Server Provider for .NET Merge Module.msm

This module 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.

The following merge modules contain language specific messages used by IBM data server driver for ODBC, CLI, and .NET. 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 9.5 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 functionality that only exists in a DB2 instance merge module. (The DB2 instance merge modules are listed below.)

Using DB2 instance Windows Installer merge modules, you can easily add IBM Data Server Runtime Client functionality 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 functionality required for database connection, SQL, and DB2 command functionality. 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 like 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 shipped 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

Part 6. Additional installation options

Chapter 7. 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.
- /v** This option allows you to pass additional command-line options and public properties to Windows Installer. You must specify this option to perform a response file installation.
- /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.
- /qn** This option allows you to perform a silent installation with no User Interface (UI).
- /qb!** This option displays a basic user interface that shows simple progress and error message handling and hides the **Cancel** button.
- /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 /L1036` 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
Dutch (Standard)	1043
English	1033
Finnish	1035
French (Standard)	1036
German	1031
Greek	1032

Table 12. Language Identifiers (continued)

Language	Identifier
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]"
```

IBM Data Server Driver for ODBC, CLI, and .NET installation command line options (Windows)

The following list describes command-line options available when you run the setup command to install IBM Data Server Driver for ODBC, CLI, and .NET on Windows operating systems. For more information on 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.

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

- To perform a response file installation, use the following command:
 `setup /u "[Full path to the response file]"`
- 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"`

Part 7. Uninstalling

Chapter 8. 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.

1. 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.
2. 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 for ODBC, CLI, and .NET. For more information, see the `db2unins` command topic.

Part 8. Appendixes

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 hard-copy 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](http://www.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 email 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 email 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/shop/publications/order. English DB2 Version 9.5 manuals in PDF format and translated versions 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
<i>Administrative API Reference</i>	SC23-5842-01	Yes
<i>Administrative Routines and Views</i>	SC23-5843-01	No
<i>Call Level Interface Guide and Reference, Volume 1</i>	SC23-5844-01	Yes
<i>Call Level Interface Guide and Reference, Volume 2</i>	SC23-5845-01	Yes
<i>Command Reference</i>	SC23-5846-01	Yes
<i>Data Movement Utilities Guide and Reference</i>	SC23-5847-01	Yes
<i>Data Recovery and High Availability Guide and Reference</i>	SC23-5848-01	Yes
<i>Data Servers, Databases, and Database Objects Guide</i>	SC23-5849-01	Yes
<i>Database Security Guide</i>	SC23-5850-01	Yes
<i>Developing ADO.NET and OLE DB Applications</i>	SC23-5851-01	Yes
<i>Developing Embedded SQL Applications</i>	SC23-5852-01	Yes
<i>Developing Java Applications</i>	SC23-5853-01	Yes
<i>Developing Perl and PHP Applications</i>	SC23-5854-01	No
<i>Developing User-defined Routines (SQL and External)</i>	SC23-5855-01	Yes
<i>Getting Started with Database Application Development</i>	GC23-5856-01	Yes
<i>Getting Started with DB2 installation and administration on Linux and Windows</i>	GC23-5857-01	Yes
<i>Internationalization Guide</i>	SC23-5858-01	Yes
<i>Message Reference, Volume 1</i>	GI11-7855-00	No
<i>Message Reference, Volume 2</i>	GI11-7856-00	No
<i>Migration Guide</i>	GC23-5859-01	Yes
<i>Net Search Extender Administration and User's Guide</i>	SC23-8509-01	Yes
<i>Partitioning and Clustering Guide</i>	SC23-5860-01	Yes
<i>Query Patroller Administration and User's Guide</i>	SC23-8507-00	Yes
<i>Quick Beginnings for IBM Data Server Clients</i>	GC23-5863-01	No

Table 13. DB2 technical information (continued)

Name	Form Number	Available in print
<i>Quick Beginnings for DB2 Servers</i>	GC23-5864-01	Yes
<i>Spatial Extender and Geodetic Data Management Feature User's Guide and Reference</i>	SC23-8508-01	Yes
<i>SQL Reference, Volume 1</i>	SC23-5861-01	Yes
<i>SQL Reference, Volume 2</i>	SC23-5862-01	Yes
<i>System Monitor Guide and Reference</i>	SC23-5865-01	Yes
<i>Troubleshooting Guide</i>	GI11-7857-01	No
<i>Tuning Database Performance</i>	SC23-5867-01	Yes
<i>Visual Explain Tutorial</i>	SC23-5868-00	No
<i>What's New</i>	SC23-5869-01	Yes
<i>Workload Manager Guide and Reference</i>	SC23-5870-01	Yes
<i>pureXML Guide</i>	SC23-5871-01	Yes
<i>XQuery Reference</i>	SC23-5872-01	No

Table 14. DB2 Connect-specific technical information

Name	Form Number	Available in print
<i>Quick Beginnings for DB2 Connect Personal Edition</i>	GC23-5839-01	Yes
<i>Quick Beginnings for DB2 Connect Servers</i>	GC23-5840-01	Yes
<i>DB2 Connect User's Guide</i>	SC23-5841-01	Yes

Table 15. Information Integration technical information

Name	Form Number	Available in print
<i>Information Integration: Administration Guide for Federated Systems</i>	SC19-1020-01	Yes
<i>Information Integration: ASNCLP Program Reference for Replication and Event Publishing</i>	SC19-1018-02	Yes
<i>Information Integration: Configuration Guide for Federated Data Sources</i>	SC19-1034-01	No
<i>Information Integration: SQL Replication Guide and Reference</i>	SC19-1030-01	Yes
<i>Information Integration: Introduction to Replication and Event Publishing</i>	SC19-1028-01	Yes

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/v9r5>.

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 85.

Displaying SQL state help from the command line processor

DB2 returns an `SQLSTATE` value for conditions that could be the result of an SQL statement. `SQLSTATE` help explains the meanings of SQL states and SQL state class codes.

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

For DB2 Version 9 topics, the DB2 Information Center URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>

For DB2 Version 8 topics, go to the Version 8 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. Clear the browser cache and then 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. Clear the browser cache and then refresh the page to display the DB2 Information Center in your preferred language.

On some browser and operating system combinations, you might have to 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

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 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. Non-Administrative and Non-Root DB2 Information Centers always run in stand-alone mode. .
2. Use the Update feature to see what updates are available. If there are updates that you would like to 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, you have to 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 Vista, the commands listed below must be run as an administrator. To launch a command prompt or graphical tool with full administrator privileges, right-click on 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 on **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:

```
/etc/init.d/db2icdv95 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 9.5 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/V9.5 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 launches to display the stand-alone Information Center.

3. Click the **Update** button (🔄). On the right hand 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 terminate the processes that were started with the help_start batch file. Do not use Ctrl-C or any other method to terminate 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 terminate the processes that were started with the help_start script. Do not use any other method to terminate the help_start script.
7. Restart the DB2 Information Center.
 - On Windows, click **Start** → **Control Panel** → **Administrative Tools** → **Services**. Then right-click on **DB2 Information Center** service and select **Start**.
 - On Linux, enter the following command:


```
/etc/init.d/db2icdv95 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 on 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 products.

DB2 documentation

Troubleshooting information can be found in the DB2 Troubleshooting Guide or the Support and Troubleshooting section of the DB2 Information Center. There you will find information on 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 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/udb/support.html>

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.

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 give you any license to these patents. You can send license inquiries, in writing, to:

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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.

This document may provide links or references to non-IBM Web sites and resources. IBM makes no representations, warranties, or other commitments whatsoever about any non-IBM Web sites or third-party resources that may be referenced, accessible from, or linked from this document. A link to a non-IBM Web site does not mean that IBM endorses the content or use of such Web site or

its owner. In addition, IBM is not a party to or responsible for any transactions you may enter into with third parties, even if you learn of such parties (or use a link to such parties) from an IBM site. Accordingly, you acknowledge and agree that IBM is not responsible for the availability of such external sites or resources, and is not responsible or liable for any content, services, products, or other materials on or available from those sites or resources. Any software provided by third parties is subject to the terms and conditions of the license that accompanies that software.

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

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

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

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

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

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

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

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

This information 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 may contain 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.

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

The following terms are trademarks or registered trademarks of the International Business Machines Corporation in the United States, other countries, or both.

pureXML	Distributed Relational Database Architecture
OpenPower	Informix
DB2	AIX
System z9	System z
POWER	OS/390
DB2 Connect	DB2 Universal Database
z/OS	Redbooks
developerWorks	System i
IBM	Rational
DRDA	Tivoli
OS/400	eServer
pSeries	ibm.com
iSeries	

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 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 Xeon, and Itanium, and Pentium, and Intel are trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- Microsoft, and Windows 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

- adding
 - databases manually 45
- AIX
 - installation requirements 14

B

- books
 - printed
 - ordering 88

C

- cataloging
 - database parameter values worksheet 56
 - databases 55
 - host databases
 - DB2 Connect 55
 - Named Pipes 51
 - TCP/IP node 53
- client configurations
 - non-supported 43
 - supported 43
- client profiles
 - configuring using the import function 48
 - creating using the export function 47
- client to server communication
 - configuring connections 41
 - TCP/IP parameter values worksheet 51
 - testing connections using the CLP 56
- clients
 - server connections 45, 49
- code directory
 - thin clients 63
- code servers
 - installing an IBM Data Server Client 63
 - installing DB2 Connect Personal Edition 63
 - thin client
 - mapping network drives 64
- command line options
 - IBM Data Server Driver for ODBC, CLI, and .NET installation 76
 - IBM Data Server Runtime Client installation 75
- command line processor (CLP)
 - cataloging a database 55
 - cataloging a node 53
 - configuring client to server connection 49
 - configuring TCP/IP
 - client 52
- commands
 - catalog database 55
 - catalog npipe 51
 - catalog tcpip 53
 - db2osconf 16
 - db2rfe - enabling root features 32, 36
 - db2setup 30
 - db2start 56
 - thnsetup 65

- communication protocols
 - Named Pipes 44
 - SSL 44
 - TCP/IP 44
 - Configuration Assistant (CA)
 - cataloging a database 41
 - configuring
 - client profiles 48
 - client to server communications 41
 - client to server connection 45
 - database connection 45
 - creating client profiles 47
 - Discovery feature 46
 - LDAP considerations 49
 - testing
 - database connections 49
 - configuring
 - client to server connection
 - command line processor (CLP) 49
 - Configuration Assistant (CA) 45
 - TCP/IP worksheet 51
 - TCP/IP
 - client 52
- ## D
- databases
 - cataloging
 - command line processor (CLP) 55
 - connections
 - configuring 45, 46
 - testing 49
 - DB2 Connect
 - installing
 - prerequisites 26
 - Personal Edition
 - installing (Windows) 63
 - thin client
 - code directory 63
 - installing 62
 - mapping network drive to code server 64
 - response files 63
 - setup 61
 - topology overview 61
 - DB2 Information Center
 - languages 89
 - updating 90
 - versions 89
 - viewing in different languages 89
 - db2osconf command
 - determining kernel configuration parameter values 16
 - db2rfe command
 - enabling root features 32, 36
 - directory structures
 - root installations compared to non-root installations 32
 - Discovery feature
 - configuring database connection 46
 - disk space requirements 13
 - documentation
 - overview 85
 - PDF 85

- documentation (*continued*)
 - printed 85
 - terms and conditions of use 92
- dropping
 - non-root instances 38

E

- examples
 - connecting to a remote database 56

F

- fix packs
 - non-root installations 37

H

- hardware
 - requirements
 - AIX 14
 - HP-UX 15
 - Linux 17
 - Solaris Operating Environment 22
 - Windows 25

- help
 - configuring language 89
 - SQL statements 88

- host databases
 - client connections 26

- HP-UX
 - installing
 - DB2 servers 15
 - IBM data server clients 15
 - kernel configuration parameters
 - modifying 16
 - recommended values 16

I

- IBM data server clients
 - cataloging
 - named pipes node 51
 - TCP/IP node 53
 - connecting to
 - host databases 26
 - IBM Data Server Client 3, 4
 - IBM Data Server Driver for ODBC, CLI, and .NET 3
 - IBM Data Server Runtime Client 3, 4
 - installing
 - on the code server 63
 - overview 6, 7
 - UNIX 30
 - Windows 27
 - overview 3
 - types 4
 - user accounts 27
- IBM Data Server Driver for ODBC, CLI, and .NET
 - installation
 - command line options 76
- IBM Data Server Runtime Client
 - installation
 - command line options 75
- import function
 - configuring client profiles 48

- installing
 - DB2 products as a non-root user 35
 - requirements
 - AIX 14
 - HP-UX 15
 - Linux 17
 - Solaris Operating Environment 22
 - Windows 25

- instances
 - removing non-root 38

K

- kernel configuration parameters
 - db2osconf command (HP-UX) 16
 - modifying on HP-UX 16
 - modifying on Linux 21
 - modifying on Solaris Operating System 24
 - recommended (HP-UX) 16

L

- LDAP (Lightweight Directory Access Protocol)
 - directory support considerations 49
- Lightweight Directory Access Protocol (LDAP)
 - directory support considerations 49
- limitations
 - non-root installations 32
- Linux
 - installation requirements 17
 - modifying kernel parameters 21
 - removing
 - DB2 non-root instances 38
- Linux library
 - libaio.so.1 17
 - libstdc++so.5 17

M

- manually adding databases
 - Configuration Assistant (CA) 45
- mapping network drives
 - thin clients 64
- memory requirements 13
- merge modules
 - DB2 instance 70
 - non-DB2 instance 69
- modifying
 - kernel parameters (HP-UX) 16
- modifying kernel parameters
 - HP-UX 16
 - Linux 21
 - Solaris Operating System 24

N

- Named Pipes
 - parameter values worksheet 50
 - supported protocol 44
- network drives
 - mapping 64
- Network File System (NFS) installation
 - on AIX 14
 - on HP-UX 15
 - on Linux 17

- Network File System (NFS) installation *(continued)*
 - on Solaris Operating Environment 22
- non-root installations
 - differences 32
 - directory structure 32
 - enabling root-based features 36
 - fix packs 37
 - installing 35
 - limitations 32
 - overview 31
 - uninstalling 38
- non-root instances
 - dropping 38
 - removing 38
- notices 95

O

- operating system requirements
 - AIX 14
 - HP-UX 15
 - Linux 17
 - Solaris Operating Environment 22
 - Windows 25
- ordering DB2 books 88

P

- parameters
 - values worksheets
 - cataloging databases 56
 - configuring client to server connections 51
 - named pipes 50
- problem determination
 - information available 92
 - tutorials 92

R

- removing
 - non-root instances 38
- requirements
 - disk 13
 - memory 13
- response files
 - creating
 - thin client 63
- root installations
 - differences 32
 - directory structure 32
- root-based features
 - non-root installation 36

S

- servers
 - client connections 45, 49
- software requirements
 - AIX 14
 - HP-UX 15
 - Linux 17
 - Solaris Operating Environment 22
 - Windows 25
- Solaris Operating Environment
 - installation requirements 22

- Solaris Operating System
 - modifying kernel parameters 24
- SQL statements
 - displaying help 88
- SSL
 - supported protocol 44
- supported client configurations 43

T

- TCP/IP
 - configuration
 - client 52
 - supported platforms 44
 - TCP/IPv6 support 44
- terms and conditions
 - use of publications 92
- testing
 - client-to-server connections 56
 - database connections 49
- thin clients
 - code directory 63
 - code server
 - mapping network drives 64
 - considerations 61
 - enabling 65
 - installation 62
 - response files 63
 - typical setup 61
- thnsetup command 65
- troubleshooting
 - online information 92
 - tutorials 92
- tutorials
 - problem determination 92
 - troubleshooting 92
 - Visual Explain 91

U

- uninstalling
 - IBM data server clients 81
 - non-root 38
- UNIX
 - installing
 - IBM data server clients 30
 - removing
 - DB2 non-root instances 38
- updates
 - DB2 Information Center 90
- user accounts
 - IBM data server clients 27

V

- Visual Explain
 - tutorial 91

W

- Windows operating systems
 - installing
 - DB2 servers (requirements) 25
 - IBM data server clients (procedure) 27
 - IBM data server clients (requirements) 25



Printed in USA

GC23-5863-01



Spine information:

DB2 Version 9.5 for Linux, UNIX, and Windows

Quick Beginnings for IBM Data Server Clients

