

IBM DB2 10.1
for Linux, UNIX, and Windows

*Getting Started with Database
Application Development*



IBM DB2 10.1
for Linux, UNIX, and Windows

*Getting Started with Database
Application Development*



Note

Before using this information and the product it supports, read the general information under Appendix C, "Notices," on page 91.

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 <http://www.ibm.com/shop/publications/order>
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at <http://www.ibm.com/planetwide/>

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

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

© Copyright IBM Corporation 2006, 2012.

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

Contents

About this book	v	Support for database application development in Fortran	35
Chapter 1. The DB2 database application development environment	1	Java software support for DB2 database products.	35
Chapter 2. Support for elements of the database application development environment	3	Support for database application development in Perl	38
Chapter 3. Supported operating systems for database application development	5	Support for database application development in PHP.	38
Installation requirements for DB2 servers and IBM data server clients (AIX)	5	Support for database application development in Ruby/Ruby on Rails	39
Installation requirements for DB2 servers and IBM data server clients (HP-UX)	7	Support for database application development in REXX	40
Installation requirements for DB2 servers and IBM data server clients (Linux)	11	Chapter 7. Supported transaction managers	41
Installation requirements for DB2 servers and IBM data server clients (Solaris)	12	Chapter 8. Supported database application development tools	43
Installation requirements for DB2 database servers and IBM data server clients (Windows)	13	Chapter 9. Configuring the database application development environment	45
Chapter 4. IBM Data Server Client support for database application development	17	Configuring the operating system for database application development	46
Chapter 5. Supported database application programming interfaces	19	Setting up the application development environment (Linux and UNIX)	46
ADO.NET development for IBM Data Servers	19	Setting up the application development environment (Windows)	47
IBM Data Server Provider for .NET	19	Configuring the IBM Data Server Client for database application development	48
ODBC .NET Data Provider	20	Configuring the development environment for application programming interfaces	49
OLE DB .NET Data Provider	21	Configuring your development environment to build and run CLI and ODBC applications	50
Introduction to DB2 Call Level Interface and ODBC	22	Installing the IBM Data Server Driver for JDBC and SQLJ	50
Introduction to embedded SQL	24	Transaction managers	54
Supported drivers for JDBC and SQLJ	25	X/Open distributed transaction processing model	54
IBM OLE DB Provider for DB2	26	Setting up the unixODBC driver manager	76
Perl DBI	27	Appendix A. DB2 database product and packaging information	79
PHP application development for IBM data servers	28	Appendix B. Overview of the DB2 technical information	81
Python, SQLAlchemy and Django Framework application development for IBM data servers	29	DB2 technical library in hardcopy or PDF format.	81
The IBM_DB Ruby driver and Rails adapter	29	Displaying SQL state help from the command line processor	84
Chapter 6. Programming languages and compilers for database application development	31	Accessing different versions of the DB2 Information Center	84
Supported .NET development software	31	Updating the DB2 Information Center installed on your computer or intranet server	84
Support for database application development in C	32	Manually updating the DB2 Information Center installed on your computer or intranet server	86
Support for database application development in C++	33	DB2 tutorials	87
Support for database application development in COBOL.	34		

DB2 troubleshooting information 88
Terms and conditions 88

Index 95

Appendix C. Notices 91

About this book

The *Getting Started with Application Development Guide* provides key concepts for supporting and configuring your database applications in the following key areas:

- database application programming interfaces
- database application programming languages
- compilers
- transaction managers
- development tools
- operating systems

The information presented in this book will help to provide you with the necessary resources for successful database application development.

Chapter 1. The DB2 database application development environment

The DB2[®] database application development environment is composed of several software elements.

These software elements include:

- **operating system**
- **IBM[®] Data Server Client**
- **database application programming interface (API)**
- **programming language**
- **transaction manager**
- **development tools**

For many of these elements, there are multiple choices available. For example: there might be several operating systems available for your hardware; DB2 database supports several APIs; and any XA-compliant transaction manager can be used to manage distributed transactions with DB2 databases.

Starting in V9.7 FP3a, the Data Server Clients periodically invoke the SYSIBM.DB2CONN_LIC stored procedure for licensing and configuration purposes.

Chapter 2. Support for elements of the database application development environment

Application development for DB2 databases is supported for a wide range of operating systems and database clients, along with numerous programming interfaces, languages, compilers, and tools.

The following documentation topics define the support for database application development on DB2 products:

- Chapter 3, “Supported operating systems for database application development,” on page 5
- Chapter 4, “IBM Data Server Client support for database application development,” on page 17
- Chapter 5, “Supported database application programming interfaces,” on page 19
- Chapter 6, “Programming languages and compilers for database application development,” on page 31
- Chapter 7, “Supported transaction managers,” on page 41
- Chapter 8, “Supported database application development tools,” on page 43

Chapter 3. Supported operating systems for database application development

Application development for DB2 databases is supported on a range of operating systems and architectures.

You can develop DB2 database applications on the following operating systems:

- AIX®
- HP-UX
- Linux
- Solaris
- Windows

There are restrictions about when you can use available 32-bit or 64-bit architectures:

- With Linux on x86 and Windows operating systems on x86, only 32-bit DB2 database instances are supported. On 64-bit Windows operating system, both the 32-bit and 64-bit instances are supported. On all other operating systems, only 64-bit DB2 database instances are supported.
- You can run 32-bit database applications with either a 32-bit or 64-bit version of one of the IBM data server clients and develop with either the 32-bit or 64-bit IBM Data Server Client
- The steps for building applications in a 32-bit environment are sometimes different from the steps for building applications in a 64-bit environment.

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. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

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

If you are planning to install the IBM DB2 pureScale® Feature, different installation prerequisites are required. See the Installing the DB2 pureScale Feature topic.

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

Table 1. AIX installation requirements

Operating System	Hardware
<p>AIX Version 6.1²</p> <ul style="list-style-type: none"> • 64-bit AIX kernel is required • AIX 6.1 Technology Level (TL) 6 and Service Pack (SP) 5 • Minimum C++ runtime level requires the xlc.rte 11.1.0.1 and xlc AIX rte 11.1.0.1 (or later) filesets. <p>AIX Version 7.1</p> <ul style="list-style-type: none"> • 64-bit AIX kernel is required • AIX 7.1 Technology Level (TL) 0 and Service Pack (SP) 3 • Minimum C++ runtime level requires the xlc.rte 11.1.0.1 and xlc AIX rte 11.1.0.1 (or later) filesets. 	<p>64-bit Common Hardware Reference Platform (CHRP) architecture, excluding POWER3 processor-based systems.¹</p> <p>All processors that are capable of running the supported AIX operating systems.</p>

- ¹ To verify that it is a CHRP architecture system, issue the command **lscfg** and look for the following output: Model Architecture: chrp. For POWER3 processor-based systems, first upgrade to POWER4 processor-based systems. POWER3 processor-based systems are not supported.
- ² In 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.
- ³ Additional considerations for POWER7[®] systems:
 - DB2 products running on POWER7 systems support Active Memory[™] Expansion.
- 32-bit clients are also supported

Note: If the minimum operating system requirements were applied using an operating system upgrade rather than using a new operating system installation, you must install I/O completion ports (IOCP) separately. You can obtain the IOCP libraries from the AIX installation CD and configure IOCP on the DB2 server. For a client only installation, IOCP is not required. For more information, see “configuring IOCP on AIX” in *Troubleshooting and Tuning Database Performance*.

Software considerations

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

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

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

- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if:
 - you want to use the DB2 Setup wizard to install a DB2 database product on Linux or UNIX operating systems
- For details regarding known AIX issues, see www.ibm.com/support/docview.wss?uid=swg21165448

Mounting considerations

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

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

Before you install DB2 database products on HP-UX operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

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

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

Note: A 64-bit HP-UX operating system is required to support DB2 products.

Table 2. HP-UX installation requirements

Operating System	Hardware
DB2 products are supported on: <ul style="list-style-type: none"> • HP-UX 11iv3 (11.31) with: <ul style="list-style-type: none"> – PHSS_37202 – PHKL_41481 – PHKL_42035 – PHKL_42335 – PHKL_41588 • HP-UX 11iv4 	Itanium based HP Integrity Series Systems Any hardware that is explicitly compatible and fully capable of running the specified operating system, all the corresponding supporting software, and any associated applications unmodified.

Kernel configuration considerations

A system restart is required if you update the kernel configuration parameters. The kernel configuration parameters are set in */etc/system*. Depending on the values of your kernel configuration parameters, modify some of them before you install the 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

- `libpam.so.0` (32-bit) is required for DB2 database servers to run 32-bit non-SQL routines.
- (Clients only) If you plan to use Kerberos authentication, your DB2 applications must be 64-bit.
- (Client and server) To use Kerberos authentication, install the HP `krb5` Kerberos support via the Kerberos packages on the operating system.
- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if you want to use the DB2 Setup wizard to install a DB2 product.
- For details regarding known HP-UX issues, see www.ibm.com/support/docview.wss?uid=swg21257602

Note: DB2 products installed on the HP-UX operating system support long host names. The length has been extended to 255 bytes, in any combination of characters or digits.

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

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

Mounting considerations

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

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

Before you install DB2 database products on Linux operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements. The `db2prereqcheck` command checks whether your system meets the prerequisites for the installation.

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

If you are planning to install the IBM DB2 pureScale Feature, different installation prerequisites are required. See the Installing the DB2 pureScale Feature topic.

DB2 database products are supported on the following hardware:

- x86 (Intel Pentium, Intel Xeon, and AMD) 32-bit Intel and AMD processors
- x64 (64-bit AMD64 and Intel EM64T processors)

- POWER[®] (System i[®], System p[®], and POWER Systems that support Linux)
- System z[®]: System z196, System z10[®], or System z9[®]

The minimum supported operating systems for Linux include:

- Red Hat Enterprise Linux (RHEL) 5 Update 6
- Red Hat Enterprise Linux (RHEL) 6
- SUSE Linux Enterprise Server (SLES) 10 Service Pack 3
- SUSE Linux Enterprise Server (SLES) 11 Service Pack 1
- Ubuntu 10.04

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

Note:

- Version 9.7 Fix Pack 2 and later fix packs (and Version 9.8 Fix Pack 2 and later fix packs) include an updated version of IBM Tivoli[®] System Automation for Multiplatforms (SA MP) Base Component that you can use in environments with SLES 11 or POWER7 systems. For more information, see “Installing IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component” or “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component”.

Multithreaded architecture limitations

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

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

Distribution Requirements

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

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

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

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

Package requirements

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

- `libpam.so.0` (32-bit) is required for DB2 database servers to run 32-bit non-SQL routines.
- `libaio.so.1` is required for DB2 database servers using asynchronous I/O.
- `libstdc++.so.6` is required for DB2 database servers and clients.

Package requirements for SLES and RHEL

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

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

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

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

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

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

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

Package requirements for SUSE Linux

Package name	Description
<code>pdksh</code> or <code>ksh93</code>	Korn Shell.
<code>openssh</code>	This package contains a set of server programs which allow users to run commands on (and from) remote computers via a secure shell. This package is not required if you use the default configuration of DB2 database systems with rsh .
<code>rsh-server</code>	This package contains a set of server programs which allow users to run commands on remote computers, login in to other computers, and copy files between computers (rsh , rexec , rlogin , and rcp). This package is not required if you configure DB2 database systems to use ssh .

Package requirements for SUSE Linux

Package name	Description
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 or ksh93	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 database systems with rsh .
/System Environment/Daemons	openssh-server	This package contains a set of server programs which allow users to run commands from a remote computer via a Secure Shell. This package is not required if you use the default configuration of DB2 database systems with rsh .
/System Environment/Daemons	rsh-server	This package contains a set of programs which allow users to run commands on a remote computer. Required for partitioned database environments. This package is not required if you configure DB2 database systems to use ssh .
/System Environment/Daemons	nfs-utils	Network File System support package. It allows access to local files from remote computers.

Software considerations

- (Client and Server) To use Kerberos authentication, install the Linux krb5 Kerberos support via the operating system Kerberos packages.
- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if:
 - you want to use the DB2 Setup wizard to install a DB2 database product on Linux or UNIX operating systems, or
 - you want to use any DB2 graphical tools on Linux for x86 and Linux on AMD 64/EM64T.
- Micro Focus does not offer support for any of its COBOL compiler products on SLES 11.

Security-enhanced Linux considerations

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

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

- check the `/etc/sysconfig/selinux` file
- run the `sestatus` command
- check the `/var/log/messages` file for SELinux notices.

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

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

If your DB2 database product installs successfully on a RHEL 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 `sqlib/samples` directory.

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

Before you install a DB2 database product on Solaris, ensure that your system meets the operating system, hardware, and software requirements. The `db2prereqcheck` command checks whether your system meets the prerequisites for the installation.

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

Table 3. Solaris installation requirements

Operating System	Hardware
Solaris 10 <ul style="list-style-type: none">• 64-bit kernel• Update 9	UltraSPARC or SPARC64 processors
Solaris 10 <ul style="list-style-type: none">• 64-bit kernel• Update 9	Solaris x64 (Intel 64 or AMD64)

- Support is only for the DB2 product to be installed on either local zone or global zone. When installed in the global zone, it can be shared by other zones.

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

- `libpam.so.0` (32-bit) is required for DB2 database servers to run 32-bit non-SQL routines.

- (Client and server) To use Kerberos authentication, install the krb5 Kerberos support via the available Kerberos packages.
- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if you want to use the DB2 Setup wizard to install a DB2 database product
- For details regarding known Solaris issues, see www.ibm.com/support/docview.wss?uid=swg21257606

Security patches can be obtained from the <http://java.sun.com> website. From the SunSolve online website, click **Patches** in the left panel.

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

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

DB2 database products support the following Solaris concepts:

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

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

Mounting considerations

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

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

Before you install a DB2 database product on Windows operating systems, ensure that the system you choose meets the necessary operating system, hardware, and software requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

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

Table 4. Windows workstation platforms

Operating System	Prerequisites	Hardware
Windows XP Professional (32-bit and 64-bit)	Windows XP Service Pack 3 or later	All Intel and AMD processors capable of running the supported Windows operating systems (32-bit and 64-bit based systems)
Windows XP Enterprise (32-bit and 64-bit)	Windows Vista Service Pack 2	
Windows XP Ultimate (32-bit and 64-bit)	Windows 7 Service Pack 1	
Windows Vista Business (32-bit and 64-bit)	IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 2.0 or later framework run time.	
Windows Vista Enterprise (32-bit and 64-bit)		
Windows Vista Ultimate (32-bit and 64-bit)		
Windows 7 Professional (32-bit and 64-bit)	64-bit IBM data server provider for .NET applications are supported.	
Windows 7 Enterprise (32-bit and 64-bit)		
Windows 7 Ultimate (32-bit and 64-bit)		

Table 5. Windows server platforms

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

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

Additional software considerations

- Windows Installer 3.0 is required. It is installed by the installer if it is not detected.
- IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 2.0 or later framework runtime. In an x64 environment, 32-bit IBM data server provider for .NET applications runs in the WOW64 emulation mode.
- If you plan to use LDAP (Lightweight Directory Access Protocol), use either a Microsoft LDAP client or the IBM Tivoli Directory Server v6 client (also known as the IBM LDAP client which is included with DB2 database products). Before installation of the Microsoft Active Directory, you must 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 installation launchpad (setup.exe), and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Internet Explorer 7.0 and later
 - Google Chrome
 - Safari 4.0

Chapter 4. IBM Data Server Client support for database application development

The Data Server Client supports application development for databases.

The Data Server Client includes the following features:

- **Precompilers for C/C++, COBOL, and Fortran**, (providing the language is supported for that platform).
- **Embedded SQL application support**, including programming libraries, include files and code samples.
- **ODBC and DB2 Call Level Interface (DB2 CLI) application support**, including programming libraries, include files, and code samples for developing applications which are easily ported to ODBC and compiled with an ODBC SDK. An ODBC SDK is available from Microsoft for Windows operating systems, and from various other vendors for many of the other supported platforms. On Windows operating systems, the ODBC and CLI driver is installed by default with the DB2 Client, supporting applications developed with the Microsoft ODBC Software Developer's Kit. For all other platforms, the ODBC and CLI driver can optionally be installed with the DB2 Client, supporting applications that can be developed with an ODBC SDK for that platform, if one exists.
- **The IBM Data Server Driver for JDBC and SQLJ**, which includes:
 - Support for JDBC 3 and JDBC 4 compliant applications
 - SQLJ support, for developing Java applications with static SQL
- **Java 5 environment** is embedded in DB2 servers to support server-side Java application artifacts including stored procedures and user-defined functions
- **ActiveX Data Objects (ADO) and Object Linking and Embedding (OLE) automation routines (UDFs and Stored Procedures)** on Windows operating systems, including code samples implemented in Microsoft Visual Basic and Microsoft Visual C++.
- **Object Linking and Embedding Database (OLE DB) table functions** on Windows operating systems.
- **C# and Visual Basic .NET applications and CLR .NET routines** on Windows operating systems.
- **Ruby and Ruby on Rails application support** , including IBM_DB gem for easy install and code samples. Ruby is a fully-integrated object-oriented programming language used to develop web applications. Ruby on Rails (RoR), also known as Rails, is a framework for developing web applications that access databases according to the Model-View-Control architectural framework
- **PHP Hypertext Preprocessor application support** , including programming libraries and code samples. PHP is a scripting language used to develop web applications. It also supports object-oriented programming
- **Data Studio** is a comprehensive suite of integrated Eclipse tools for the database developer and development DBA. Data Studio reduces the time to perform day-to-day administration tasks, create, deploy, and debug SQL and Java stored procedures, deploy data-centric web services, and create queries for relational and XML data using SQL and XQuery for DB2 and Informix® data servers.
- **Interactive SQL** through the Command Line Processor (CLP) to prototype SQL statements or to perform ad hoc queries against the database.

- **A set of documented APIs** for other application development tools to implement precompiler support for DB2 directly within their products. For example, IBM COBOL on AIX uses this interface.

Chapter 5. Supported database application programming interfaces

You can use any of the following programming interfaces to manage or access DB2 databases.

You can:

- Use DB2 APIs to perform administrative functions such as backing up and restoring databases.
- Embed static and dynamic SQL statements in your applications programmed in C, C++, COBOL, FORTRAN, and REXX.
- Code DB2 Call Level Interface (DB2 CLI) function calls in your applications to invoke dynamic SQL statements.
- Develop Java applications and applets using the Java Database Connectivity application programming interface (JDBC API) or embedded SQL for Java (SQLJ).
- Develop Microsoft Visual Basic and Visual C++ applications using ActiveX Data Object (ADO).
- Develop ADO.NET applications using IBM Data Server Provider for .NET (DB2 .NET).
- Develop applications using IBM or third-party tools such as Excel, Perl, and Open Database Connectivity (ODBC) end-user tools such as Lotus® Approach®, and its programming language, LotusScript®.
- Develop web applications using PHP scripting language and Ruby on Rails (RoR) which is a framework for developing web applications that access databases.
- Develop web applications using Python, which is a general purpose, high-level scripting language that is well-suited for rapid application development.
- Starting with Version 9.7 Fix Pack 1, use Oracle Call Interface (OCI) applications to access DB2 databases. The IBM Data Server Driver for DB2CI provides support for calls to a number of OCI APIs.

ADO.NET development for IBM Data Servers

The IBM Data Server Provider for .NET is an extension of the ADO.NET interface that allows .NET applications to access an IBM data server through a secure connection, execute commands, and retrieve result sets.

Reference documentation is included with the IBM Data Server Provider for .NET when installing the IBM Data Server Client, presenting detailed information about all the IBM Data Server Provider for .NET objects and their members. During the DB2 installation process, this documentation is registered with Microsoft Visual Studio .NET. To view the IBM Data Server Provider for .NET documentation from Microsoft Visual Studio .NET, select the Help menu option, and Contents. Once the help viewer opens, filter by *IBM Database Documentation*.

IBM Data Server Provider for .NET

The IBM Data Server Provider for .NET extends data server support for the ADO.NET interface. The provider delivers high-performing, secure access to IBM data servers.

Two providers are included in the IBM Data Server Provider for .NET client package. These providers are sometimes referred to as the Common .NET Providers.

The DB2 .NET Provider (IBM.Data.DB2.dll)

With the DB2 .NET Provider your .NET applications can access the following database management systems:

- DB2 Database for Linux, UNIX, and Windows, Version 9.1, Version 9.5, Version 9.7, Version 9.8, and Version 10.1
- DB2 Universal Database™ Version 8 for Windows, UNIX, and Linux
- DB2 for z/OS® Version 8, Version 9, and Version 10 , through DB2 Connect™
- IBM DB2 for IBM i Version 5 Release 4, Version 6 Release 1 and Version 7 Release 1, through DB2 Connect (for IBM DB2 Version 9.7 Fix Pack 4 and higher versions)
- IBM DB2 for IBM i Version 5 Release 4 and Version 6 Release 1, through DB2 Connect (for IBM DB2 Version 9.7 Fix Pack 3 and earlier versions)
- IBM Informix Version 11.10, Version 11.50, and Version 11.70

The remainder of these topics discuss the Common DB2 .NET Provider.

The Informix database server .NET Provider (IBM.Data.Informix.dll)

With the Informix database server .NET Provider your .NET applications can access the following database management systems:

- IBM Informix, Version 11.10 and Version 11.50

For more information about this provider, see the: IBM Informix Dynamic Server Information Center.

To develop and run applications that use Data Server Provider for .NET you need the .NET Framework.

In addition to the IBM Data Server Provider for .NET, the IBM Database Add-Ins for Visual Studio enable you to quickly and easily develop .NET applications for IBM data servers using Microsoft Visual Studio. You can also use the Add-Ins to create database objects such as indexes and tables, and develop server-side objects, such as stored procedures and user-defined functions.

ODBC .NET Data Provider

The ODBC .NET Data Provider makes ODBC calls to a DB2 data source using the CLI Driver. Therefore, the connection string keywords supported by the ODBC .NET Data Provider are the same as those supported by the CLI driver. This provider is no longer tested. Users are recommended to use the IBM Data Server Provider for .NET.

Also, the ODBC .NET Data Provider has the same restrictions as the CLI driver. There are additional restrictions for the ODBC .NET Data Provider, which are identified in the topic: “ODBC .NET Data Provider restrictions” in *Developing ADO.NET and OLE DB Applications*.

In order to use the ODBC .NET Data Provider, you must have the .NET Framework Version 2.0, 3.0, or 3.5 installed. For DB2 Universal Database for AS/400® V5R4 and earlier, the following fix is required on the server: APAR II13348.

The supported connection keywords for the ODBC .NET Data Provider are listed in the table 1:

Table 6. Useful **ConnectionString** keywords for the ODBC .NET Data Provider

Keyword	Value	Meaning
DSN	database alias	The DB2 database alias as cataloged in the database directory
UID	user ID	The user ID used to connect to the DB2 server
PWD	password	The password for the user ID used to connect to the DB2 server

Note: For the full list of **ConnectionString** keywords, see the Microsoft documentation.

The following code is an example of creating an `OdbcConnection` to connect to the SAMPLE database:

```
[Visual Basic .NET]
Dim con As New OdbcConnection("DSN=sample;UID=userid;PWD=password;")
con.Open()

[C#]
OdbcConnection con = new OdbcConnection("DSN=sample;UID=userid;PWD=password;");
con.Open()
```

OLE DB .NET Data Provider

The OLE DB .NET Data Provider uses the IBM DB2 OLE DB Driver, which is referred to in a `ConnectionString` object as `IBMDADB2`.

The connection string keywords supported by the OLE DB .NET Data Provider are the same as those supported by the IBM OLE DB Provider for DB2. This provider is no longer tested. Users are recommended to use the IBM Data Server Provider for .NET.

Also, the OLE DB .NET Data Provider has the same restrictions as the IBM DB2 OLE DB Provider. There are additional restrictions for the OLE DB .NET Data Provider, which are identified in the topic: "OLE DB .NET Data Provider restrictions" in *Developing ADO.NET and OLE DB Applications*.

In order to use the OLE DB .NET Data Provider, you must have the .NET Framework Version 2.0, 3.0, or 3.5 installed.

For DB2 Universal Database for AS/400 R520, R530 and R540, the following fix is required on the server: APAR ii13348.

All the supported connection keywords for the OLE DB .NET Data Provider are shown in table 1:

Table 7. Useful **ConnectionString** keywords for the OLE DB .NET Data Provider

Keyword	Value	Meaning
PROVIDER	IBMDADB2	Specifies the IBM OLE DB Provider for DB2 (required)

Table 7. Useful **ConnectionString** keywords for the OLE DB .NET Data Provider (continued)

Keyword	Value	Meaning
DSN or Data Source	database alias	The DB2 database alias as cataloged in the database directory
UID	user ID	The user ID used to connect to the DB2 data server
PWD	password	The password for the user ID used to connect to the DB2 data server

Note: For the full list of **ConnectionString** keywords, see the Microsoft documentation.

Example of creating an `OleDbConnection` to connect to the SAMPLE database is:

```
[Visual Basic .NET]
Dim con As New OleDbConnection("Provider=IBMDADB2;" +
    "Data Source=sample;UID=userid;PWD=password;")
con.Open()

[C#]
OleDbConnection con = new OleDbConnection("Provider=IBMDADB2;" +
    "Data Source=sample;UID=userid;PWD=password;");
con.Open()
```

Introduction to DB2 Call Level Interface and ODBC

DB2 Call Level Interface (CLI) is IBM's callable SQL interface to the DB2 family of database servers. It is a 'C' and 'C++' application programming interface for relational database access that uses function calls to pass dynamic SQL statements as function arguments.

You can use the CLI interface to access the following IBM data server databases:

- DB2 Version 9 for Linux, UNIX, and Windows
- DB2 Universal Database Version 8 (and later) for OS/390® and z/OS
- DB2 for IBM i 5.4 and later
- IBM Informix Version 11.50 (starting in DB2 Version 9.7 Fix Pack 1 and later), Version 11.70 (starting in DB2 Version 9.7 Fix Pack 3 and later)

CLI is an alternative to embedded dynamic SQL, but unlike embedded SQL, it does not require host variables or a precompiler. Applications can be run against a variety of databases without having to be compiled against each of these databases. Applications use procedure calls at run time to connect to databases, issue SQL statements, and retrieve data and status information.

The CLI interface provides many features not available in embedded SQL. For example:

- CLI provides function calls that support a way of querying database catalogs that is consistent across the DB2 family. This reduces the need to write catalog queries that must be tailored to specific database servers.
- CLI provides the ability to scroll through a cursor:
 - Forward by one or more rows
 - Backward by one or more rows

- Forward from the first row by one or more rows
- Backward from the last row by one or more rows
- From a previously stored location in the cursor.
- Stored procedures called from application programs that were written using CLI can return result sets to those programs.

CLI is based on the Microsoft Open Database Connectivity (ODBC) specification, and the International Standard for SQL/CLI. These specifications were chosen as the basis for the DB2 Call Level Interface in an effort to follow industry standards and to provide a shorter learning curve for those application programmers already familiar with either of these database interfaces. In addition, some DB2 specific extensions have been added to help the application programmer specifically exploit DB2 features.

The CLI driver also acts as an ODBC driver when loaded by an ODBC driver manager. It conforms to ODBC 3.51.

CLI Background information

To understand CLI or any callable SQL interface, it is helpful to understand what it is based on, and to compare it with existing interfaces.

The X/Open Company and the SQL Access Group jointly developed a specification for a callable SQL interface referred to as the *X/Open Call Level Interface*. The goal of this interface is to increase the portability of applications by enabling them to become independent of any one database vendor's programming interface. Most of the X/Open Call Level Interface specification has been accepted as part of the ISO Call Level Interface International Standard (ISO/IEC 9075-3:1995 SQL/CLI).

Microsoft developed a callable SQL interface called Open Database Connectivity (ODBC) for Microsoft operating systems based on a preliminary draft of X/Open CLI.

The ODBC specification also includes an operating environment where database-specific ODBC drivers are dynamically loaded at run time by a driver manager based on the data source (database name) provided on the connect request. The application is linked directly to a single driver manager library rather than to each DBMS's library. The driver manager mediates the application's function calls at run time and ensures they are directed to the appropriate DBMS-specific ODBC driver. Because the ODBC driver manager only knows about the ODBC-specific functions, DBMS-specific functions cannot be accessed in an ODBC environment. DBMS-specific dynamic SQL statements are supported through a mechanism called an escape clause.

ODBC is not limited to Microsoft operating systems; other implementations are available on various platforms.

The CLI load library can be loaded as an ODBC driver by an ODBC driver manager. For ODBC application development, you must obtain an ODBC Software Development Kit. For the Windows platform, the ODBC SDK is available as part of the Microsoft Data Access Components (MDAC) SDK, available for download from <http://www.microsoft.com/downloads>. For non-Windows platforms, the ODBC SDK is provided by other vendors. When developing ODBC applications that may connect to DB2 servers, use the Call Level Interface Guide and Reference Volume 1 and the Call Level Interface Guide and Reference Volume 2 (for information about

DB2 specific extensions and diagnostic information), in conjunction with the ODBC Programmer's Reference and SDK Guide available from Microsoft.

Applications written using CLI APIs link directly to the CLI library. CLI includes support for many ODBC and ISO SQL/CLI functions, as well as DB2 specific functions.

The following DB2 features are available to both ODBC and CLI applications:

- double byte (graphic) data types
- stored procedures
- Distributed Unit of Work (DUOW), two phase commit
- compound SQL
- user defined types (UDT)
- user defined functions (UDF)

Introduction to embedded SQL

Embedded SQL database applications connect to databases and execute embedded SQL statements. Embedded SQL statements are embedded within a host language application. Embedded SQL database applications support the embedding of SQL statements to be executed statically or dynamically.

You can develop embedded SQL applications for DB2 in the following host programming languages: C, C++, COBOL, FORTRAN, and REXX.

Note: Support for embedded SQL in FORTRAN and REXX has been deprecated and will remain at the DB2 Universal Database , Version 5.2 level.

Building embedded SQL applications involves two prerequisite steps before application compilation and linking.

- Preparing the source files containing embedded SQL statements using the DB2 precompiler.

The PREP (PRECOMPILE) command is used to invoke the DB2 precompiler, which reads your source code, parses and converts the embedded SQL statements to DB2 runtime services API calls, and finally writes the output to a new modified source file. The precompiler produces access plans for the SQL statements, which are stored together as a package within the database.

- Binding the statements in the application to the target database.

Binding is done by default during precompilation (the PREP command). If binding is to be deferred (for example, running the BIND command later), then the BINDFILE option needs to be specified at PREP time in order for a bind file to be generated.

Once you have precompiled and bound your embedded SQL application, it is ready to be compiled and linked using the host language-specific development tools.

To aid in the development of embedded SQL applications, you can refer to the embedded SQL template in C. Examples of working embedded SQL sample applications can also be found in the %DB2PATH%\SQLLIB\samples directory.

Note: %DB2PATH% refers to the DB2 installation directory

Static and dynamic SQL

SQL statements can be executed in one of two ways: statically or dynamically.

Statically executed SQL statements

For statically executed SQL statements, the syntax is fully known at precompile time. The structure of an SQL statement must be completely specified for a statement to be considered static. For example, the names for the columns and tables referenced in a statement must be fully known at precompile time. The only information that can be specified at run time are values for any host variables referenced by the statement. However, host variable information, such as data types, must still be precompiled. You precompile, bind, and compile statically executed SQL statements before you run your application. Static SQL is best used on databases whose statistics do not change a great deal.

Dynamically executed SQL statements

Dynamically executed SQL statements are built and executed by an application at run time. An interactive application that prompts the end user for key parts of an SQL statement, such as the names of the tables and columns to be searched, is a good example of a situation suited for dynamic SQL.

Supported drivers for JDBC and SQLJ

The DB2 product includes support for two types of JDBC driver architecture.

According to the JDBC specification, there are four types of JDBC driver architectures:

Type 1

Drivers that implement the JDBC API as a mapping to another data access API, such as Open Database Connectivity (ODBC). Drivers of this type are generally dependent on a native library that limits their portability. The DB2 database system does not provide a type 1 driver.

Type 2

Drivers that are written partly in the Java programming language and partly in native code. The drivers use a native client library specific to the data source to which they connect. Because of the native code, their portability is limited.

Type 3

Drivers that use a pure Java client and communicate with a data server using a data-server-independent protocol. The data server then communicates the client's requests to the data source. The DB2 database system does not provide a type 3 driver.

Type 4

Drivers that are pure Java and implement the network protocol for a specific data source. The client connects directly to the data source.

DB2 Database for Linux, UNIX, and Windows supports the following driver:

Driver name	Packaged as	Driver type
IBM Data Server Driver for JDBC and SQLJ	<ul style="list-style-type: none">db2jcc.jar and sqlj.zip for JDBC 3.0 supportdb2jcc4.jar and sqlj4.zip for support of some JDBC 4.0 or later functions	Type 2 and Type 4

IBM Data Server Driver for JDBC and SQLJ (type 2 and type 4)

The IBM Data Server Driver for JDBC and SQLJ is a single driver that includes JDBC type 2 and JDBC type 4 behavior. When an application loads the IBM Data Server Driver for JDBC and SQLJ, a single driver instance is loaded for type 2 and type 4 implementations. The application can make type 2 and type 4 connections using this single driver instance. The type 2 and type 4 connections can be made concurrently. IBM Data Server Driver for JDBC and SQLJ type 2 driver behavior is referred to as *IBM Data Server Driver for JDBC and SQLJ type 2 connectivity*. IBM Data Server Driver for JDBC and SQLJ type 4 driver behavior is referred to as *IBM Data Server Driver for JDBC and SQLJ type 4 connectivity*.

Two versions of the IBM Data Server Driver for JDBC and SQLJ are available. IBM Data Server Driver for JDBC and SQLJ version 3.5x is JDBC 3.0-compliant. IBM Data Server Driver for JDBC and SQLJ version 4.x is compliant with JDBC 4.0 or later.

The IBM Data Server Driver for JDBC and SQLJ supports these JDBC and SQLJ functions:

- Version 3.5x supports all of the methods that are described in the JDBC 3.0 specifications.
- Version 4.x supports all of the methods that are described in the JDBC 4.0 or later specifications.
- SQLJ application programming interfaces, as defined by the SQLJ standards, for simplified data access from Java applications.
- Connections that are enabled for connection pooling. WebSphere® Application Server or another application server does the connection pooling.
- Connections to a data server from Java user-defined functions and stored procedures use IBM Data Server Driver for JDBC and SQLJ type 2 connectivity only. Applications that call user-defined functions or stored procedures can use IBM Data Server Driver for JDBC and SQLJ type 2 connectivity or IBM Data Server Driver for JDBC and SQLJ type 4 connectivity to connect to a data server. The IBM Data Server Driver for JDBC and SQLJ is the default driver for Java routines.
- Support for distributed transaction management. This support implements the Java 2 Platform, Enterprise Edition (J2EE) Java Transaction Service (JTS) and Java Transaction API (JTA) specifications, which conform to the X/Open standard for distributed transactions (*Distributed Transaction Processing: The XA Specification*, available from <http://www.opengroup.org>).

IBM OLE DB Provider for DB2

The IBM OLE DB Provider for DB2 allows DB2 to act as a resource manager for the OLE DB provider. This support gives OLE DB-based applications the ability to extract or query DB2 data using the OLE interface.

Microsoft OLE DB is a set of OLE/COM interfaces that provides applications with uniform access to data stored in diverse information sources. The OLE DB architecture defines OLE DB consumers and OLE DB providers. An OLE DB consumer is any system or application that uses OLE DB interfaces; an OLE DB provider is a component that exposes OLE DB interfaces.

The IBM OLE DB Provider for DB2, whose provider name is IBM DADB2, enables OLE DB consumers to access data on a DB2 database server. If DB2 Connect is

installed, these OLE DB consumers can also access data on a host DBMS such as DB2 for z/OS, DB2 Server for VM and VSE, or DB2 Universal Database for AS/400.

The IBM OLE DB Provider for DB2 offers the following features:

- Support level 0 of the OLE DB provider specification, including some additional level 1 interfaces.
- A free threaded provider implementation, which enables the application to create components in one thread and use those components in any other thread.
- An Error Lookup Service that returns DB2 error messages.

Note that the IBM OLE DB Provider resides on the client and is different from the OLE DB table functions, which are also supported by DB2 database systems.

Subsequent sections of this document describe the specific implementation of the IBM OLE DB Provider for DB2. For more information about the Microsoft OLE DB 2.0 specification, refer to the Microsoft OLE DB 2.0 Programmer's Reference and Data Access SDK, available from Microsoft Press.

Version Compliance

The IBM OLE DB Provider for DB2 complies with Version 2.7 or later of the Microsoft OLE DB specification.

System Requirements

Refer to the announcement letter for the IBM OLE DB Provider for DB2 data servers to see the supported Windows operating systems.

To install the IBM OLE DB Provider for DB2, you must first be running on one of the supported operating systems listed previously. You also need to install a full DB2 product, IBM Data Server Driver for ODBC and CLI, or IBM Data Server Driver Package.

Perl DBI

DB2 supports the Perl Database Interface (DBI) specification for data access through the DBD::DB2 driver.

The DB2 Perl DBI website is located at:

<http://www.ibm.com/software/data/db2/perl/>

and contains the latest DBD::DB2 driver, and related information.

Perl is an interpreted language and the Perl DBI Module uses dynamic SQL. These properties make Perl an ideal language for quickly creating and revising DB2 applications. The Perl DBI Module uses an interface that is quite similar to the CLI and JDBC interfaces, which makes it easy to port Perl applications to CLI and JDBC, or to port CLI and JDBC applications to Perl.

PHP application development for IBM data servers

PHP: Hypertext Preprocessor (PHP) is an interpreted programming language that is widely used for developing web applications. PHP has become a popular language for web development because it is easy to learn, focuses on practical solutions, and supports the most commonly required functionality in web applications.

PHP is a modular language that enables you to customize the available functionality through the use of extensions. These extensions can simplify tasks such as reading, writing, and manipulating XML, creating SOAP clients and servers, and encrypting communications between server and browser. The most popular extensions for PHP, however, provide read and write access to databases so that you can easily create a dynamic database-driven website.

IBM provides the listed PHP extensions for accessing IBM data server databases:

ibm_db2

A procedural application programming interface (API) that, in addition to the normal create, read, update, and write database operations, also offers extensive access to the database metadata. You can compile the `ibm_db2` extension with either PHP 4 or PHP 5. This extension is written, maintained, and supported by IBM.

pdo_ibm

A driver for the PHP Data Objects (PDO) extension that offers access to IBM data server databases through the standard object-oriented database interface introduced in PHP 5.1.

These extensions are included as part of the IBM Data Server Driver Package (DS Driver) of Version 1.7.0. This version or a later version is supported to connect to IBM DB2 Version 9.7 for Linux, UNIX, and Windows. You can check the version of `ibm_db2` extension by issuing the `php --re ibm_db2` command.

The most recent versions of `ibm_db2` and `pdo_ibm` are also available from the PHP Extension Community Library (PECL) at <http://pecl.php.net/>.

PHP applications can access the listed IBM data server databases:

- IBM DB2 Version 9.1 for Linux, UNIX, and Windows, Fix Pack 2 and later
- IBM DB2 Universal Database (DB2 UDB) Version 8 for Linux, UNIX, and Windows, Fixpak 15 and later
- Remote connections to IBM DB2 for IBM i V5R3
- Remote connections to IBM DB2 for IBM i Version 5.4 and later
- Remote connections to IBM DB2 for z/OS, Version 8 and later

A third extension, Unified ODBC, has historically offered access to DB2 database systems. For new applications, however, you can use either `ibm_db2` and `pdo_ibm` because they offer significant performance and stability benefits over Unified ODBC. The `ibm_db2` extension API makes porting an application that was previously written for Unified ODBC almost as easy as globally changing the `odbc_` function name to `db2_` throughout the source code of your application.

Python, SQLAlchemy and Django Framework application development for IBM data servers

Python is a general purpose, high level scripting language that is well suited for rapid application development. Python emphasizes code readability and supports a variety of programming paradigms, including procedural, object-oriented, aspect-oriented, meta, and functional programming. The Python language is managed by the Python Software Foundation.

The listed extensions are available for accessing IBM data server databases from a Python application:

ibm_db

This API is defined by IBM and provides the best support for advanced features. In addition to issuing SQL queries, calling stored procedures, and using pureXML[®], you can access metadata information.

ibm_db_dbi

This API implements Python Database API Specification v2.0. Because the `ibm_db_dbi` API conforms to the specification, it does not offer some of the advanced features that the `ibm_db` API supports. If you have an application with a driver that supports Python Database API Specification v2.0, you can easily switch to `ibm_db`. The `ibm_db` and `ibm_db_dbi` APIs are packaged together.

ibm_db_sa

This adaptor supports SQLAlchemy, which offers a flexible way to access IBM data servers. SQLAlchemy is a popular open source Python SQL toolkit and object-to-relational mapper (ORM).

ibm_db_django

This adaptor provides access to IBM data servers from Django. Django is a popular web framework used to build high-performing, elegant web applications quickly.

Python applications can access the listed IBM data servers:

- IBM DB2 Version 9.1 for Linux, UNIX, and Windows, Fix Pack 2 and later
- IBM DB2 Universal Database (DB2 UDB) Version 8 for Linux, UNIX, and Windows, Fixpak 15 and later
- Remote connections to IBM DB2 for IBM i V5R3, with PTF SI27358 (includes SI27250)
- Remote connections to IBM DB2 for IBM i 5.4 and later, with PTF SI27256
- Remote connections to IBM DB2 for z/OS, Version 8 and Version 9
- IBM Informix Dynamic Server v11.10 and later

The IBM_DB Ruby driver and Rails adapter

With the introduction of support for the Ruby on Rails framework, Rails applications can now access data on IBM data servers.

Collectively known as the `IBM_DB` gem, the `IBM_DB` Ruby driver and Rails adapter allows Ruby applications to access the listed database management systems:

- DB2 for Linux, UNIX, and Windows Version 9 and later
- DB2 Universal Database (DB2 UDB) Version 8 for Linux, UNIX, and Windows

- DB2 UDB Version 5, Release 1 (and later) for AS/400 and iSeries®, through DB2 Connect
- DB2 for z/OS, Version 8 and Version 9, through DB2 Connect
- Informix Dynamic Server, Version 11.10 and later

Note: Client applications must use IBM Data Server Driver Version 9.5 or later when accessing Informix Dynamic Server Version 11.10. Previous versions are not supported. Client applications can also use IBM Data Server Runtime Client or IBM Data Server Client.

The IBM_DB Ruby driver can be used to connect to and access data from the IBM data servers mentioned previously. The IBM_DB Ruby adapter allows any database-backed Rails application to interface with IBM data servers.

For more information about IBM Ruby projects and the RubyForge open source community, see <http://rubyforge.org/projects/rubyibm/>

For a list of installation requirements for DB2 database products, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.qb.server.doc/doc/r0025127.html>

For a list of installation requirements for IBM Informix Dynamic Server, see http://publib.boulder.ibm.com/infocenter/idshelp/v111/topic/com.ibm.expr.doc/ids_in_004x.html

For information about downloading an IBM Data Server Driver Package (DS Driver), see <http://www.ibm.com/software/data/support/data-server-clients/index.html>.

Chapter 6. Programming languages and compilers for database application development

Application development for DB2 databases is supported on a range of programming languages and compilers.

You can develop DB2 database applications in the following languages:

- C++
- C#, VB .NET, and other .NET languages
- C
- COBOL
- Fortran
- Java
- Perl
- PHP
- Python
- REXX
- Ruby and Ruby on Rails

The APIs and programming languages for developing external routines are generally compatible with those for database application development. There are some minor differences in prerequisites and support. To review the supported APIs and programming language for external routine development see: *Supported APIs and programming languages for external routine development*.

Supported .NET development software

To develop and deploy .NET applications that run against IBM data servers, you will need to use supported development software and operating systems.

Supported operating systems for developing and deploying .NET Framework 2.0, 3.0, 3.5 and 4.0 applications

- Windows XP, Service Pack 2 (32-bit and 64-bit editions)
- Windows Server 2003 (32-bit and 64-bit editions)
- Windows Vista (32-bit and 64-bit editions)
- Windows Server 2008 (32-bit and 64-bit editions)
- Windows Server 2008 R2 (64-bit edition)
- Windows 7 (32-bit and 64-bit editions)

Supported development software for .NET Framework applications

In addition to an IBM data server client or Driver Package, you will need one of the following supported tools to develop .NET Framework applications.

- Visual Studio 2008
- Visual Studio 2010

Supported deployment software for .NET Framework applications

In addition to an IBM data server client or Driver Package, you will need one of the following supported packages to deploy .NET Framework applications. In most cases one of these are included with a Windows installation.

- .NET Framework Version 2.0 Redistributable Package
- .NET Framework Version 3.0 Redistributable Package
- .NET Framework Version 3.5 Redistributable Package
- .NET Framework Version 4.0 Redistributable Package

Support for database application development in C

Application development for DB2 databases is supported on a range of C language compilers.

The following table lists C compilers that are supported for DB2 database application development.

Unless otherwise noted, the compilers listed in the following table support both 32-bit and 64-bit application development.

Table 8. Supported C compilers

Operating System	Supported Compiler(s)
AIX	<ul style="list-style-type: none">• IBM XL C/C++ Enterprise Edition Version 9.0 for AIX• IBM XL C/C++ Enterprise Edition Version 10.1 for AIX• IBM XL C/C++ Enterprise Edition Version 11.0 for AIX
HP-UX on IPF (IA64)	<ul style="list-style-type: none">• For HP-UX 11i v3, use HP aC++ Compiler Version A.06.12
Linux on x86	<ul style="list-style-type: none">• GNU/Linux gcc versions 3.3 and 3.4• GNU/Linux gcc version 4.1• Intel C Compiler Version 10.1
Linux on x64	<ul style="list-style-type: none">• GNU/Linux gcc versions 3.3 and 3.4• GNU/Linux gcc version 4.1• Intel C Compiler Version 10.1
Linux on POWER (PowerPC®, Power Systems™ Servers, pSeries®)	<ul style="list-style-type: none">• GNU/Linux gcc versions 3.3 and 3.4• GNU/Linux gcc version 4.1• IBM XL C/C++ Advanced Edition Version 9.0 for Linux• IBM XL C/C++ Advanced Edition Version 10.1 for Linux• IBM XL C/C++ Advanced Edition Version 11.0 for Linux
Linux on System z (s/390x)	<ul style="list-style-type: none">• GNU/Linux gcc versions 3.3 and 3.4• GNU/Linux gcc version 4.1
Solaris on UltraSPARC	<ul style="list-style-type: none">• Sun Studio 10• Sun Studio 11• Sun Studio 12• Forte C++ Version 6.2• Sun ONE Studio 7, Compiler Collection

Table 8. Supported C compilers (continued)

Operating System	Supported Compiler(s)
Solaris on x64	<ul style="list-style-type: none"> • Sun Studio 10 • Sun Studio 11 • Sun Studio 12
Windows on x86	<ul style="list-style-type: none"> • Intel Proton Compiler for Windows 32-bit applications, Version 9.0.021 or later • Microsoft Visual C++ .NET or later
Windows on x64 (x64, AMD64/EM64T)	<ul style="list-style-type: none"> • Microsoft Platform SDK for Windows Server 2003 SP1 (Build 1830 or later), C/C++ compiler for x64 • Intel Proton Compiler for Windows x64, Version 9.0.024 or later • Microsoft Visual C++ 2008, 2010

Support for database application development in C++

Application development for DB2 databases is supported on a range of C++ language compilers.

The following table lists C++ compilers that are supported for DB2 database application development.

Unless otherwise noted, the compilers listed in the following table support both 32-bit and 64-bit application development.

Table 9. Supported C++ compilers

Operating System	Supported Compiler(s)
AIX	<ul style="list-style-type: none"> • IBM XL C/C++ Enterprise Edition Version 9.0 for AIX • IBM XL C/C++ Enterprise Edition Version 10.1 for AIX • IBM XL C/C++ Enterprise Edition Version 11.0 for AIX
HP-UX on IPF (IA64)	<ul style="list-style-type: none"> • For HP-UX 11i v3, use HP aC++ Compiler Version A.06.12
Linux on x86	<ul style="list-style-type: none"> • GNU/Linux g++ version 4.1.1 and 4.1.0 on RHEL5 and SLES10 • Intel C++ Compiler Version 9.1 and later <p>Note: If you experience problems with the g++ version 4.1.1 compiler on RHEL5, revert back to the g++ version 3.4 compiler which is available as a compat package on RHEL 5.</p>
Linux on x64	<ul style="list-style-type: none"> • GNU/Linux g++ versions 4.1.1 and 4.1.0 on RHEL5 and SLES10 <p>Note: If you experience problems with the g++ version 4.1.1 compiler on RHEL5, revert back to the g++ version 3.4 compiler which is available as a compat package on RHEL 5.</p>

Table 9. Supported C++ compilers (continued)

Operating System	Supported Compiler(s)
Linux on POWER (PowerPC, Power Systems Servers, pSeries)	<ul style="list-style-type: none"> • GNU/Linux g++ version 4.1.1 and 4.1.0 on RHEL5 and SLES10 • IBM XL C/C++ Enterprise Edition Version 9.0 for Linux on pSeries <p>Note: If you experience problems with the g++ version 4.1.1 compiler on RHEL 5, revert back to the g++ version 3.4 compiler which is available as a compat package on RHEL 5.</p>
Linux on zSeries® (s/390x)	<ul style="list-style-type: none"> • GNU/Linux g++ versions 3.4 and 3.3 on SLES9 • GNU/Linux g++ version 4.1.1 and 4.1.0 on RHEL5 and SLES10 <p>Note: If you experience problems with the g++ version 4.1.1 compiler on RHEL5, revert back to the g++ version 3.4 compiler which is available as a compat package on RHEL 5.</p>
Solaris on UltraSPARC	<ul style="list-style-type: none"> • Sun Studio 10 • Sun Studio 11 • Sun Studio 12 • Forte C++ Version 6.2 • Sun ONE Studio 7, Compiler Collection
Solaris on x64	<ul style="list-style-type: none"> • Sun Studio 10 • Sun Studio 11 • Sun Studio 12
Windows on x86	<ul style="list-style-type: none"> • Intel Proton Compiler for Windows 32-bit applications, Version 9.0.021 or later • Microsoft Visual C++ .NET
Windows on x64	<ul style="list-style-type: none"> • Microsoft Windows Server 2003 SP1 SDK (Build 1830 or later), C/C++ compiler for x64 • Intel Proton Compiler for Windows x64, Version 9.0.024 or later • Microsoft Visual C++ 2008, 2010

Support for database application development in COBOL

This reference table outlines the current level of compiler support provided by DB2 for application development with the COBOL programming language.

The following table lists COBOL compilers that are supported for DB2 database application development.

Unless otherwise noted, the compilers listed in the following table support both 32-bit and 64-bit application development.

Table 10. Supported COBOL compilers

Operating system	Supported compilers
AIX	<ul style="list-style-type: none"> • IBM COBOL for AIX Version 3.1 • IBM COBOL for AIX Version 4.1 • Micro Focus Server Express®, Version 5.1 • Micro Focus Server Express, Version 6.0
HP-UX on IPF (IA64)	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0
Linux on x86	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0
Linux on x64	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0
Linux on zSeries (s/390x)	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0
Solaris on UltraSPARC	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0
Windows on x86	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0 • IBM Rational® Developer for zSeries v7.6.0
Windows on x64	<ul style="list-style-type: none"> • Micro Focus Server Express, Version 5.1 • Micro Focus Server Express, Version 6.0

Support for database application development in Fortran

Application development for DB2 databases is supported on a range of Fortran language compilers.

The following table lists Fortran compilers that are supported for DB2 database application development.

Unless otherwise noted, the compilers listed in the following table support both 32-bit and 64-bit application development.

Table 11. Supported Fortran compilers

Operating System	Supported Compiler(s)
AIX	<ul style="list-style-type: none"> • IBM XL Fortran for AIX Version 12.1
HP-UX on IPF (IA64)	<ul style="list-style-type: none"> • HP-UX F90 B.11.23
Solaris on UltraSPARC	<ul style="list-style-type: none"> • SPARCCompiler Fortran Versions 4.2 and 5.0 (32-bit only)

Java software support for DB2 database products

You require the appropriate level of IBM Software Development Kit (SDK) for Java to use Java-based tools and to create and run Java applications, including stored procedures and user-defined functions.

If the IBM SDK for Java is required by a component being installed and the SDK for Java is not already installed in that path, the SDK for Java will be installed if you use either the DB2 Setup wizard or a response file to install the product.

The SDK for Java is not installed with IBM Data Server Runtime Client or IBM Data Server Driver Package.

The following table lists the installed SDK for Java levels for DB2 database products according to operating system platform:

Operating System Platform	SDK for Java level
AIX	SDK 7
HP-UX for Itanium-based systems	SDK 6
Linux on x86	SDK 7
Linux on AMD64/EM64T	SDK 7
Linux on zSeries	SDK 7
Linux on POWER	SDK7
Solaris Operating System	SDK 7
Windows x86	SDK 7
Windows x64	SDK 7

Note:

1. The SDK for Java software can be downloaded from the developerWorks® Web page at: <http://www.ibm.com/developerworks/java/jdk/index.html> . For a list of the supported levels of the SDK for Java, see the table later in this section entitled DB2 Database for Linux, UNIX, and Windows support for SDKs for Java.

Note: For Windows operating system platforms, use the IBM Development Package for Eclipse downloads.

2. DB2 GUI tools only run on Linux on x86, Linux on AMD64/EM64T, Windows x86, and Windows x64.
3. On Windows x86 and Linux on x86:
 - the 32-bit SDK is installed
 - 32-bit applications and Java external routines are supported
4. On all supported platforms (except Windows x86, and Linux on x86):
 - 32-bit applications are supported
 - 32-bit Java external routines are not supported
 - 64-bit applications and Java external routines are supported

Supported Java application development software

The following table lists the supported levels of the SDK for Java. The listed levels and forward-compatible later versions of the same levels are supported.

Because there are frequent SDK for Java fixes and updates, not all levels and versions have been tested. If your database application has problems that are related to the SDK for Java, try the next available version of your SDK for Java at the given level.

Versions of SDK for Java, other than IBM SDK, are supported only for building and running stand-alone Java applications. For building and running new Java stored procedures and user-defined functions, only the IBM SDK for Java that is included with the DB2 Database for Linux, UNIX, and Windows product is supported. For running Java stored procedures and user-defined functions that were built by prior DB2 releases, refer to Table 1, column “Java Stored Procedures and User Defined Functions” for details.

Table 12. DB2 Database for Linux, UNIX, and Windows supported levels of SDKs for Java

	Java applications using JDBC driver db2java.zip or db2jcc.jar	Java applications using JDBC driver db2jcc4.jar	Java Stored Procedures and User Defined Functions	DB2 Graphical Tools
AIX	1.4.2 to 7	7	1.4.2 ⁶ to 7 ⁵	N/A
HP-UX for Itanium-based systems	1.4.2 to 6 ¹	6 ¹	1.4.2 ⁶ to 6	N/A
Linux on POWER	1.4.2 to 7 ^{3,4}	7 ^{3,4}	1.4.2 ⁶ to 7	N/A
Linux on x86	1.4.2 to 7 ^{2,3,4}	7 ^{2,3,4}	1.4.2 ⁶ to 7	5 to 7
Linux on AMD64 and Intel EM64T processors	1.4.2 to 7 ^{2,3,4}	7 ^{2,3,4}	1.4.2 ⁶ to 7	N/A
Linux on zSeries	1.4.2 to 7 ^{3,4}	7 ^{3,4}	1.4.2 ⁶ to 7	N/A
Solaris operating system	1.4.2 to 7 ²	7 ²	1.4.2 ⁶ to 7	N/A
Windows on x86	1.4.2 to 7 ²	7 ²	1.4.2 ⁶ to 7	5 to 7
Windows on x64, for AMD64 and Intel EM64T processors	1.4.2 to 7 ²	7 ²	1.4.2 ⁶ to 7	5 to 7

Note:

1. The same levels of the SDK for Java that are available from Hewlett-Packard are supported for building and running stand-alone client applications that run under the IBM Data Server Driver for JDBC and SQLJ.
2. The same levels of the SDK for Java that are available from Oracle are supported for building and running stand-alone applications with the IBM Data Server Driver for JDBC and SQLJ. However, if you set the IBM Data Server Driver for JDBC and SQLJ property securityMechanism for a type of security that uses encryption, the SDK for Java must support the type of encryption that you use. For example, the SDK for Java that you use might support 256-bit AES (strong) encryption, but not 56-bit DES (weak) encryption. You can specify the encryption algorithm by setting the IBM Data Server Driver for JDBC and SQLJ property encryptionAlgorithm. To use 256-bit AES encryption, set encryptionAlgorithm to 2. When you use 256-bit AES encryption with the SDK for Java from Oracle, you might need to install the JCE Unlimited Strength Jurisdiction Policy File, which is available from Oracle.
3. A minimum level of SDK for Java 1.4.2 SR6 is required for SUSE Linux Enterprise Server (SLES) 10. A minimum level of SDK for Java 1.4.2 SR7 is required for Red Hat Enterprise Linux (RHEL) 5.
4. SDK for Java 6 support on Linux requires SDK for Java 6 SR3 or later.
5. If SDK for Java 6 SR2 or later is used, set DB2LIBPATH=*java_home/jre/lib/ppc64*.

6. Support for Java stored procedures and user-defined functions built by IBM SDK for Java 1.4.2 was deprecated in Version 9.7 and might be removed in a future release. IBM SDK for Java 1.4.2 has an End of Service date of September 2011. It is recommended to remove SDK for Java 1.4.2 dependency well before this date. Removing this dependency can be done by rebuilding Java stored procedures and user-defined functions with the SDK for Java included in DB2 Version 9.1, DB2 Version 9.5, DB2 Version 9.7 or DB2 Version 10.1 .

Support for database application development in Perl

Application development for DB2 databases is supported for a range of Perl language versions.

The following table lists versions of the Perl programming language that are supported for DB2 database application development.

Unless otherwise noted, the versions listed in the following table support both 32-bit and 64-bit application development.

Table 13. Supported versions of the Perl programming language

Operating System	Supported Version(s)
AIX	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
HP-UX on IPF (IA64)	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Linux on x86	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Linux on x64	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Linux on POWER (PowerPC, Power Systems Servers, pSeries)	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Linux on zSeries (s/390x)	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Solaris on UltraSPARC	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Solaris on x64	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Windows on x86	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later
Windows on x64	<ul style="list-style-type: none"> • Perl 5.8 or later • DBI 1.41 or later

Support for database application development in PHP

Application development for DB2 databases is supported for a range of PHP language versions.

The following table lists versions of the PHP scripting language that are supported for DB2 database application development.

Though the base support is listed in the following table, using the latest version of PHP 5 is recommended.

Unless otherwise noted, the version listed in the following table support both 32-bit and 64-bit application development.

Table 14. Supported versions of the PHP scripting language

Operating System	Supported Version(s)
AIX	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
HP-UX on IPF (IA64)	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Linux on x86 (32-bit DB2 instances and language only)	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Linux on AMD64/EM64T	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Linux on POWER (PowerPC, System i, pSeries)	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Linux on zSeries (s/390x)	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Solaris on UltraSPARC	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Solaris on x64	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Windows on x86 (32-bit DB2 instances only)	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver
Windows on x64 (AMD64/EM64T)	<ul style="list-style-type: none"> • PHP 4.3.4 or later for IBM_DB2 driver • PHP 5.1 or later for PDO_IBM driver

Support for database application development in Ruby/Ruby on Rails

Application development for DB2 databases is supported for a range of Ruby and Ruby on Rails language versions.

The following table lists versions of Ruby/Ruby on Rails that are supported for DB2 database application development.

Unless otherwise noted, the version listed in the following table support both 32-bit and 64-bit application development.

Table 15. Supported versions for Ruby/Ruby on Rails

Operating System	Supported Versions
AIX	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.

Table 15. Supported versions for Ruby/Ruby on Rails (continued)

Operating System	Supported Versions
HP-UX on IPF (IA64)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Linux on x86 (32-bit DB2 instances and language only)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Linux on AMD64/EM64T (x64)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Linux on POWER (PowerPC, System i, pSeries)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Linux on zSeries (s/390x)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Solaris on UltraSPARC	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Solaris on x64	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Windows on x86 (32-bit application development only)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.
Windows on x64 (AMD64/EM64T, 32-bit application development only)	<ul style="list-style-type: none"> • Ruby 1.8.5 and later • Ruby on Rails 1.2.1 and later.

Support for database application development in REXX

Application development for DB2 databases is supported for a range of REXX programming language versions.

The following table lists versions of the REXX programming language that are supported for DB2 database application development.

Only 32-bit versions of REXX are supported.

Table 16. Supported versions of the REXX programming language

Operating System	Supported Version(s)
AIX ¹	<ul style="list-style-type: none"> • REXXSAA 4.00
Linux on x86 (32-bit DB2 instances and compilers only) ¹	<ul style="list-style-type: none"> • Object REXX Interpreter for Linux Version 2.1
Linux on zSeries (s/390x) ¹	<ul style="list-style-type: none"> • Object REXX 2.2.0 for Linux/390

- ¹Only 32-bit versions of REXX are supported.

Chapter 7. Supported transaction managers

A DB2 database can serve as a resource manager for any XA-compliant transaction manager.

Here is a list of some common XA-compliant transaction managers that can be used to coordinate distributed transactions involving DB2 databases:

- IBM TXSeries CICS®
- IBMWebSphere MQ
- BEA Tuxedo
- BEA WebLogic
- Microsoft Distributed Transaction Coordinator (DTC)

Chapter 8. Supported database application development tools

These integrated development environments and development tools facilitate DB2 database application development.

IBM integrated database application development environments

- IBM Data Studio (complimentary IBM component)
- Rational Application Developer (separately available IBM product)
- Rational Software Architect (separately available IBM product)
- InfoSphere® Data Architect (separately available IBM product)

Database application plug-ins for integrated development environments

- IBM Database Add-Ins for Visual Studio

Tools for developing SQL, SQL/XML, and XQuery statements

- Command line processor (CLP)
- IBM Data Studio

Tools for database application monitoring and performance tuning

- Snapshot and Event Monitors
- Static Profiling for CLI, ODBC, and JDBC applications
- IBM Data Studio (SQL PL Profiling)

Tools for developing web applications

- IBM Data Studio - flexible web services data access (no programming required)
- Rational Application Developer - comprehensive Java development environment for web applications
- Zend Studio - PHP development and deployment environment for web applications
- Ruby on Rails - rapid development of web applications using the Ruby on Rails framework and the new Rails Adapter to access data from IBM data servers
- Microsoft Visual Studio with IBM Database Add-Ins - development environment for .NET based web applications

IBM Database Enterprise Developer Edition

The IBM Database Enterprise Developer Edition product enables an application developer to design, build, and prototype applications for deployment on any of the IBM Information Management client or server platforms. Using the software that comes with this product, you can develop and test applications that run on one operating system and access databases on the same or on a different operating system. For example, you can create an application that runs on a Windows operating system but accesses a database on a UNIX operating system such as AIX.

Chapter 9. Configuring the database application development environment

To configure the database application development environment, you must configure the operating system, Data Server Client, development environment, and transaction manager.

Before you begin

To configure for DB2 database application development, you need (if necessary) :

- A supported operating system
See Chapter 3, “Supported operating systems for database application development,” on page 5.
- The IBM Data Server Client installed
See “Methods for installing IBM data server clients” in *Installing IBM Data Server Clients*.
- API drivers and, if required, driver managers installed
See Chapter 5, “Supported database application programming interfaces,” on page 19.
- Compilers or interpreters required for the programming languages you will be using installed
See Chapter 6, “Programming languages and compilers for database application development,” on page 31.
- A transaction manager installed
See Chapter 7, “Supported transaction managers,” on page 41.
- Development tools installed
See Chapter 8, “Supported database application development tools,” on page 43.

About this task

The DB2 database application development environment is composed of several software elements:

- operating system
- Data Server Client
- database application programming interface (API)
- programming language
- transaction manager
- development tools

Each of these elements requires some configuration for DB2 database application development.

Procedure

To configure for DB2 database application development, do the following steps:

1. Configure the operating system.
See “Configuring the operating system for database application development” on page 46.

2. Configure the Data Server Client.
See “Configuring the IBM Data Server Client for database application development” on page 48.
3. Configure the development environment.
See “Configuring the development environment for application programming interfaces” on page 49.
4. Configure the transaction manager.
See “Configuration considerations for XA transaction managers” on page 69.

Configuring the operating system for database application development

Part of configuring for DB2 database application development is configuring operating system local and global environment variables.

Before you begin

To configure environment variables for DB2 database application development, you must have:

- A supported operating system
See Chapter 3, “Supported operating systems for database application development,” on page 5.
- A IBM Data Server Client installed
See “Methods for installing IBM data server clients” in *Installing IBM Data Server Clients*.

Procedure

To configure environment variables for DB2 database application development:

Perform one of the following tasks:

- For UNIX operating systems, “Setting up the application development environment (Linux and UNIX).”
- For Windows operating systems, “Setting up the application development environment (Windows)” on page 47

Setting up the application development environment (Linux and UNIX)

You must set environment variables for database manager instances on Linux and UNIX operating systems. This can be accomplished by launching the scripts that are provided for each instance.

About this task

Each database manager instance has two scripts to set up the environment for that instance: `db2profile` and `db2cshrc`. Setting up your environment is not required when using type 4 connectivity in your JDBC applications.

Procedure

1. Run the correct script to set up your environment for the instance that you will use to develop your application.

- For bash or Korn shell:
 - `.$HOME/sqllib/db2profile`
 - For C shell:
 - `source $HOME/sqllib/db2cshrc`
2. If you are developing applications using ODBC, DB2 CLI, JDBC, or SQLJ, perform the following tasks:
- “Setting up the ODBC environment” in *Call Level Interface Guide and Reference Volume 1*.
 - “Installing the IBM Data Server Driver for JDBC and SQLJ” on page 50.

UNIX environment variable settings

Environment variables for UNIX operating systems are necessary for proper interaction with your database environment.

Depending on the UNIX platform you are on, values for the following environment variables are set, either in `db2profile` (for bash or korn shell) or `db2cshrc` (for C shell), and a call to these files are put in the instance owner's `.bashrc` (bash shell), and either in the `.profile` (bash or korn shell) file or in the `.login` (C shell) file.

AIX:

- `PATH`, includes several DB2 directories including `sqllib/bin`
- `LIBPATH`, includes the directory `sqllib/lib` (see following note)

HP-UX:

- `PATH`, includes several DB2 directories including `sqllib/bin`
- `SHLIB_PATH` (32-bit and 64-bit) or `LD_LIBRARY_PATH` (64-bit), includes the directory `sqllib/lib` (see following note)

Linux and Solaris:

- `PATH`, includes several DB2 directories including `sqllib/bin`
- `LD_LIBRARY_PATH`, includes the directory `sqllib/lib` (see following note)

The blank files `sqllib/userprofile` and `sqllib/usercshrc` are created during instance creation so users can add their own instance environmental settings. If you do not want the new environment settings in the `db2profile` or `db2cshrc` scripts, you can override them using the corresponding "user" script, which is called at the end of the `db2profile` or `db2cshrc` script. During an instance upgrade (`db2iupgrade`), the user scripts are copied over so that your environment modifications will still be in use.

Setting up the application development environment (Windows)

To set up an application development environment on Windows operating systems, you must follow these steps.

About this task

When you install the IBM Data Server Client on Windows operating systems, the install program updates the configuration registry with the environment variables **INCLUDE**, **LIB**, and **PATH**. The system-wide environment variable, **DB2INSTANCE**, is set by install to the default instance created, called **DB2**. **DB2PATH** is set inside a **DB2** command window when the window is opened.

You can override these environment variables to set the values for the system or the currently logged-on user. Exercise caution when changing these environment variables. Do not change the **DB2PATH** environment variable. **DB2INSTANCE** is defined as a system-level environment variable. If **DB2INSTANCE** is not set, you do not have to make use of the **DB2INSTDEF** DB2 registry variable which defines the default instance name to use.

Procedure

1. To override the environment variable settings, use the Windows Control Panel. When using the variable `%DB2PATH%` in a command, put the full path in quotation marks, as in:

```
set LIB="%DB2PATH%;%LIB%"
```

The default installation value for this variable is `\Program Files\IBM\SQLLIB`, which contains a space, so put the full path in quotation marks.

2. In addition, you must take the following specific steps for running DB2 applications:
 - When building C or C++ programs, you must ensure that the **INCLUDE** environment variable contains `%DB2PATH%\INCLUDE` as the first directory. To do this, update the environment for your compiler:
 - a. Launch the shortcut to the Visual Studio Command Prompt.
 - b. In the Visual Studio Command window, run `DB2CMD.exe` to open the DB2 Command Window.
 - c. In the DB2 Command Window, set your **INCLUDE** path as follows :

```
set INCLUDE=%DB2PATH%\INCLUDE;%INCLUDE%
```
 - When building Micro Focus COBOL programs, set the **COBCPY** environment variable to point to `%DB2PATH%\INCLUDE\cobl_mf`.
 - When building IBM COBOL programs, set the **SYSLIB** environment variable to point to `%DB2PATH%\INCLUDE\cobl_a`.
 - Ensure the **LIB** environment variable points to `%DB2PATH%\lib` by using:

```
set LIB="%DB2PATH%\lib";%LIB%
```

To enable cross-developing 64-bit applications from a 32-bit environment, see “Upgrading 32-bit database applications to run on 64-bit instances” in *Upgrading to DB2 Version 10.1*.
 - To run an application against a remote server, both the client and the server must be properly configured. For instance, the **DB2COMM** environment variable must be set at the server of the remote database.

Configuring the IBM Data Server Client for database application development

The Data Server Client supports database application development; part of configuring for application development is configuring the client.

Before you begin

To configure the Data Server Client for database application development, you must have:

- Installed a Data Server Client
See “Installing IBM data server clients (Linux and UNIX)” and “Installing IBM data server clients (Windows)” in *Installing IBM Data Server Clients*.
- Optional: Created sample databases with which to test you applications

See “The SAMPLE database” in *SQL Reference Volume 1*.

- Completed basic configuration steps for your Data Server Client

See “Client-to-server communications configuration overview” in *Installing IBM Data Server Clients*.

About this task

For information about configuring the Data Server Client for database application development, see “Configuring the development environment for application programming interfaces.”

Configuring the development environment for application programming interfaces

Part of configuring for DB2 database application development is configuring the development environment for any application programming interfaces (APIs) that you will use in your applications.

Before you begin

To configure the application development environment for the application programming interfaces (APIs) that you will use, you must have:

- An IBM data server client installed
- The API drivers installed

See Chapter 5, “Supported database application programming interfaces,” on page 19.

- For the ODBC API only: an ODBC driver manager installed

See “Setting up the Windows CLI environment” in *Call Level Interface Guide and Reference Volume 1*.

Procedure

To configure the development environment for APIs, do the following steps:

1. For ADO .NET, see “ADO.NET development for IBM Data Servers” in *Developing ADO.NET and OLE DB Applications* for details.
2. For CLI and ODBC, perform any of the following tasks that apply to your environment:
 - “Setting up your application development environment to run CLI and ODBC applications” in *Call Level Interface Guide and Reference Volume 1*
 - “Setting up the UNIX ODBC Driver Manager” in *Call Level Interface Guide and Reference Volume 1*
3. For Embedded SQL, perform the following task: “Setting up the embedded SQL development environment” in *Developing Embedded SQL Applications*.
4. For JDBC and SQLJ, perform the following task: “Installing the IBM Data Server Driver for JDBC and SQLJ” in *Developing Java Applications*.
5. For OLE DB, see “IBM OLE DB Provider for DB2” in *Getting Started with Database Application Development* for details.
6. For Perl, see “Perl DBI” in *Developing Perl, PHP, Python, and Ruby on Rails Applications* for details.
7. For PHP, perform one of the following tasks:
 - Set up the PHP environment on Linux or UNIX.

- Set up the PHP environment on Windows.
8. For Ruby/Ruby on Rails, perform any of the tasks that are listed in “The IBM_DB Ruby driver and Rails adapter” in *Getting Started with Database Application Development* that apply to your environment.
 9. For Python, perform the task “Setting up the Python environment for IBM data servers” in *Developing Perl, PHP, Python, and Ruby on Rails Applications*.

Configuring your development environment to build and run CLI and ODBC applications

You can run CLI and ODBC applications against a DB2 database server using the IBM Data Server Client, the IBM Data Server Runtime Client, or the IBM Data Server Driver for ODBC and CLI. However, to compile CLI or ODBC applications, you need the IBM Data Server Client.

Procedure

In order for a CLI application to successfully access a DB2 database:

1. Ensure the CLI/ODBC driver was installed during the DB2 client install.
2. For the IBM Data Server Client and Runtime Client only: If the database is being accessed from a remote client, catalog the database and hostname of the machine the database is on.

On Windows operating systems, you can use the CLI/ODBC Settings GUI to catalog the DB2 database.

3. Optional: Explicitly bind the CLI /ODBC bind files to the database with the command:

```
db2 bind ~/sqllib/bnd/@db2cli.lst blocking all sqlerror continue \
      messages cli.msg grant public
```

On Windows operating systems, you can use the CLI/ODBC Settings GUI to bind the CLI/ODBC bind files to the database.

4. Optional: Change the CLI /ODBC configuration keywords by editing the `db2cli.ini` file. For information about the location of the `db2cli.ini` file, see “`db2cli.ini` initialization file” in *Call Level Interface Guide and Reference Volume 1*.

On Windows operating systems, you can use the CLI/ODBC Settings GUI to set the CLI/ODBC configuration keywords.

Results

Once you have completed steps 1 to 4, proceed to setting up your Windows CLI environment, or setting up your Linux or UNIX ODBC environment if you are running ODBC applications on Linux or UNIX.

Installing the IBM Data Server Driver for JDBC and SQLJ

After you install the IBM Data Server Driver for JDBC and SQLJ, you can prepare and run JDBC or SQLJ applications.

Before you begin

Before you install the IBM Data Server Driver for JDBC and SQLJ, you need the following software.

- An SDK for Java, 1.4.2 or later.

For all DB2 products except the IBM Data Server Runtime Client and the IBM Data Server Driver Package, the DB2 Database for Linux, UNIX, and Windows installation process automatically installs the SDK for Java, Version 5.

If you want to use JDBC 4.0 functions, you must install an SDK for Java, 6 or later.

If you want to use JDBC 4.1 functions, you must install an SDK for Java, 7 or later.

If you plan to run JDBC or SQLJ applications on your system, but not to prepare them, you need a Java runtime environment only.

Important: Support for the SDK for Java 1.4.2 is deprecated for Java routines, and might be discontinued in a future release.

- JVM native threads support

Any JVMs that run Java applications that access DB2 databases must include native threads support. You can specify native threads as the default thread support for some JVMs by setting the **THREADS_FLAG** environment variable to "native". Refer to the documentation for your Java environment for instructions on making native threads the default on your system.

- Unicode support for System i servers

If any SQLJ or JDBC programs will use IBM Data Server Driver for JDBC and SQLJ type 4 connectivity to connect to a DB2 for i server, the System i operating system must support the Unicode UTF-8 encoding scheme. The following table lists the System i PTFs that you need for Unicode UTF-8 support:

Table 17. System i PTFs for Unicode UTF-8 support

System i version	PTF numbers
V5R3 or later	None (support is included)

- Java support for HP-UX clients and servers

HP-UX servers: The IBM Data Server Driver for JDBC and SQLJ does not support databases that are in the HP-UX default character set, Roman8. Therefore, when you create a database on an HP-UX server that you plan to access with the IBM Data Server Driver for JDBC and SQLJ, you must create the database with a different character set.

HP-UX clients and servers: The Java environment on an HP-UX system requires special setup to run stored procedures under the IBM Data Server Driver for JDBC and SQLJ.

Restrictions

If you install the IBM Data Server Driver for JDBC and SQLJ on a Windows 64-bit operating system, you cannot use IBM Data Server Driver for JDBC and SQLJ type 2 connectivity to connect to a DB2 Database for Linux, UNIX, and Windows instance from a 32-bit Java application.

Procedure

To install the IBM Data Server Driver for JDBC and SQLJ:

1. During the DB2 Database for Linux, UNIX, and Windows installation process, select Java support on UNIX or Linux, or JDBC support on Windows. These selections are defaults. If you have already installed DB2 Database for Linux, UNIX, and Windows without JDBC support, you can run the installation process in Custom mode to add JDBC support.

Selection of Java support or JDBC support causes the installation process to perform the following actions:

- Installs the IBM Data Server Driver for JDBC and SQLJ class files.
The files are placed in the sql11b\java directory for Windows systems, or the sql11b/java directory for UNIX or Linux systems.

The files names are:

db2jcc.jar or db2jcc4.jar

Include db2jcc.jar in the **CLASSPATH** if you plan to use the version of the IBM Data Server Driver for JDBC and SQLJ that includes only JDBC 3.0 or earlier functions.

Include db2jcc4.jar in the **CLASSPATH** if you plan to use the version of the IBM Data Server Driver for JDBC and SQLJ that includes JDBC 4.0 or later functions, and JDBC 3.0 or earlier functions.

sqlj.zip or sqlj4.zip

Include sqlj.zip in the **CLASSPATH** if you plan to prepare SQLJ applications that include only JDBC 3.0 or earlier functions.

Include sqlj4.zip in the **CLASSPATH** if you plan to prepare SQLJ applications that include JDBC 4.0 or later functions, and JDBC 3.0 or earlier functions.

- Modifies the **CLASSPATH** to include the IBM Data Server Driver for JDBC and SQLJ class files.

Important: This step is performed automatically only for the db2jcc.jar and sqlj.zip file. If you are using the db2jcc4.jar file or the sqlj4.zip file, you must modify the **CLASSPATH** manually. Change db2jcc.jar to db2jcc4.jar or sqlj.zip to sqlj4.zip in the **CLASSPATH**.

You also must make this change in every DB2 command line window that you open.

Important: Include db2jcc.jar or db2jcc4.jar in the **CLASSPATH**. Do not include both files.

Important: Include sqlj.zip or sqlj4.zip in the **CLASSPATH**. Do not include both files. Do not include db2jcc.jar with sqlj4.zip, or db2jcc4.jar with sqlj.zip.

- If IBM Data Server Driver for JDBC and SQLJ client license files exist, the installation process installs them and modifies the **CLASSPATH** to include them.
The files are placed in the sql11b\java directory for Windows systems, or the sql11b/java directory for UNIX or Linux systems. The file names are:

Table 18. IBM Data Server Driver for JDBC and SQLJ license files

License file	Server to which license file permits a connection	Product that includes license file
db2jcc_license_cisuz.jar	DB2 for z/OS DB2 for i	All DB2 Connect products

Client license files are not required for connections to DB2 Database for Linux, UNIX, and Windows, Cloudscape, or IBM Informix databases from the IBM Data Server Driver for JDBC and SQLJ version 3.50 or later. Client

license files are not required for direct connections to DB2 for z/OS if DB2 Connect Unlimited Edition for System z server license activation is performed.

Recommendation: If you connect to DB2 for z/OS data servers directly, rather than through a DB2 Connect gateway, and you use DB2 Connect Unlimited Edition for System z 9.7 Fix Pack 3 or later, you must activate the server license key for DB2 Connect Unlimited Edition for System z. You must activate the license key on each DB2 for z/OS subsystem or data sharing group to which you will connect. See the DB2 Connect information for details on server license activation.

- Installs IBM Data Server Driver for JDBC and SQLJ native libraries for support of IBM Data Server Driver for JDBC and SQLJ type 2 connectivity. The files are placed in the `sql11b\bin` directory for Windows systems, or the `sql11b/lib` directory for UNIX or Linux systems.

The file names are:

libdb2jcct2.so

For AIX, HP-UX on IPF, Linux, and Solaris

db2jcct2.dll

For Windows

2. Customize the driver-wide configuration properties, if any of the defaults are inappropriate.
3. Configure TCP/IP.

Servers must be configured for TCP/IP communication in the following cases:

- JDBC or SQLJ applications that use IBM Data Server Driver for JDBC and SQLJ type 4 connectivity.
- JDBC or SQLJ applications that use IBM Data Server Driver for JDBC and SQLJ type 2 connectivity, and specify the *server* and *port number* in the connection URL.

Ensure that the TCP/IP listener is running. To activate the TCP/IP listener:

- a. Set the environment variable **DB2COMM** to TCP/IP:

```
db2set DB2COMM=TCPIP
```

- b. Update the database manager configuration file with the TCP/IP service name as specified in the services file:

```
db2 update dbm cfg using svcname TCP/IP-service-name
```

The port number used for applets and SQLJ programs must be the same as the TCP/IP **svcname** number used in the database manager configuration file.

- c. Execute the **db2stop** and **db2start** commands for the service name setting to take effect.
4. On DB2 Database for Linux, UNIX, and Windows servers on which you plan to run Java stored procedures or user-defined functions, update the database manager configuration to include the path where the SDK for Java is located. You can do this by entering commands similar to these on the server command line:

- For database systems on UNIX or Linux:

```
db2 update dbm cfg using JDK_PATH /home/db2inst/jdk15
```

`/home/db2inst/jdk15` is the path where the SDK for Java is installed.

- For database systems on Windows:

```
db2 update dbm cfg using JDK_PATH c:\Program Files\jdk15
```

c:\Program Files\jdk15 is the path where the SDK for Java is installed.

To verify the correct value for the **jdk_path** field in the DB2 database manager configuration, enter the following command on the database server:

```
db2 get dbm cfg
```

You might want to redirect the output to a file for easier viewing. The **jdk_path** field appears near the beginning of the output.

5. If you plan to call SQL procedures that are on DB2 Database for Linux, UNIX, and Windows servers from Java programs, and the date and time format that is associated with the territory code of the database servers is **not** the USA format, take the following actions:
 - a. Set the **DB2_SQLROUTINE_PREPOPTS** registry variable on the database servers to indicate that the default datetime format is ISO:

```
db2set DB2_SQLROUTINE_PREPOPTS="DATETIME ISO"
```
 - b. Redefine any existing SQL procedures that you plan to call from Java programs.

These steps are necessary to ensure that the calling application receives date and time values correctly.

6. If you plan to access DB2 for z/OS database servers with your Java applications, follow the instructions in "Special setup for accessing DB2 for z/OS servers from Java programs" in *Developing Java Applications*.

Transaction managers

Transaction managers are software components used to assign identifiers to specific database transactions, monitor transaction processing, and to validate or act when a transaction completes or fails.

Transaction management guards against hardware and software errors that might leave a transaction partially completed, with a database left in an unknown, inconsistent state. The transaction manager ensures that all operations in any uncommitted (i.e., not completely processed) transactions are cancelled.

The DB2 Database for Linux, UNIX, and Windows transaction manager (TM), provided with the DB2 product and DB2 Connect, performs all of these functions. The database manager provides transaction manager functions used to coordinate the updating of several databases within a single unit of work. The database client automatically coordinates the unit of work, and uses a transaction manager database to register each transaction and track its completion status. You can use the DB2 transaction manager with DB2 databases.

If you have resources other than DB2 databases that you want to participate in a two-phase commit transaction, you can use an XA-compliant transaction manager.

X/Open distributed transaction processing model

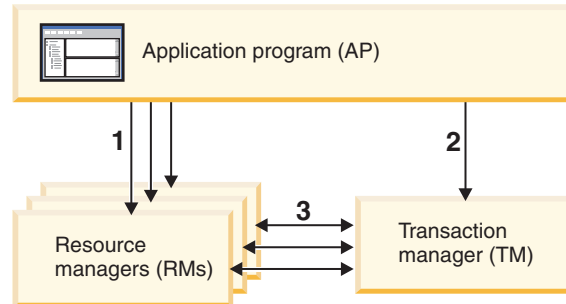
The X/Open Distributed Transaction Processing (DTP) model includes a number of interrelated components that control how distributed transactions are processed.

These components include the:

- Application program (AP)

- Transaction manager (TM)
- Resources managers (RM).

Figure 1 illustrates this model, and shows the relationship among these components.



Legend

- 1 - AP uses resources from a set of RMs
- 2 - AP defines transaction boundaries through TM interfaces
- 3 - TM and RMs exchange transaction information

Figure 1. X/Open distributed transaction processing (DTP) model

Application program (AP)

The application program (AP) defines transaction boundaries, and defines the application-specific actions that make up the transaction.

For example, a CICS application program might want to access resource managers (RMs), such as a database and a CICS Transient Data Queue, and use programming logic to manipulate the data. Each access request is passed to the appropriate resource managers through function calls specific to that RM. In the case of DB2 products, these could be function calls generated by the DB2 database precompiler for each SQL statement, or database calls coded directly by the programmer using the APIs.

A transaction manager (TM) product usually includes a transaction processing (TP) monitor to run the user application. The TP monitor provides APIs to allow an application to start and end a transaction, and to perform application scheduling and load balancing among the many users who want to run the application. The application program in a distributed transaction processing (DTP) environment is really a combination of the user application and the TP monitor.

To facilitate an efficient online transaction processing (OLTP) environment, the TP monitor pre-allocates a number of server processes at startup, and then schedules and reuses them among the many user transactions. This conserves system resources, by allowing more concurrent users to be supported with a smaller number of server processes and their corresponding RM processes. Reusing these processes also avoids the overhead of starting up a process in the TM and RMs for each user transaction or program. (A program invokes one or more transactions.)

This also means that the server processes are the actual "user processes" to the TM and the RMs. This has implications for security administration and application programming.

The following types of transactions are possible from a TP monitor:

- Non-XA transactions

These transactions involve RMs that are not defined to the TM, and are therefore not coordinated under the two-phase commit protocol of the TM. This might be necessary if the application needs to access an RM that does not support the XA interface. The TP monitor simply provides efficient scheduling of applications and load balancing. Since the TM does not explicitly "open" the RM for XA processing, the RM treats this application as any other application that runs in a non-DTP environment.

- Global transactions

These transactions involve RMs that are defined to the TM, and are under the TM's two-phase commit control. A global transaction is a unit of work that could involve one or more RMs. A *transaction branch* is the part of work between a TM and an RM that supports the global transaction. A global transaction could have multiple transaction branches when multiple RMs are accessed through one or more application processes that are coordinated by the TM.

Loosely coupled global transactions exist when each of a number of application processes accesses the RMs as if they are in a separate global transaction, but those applications are under the coordination of the TM. Each application process will have its own transaction branch within an RM. When a commit or rollback is requested by any one of the APs, TM, or RMs, the transaction branches are completed altogether. It is the application's responsibility to ensure that resource deadlock does not occur among the branches. (Note that the transaction coordination performed by the DB2 transaction manager for applications prepared with the SYNCPOINT(TWOPHASE) option is roughly equivalent to these loosely coupled global transactions.

Tightly coupled global transactions exist when multiple application processes take turns to do work under the same transaction branch in an RM. To the RM, the two application processes are a single entity. The RM must ensure that resource deadlock does not occur within the transaction branch.

Transaction manager (TM)

The transaction manager (TM) assigns identifiers to transactions, monitors their progress, and takes responsibility for transaction completion and failure. The transaction branch identifiers (known as XIDs) are assigned by the TM to identify both the global transaction, and the specific branch within an RM. This is the correlation token between the log in a TM and the log in an RM. The XID is needed for two-phase commit, or rollback, to perform the *resynchronization* operation (also known as a *resync*) on system startup, or to let the administrator perform a *heuristic* operation (also known as *manual intervention*), if necessary.

After a TP monitor is started, it asks the TM to open all the RMs that a set of application servers have defined. The TM passes `xa_open` calls to the RMs, so that they can be initialized for DTP processing. As part of this startup procedure, the TM performs a resync to recover all *indoubt transactions*. An indoubt transaction is a global transaction that was left in an uncertain state. This occurs when the TM (or at least one RM) becomes unavailable after successfully completing the first phase (that is, the prepare phase) of the two-phase commit protocol. The RM will not know whether to commit or roll back its branch of the transaction until the TM

can reconcile its own log with the RM logs when they become available again. To perform the resync operation, the TM issues a **xa_recover** call one or more times to each of the RMs to identify all the indoubt transactions. The TM compares the replies with the information in its own log to determine whether it should inform the RMs to **xa_commit** or **xa_rollback** those transactions. If an RM has already committed or rolled back its branch of an indoubt transaction through a heuristic operation by its administrator, the TM issues an **xa_forget** call to that RM to complete the resync operation.

When a user application requests a commit or a rollback, it must use the API provided by the TP monitor or TM, so that the TM can coordinate the commit and rollback among all the RMs involved. For example, when a WebSphere application issues a request to commit a transaction, the WebSphere XA TM will in turn issue XA calls, such as **xa_end**, **xa_prepare**, **xa_commit**, or **xa_rollback** to request the RM to commit or roll back the transaction. The TM could choose to use one-phase instead of two-phase commit if only one RM is involved, or if an RM replies that its branch is read-only.

Resource managers (RM)

A resource manager (RM) provides access to shared resources, such as databases.

The DB2 system, as resource manager of a database, can participate in a *global transaction* that is being coordinated by an XA-compliant TM. As required by the XA interface, the database manager provides a *db2xa_switch* external C variable of type *xa_switch_t* to return the XA switch structure to the TM. This data structure contains the addresses of the various XA routines to be invoked by the TM, and the operating characteristics of the RM.

There are two methods by which the RM can register its participation in each global transaction: *static registration* and *dynamic registration*:

- Static registration requires the TM to issue (for every transaction) the **xa_start**, **xa_end**, and **xa_prepare** series of calls to all the RMs defined for the server application, regardless of whether a given RM is used by the transaction. This is inefficient if not every RM is involved in every transaction, and the degree of inefficiency is proportional to the number of defined RMs.
- Dynamic registration (used by DB2) is flexible and efficient. An RM registers with the TM using an **ax_reg** call only when the RM receives a request for its resource. Note that there is no performance disadvantage with this method, even when there is only one RM defined, or when every RM is used by every transaction, because the **ax_reg** and the **xa_start** calls have similar paths in the TM.

The XA interface provides two-way communication between a TM and an RM. It is a system-level interface between the two DTP software components, not an ordinary application program interface to which an application developer codes. However, application developers should be familiar with the programming restrictions that the DTP software components impose.

Although the XA interface is invariant, each XA-compliant TM may have product-specific ways of integrating an RM. For information about integrating your DB2 product as a resource manager with a specific transaction manager, see the appropriate TM product documentation.

xa_open string formats

xa_open string format for DB2 Database for Linux, UNIX, and Windows and DB2 Connect Version 8 FixPak 3 and later

This is the format for the xa_open string:

```
parm_id1 = parm_value,parm_id2 = parm_value, ...
```

It does not matter in what order these parameters are specified. Valid values for *parm_id* are described in the following section.

Note: Unless explicitly stated, these parameters are not case sensitive and have no default value.

AXLIB

Library that contains the **ax_reg** and **ax_unreg** functions for the TP monitor.

This value is used by DB2 Database for Linux, UNIX, and Windows to obtain the addresses of the required **ax_reg** and **ax_unreg** functions. It can be used to override assumed values based on the TPM parameter, or it can be used by TP monitors that do not appear on the list for TPM. On AIX, if the library is an archive library, the archive member should be specified in addition to the library name. For example: AXLIB=/usr/mqm/lib/libmqmax_r.a(libmqmax_r.o). This parameter is optional.

CHAIN_END

xa_end chaining flag. Valid values are T, F, or no value. XA_END chaining is an optimization that can be used by DB2 Database for Linux, UNIX, and Windows to reduce network flows. If the TP monitor environment is such that it can be guaranteed that **xa_prepare** will be invoked within the same thread or process immediately following the call to **xa_end**, and if CHAIN_END is on, the xa_end flag will be chained with the **xa_prepare** command, thus eliminating one network flow. A value of T means that CHAIN_END is on; a value of F means that CHAIN_END is off; no specified value means that CHAIN_END is on. This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

CREG

xa_start chaining flag. Valid values are T, or F, or no value. **xa_start** chaining is an optimization that is used by DB2 Database for Linux, UNIX, and Windows to reduce network flows. The parameter is only valid if the TP monitor is using static registration (see SREG). The TP monitor environment is such that it can guarantee that an SQL statement will be invoked immediately after the call to the XA API **xa_start**. If CREG is set to T, the SQL statement is chained to the **xa_start** request, thus eliminating one network flow. This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

CT Connect Timeout. Valid values are 0 - 32767. CT specifies the amount of time, in seconds, that an application will wait when attempting to establish a connection with the server. If a connection is not established in the amount of time specified, an error will be returned. Specifying a value of 0 means that the application will attempt to wait until a connection is established regardless of how long it takes. However, it is possible that the connection attempt will be terminated by the default TCP/IP timeout setting. If this parameter is not specified, the default value of 0 is used.

DB Database alias. Database alias used by the application to access the database. This parameter must be specified.

HOLD_CURSOR

Specifies whether cursors are held across transaction commits. Valid values are T, F, or no value. TP monitors typically reuse threads or processes for multiple applications. To ensure that a newly loaded application does not inherit cursors opened by a previous application, cursors are closed after a commit. If HOLD_CURSORS is on, cursors with hold attributes are not closed, and will persist across transaction commit boundaries. When using this option, the global transaction must be committed or rolled back from the same thread of control. If HOLD_CURSOR is off, the opening of any cursors with hold attributes will be rejected. A value of T means that HOLD_CURSOR is on; a value of F means that HOLD_CURSOR is off; This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

PWD

Password. A password that is associated with the user ID. Required if a user ID is specified. This parameter is case-sensitive.

SINGLE_PROCESS

Specifies whether direct XA support is enabled for clients accessing DB2 for z/OS. Valid values are T, F, or no value. For instance-based clients (IBM data server clients), XA support must be explicitly enabled either through this setting or the **enableDirectXA** setting in the db2dsdriver configuration file. For instance-less clients (IBM data server drivers), XA support is enabled by default for Microsoft Distributed Transaction Coordinator or Microsoft Component Services (COM+). For all other supported transaction managers, XA support for instance-less clients must be explicitly enabled through the SINGLE_PROCESS keyword.

SREG

Static Registration. Valid values are T, or F, or no value. DB2 Database for Linux, UNIX, and Windows supports two methods of registering a global transaction. The first is Dynamic Registration, where DB2 Database for Linux, UNIX, and Windows calls the TP's **ax_reg** function to register the transaction (see AXLIB). The second method is Static Registration, where the TP calls the XA API **xa_start** to initiate a global transaction. Note that both dynamic and static registration are mutually exclusive. If this parameter is not specified, the default value of F is used.

SUSPEND_CURSOR

Specifies whether cursors are to be kept when a transaction thread of control is suspended. Valid values are T, F, or no value. TP monitors that suspend a transaction branch can reuse the suspended thread or process for other transactions. If SUSPEND_CURSOR is off, all cursors except cursors with hold attributes are closed. On resumption of the suspended transaction, the application must obtain the cursors again. If SUSPEND_CURSOR is on, any open cursors are not closed, and are available to the suspended transaction on resumption. A value of T means that SUSPEND_CURSOR is on; a value of F means that SUSPEND_CURSOR is off; This parameter can be used to override the setting derived from a specified TPM value. If this parameter is not specified, the default value of F is used.

TCTX

Specifies whether the transaction should use a trusted connection. Valid values are TRUE or FALSE. If this parameter is set to TRUE it tells the transaction manager to try to open a trusted connection.

TOC

The entity (“Thread of Control”) to which all DB2 XA Connections are bound. Valid values are T, or P, or not set. TOC is the entity where all DB2 XA Connections are bound. All DB2 XA Connections formed within an entity must be unique. That is, they cannot have two connections to the same database within the entity. The TOC has two parameters: T (OS Thread) and P (OS Process). When set to a value of T, all DB2 XA Connections formed under a particular OS Thread are unique to that thread only. Multiple threads cannot share DB2 XA Connections. Each OS thread must form its own set of DB2 XA Connections. When set to a value of P, all DB2 XA Connections are unique to the OS Process and all XA Connections can be shared between OS threads. If this parameter is not specified, the default value of T is used.

TPM

Transaction processing monitor name. Name of the TP monitor being used. For supported values, see the next table. This parameter can be specified to allow multiple TP monitors to use a single DB2 instance. The specified value overrides the value specified in the **tp_mon_name** database manager configuration parameter. This parameter is optional.

UID

User ID. Specifies the user ID that has authority to connect to the database. Required if a password is specified. This parameter is case-sensitive.

UREGNM

User Registry Name. When an identity mapping service is being used, this parameter gives the name of the registry to which the user name given in the UID parameter belongs.

TPM and tp_mon_name values

The xa_open string TPM parameter and the **tp_mon_name** database manager configuration parameter are used to indicate to DB2 Database for Linux, UNIX, and Windows which TP monitor is being used. The **tp_mon_name** value applies to the entire DB2 instance. The TPM parameter applies only to the specific XA resource manager. The TPM value overrides the **tp_mon_name** parameter. Valid values for the TPM and **tp_mon_name** parameters are as follows:

Table 19. Valid Values for TPM and tp_mon_name

TPM Value	TP Monitor Product	Internal Settings
CICS	IBM TXSeries CICS	AXLIB=libEncServer (for Windows) =/usr/lpp/cics/lib/libEncServer (for UNIX based systems) HOLD_CURSOR=T CHAIN_END=T SUSPEND_CURSOR=F TOC=T

Table 19. Valid Values for TPM and tp_mon_name (continued)

TPM Value	TP Monitor Product	Internal Settings
MQ	IBM MQSeries®	AXLIB=mqmax (for Windows) =/usr/mqm/lib/libmqmax_r.a (for AIX threaded applications) =/usr/mqm/lib/libmqmax.a (for AIX non-threaded applications) =/opt/mqm/lib/libmqmax.so (for Solaris) =/opt/mqm/lib/libmqmax_r.sl (for HP threaded applications) =/opt/mqm/lib/libmqmax.sl (for HP non-threaded applications) =/opt/mqm/lib/libmqmax_r.so (for Linux threaded applications) =/opt/mqm/lib/libmqmax.so (for Linux non-threaded applications) HOLD_CURSOR=F CHAIN_END=F SUSPEND_CURSOR=F TOC=P
CB	IBM Component Broker	AXLIB=somtrx1i (for Windows) =libsomtrx1 (for UNIX based systems) HOLD_CURSOR=F CHAIN_END=T SUSPEND_CURSOR=F TOC=T
SF	IBM San Francisco	AXLIB=ibmsfDB2 HOLD_CURSOR=F CHAIN_END=T SUSPEND_CURSOR=F TOC=T
TUXEDO	BEA Tuxedo	AXLIB=libtux HOLD_CURSOR=F CHAIN_END=F SUSPEND_CURSOR=F TOC=T
MTS	Microsoft Transaction Server	It is not necessary to configure DB2 Database for Linux, UNIX, and Windows for MTS. MTS is automatically detected by the ODBC driver in DB2 Database for Linux, UNIX, and Windows.
JTA	Java Transaction API	It is not necessary to configure DB2 Database for Linux, UNIX, and Windows for Enterprise Java Servers (EJS) such as IBM WebSphere. The JDBC driver in DB2 Database for Linux, UNIX, and Windows automatically detects this environment. Therefore this TPM value is ignored.

xa_open string format for earlier versions

Earlier versions of DB2 Database for Linux, UNIX, and Windows used the xa_open string format described here. This format is still supported for compatibility reasons. Applications should be migrated to the new format when possible.

Each database is defined as a separate resource manager (RM) to the transaction manager (TM), and the database must be identified with an `xa_open` string that has the following syntax:

```
"database_alias,userid,password"
```

The `database_alias` is required to specify the alias name of the database. The alias name is the same as the database name unless you explicitly cataloged an alias name after you created the database. The `userid` and `password` are optional and, depending on the authentication method, are used to provide authentication information to the database.

Examples

1. You are using IBM TXSeries CICS on Windows. The TXSeries documentation indicates that you need to configure `tp_mon_name` with a value of `libEncServer:C`. This is still an acceptable format; however, with DB2 Database for Linux, UNIX, and Windows or DB2 Connect Version 8 FixPak 3 and later, you have the option of:

- Specifying a `tp_mon_name` of CICS (recommended for this scenario):

```
db2 update dbm cfg using tp_mon_name CICS
```

For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:

```
db=dbalias,uid=userid,pwd=password
```

- For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:

```
db=dbalias,uid=userid,pwd=password,tpm=cics
```

2. You are using IBM MQSeries on Windows. The MQSeries documentation indicates that you need to configure `tp_mon_name` with a value of `mqmax`. This is still an acceptable format; however, with DB2 Database for Linux, UNIX, and Windows or DB2 Connect Version 8 FixPak 3 and later, you have the option of:

- Specifying a `tp_mon_name` of MQ (recommended for this scenario):

```
db2 update dbm cfg using tp_mon_name MQ
```

For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:

```
uid=userid,db=dbalias,pwd=password
```

- For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:

```
uid=userid,db=dbalias,pwd=password,tpm=mq
```

3. You are using both IBM TXSeries CICS and IBM MQSeries on Windows. A single DB2 instance is being used. In this scenario, you would configure as follows:

- a. For each database defined to CICS in the Region—> Resources—> Product—> XAD—> Resource manager initialization string, specify:

```
pwd=password,uid=userid,tpm=cics,db=dbalias
```

- b. For each database defined as a resource in the queue manager properties, specify an `XaOpenString` as:

```
db=dbalias,uid=userid,pwd=password,tpm=mq
```

4. You are developing your own XA-compliant transaction manager (XA TM) on Windows, and you want to tell DB2 Database for Linux, UNIX, and Windows

that library "myaxlib" has the required functions **ax_reg** and **ax_unreg**. Library "myaxlib" is in a directory specified in the PATH statement. You have the option of:

- Specifying a **tp_mon_name** of myaxlib:
`db2 update dbm cfg using tp_mon_name myaxlib`

and, for each database defined to the XA TM, specifying an xa_open string:

```
db=dbalias,uid=userid,pwd=password
```

- For each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password,axlib=myaxlib`

5. You are developing your own XA-compliant transaction manager (XA TM) on Windows, and you want to tell DB2 Database for Linux, UNIX, and Windows that library "myaxlib" has the required functions **ax_reg** and **ax_unreg**. Library "myaxlib" is in a directory specified in the PATH statement. You also want to enable XA END chaining. You have the option of:

- For each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password,axlib=myaxlib,chain_end=T`
- For each database defined to the XA TM, specifying an xa_open string:
`db=dbalias,uid=userid,pwd=password,axlib=myaxlib,chain_end`

Updating host or System i database servers with an XA-compliant transaction manager

Host and System i database servers may be updatable depending upon the architecture of the XA Transaction Manager.

Before you begin

- To support commit sequences from different processes, the DB2 Connect connection concentrator must be enabled.
- DB2 Connect with the DB2 sync point manager (SPM) must be configured.

Procedure

To enable the DB2 Connect connection concentrator, set the database manager configuration parameter **max_connections** to a value greater than **max_coordagents**.

Resolving indoubt transactions manually

An XA-compliant transaction manager (Transaction Processing Monitor) uses a two-phase commit process similar to that used by the DB2 transaction manager. The principal difference between the two environments is that the TP monitor provides the function of logging and controlling the transaction, instead of the DB2 transaction manager and the transaction manager database.

Errors similar to those that occur for the DB2 transaction manager can occur when using an XA-compliant transaction manager. Similar to the DB2 transaction manager, an XA-compliant transaction manager attempts to resynchronize indoubt transactions.

If you cannot wait for the transaction manager to automatically resolve indoubt transactions, you can manually resolve them. This manual process is sometimes referred to as *making a heuristic decision*.

About this task

The **LIST INDOUBT TRANSACTIONS** command (using the **WITH PROMPTING** option), or the related set of APIs (db2XaListIndTrans, sqlxphcm, sqlxhfrg, sqlxphrl), allows you to query, commit, and roll back indoubt transactions. In addition, it also allows you to “forget” transactions that were heuristically committed or rolled back, by removing the log records and releasing the log space.

Manually resolve indoubt transactions by using these commands (or related APIs) with *extreme caution*, and only as a last resort. The best strategy is to wait for the transaction manager to drive the resynchronization process. You could experience data integrity problems if you manually commit or roll back a transaction in one of the participating databases, and the opposite action is taken against another participating database. Recovering from data integrity problems requires you to understand the application logic, to identify the data that was changed or rolled back, and then to perform a point-in-time recovery of the database, or manually undo or reapply the changes. If you cannot wait for the transaction manager to initiate the resynchronization process, and you must release the resources tied up by an indoubt transaction, heuristic operations are necessary. This situation could occur if the transaction manager will not be available for an extended period of time to perform the resynchronization, and the indoubt transaction is tying up resources that are urgently needed. An indoubt transaction ties up the resources that were associated with this transaction before the transaction manager or resource managers became unavailable. For the database manager, these resources include locks on tables and indexes, log space, and storage taken up by the transaction. Each indoubt transaction also decreases (by one) the maximum number of concurrent transactions that can be handled by the database. Moreover, an offline backup cannot be taken unless all indoubt transactions have been resolved.

The heuristic forget function is required in the following situations:

- When a heuristically committed or rolled back transaction causes a log full condition, indicated in output from the **LIST INDOUBT TRANSACTIONS** command
- When an offline backup is to be taken

The heuristic forget function releases the log space occupied by an indoubt transaction. The implication is that if a transaction manager eventually performs a resynchronization operation for this indoubt transaction, it could potentially make the wrong decision to commit or roll back other resource managers, because there is no log record for the transaction in this resource manager. In general a “missing” log record implies that the resource manager has rolled back the transaction.

Procedure

To resolve indoubt transactions manually:

1. Connect to the database for which you require all transactions to be complete.
2. List the indoubt transactions:
 - a. For DB2 database servers, use the **LIST INDOUBT TRANSACTIONS WITH PROMPTING** command. The *xid* represents the global transaction ID, and is identical to the *xid* used by the transaction manager and by other resource managers participating in the transaction.
 - b. For host or System i database servers, you can use one of the following: You can obtain indoubt information directly from the host or System i server. To obtain indoubt information directly from DB2 for z/OS, invoke the **DISPLAY THREAD TYPE(INDOUBT)** command. Use the **RECOVER**

command to make a heuristic decision. To obtain indoubt information directly from DB2 for i, invoke the **wrkcmdfn** command. You can obtain indoubt information from the DB2 Connect server used to access the host or System i database server. You can obtain indoubt information from the DB2 Connect server used to access the host or System i database server.

3. For each indoubt transaction, use the information shown about the application and the operating environment to determine the other participating resource managers.
4. Determine the actions to take with each indoubt transaction:
 - If the transaction manager is available, and the indoubt transaction in a resource manager was caused by the resource manager not being available in the second commit phase, or for an earlier resynchronization process, do the following:
 - a. Check the transaction manager's log to determine what action has been taken against the other resource managers.
 - b. Take the same action against the database; that is, use the **LIST INDOUBT TRANSACTIONS WITH PROMPTING** command, to either heuristically commit or heuristically roll back the transaction.
 - If the transaction manager is *not* available, use the status of the transaction in the other participating resource managers to determine what action to take:
 - If at least one of the other resource managers has committed the transaction, heuristically commit the transaction in all the resource managers.
 - If at least one of the other resource managers has rolled back the transaction, heuristically roll back the transaction.
 - If the transaction is in the “prepared” (indoubt) state in all of the participating resource managers, heuristically roll back the transaction.
 - If one or more of the other resource managers is not available, heuristically roll back the transaction.

What to do next

To obtain indoubt transaction information from DB2 Database for Linux, UNIX, and Windows, connect to the database and issue the **LIST INDOUBT TRANSACTIONS WITH PROMPTING** command, or call the `db2XaListIndTrans` API from a client application.

Indoubt transaction management APIs

There might be times where it is useful to be able to query, commit, and roll back indoubt transactions without waiting for the Transaction Manager (TM) to perform a *re-sync* action. This situation might arise, for example, if the communication line is broken, and an indoubt transaction is tying up needed resources.

A set of APIs is provided for tool writers to perform heuristic functions on indoubt transactions when the resource owner (such as the database administrator) cannot wait for the TM to perform the *re-sync* operation. For the database manager, these resources include locks on tables and indexes, log space, and storage used by the transaction. Each indoubt transaction also decreases by one the maximum number of concurrent transactions that could be processed by the database manager.

The heuristic APIs have the capability to query, commit, and roll back indoubt transactions, and to cancel transactions that have been heuristically committed or rolled back, by removing the log records and releasing log pages.

Attention: The heuristic APIs should be used with caution and only as a last resort. The TM should drive the re-sync events. If the TM has an operator command to start the re-sync action, it should be used. If the user cannot wait for a TM-initiated re-sync, heuristic actions are necessary.

Although there is no set way to perform these actions, the following guidelines may be helpful:

- Use the `db2XaListIndTrans` function to display the indoubt transactions. They have a status = 'P' (prepared), and are not connected. The *gtrid* portion of an *xid* is the global transaction ID that is identical to that in other resource managers (RM) that participate in the global transaction.
- Use knowledge of the application and the operating environment to identify the other participating RMs.
- If the transaction manager is CICS, and the only RM is a CICS resource, perform a heuristic rollback.
- If the transaction manager is not CICS, use it to determine the status of the transaction that has the same *gtrid* as does the indoubt transaction.
- If at least one RM has committed or rolled back, perform a heuristic commit or a rollback.
- If they are all in the prepared state, perform a heuristic rollback.
- If at least one RM is not available, perform a heuristic rollback.

If the transaction manager is available, and the indoubt transaction is due to the RM not being available in the second phase, or in an earlier re-sync, the DBA should determine from the TM's log what action has been taken against the other RMs, and then do the same. The *gtrid* is the matching key between the TM and the RMs.

Do not execute `sqlxhfrg` unless a heuristically committed or rolled back transaction happens to cause a log full condition. The `forget` function releases the log space occupied by this indoubt transaction. If a transaction manager eventually performs a re-sync action for this indoubt transaction, the TM could make the wrong decision to commit or to roll back other RMs, because no record was found in this RM. In general, a missing record implies that the RM has rolled back.

Security considerations for XA transaction managers

The TP monitor pre-allocates a set of server processes and runs the transactions from different users under the IDs of the server processes. To the database, each server process appears as a big application that has many units of work, all being run under the same ID associated with the server process.

For example, in an AIX environment using CICS, when a TXSeries CICS region is started, it is associated with the AIX user name under which it is defined. All the CICS Application Server processes are also being run under this TXSeries CICS "master" ID, which is usually defined as "cics". CICS users can invoke CICS transactions under their DCE login ID, and while in CICS, they can also change their ID using the `CESN` signon transaction. In either case, the end user's ID is not available to the RM. Consequently, a CICS Application Process might be running transactions on behalf of many users, but they appear to the RM as a single program with many units of work from the same "cics" ID. Optionally, you can specify a user ID and password on the `xa_open` string, and that user ID will be used, instead of the "cics" ID, to connect to the database.

There is not much impact on static SQL statements, because the binder's privileges, not the end user's privileges, are used to access the database. This does mean, however, that the EXECUTE privilege of the database packages must be granted to the server ID, and not to the end user ID.

For dynamic statements, which have their access authentication done at run time, access privileges to the database objects must be granted to the server ID and not to the actual user of those objects. Instead of relying on the database to control the access of specific users, you must rely on the TP monitor system to determine which users can run which programs. The server ID must be granted all privileges that its SQL users require.

To determine who has accessed a database table or view, you can perform the following steps:

1. From the SYSCAT.PACKAGEDEP catalog view, obtain a list of all packages that depend on the table or view.
2. Determine the names of the server programs (for example, CICS programs) that correspond to these packages through the naming convention used in your installation.
3. Determine the client programs (for example, CICS transaction IDs) that could invoke these programs, and then use the TP monitor's log (for example, the CICS log) to determine who has run these transactions or programs, and when.

XA function supported by DB2 Database for Linux, UNIX, and Windows

DB2 Database for Linux, UNIX, and Windows supports the XA91 specification defined in *X/Open CAE Specification Distributed Transaction Processing: The XA Specification*, with the following exceptions:

- Asynchronous services

The XA specification allows the interface to use asynchronous services, so that the result of a request can be checked at a later time. The database manager requires that the requests be invoked in synchronous mode.

- Registration

The XA interface allows two ways to register an RM: static registration and dynamic registration. DB2 supports both dynamic and static registration. DB2 provides two switches to control the type of registration used.

- `db2xa_switch_std` for dynamic registration
- `db2xa_switch_static_std` for static registration

- Association migration

The DB2 product does not support transaction migration between threads of control.

XA switch usage and location

As required by the XA interface, the database manager provides a `db2xa_switch_std` and a `db2xa_switch_static_std` external C variable of type `xa_switch_t` to return the XA switch structure to the TM. Other than the addresses of various XA functions, the following fields are returned:

Field **Value**

name The product name of the database manager. For example, IBM DB2 Version 9.7 for AIX.

flags For `db2xa_switch_std` TMREGISTER | TMNOMIGRATE is set

Explicitly states that the DB2 product uses dynamic registration, and that the TM should not use association migration. Implicitly states that asynchronous operation is not supported.

For *db2xa_switch_static_std* TMNOMIGRATE is set

Explicitly states that the DB2 product uses static registration, and that the TM should not use association migration. Implicitly states that asynchronous operation is not supported.

version

Must be zero.

Using the DB2 Database for Linux, UNIX, and Windows XA switch

The XA architecture requires that a Resource Manager (RM) provide a *switch* that gives the XA Transaction Manager (TM) access to the RM's **xa_** routines. An RM switch uses a structure called *xa_switch_t*. The switch contains the RM's name, non-NULL pointers to the RM's XA entry points, a flag, and a version number.

Linux and UNIX

The switch for DB2 Database for Linux, UNIX, and Windows can be obtained through either of the following two ways:

- Through one additional level of indirection. In a C program, this can be accomplished by defining the macro:

```
#define db2xa_switch_std (*db2xa_switch_std)
#define db2xa_switch_static_std (*db2xa_switch_std)
```

prior to using *db2xa_switch_std* or *db2xa_switch_static_std*.

- By calling **db2xacic_std** or **db2xacicst_std**

DB2 provides these APIs, which return the address of the *db2xa_switch_std* or *db2xa_switch_static_std* structure. This function is prototyped as:

```
struct xa_switch_t * SQL_API_FN db2xacic_std( )
struct xa_switch_t * SQL_API_FN db2xacicst_std( )
```

With either method, you must link your application with `libdb2`.

Windows

The pointer to the *xa_switch* structure, *db2xa_switch_std*, or *db2xa_switch_static_std* is exported as DLL data. This implies that a Windows application using this structure must reference it in one of three ways:

- Through one additional level of indirection. In a C program, this can be accomplished by defining the macro:

```
#define db2xa_switch_std (*db2xa_switch_std)
#define db2xa_switch_static_std (*db2xa_switch_std)
```

prior to using *db2xa_switch_std* or *db2xa_switch_static_std*.

- If using the Microsoft Visual C++ compiler, *db2xa_switch_std* or *db2xa_switch_static_std* can be defined as:

```
extern __declspec(dllimport) struct xa_switch_t db2xa_switch_std
extern __declspec(dllimport) struct xa_switch_t db2xa_switch_static_std
```

- By calling **db2xacic_std** or **db2xacicst_std**

DB2 provides this API, which returns the address of the *db2xa_switch_std* or *db2xa_switch_static_std* structure. This function is prototyped as:

```

struct xa_switch_t * SQL_API_FN db2xacic_std( )
struct xa_switch_t * SQL_API_FN db2xacicst_std( )

```

With any of these methods, you must link your application with `db2api.lib`.

Example C Code

The following code illustrates the different ways in which the `db2xa_switch_std` or `db2xa_switch_static_std` can be accessed via a C program. Be sure to link your application with the appropriate library.

```

#include <stdio.h>
#include <xa.h>

struct xa_switch_t * SQL_API_FN db2xacic_std( );

#ifdef DECLSPEC_DEFN
extern __declspec(dllimport) struct xa_switch_t db2xa_switch_std;
#else
#define db2xa_switch_std (*db2xa_switch_std)
extern struct xa_switch_t db2xa_switch_std;
#endif

main( )
{
    struct xa_switch_t *foo;
    printf ( "switch_std.name );
    foo = db2xacic_std();
    printf ( "name );
    return ;
}

```

XA interface problem determination

When an error is detected during an XA request from the TM, the application program may not be able to get the error code from the TM. If your program abends, or gets a cryptic return code from the TP monitor or the TM, you should check the First Failure Service Log, which reports XA error information when diagnostic level 3 or greater is in effect.

You should also consult the console message, TM error file, or other product-specific information about the external transaction processing software that you are using.

The database manager writes all XA-specific errors to the First Failure Service Log with SQLCODE -998 (transaction or heuristic errors) and the appropriate reason codes. Following are some of the more common errors:

- Invalid syntax in the `xa_open` string.
- Failure to connect to the database specified in the open string as a result of one of the following:
 - The database has not been cataloged.
 - The database has not been started.
 - The server application's user name or password is not authorized to connect to the database.
- Communications error.

Configuration considerations for XA transaction managers

You should consider the following configuration parameters when you are setting up your TP monitor environment: `tp_mon_name`, `tm_database`, `maxappls`, and `autorestart`.

- **tp_mon_name**

This database manager configuration parameter identifies the name of the TP monitor product being used (for example, TXSeries CICS).

- **tm_database**

Because DB2 Database for Linux, UNIX, and Windows does *not* coordinate transactions in the XA environment, this database manager configuration parameter is not used for XA-coordinated transactions.

- **maxappls**

This database configuration parameter specifies the maximum number of active applications allowed. The value of this parameter must be equal to or greater than the sum of the connected applications, plus the number of these applications that may be concurrently in the process of completing a two-phase commit or rollback. This sum should then be increased by the anticipated number of indoubt transactions that might exist at any one time.

For a TP monitor environment (for example, TXSeries CICS), you may need to increase the value of the **maxappls** parameter. This would help to ensure that all TP monitor processes can be accommodated.

- **autorestart**

This database configuration parameter specifies whether the **RESTART DATABASE** routine will be invoked automatically when needed. The default value is YES (that is, enabled).

A database containing indoubt transactions requires a restart database operation to start up. If **autorestart** is not enabled when the last connection to the database is dropped, the next connection will fail and require an explicit **RESTART DATABASE** invocation. This condition will exist until the indoubt transactions have been removed, either by the transaction manager's resync operation, or through a heuristic operation initiated by the administrator. When the **RESTART DATABASE** command is issued, a message is returned if there are any indoubt transactions in the database. The administrator can then use the **LIST INDOUBT TRANSACTIONS** command and other command line processor commands to find get information about those indoubt transactions.

Resource manager setup: Each database is defined as a separate resource manager (RM) to the transaction manager (TM), and the database must be identified with an xa_open string.

When setting up a database as a resource manager, you do not need the xa_close string. If provided, this string will be ignored by the database manager.

Database connection considerations

Automatic client reroute (ACR)

Whenever a server crashes, each client that is connected to that server gets a communication error which terminates the connection and concludes in an application error. In application environments where availability is important, the user will either have a redundant setup or will fail the server over to a standby node. In either case, the DB2 Database for Linux, UNIX, and Windows client code will attempt to re-establish the connection to either the original database (which might be running on a failover node where the IP address fails over as well), or to a new database on a different server. The application is then notified using an SQLCODE to indicate that the connection has been rerouted and that the specific transaction being run has been rolled back. At that point, the application can choose to rerun that transaction or continue on.

Data consistency between the failed primary database and the "failed to" standby database when using ACR is very dependent upon the state of the database logs in the database to which the connection has been rerouted. For the purposes of this discussion, we will call this database the "standby database" and the server on which this standby database resides the "standby server". If the standby database is an exact copy of the failed primary database at the point in time of the failure then the data at the standby database will be consistent and there will be no data integrity issues. However, if the standby database is not an exact copy of the failed primary database then there might be data integrity issues resulting from inconsistent transaction outcomes for transactions which have been prepared by the XA Transaction Manager but yet to be committed. These are known as indoubt transactions. The Database Administrator and application developers who are using the ACR function must be aware of the risk of data integrity problems when using this capability.

The following sections describe the various DB2 Database for Linux, UNIX, and Windows environments and the risks of data integrity problems in each.

High availability disaster recovery (HADR)

The DB2 High Availability Disaster Recovery feature (HADR) can be used to control the level of log duplication between the primary and standby databases when the application regains connectivity after a primary database failure. The database configuration parameter which controls the level of log duplication is called **hadr_syncmode**. There are four possible values for this parameter:

- SYNC

This mode provides the greatest protection against transaction loss at the cost of longest transaction response time among the four modes. As the name of this mode suggests, SYNC is used to synchronize the writing of the transaction log in the primary database and in the standby database. Synchronization is accomplished when the primary database has written its own log files and it has received acknowledgement from the standby database that the logs have also been written on the standby database.

If an XA Transaction Manager is being used to coordinate transactions involving DB2 resources, then it is strongly recommended that SYNC mode be used. SYNC mode will guarantee data integrity as well as transaction resynchronization integrity when a client is rerouted to the standby database since it is an exact replica of the primary database.

- NEARSYNC

This mode provides slightly less protection against transaction loss, in exchange for a shorter transaction response time when compared with SYNC mode. The primary database considers log write successful only when logs have been written to its own log files and it has received acknowledgement from the standby database that the logs have also been written to main memory on the standby database. If the standby database crashes before it can copy the logs from memory to disk, the logs are lost on the standby database in the short term.

Given the possibility that database logs are lost, and the situation where the standby database is not an exact replica of the primary database, it is possible that data integrity will be compromised. The compromise occurs if the given transaction was indoubt and then the primary database crashes. Assume the transaction outcome is COMMIT. When the XA TM issues the subsequent XA_COMMIT request, it will fail since the primary database has crashed. Since the XA_COMMIT request has failed, the XA TM will need to recover this

transaction on this database by issuing an XA_RECOVER request. The standby database will respond by returning the list of all its transactions which are indoubt. If the standby database were to crash and restart before the “in memory,” database logs were written to disk, and before the XA_RECOVER request was issued by the XA TM, the standby database would have lost the log information about the transaction and could not return it in response to the XA_RECOVER request. The XA TM would then assume the database committed this transaction. But, what has really occurred is the data manipulation will have been lost and the appearance that the transaction was rolled back. This results in a data integrity issue since all other resources involved in this transaction were COMMITTED by the XA TM.

Using NEARSYNC is a good compromise between data integrity and transaction response time since the likelihood of both the primary and standby databases crashing should be low. However, a database administrator still needs to understand that there is a possibility of data integrity problems.

- ASYNC

Compared with the SYNC and NEARSYNC modes, the ASYNC mode results in shorter transaction response times but might cause greater transaction losses if the primary database fails. The primary database considers log write successful only when logs have been written to its own log files and the logs have been delivered to the TCP layer on the primary database's host machine. The primary database does not wait for acknowledgement of any kind from the standby database. The logs might be still on their way to the standby database when the primary database considers relevant transactions committed.

If the same scenario as described in NEARSYNC occurs, the likelihood of loss of transaction information is higher than with NEARSYNC. Therefore, the likelihood of data integrity issues is higher than with NEARSYNC and SYNC.

- SUPERASYNC

This mode has the shortest transaction response time but has also the highest probability of transaction losses if the primary system fails. This mode is useful when you do not want transactions to be blocked or experience elongated response times due to network interruptions or congestion. The primary database considers log write successful only when logs have been written to its own log files. The primary database does not wait for acknowledgement of any kind from the standby database. The logs might be still on their way to the standby database when the primary database considers relevant transactions committed.

If the same scenario as described in NEARSYNC occurs, the likelihood of loss of transaction information is higher than with NEARSYNC. Therefore, the likelihood of data integrity issues is higher than with NEARSYNC and SYNC.

Partitioned database environments

The use of automatic client reroute (ACR) in partitioned database environments can also lead to data integrity issues. If the standby database is defined to be a different database partition of the same database, then recovery of indoubt transactions in scenarios as described in the High Availability Disaster Recovery NEARSYNC section, can result in data integrity problems. This occurs because the database partitions do not share database transaction logs. Therefore the standby database (database partition B) will have no knowledge of indoubt transactions that exist at the primary database (database partition A).

Non-partitioned database environments

The use of ACR in non-partitioned database environments can also lead to data integrity issues. Assuming disk failover technology, such as IBM PowerHA[®] SystemMirror for AIX, Microsoft Cluster Service (MSCS), or HP's Service Guard, is not in use then the standby database will not have the database transaction logs that existed on the primary database when it failed. Therefore, the recovery of indoubt transactions in scenarios as described in the High Availability Disaster Recovery NEARSYNC section, can result in data integrity problems.

Transactions accessing partitioned databases

In a partitioned database environment, user data might be distributed across database partitions. An application accessing the database connects and sends requests to the database partition identified as the coordinator node. Different applications can connect to different database partitions, and the same application can choose different database partitions for different connections.

For transactions against a database in a partitioned database environment, all access must be through the *same* database partition. That is, the same database partition must be used from the start of the transaction until (and including) the time that the transaction is committed.

Any transaction against the partitioned database must be committed before disconnecting.

Global variables and XA transactions:

In an XA environment, session global variable objects will not change. Session global variables, within an XA transaction, are not available to other sessions joining the transaction.

Introduction

If a global variable is created within a session, then it cannot be used by other sessions until the unit of work has committed. However, the newly created variable can be used within the session that created it before the unit of work is committed. This behavior is consistent with other created objects such as tables.

The length attribute and data type of a global variable does not vary once it is created. Furthermore, the global variable cannot be altered at all (for example, its default value).

The setting of a global variable's value is non-transactional. Hence, an application cannot rollback the setting of the value of a global variable. Note the following:

- If the creation of a global variable is rolled back then the variable no longer exists.
- If the drop of a global variable is rolled back, the value of the global variable will be what it was before the drop.

Session global variables can have a default value defined for them. This value is automatically assigned to the variable as soon as it is first referenced in a session.

XA interaction

Session global variables are scoped to a per session level within an XA transaction and are not available to other sessions that join the transaction.

Global variables are non-transactional state objects. Hence, in an XA environment (for both tightly coupled transactions and loosely coupled transactions), global variable objects (both existence and value) will not change at XA_end (nor for other XA APIs, like XA_start, XA_rollback).

Note: In contrast, DB2 for Linux, UNIX and Windows has implemented declared global temporary tables as transactional state objects. Hence, when the following XA APIs are executed, declared global temporary tables are dropped.

- XA_rollback
- XA_end(SUCCESS)
- XA_start(NOFLAGS)
- XA_start(JOIN)

Configuring IBM WebSphere Application Server:

About this task

IBM WebSphere Application Server is a Java-based application server. It can use the DB2 Database for Linux, UNIX, and Windows XA support via the Java Transaction API (JTA) provided by the DB2 JDBC driver. Refer to IBM WebSphere documentation regarding how to use the Java Transaction API with WebSphere Application Server. WebSphere Application Server documentation can be viewed online at <http://www.ibm.com/software/webservers/appserv/was/library/>.

Configuring IBM WebSphere MQ:

About this task

IBM WebSphere MQ is a message oriented business integration server. It can use the DB2 Database for Linux, UNIX, and Windows XA support to integrate database updates within WebSphere MQ units of work. Refer to IBM WebSphere MQ documentation regarding how to use the transactional support in WebSphere MQ. WebSphere MQ documentation can be viewed online at <http://publib.boulder.ibm.com/infocenter/wmqv7/v7r0/index.jsp>.

Configuring IBM TXSeries CICS:

About this task

For information about how to configure IBM TXSeries CICS to use DB2 Database for Linux, UNIX, and Windows as a resource manager, refer to your *IBM TXSeries CICS Administration Guide*. TXSeries documentation can be viewed online at <http://www-01.ibm.com/software/htp/cics/txseries/library/>. Host and System i database servers can participate in CICS-coordinated transactions.

Configuring BEA Tuxedo:

What follows is a description of the process to configure BEA Tuxedo for use with DB2 Database for Linux, UNIX, and Windows. There are some differences that are noted based on whether Tuxedo is working with a 64-bit instance of DB2 Database for Linux, UNIX, and Windows or a 32-bit instance of DB2 Database for Linux, UNIX, and Windows.

About this task

Note: There are new names for the XA switch data structures: *db2xa_switch_std* and *db2xa_switch_static_std*. There are also new names for the APIs: *db2xacic* and *db2xacicst*. The old switch data structure and API names can be used but only when working with a 32-bit instance of DB2 Database for Linux, UNIX, and Windows.

Procedure

To configure Tuxedo to use DB2 Database for Linux, UNIX, and Windows as a resource manager:

1. Install Tuxedo as specified in the documentation for that product. Ensure that you perform all basic Tuxedo configuration, including the log files and environment variables. You also require a compiler and the Data Server Client. Install these if necessary.
2. At the Tuxedo server ID, set the **DB2INSTANCE** environment variable to reference the instance that contains the databases that you want Tuxedo to use. Set the **PATH** variable to include the DB2 database program directories. Confirm that the Tuxedo server ID can connect to the DB2 databases.
3. Update the **tp_mon_name** database manager configuration parameter with the value TUXEDO.
4. Add a definition for the DB2 database product to the Tuxedo resource manager definition file. In the examples that follow, **UDB_XA** is the locally-defined Tuxedo resource manager name for DB2 V9.1, and *db2xa_switch_std* is the DB2 database-defined name for a structure of type *xa_switch_t*:

- For AIX. In the file `${TUXDIR}/udataobj/rm`, add the definition:

```
# DB2 UDB
UDB_XA:db2xa_switch_std:-L${DB2DIR} /lib -ldb2
```

where `{TUXDIR}` is the directory where you installed Tuxedo, and `{DB2DIR}` is the DB2 instance directory.

- For Windows. In the file `%TUXDIR%\udataobj\rm`, add the definition:

```
# DB2 UDB
UDB_XA;db2xa_switch_std;%DB2DIR%\lib\db2api.lib
```

where `%TUXDIR%` is the directory where you installed Tuxedo, and `%DB2DIR%` is the DB2 instance directory.

5. Build the Tuxedo transaction monitor server program for DB2:

- For AIX:

```
${TUXDIR}/bin/buildtms -r UDB_XA -o ${TUXDIR}/bin/TMS_UDB
```

where `{TUXDIR}` is the directory where you installed Tuxedo.

- For Windows:

```
XA -o UDB
```

6. Build the application servers. In the examples that follow, the `-r` option specifies the resource manager name, the `-f` option (used one or more times) specifies the files that contain the application services, the `-s` option specifies the application service names for this server, and the `-o` option specifies the output server file name:

- For AIX:

```
${TUXDIR}/bin/buildserver -r UDB_XA -f svcfile.o -s SVC1,SVC2
-o UDBserver
```

where {TUXDIR} is the directory where you installed Tuxedo.

- For Windows:

```
XA -f svcfile.o -s SVC1,SVC2
-o UDBserver
```

where %TUXDIR% is the directory where you installed Tuxedo.

7. Set up the Tuxedo configuration file to reference the DB2 server. In the *GROUPS section of the UDBCONFIG file, add an entry similar to:

```
UDB_GRP LMID=simp GRPNO=3
TMSNAME=TMS_UDB TMSCOUNT=2
OPENINFO="UDB_XA:db=sample,uid=db2_user,pwd=db2_user_pwd"
```

where the TMSNAME parameter specifies the transaction monitor server program that you built previously, and the OPENINFO parameter specifies the resource manager name. This is followed by the database name, and the DB2 database user ID and password, which are used for authentication. The application servers that you built previously are referenced in the *SERVERS section of the Tuxedo configuration file.

8. If the application is accessing data residing on DB2 for z/OS and OS/390, DB2 for IBM i, or DB2 for VM&VSE, the DB2 Connect XA concentrator will be required.
9. Start Tuxedo:

```
tmbboot -y
```

After the command completes, Tuxedo messages should indicate that the servers are started. In addition, if you issue the DB2 command **LIST APPLICATIONS ALL**, you should see two connections (in this situation) specified by the TMSCOUNT parameter in the UDB_GRP in the Tuxedo configuration file, UDBCONFIG.

Setting up the unixODBC driver manager

To run ODBC applications on Linux or UNIX operating systems, you must configure the unixODBC Driver Manager.

Procedure

To set up the unixODBC Driver Manager for use with CLI and ODBC applications:

1. Download the latest unixODBC source code from <http://www.unixodbc.org>.

2. Untar the source files. For example:

```
gzip -d unixODBC-2.2.11.tar.gz
tar xf unixODBC-2.2.11.tar
```

3. For AIX only: Configure the C compiler to be thread-enabled:

```
export CC=xlc_r
export CCC=xlc_r
```

4. To compile a 64-bit version of the driver manager using the xlc_r compilers, set the environment variables **OBJECT_MODE** and **CFLAGS**:

```
export OBJECT_MODE=64
export CFLAGS=-q64 -DBUILD_REAL_64_BIT_MODE
```

5. Install the driver manager in either your home directory or the default /usr/local prefix:

- (Home directory) Issue the following command in the directory where you untarred the source files:

```
./configure --prefix=$HOME -DBUILD_REAL_64_BIT_MODE --enable-gui=no
--enable-drivers=no
```

- (/usr/local as root) Issue the following command:

```
./configure --enable-gui=no --enable-drivers=no
```

6. Optional: Examine all configuration options by issuing the following command:

```
./configure --help
```

7. Build and install the driver manager:

```
make
make install
```

Libraries will be copied to the [prefix]/lib directory, and executables will be copied to the [prefix]/bin directory.

8. For AIX only: Extract the shared library from the ODBC driver for DB2 to yield shr.o on 32-bit operating systems and shr_64.o on 64-bit operating systems. To avoid confusion, rename the files db2.o and db2_64.o. These steps are necessary on AIX because the unixODBC Driver Manager loads the driver dynamically.

- On 32-bit operating systems, issue the following commands:

```
cd INSTHOME/sql/lib/lib
ar -x libdb2.a
mv shr.o db2.o
```

where *INSTHOME* is the home directory of the instance owner.

- On 64-bit operating systems, issue the following commands:

```
cd INSTHOME/sql/lib/lib
ar -x -X 64 libdb2.a
mv shr_64.o db2_64.o
```

where *INSTHOME* is the home directory of the instance owner.

Ensure that your INI file references the correct library.

9. Optional: For AIX only: Extract libodbc.a, libodbcinst.a, and libodbccr.a if you will be dynamically loading the driver manager:

```
ar -x libodbc.a
ar -x libodbcinst.a
ar -x libodbccr.a
```

This produces libodbc.so.1, libodbcinst.so.1, and libodbccr.so.1 in the [prefix]/lib/so directory.

10. Build the application and ensure it is linked to the unixODBC Driver Manager by including the -L[prefix]/lib -lodbc option in the compile and link command.
11. Specify the paths for at least the user INI file (odbc.ini) or the system INI file (odbcinst.ini), and set the **ODBCHOME** environment variable to the directory where the system INI file was created.

Important: Provide absolute paths when specifying the paths of the user and system INI files. Do not use relative paths or environment variables.

Note: If compiling 64-bit applications for the ODBC Driver, use the -DODBC64 option to enable the 64-bit definitions in the driver manager.

Appendix A. DB2 database product and packaging information

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

The following tables list the DB2 database products and features that are available on each operating system.

Table 20. Availability of DB2 database products and features on UNIX operating systems

DB2 database products and features	AIX	HP-UX on IA-64 hardware	Solaris on UltraSPARC hardware	Solaris on x86-64 ("x64") hardware
DB2 Advanced Enterprise Server Edition	Yes	Yes	Yes	Yes
DB2 Enterprise Server Edition	Yes	Yes	Yes	Yes
DB2 Workgroup Server Edition	Yes	Yes	Yes	Yes
DB2 Personal Edition	No	No	No	No
DB2 Express Edition	No	No	No	Yes
DB2 Express-C	No	No	No	Yes
Data Server Client and Data Server Runtime Client	Yes	Yes	Yes	Yes
DB2 Advanced Access Control Feature	Yes	Yes	Yes	Yes
IBM DB2 High Availability Feature for Express Edition	No	No	No	Yes
IBM Homogeneous Replication Feature for DB2 Enterprise Server Edition	Yes	Yes	Yes	Yes
IBM DB2 Performance Optimization Feature for Enterprise Server Edition	Yes	Yes	Yes	Yes
DB2 Storage Optimization Feature	Yes	Yes	Yes	Yes

Table 21. Availability of DB2 database products and features on Linux operating systems

DB2 database products and features	Linux on x86-32 hardware	Linux on x64 hardware	Linux on IBM Power Systems	Linux on System z
DB2 Advanced Enterprise Server Edition	No ¹	Yes	Yes	Yes
DB2 Enterprise Server Edition	No ¹	Yes	Yes	Yes
DB2 Workgroup Server Edition	Yes	Yes	Yes	No
DB2 Personal Edition	Yes	Yes	No	No
DB2 Express Edition	Yes	Yes	Yes	No
DB2 Express-C	Yes	Yes	Yes	No
Data Server Client and Data Server Runtime Client	Yes	Yes	Yes	Yes
DB2 Advanced Access Control Feature	No ¹	Yes	Yes	Yes
DB2 Geodetic Data Management Feature	No ¹	Yes	No	Yes
IBM DB2 High Availability Feature for Express Edition	Yes	Yes	Yes	No
IBM Homogeneous Replication Feature for DB2 Enterprise Server Edition	No ¹	Yes	Yes	Yes

Table 21. Availability of DB2 database products and features on Linux operating systems (continued)

DB2 database products and features	Linux on x86-32 hardware	Linux on x64 hardware	Linux on IBM Power Systems	Linux on System z
IBM DB2 Performance Optimization Feature for Enterprise Server Edition	No ¹	Yes	Yes	Yes
DB2 Storage Optimization Feature	No ¹	Yes	Yes	Yes
Note:				
1. Only test and development are supported in DB2 Enterprise Server Edition for Linux on x86-32 hardware. However, you can use DB2 Enterprise Server Edition for Linux on x86-32 as a try and buy or by installing the IBM Database Enterprise Developer Edition license certificate onto a trial copy.				

Table 22. Availability of DB2 database products and features on Windows operating systems

DB2 database products and features	Windows XP, Windows Vista, Windows 7 (32-bit editions)	Windows XP, Windows Vista, Windows 7 (64-bit editions) on x86 64-bit ("x64") hardware	Windows Server 2003, Windows Server 2008 (32-bit editions)	Windows Server 2003, Windows Server 2008 (64-bit editions) on x64 hardware
DB2 Advanced Enterprise Server Edition	No ¹	No ¹	Yes	Yes
DB2 Enterprise Server Edition	No ¹	No ¹	Yes	Yes
DB2 Workgroup Server Edition	Yes	Yes	Yes	Yes
DB2 Personal Edition	Yes	Yes	Yes	Yes
DB2 Express Edition	Yes	Yes	Yes	Yes
DB2 Express-C	Yes	Yes	Yes	Yes
Data Server Client and Data Server Runtime Client	Yes	Yes	Yes	Yes
DB2 Advanced Access Control Feature	No ¹	No ¹	Yes	Yes
DB2 Geodetic Data Management Feature	No ¹	No ¹	Yes	Yes
IBM DB2 High Availability Feature for Express Edition	Yes	Yes	Yes	Yes
IBM Homogeneous Replication Feature for DB2 Enterprise Server Edition	No ¹	No ¹	Yes	Yes
IBM DB2 Performance Optimization Feature for Enterprise Server Edition	No ¹	No ¹	Yes	Yes
DB2 Storage Optimization Feature	No ¹	No ¹	Yes	Yes
Note:				
1. Only test and development are supported in DB2 Advanced Enterprise Server Edition, DB2 Enterprise Server Edition, and IBM Database Enterprise Developer Edition on Windows Vista Ultimate, Enterprise and Business editions and Windows XP Professional edition.				

For information about DB2 database product and packaging information, see <http://www-1.ibm.com/support/docview.wss?rs=73&uid=swg21219983>.

For information about DB2 Express-C, see: www.ibm.com/software/data/db2/express

Appendix B. Overview of the DB2 technical information

DB2 technical information is available in multiple formats that can be accessed in multiple ways.

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

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

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

You can access additional DB2 technical information such as technotes, white papers, and IBM Redbooks® publications online at [ibm.com](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/e-business/linkweb/publications/servlet/pbi.wss. English and translated DB2 Version 10.1 manuals in PDF format can be downloaded from www.ibm.com/support/docview.wss?rs=71&uid=swg2700947.

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

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

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

Table 23. DB2 technical information

Name	Form Number	Available in print	Last updated
<i>Administrative API Reference</i>	SC27-3864-00	Yes	April, 2012
<i>Administrative Routines and Views</i>	SC27-3865-00	No	April, 2012
<i>Call Level Interface Guide and Reference Volume 1</i>	SC27-3866-00	Yes	April, 2012
<i>Call Level Interface Guide and Reference Volume 2</i>	SC27-3867-00	Yes	April, 2012
<i>Command Reference</i>	SC27-3868-00	Yes	April, 2012
<i>Database Administration Concepts and Configuration Reference</i>	SC27-3871-00	Yes	April, 2012
<i>Data Movement Utilities Guide and Reference</i>	SC27-3869-00	Yes	April, 2012
<i>Database Monitoring Guide and Reference</i>	SC27-3887-00	Yes	April, 2012
<i>Data Recovery and High Availability Guide and Reference</i>	SC27-3870-00	Yes	April, 2012
<i>Database Security Guide</i>	SC27-3872-00	Yes	April, 2012
<i>DB2 Workload Management Guide and Reference</i>	SC27-3891-00	Yes	April, 2012
<i>Developing ADO.NET and OLE DB Applications</i>	SC27-3873-00	Yes	April, 2012
<i>Developing Embedded SQL Applications</i>	SC27-3874-00	Yes	April, 2012
<i>Developing Java Applications</i>	SC27-3875-00	Yes	April, 2012
<i>Developing Perl, PHP, Python, and Ruby on Rails Applications</i>	SC27-3876-00	No	April, 2012
<i>Developing User-defined Routines (SQL and External)</i>	SC27-3877-00	Yes	April, 2012
<i>Getting Started with Database Application Development</i>	GI13-2046-00	Yes	April, 2012

Table 23. DB2 technical information (continued)

Name	Form Number	Available in print	Last updated
<i>Getting Started with DB2 Installation and Administration on Linux and Windows</i>	GI13-2047-00	Yes	April, 2012
<i>Globalization Guide</i>	SC27-3878-00	Yes	April, 2012
<i>Installing DB2 Servers</i>	GC27-3884-00	Yes	April, 2012
<i>Installing IBM Data Server Clients</i>	GC27-3883-00	No	April, 2012
<i>Message Reference Volume 1</i>	SC27-3879-00	No	April, 2012
<i>Message Reference Volume 2</i>	SC27-3880-00	No	April, 2012
<i>Net Search Extender Administration and User's Guide</i>	SC27-3895-00	No	April, 2012
<i>Partitioning and Clustering Guide</i>	SC27-3882-00	Yes	April, 2012
<i>pureXML Guide</i>	SC27-3892-00	Yes	April, 2012
<i>Spatial Extender User's Guide and Reference</i>	SC27-3894-00	No	April, 2012
<i>SQL Procedural Languages: Application Enablement and Support</i>	SC27-3896-00	Yes	April, 2012
<i>SQL Reference Volume 1</i>	SC27-3885-00	Yes	April, 2012
<i>SQL Reference Volume 2</i>	SC27-3886-00	Yes	April, 2012
<i>Text Search Guide</i>	SC27-3888-00	Yes	April, 2012
<i>Troubleshooting and Tuning Database Performance</i>	SC27-3889-00	Yes	April, 2012
<i>Upgrading to DB2 Version 10.1</i>	SC27-3881-00	Yes	April, 2012
<i>What's New for DB2 Version 10.1</i>	SC27-3890-00	Yes	April, 2012
<i>XQuery Reference</i>	SC27-3893-00	No	April, 2012

Table 24. DB2 Connect-specific technical information

Name	Form Number	Available in print	Last updated
<i>DB2 Connect Installing and Configuring DB2 Connect Personal Edition</i>	SC27-3861-00	Yes	April, 2012
<i>DB2 Connect Installing and Configuring DB2 Connect Servers</i>	SC27-3862-00	Yes	April, 2012
<i>DB2 Connect User's Guide</i>	SC27-3863-00	Yes	April, 2012

Displaying SQL state help from the command line processor

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

Procedure

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

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

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

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

Accessing different versions of the DB2 Information Center

Documentation for other versions of DB2 products is found in separate information centers on ibm.com[®].

About this task

For DB2 Version 10.1 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1>.

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

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

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

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

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

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

A locally installed DB2 Information Center must be updated periodically.

Before you begin

A DB2 Version 10.1 Information Center must already be installed. For details, see the “Installing the DB2 Information Center using the DB2 Setup wizard” topic in *Installing DB2 Servers*. All prerequisites and restrictions that applied to installing the Information Center also apply to updating the Information Center.

About this task

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

- Automatic updates update existing Information Center features and languages. One benefit of automatic updates is that the Information Center is unavailable for a shorter time compared to during a manual update. In addition, automatic updates can be set to run as part of other batch jobs that run periodically.
- Manual updates can be used to update existing Information Center features and languages. Automatic updates reduce the downtime during the update process, however you must use the manual process when you want to add features or languages. For example, a local Information Center was originally installed with both English and French languages, and now you want to also install the German language; a manual update will install German, as well as, update the existing Information Center features and languages. However, a manual update requires you to manually stop, update, and restart the Information Center. The Information Center is unavailable during the entire update process. In the automatic update process the Information Center incurs an outage to restart the Information Center after the update only.

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

Procedure

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

1. On Linux operating systems,
 - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `/opt/ibm/db2ic/V10.1` directory.
 - b. Navigate from the installation directory to the `doc/bin` directory.
 - c. Run the `update-ic` script:
`update-ic`
2. On Windows operating systems,
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `<Program Files>\IBM\DB2 Information Center\Version 10.1` 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 `update-ic.bat` file:
`update-ic.bat`

Results

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

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

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

About this task

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

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

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

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

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

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

Procedure

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

1. Stop the *DB2 Information Center*.
 - On Windows, click **Start > Control Panel > Administrative Tools > Services**. Then right-click **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:

```
/etc/init.d/db2icdv10 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 10.1* 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/V10.1` directory.
 - b. Navigate from the installation directory to the `doc/bin` directory.
 - c. Run the `help_start` script:

```
help_start
```

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

3. Click the **Update** button (🔧). (JavaScript must be enabled in your browser.) On the right panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check that 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 `doc\bin` directory within the installation directory, and run the `help_end.bat` file:

```
help_end.bat
```

Note: The `help_end` batch file contains the commands required to safely stop the processes that were started with the `help_start` batch file. Do not use `Ctrl-C` or any other method to stop `help_start.bat`.
 - On Linux, navigate to the `doc/bin` directory within the installation directory, and run the `help_end` script:

```
help_end
```

Note: The `help_end` script contains the commands required to safely stop the processes that were started with the `help_start` script. Do not use any other method to stop the `help_start` script.
7. Restart the *DB2 Information Center*.
 - On Windows, click **Start > Control Panel > Administrative Tools > Services**. Then right-click **DB2 Information Center** service and select **Start**.
 - On Linux, enter the following command:

```
/etc/init.d/db2icdv10 start
```

Results

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

DB2 tutorials

The DB2 tutorials help you learn about various aspects of DB2 database 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/db2luw/v10r1/>.

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

DB2 tutorials

To view the tutorial, click the title.

“pureXML” in *pureXML Guide*

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

DB2 troubleshooting information

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

DB2 documentation

Troubleshooting information can be found in the *Troubleshooting and Tuning Database Performance* or the Database fundamentals section of the *DB2 Information Center*, which contains:

- Information about how to isolate and identify problems with DB2 diagnostic tools and utilities.
- Solutions to some of the most common problem.
- Advice to help solve other problems you might encounter with your DB2 database products.

IBM Support Portal

See the IBM Support Portal 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 IBM Support Portal at http://www.ibm.com/support/entry/portal/Overview/Software/Information_Management/DB2_for_Linux,_UNIX_and_Windows

Terms and conditions

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

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal use: You may reproduce these publications for your personal, noncommercial 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.

Rights: 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.

IBM Trademarks: IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml

Appendix C. Notices

This information was developed for products and services offered in the U.S.A. Information about non-IBM products is based on information available at the time of first publication of this document and is subject to change.

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

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

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

Any references in this information to websites not owned by IBM are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

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

IBM Canada Limited
U59/3600
3600 Steeles Avenue East
Markham, Ontario L3R 9Z7
CANADA

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

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

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

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

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

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating

platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

© (*your company name*) (*year*). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. *_enter the year or years_*. All rights reserved.

Trademarks

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

The following terms are trademarks or registered trademarks of other companies

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle, its affiliates, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Celeron, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Index

Special characters

.NET
application development software 31

A

ActiveX Data Object (ADO) specification
IBM Data Server Provider for .NET 20
ActiveX data objects
DB2 Data Server Client support 17
ADO.NET applications
developing 19
AIX
installing
DB2 server products 5
IBM data server clients 5
APIs
heuristic 65
overview 19
application development
application programming interfaces (APIs)
configuring 49
summary 19
configuring
environment 45
environment (Windows) 47
environment for using application programming
interfaces (APIs) 49
IBM Data Server Client 48
operating systems 46
elements supported 3
embedded SQL overview 24
environment
configuring 45
overview 1
IBM Data Server Client
configuring 48
details 17
IBM Data Server Provider for .NET 20
Linux
configuring 46
operating systems
configuring 46
Linux 46
supported 5
UNIX 46
Perl DBI 27
programming languages 31
tools 3, 43
transaction managers 41
UNIX
configuring 46
Windows
configuring 47
applications
development tools 43
programming interfaces 19
automatic client reroute
details 70

B

BEA Tuxedo
configuring 75

C

C language
versions supported 32
C/C++ language
versions supported 33
call level interface (CLI)
environment setup 50
IBM Data Server Client support 17
overview 22
client reroute
automatic 70
clients
configuring 48
COBOL language
versions supported 34
command-line processor (CLP)
IBM Data Server Client 17
connection keywords
ODBC .NET Data Provider 20
OLE DB .NET Data Provider 21
connections
failures
automatic client reroute 70
core level functions
ODBC 22

D

Data Studio
development tools 43
database applications
programming interfaces 19
DB2 Information Center
updating 84, 86
versions 84
DB2 products
general information 79
packaging 79
Distributed Transaction Processing (DTP)
application program (AP) 54
configuring 70
databases
connections 70
updates (host and System i) 63
error handling 64
resource manager (RM) 54
security 66
transaction manager (TM) 54
X/Open Distributed Transaction Processing (DTP)
model 54
documentation
overview 81
PDF files 81
printed 81
terms and conditions of use 88

DTP (Distributed Transaction Processing)
see Distributed Transaction Processing (DTP) 54

E

embedded SQL applications
IBM Data Server Client support 17
overview 24
environment variables
UNIX
overview 47

F

FORTRAN language
supported versions 35

G

global variables
XA transactions 73

H

hardware
requirements
DB2 server products (AIX) 5
DB2 server products (HP-UX) 7
DB2 server products (Linux) 8
DB2 server products (Solaris Operating Environment) 12
DB2 server products (Windows) 13
IBM data server clients (AIX) 5
IBM data server clients (HP-UX) 7
IBM data server clients (Linux) 8
IBM data server clients (Solaris Operating Environment) 12
IBM data server clients (Windows) 13
help
SQL statements 84
heuristic decisions
resolving indoubt transactions 64
heuristic operations
resolving indoubt transactions
manually 64
host databases
updating with XA transaction managers 63
HP-UX
installing
DB2 servers 7
IBM data server clients 7

I

IBM data server clients
application development 3
IBM Data Server Driver for JDBC and SQLJ
installing 50
levels for DB2 versions 36
IBM Data Server Provider for .NET
documentation 19
overview 20
IBM Database Add-Ins for Visual Studio
development tools 43

IBM OLE DB Provider
consumers 26
installing 26
overview 26
providers 26
system requirements 26
versions 26
IBM TXSeries CICS
configuring 74
ibm_db API
details 29
IBM_DB Ruby driver and Rails adapter
details 29
ibm_db_dbi API
details 29
ibm_db_sa adaptor
details 29
ibm_db2 API
details 28
indoubt transactions
recovering
X/Open Distributed Transaction Processing (DTP)
model 54
resolving 64
installation
IBM Data Server Driver for JDBC and SQLJ 50
requirements
AIX 5
HP-UX 7
Linux 8
Solaris Operating Environment 12
Windows 13

J

Java
DB2 product support 36
IBM Data Server Client support 17
JDBC
drivers
details 25, 36
IBM Data Server Client support 17
IBM Data Server Driver for JDBC and SQLJ installation 50

L

Linux
configuring
application development 46
installing
DB2 servers 8
IBM data server clients 8
libraries
libaio.so.1 8
libstdc so.5 8
LIST INDOUBT TRANSACTIONS command
resolving indoubt transactions manually 64

M

Microsoft Visual Studio .NET
viewing IBM Data Server Provider for .NET
documentation 19

N

notices 91

O

ODBC

- CLI 22
- core level functions 22
- driver managers
 - unixODBC 76

ODBC .NET Data Provider

- overview 20

OLE

- IBM Data Server Client support 17

OLE DB

- table functions 26

OLE DB .NET Data Provider

- overview 21

operating systems

- application development 3, 5
- configuring 46
- installation requirements
 - AIX 5
 - HP-UX 7
 - Linux 8
 - Solaris Operating Environment 12
 - Windows 13

P

partitioned databases

- transactions
 - accessing databases 70

pdo_ibm

- details 28

Perl

- Database Interface (DBI) specification 27
- versions 38

PHP

- application development 28
- extensions for IBM data servers 28
- versions 39

precompilers

- IBM Data Server Client support 17

problem determination

- information available 88
- tutorials 88

programming languages

- C 32
- C++ 33
- COBOL 34
- FORTRAN 35
- Perl 38
- PHP 39
- REXX 40
- Ruby on Rails 39
- summary 31

Python

- application development 29
- extensions for IBM data servers 29

R

Rails adapter

- details 29

Rational Application Developer (RAD)

- supported development tools 43

Rational Data Architect

- supported development tools 43

Rational Software Architect

- supported development tools 43

resolution

- indoubt transactions 64

resource managers (RM)

- overview 54
- setting up database as 70

REXX language

- versions supported 40

Ruby driver

- details 29

Ruby on Rails

- development tools 43

Ruby/Ruby on Rails language

- supported versions 39

S

SAMPLE database

- connecting
 - ODBC .NET Data Provider 20
 - OLE DB .NET Data Provider 21

SDKs

- product levels 36

software requirements

- AIX 5
- HP-UX 7
- Linux 8
- Solaris Operating Environment 12
- Windows 13

Solaris operating systems

- installation requirements
 - DB2 servers 12
 - IBM data server clients 12

SQL Access Group 22

SQL Editor 43

SQL statements

- dynamic SQL 24
- help
 - displaying 84
- static SQL 24

SQLAlchemy

- adapter for IBM data servers 29

SQLJ

- drivers 25
- IBM Data Server Client support 17

static profiling

- development tools 43

System i

- databases
 - updating with XA transaction managers 63

system requirements

- IBM OLE DB Provider for DB2 26

T

terms and conditions

- publications 88

threads

- IBM OLE DB Provider for DB2 26

TPM values 58

TPMONNAME values 58

- transaction managers
 - BEA Tuxedo 75
 - distributed transaction processing 54
 - IBM TXSeries CICS 74
 - IBM WebSphere Application Server 74
 - IBM WebSphere MQ 74
 - overview 54
 - problem determination 69
 - XA architecture 67
 - XA-compliant 41
- transaction processing monitors
 - BEA Tuxedo 75
 - configuration considerations 70
 - IBM TXSeries CICS 74
 - security 66
- transactions
 - accessing partitioned databases 70
 - global 54
 - loosely coupled
 - X/Open Distributed Transaction Processing (DTP) model 54
 - non-XA 54
 - tightly coupled 54
 - two-phase commit 54
- troubleshooting
 - online information 88
 - tutorials 88
- tutorials
 - list 87
 - problem determination 88
 - pureXML 87
 - troubleshooting 88
- Tuxedo
 - configuring 75
- TXSeries CICS 74

U

- UNIX
 - application development
 - environment variable settings 47
 - configuring
 - application development 46
- unixODBC driver manager
 - setting up 76
- updates
 - DB2 Information Center 84, 86

W

- WebSphere Application Server
 - overview 74
- WebSphere MQ
 - overview 74
- Windows
 - application development environment configuration 47
 - installing
 - DB2 servers (requirements) 13
 - IBM data server clients (requirements) 13

X

- X/Open Company 22
- X/Open distributed transaction processing (DTP) model
 - details 54
- X/Open SQL CLI 22

- XA
 - interface 54
 - specification 67
 - switch 67
 - transactions 73
- XA transaction managers
 - configuration considerations 70
 - list 41
 - security considerations 66
 - troubleshooting 69
 - updating host databases 63
 - updating System i databases 63

Z

- Zend Core 43



Printed in USA

GI13-2046-00



Spine information:

IBM DB2 10.1 for Linux, UNIX, and Windows

Getting Started with Database Application Development

