

---

## Product Overview

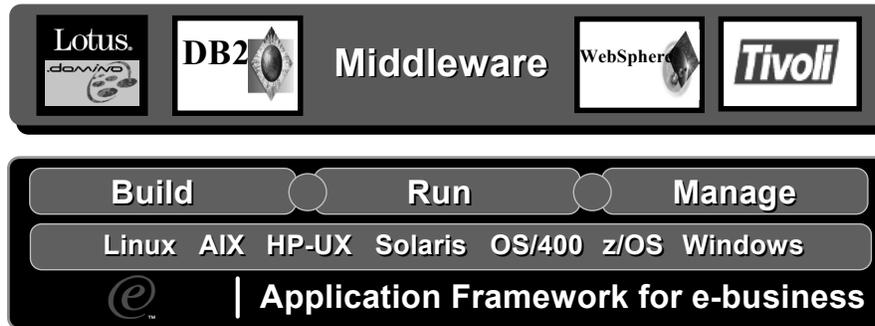


- ◆ DB2 UNIVERSAL DATABASE
- ◆ DB2 CONNECTIVITY
- ◆ DB2 APPLICATION DEVELOPMENT
- ◆ DB2 ADMINISTRATION

*I*n this chapter, you will be introduced to the DB2 Universal Database (DB2 UDB) family of products for UNIX and Intel platforms. DB2 has the ability to store all kinds of electronic information. This includes traditional relational data as well as structured and unstructured binary information, documents and text in many languages, graphics, images, multimedia (audio and video), information specific to operations like engineering drawings, maps, insurance claims forms, numerical control streams, or any type of electronic information. This chapter will illustrate some of the ways to access data in a DB2 database using some of the interfaces provided with the products. A description of each of the DB2 products will be provided to illustrate some of the features and functions.

The DB2 database is an important part of IBM's e-business software portfolio. The e-business Application Framework provides an open blueprint on how to build e-business applications. Popular IBM e-business tools include WebSphere Studio for developing Java programs or components and Tivoli software for distributed systems management. As for application server software, IBM offers several types of servers depending on the business requirement, from WebSphere Message

Queuing software to Java-based transaction processing with Websphere Application Server. The most popular IBM software servers are its database servers, specifically the DB2 Family.



**Fig. 1–1** *IBM Application Framework for e-business*

The DB2 Family executes on pervasive devices, Intel, UNIX, iSeries, and main-frame platforms. Supported operating environments include: Linux, Windows 2000/NT/XP/.NET, AIX, HP-UX, Sun Solaris, OS/400, VSE/VM, and OS/390. The DB2 code base is optimized for each platform to ensure maximum performance. The SQL API is common to all platforms, which allows applications written on one platform to access data on any platform. Internally, DB2 running on the OS/400, VM/VSE, or OS/390 environments differ from DB2 on the Intel and UNIX platforms, but it is the common SQL API that enables applications to work together. The DB2 code base on Intel and UNIX platforms are identical.

DB2 V8.1 for Linux, Windows, and UNIX provides seamless database connectivity using the most popular network communications protocols, including NetBIOS, TCP/IP, Named Pipes, and APPC.

## DB2 and e-business

As a core component of IBM's e-business Application Framework, DB2 is a catalyst for delivering applications that change the way a company performs its operations. The transformation process of taking a business to an e-business, often centers around applications that impact the entire operations of a company, they can include: electronic commerce, Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), Purchasing and Logistics Management (PLM), and Supply Chain Management (SCM). It is the process of re-engineering the company to integrate these new applications that enables a company to leverage these applications to their fullest potential..

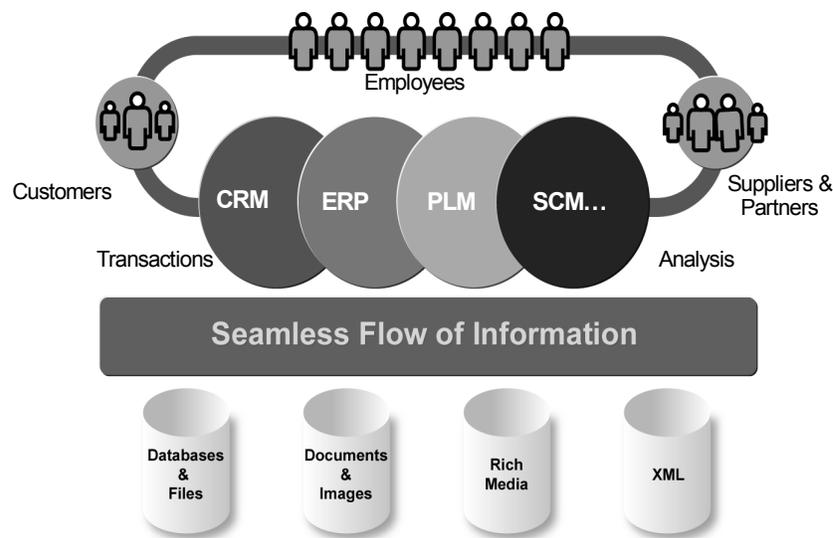


Fig. 1–2 *e-business demands for information*

## DB2 and Linux

Many companies are investing on Linux platforms because of the platform independence this environment provides. Today, the most popular hardware platform that is being used for Linux is the Intel platform in a server-based environment, however, many companies are also using Linux to consolidate workloads onto more powerful hardware, including IBM's largest enterprise zSeries servers which can run both DB2 UDB and DB2 Connect. Accompanying the major hardware vendors supporting Linux, are the major software vendors, who are also porting their applications to support Linux.

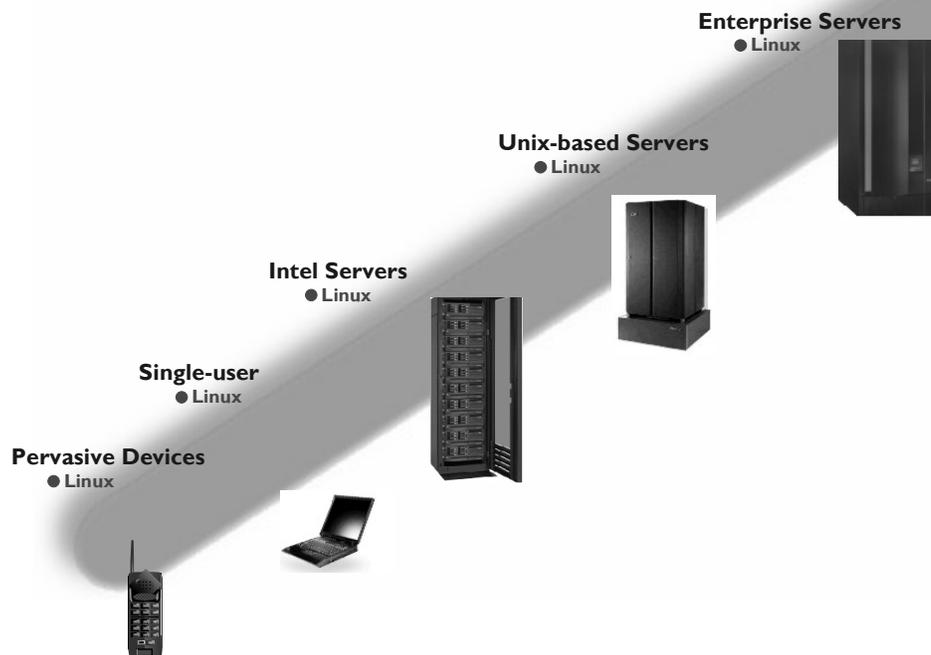


Fig. 1-3 DB2 UDB and Linux

## DB2 Universal Database

In the distributed environment, DB2 offers several packaging options:

- **Enterprise Server Edition** - This offering is often used to build e-business applications and support departmental applications to large enterprise data warehouses. It offers the most connectivity options and can share data with third-party databases and DB2 on heterogeneous platforms.
- **Workgroup Server Edition** - This offering is often used in smaller departmental applications or for applications that do not need access to remote databases on the OS/400, VM/VSE, or OS/390 platform.
- **Personal Edition** - This full-function database offering is for single-users and will not accept remote database requests. This offering is available on Linux and Windows.
- **Everyplace** - This is a mobile computing offering that gives mobile workers access to DB2 data sources in the enterprise through handheld devices such as personal digital assistants (PDAs) and handheld personal computers (HPCs). DB2 Everyplace can be run on a number of pervasive operating environments, including: Palm OS, Linux, Windows CE, Neutrino, PocketPC and Symbian.



**Note:** DB2 Enterprise Server Edition is a merge of the products previously known as DB2 UDB Enterprise Edition and DB2 UDB Enterprise-Extended Edition. The product previously known as DB2 Satellite Edition has been merged into the new release of DB2 Personal Edition.

## DB2 Enterprise Server Edition

DB2 Enterprise Server Edition is a relational database management system that is the foundation of many mission-critical systems and the primary focus of this certification guide. It is fully web-enabled, scalable from single processors to symmetric multiprocessors and to massively parallel systems; it supports unstructured data such as image, audio, video, text, spatial, and XML with its object relational capabilities.

Applications for DB2 Enterprise Server Edition can scale upward and execute on massively parallel systems or can scale downward with applications executing on single-user database systems.

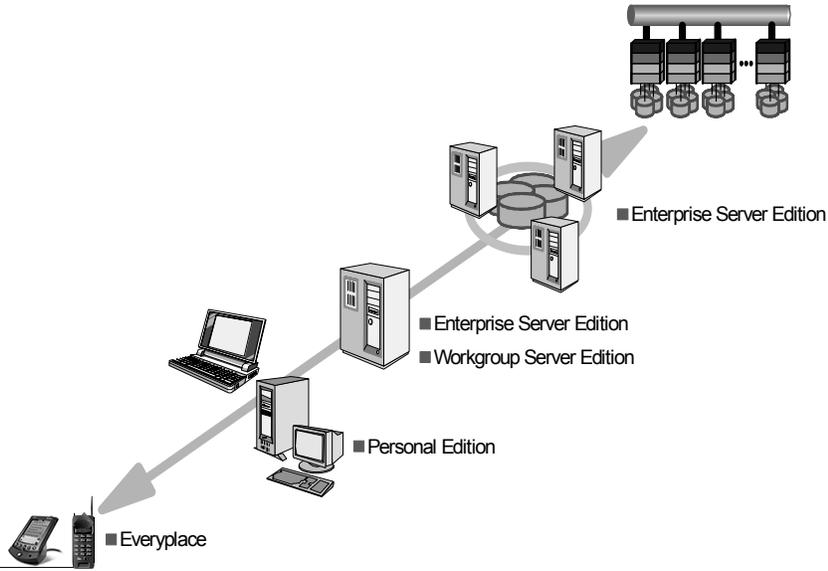


Fig. 1-4 DB2 Universal Database

DB2 Enterprise Server Edition provides an option that enables DB2 to partition data across clusters or massively parallel computers. To the end-user or application developer, the database appears to be on a single system. SQL statements are processed in parallel, thus increasing the execution speed for any given query.

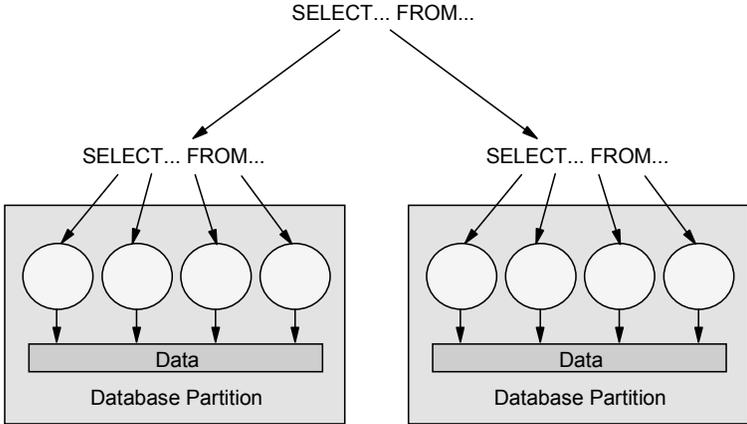


Fig. 1-5 DB2 processing a query in parallel

This option is often selected to support large database applications such as data warehousing applications. By providing for inter and intra-parallelism, databases can scale to multiple terabytes.

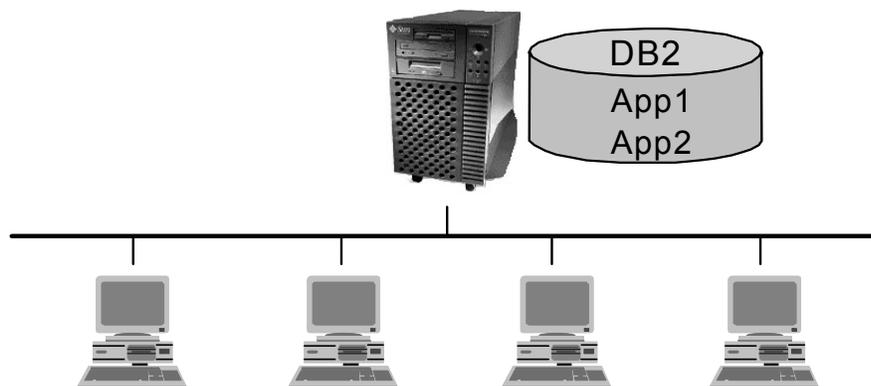
DB2 Enterprise Server Edition is available on the Linux, UNIX, and Windows platforms. This product is covered in detail in the *Advanced DBA Certification Guide and Reference for DB2 UDB V8 for Linux, Unix and Windows*.



**Note:** The Database Partitioning Feature (DPF) allows DB2 Enterprise Server Edition to partition workloads on a single physical server or across a cluster of servers.

## DB2 Workgroup Server Edition

DB2 Workgroup Server Edition is designed for use in a LAN environment. It provides support for both remote and local clients. A server with DB2 Workgroup Server Edition installed can be connected to a network and participate in a distributed environment as show in Figure 1–6.



**Fig. 1–6** *DB2 Workgroup Server Edition with remote clients*

In Fig. 1–6, App1 and App2 are local database applications. Remote clients can also execute App1 and App2 if the necessary setup has been performed. A DB2 application does not contain any specific information regarding the physical location of the database. DB2 client applications communicate with DB2 Workgroup Server Edition using a supported distributed protocol. Depending on the client and server operating systems involved, DB2 Workgroup Server Edition supports the TCP/IP, NetBIOS, Named Pipes, and APPC protocols.

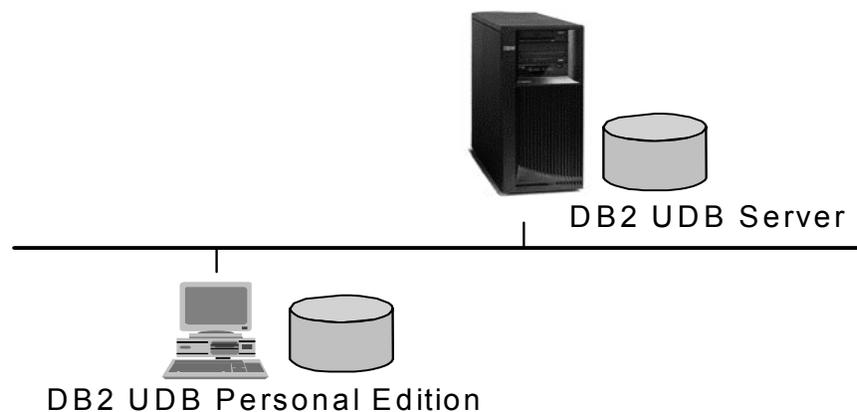
DB2 Workgroup Server Edition is available on the Linux, UNIX, and Windows platforms.

## DB2 Personal Edition

DB2 Personal Edition is a full-function database that enables a single-user to create databases on the workstation on which it was installed. It can be used as a remote client to a DB2 server, since it also contains the DB2 client components. Applications written to execute on the Personal Edition offering can also be used to access a DB2 Server, with no programming changes.

The DB2 Personal Edition product is often used by end-users requiring access to local and remote DB2 databases or developers prototyping applications that will be accessing other DB2 databases.

Fig. 1–7 shows an example of a DB2 Personal Edition installation. In this example, the user can access a local database on his or her desktop machine and access remote databases located on the database server. From the desktop, the user can make changes to the database throughout the day and replicate those changes as a client to the remote database on the DB2 server.



**Fig. 1–7** *DB2 Personal Edition*

DB2 Personal Edition includes graphical tools (via the DB2 Administration Client component) that enable a user to administer, tune for performance, access remote DB2 servers, process SQL queries, and manage other servers from a single workstation.

The Satellite Administration Center provides utilities that minimize the cost of managing a large number of systems by administering collections of systems run-

ning the same application and database in groups. Scripts stored at a central administration control point are used to accomplish administrative tasks on the systems in a group. This model also allows for centralized problem determination.

This product is available on Linux and Windows.

## DB2 Everyplace

DB2 Everyplace is a tiny “fingerprint” database of about 180K. It is designed for low-cost, low-power, small form-factor devices such as personal digital assistants (PDAs), handheld personal computers (HPCs), or embedded devices. DB2 Everyplace runs on devices that use the Palm Computing® Platform, Linux, Windows CE, PocketPC, Symbian, and QNX Neutrino. DB2 Everyplace provides a local data store on the mobile or embedded device for storing relational data from elsewhere in the enterprise. Relational data can be synchronized to the handheld device from other DB2 data sources such as DB2 Universal Database for Linux, UNIX, Windows, DB2 for OS/390, and DB2 for OS/400. DB2 Everyplace with DB2 Everyplace Synchronization Server will also synchronize data from other JDBC-compliant data sources such as Oracle and Microsoft.

The DB2 Everyplace Sync Server mobilizes professionals with e-business information anywhere, anytime. It extends the power of DB2 to a wide range of handheld devices, such as those running the Palm Pilot.

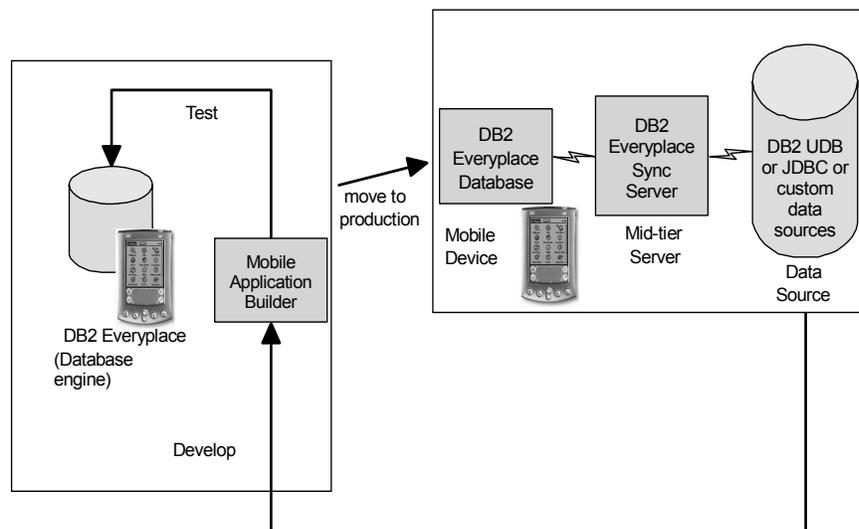


Fig. 1–8 DB2 Everyplace

The DB2 Everyplace Mobile Application Builder supports building applications for small handheld devices that access DB2 Everyplace databases. Some of its capabilities include:

- Supports visual construction of forms for different devices
- Supports the lightweight DB2 Everywhere database on the device
- Provides scripting capabilities for user-defined logic
- Integrates with other tools for application testing and debugging

## DB2 Connectivity

DB2 provides a variety of options for connecting to other DB2 and non-DB2 databases.

- **DB2 Clients** - Client code is required on workstations for remote users to access a DB2 database.
- **DB2 Connect** - This product provides support for applications executing on UNIX and Intel platforms to transparently access DB2 databases on the OS/400, VM/VSE, and OS/390 environments. Note that DB2 Connect is not required to access DB2 for any of the UNIX or Intel platforms.
- **DB2 Replication** - This product provides replication capabilities for DB2 databases and is integrated in DB2 on the UNIX and Intel platforms.
- **DB2 Relational Connect** - This product allows DB2 clients to access and join tables between DB2 and other heterogeneous databases, such as Oracle.
- **DB2 Life Sciences Data Connect** - This product allows DB2 to support a federated database environment and provides access to data commonly found in the Life Sciences industry, including: Excel spreadsheets, BLAST (Basic Local Aligned Search Tool), flat files, Documentum documents, and XML tagged files.
- **WebSphere Application Server** - This Application Server is shipped with DB2 Universal Developer's Edition and allows developers to use Java as platform in a transaction processing environment.

The *DB2 UDB Quick Beginnings* manual for each platform provides operating system requirements for implementing DB2 distributed configurations.

## DB2 Universal Database Clients

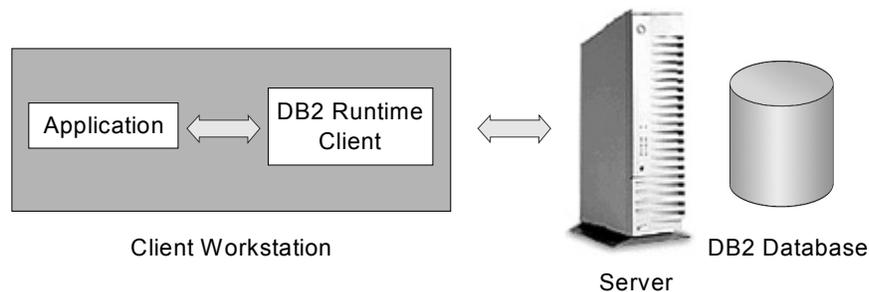
A DB2 client can be configured to establish communications with a DB2 server using various communication protocols. The supported protocols vary according to operating system:

- TCP/IP - used in all environments
- NetBIOS - used in Windows environments
- APPC - used in IBM operating environments
- Named Pipe - used in Windows environments

A DB2 client has a number of options for what client code needs to be installed, which will be depend on the requirements of the client. The options include:

- DB2 Runtime Client
- DB2 Administration Client
- DB2 Application Development Client
- DB2 Thin Client

Once a DB2 application has been developed, only the DB2 Runtime Client component needs to be installed on each workstation executing the application. Fig. 1–9 shows the relationship between the application, DB2 Runtime Client, and the DB2 database server. If the application and database are installed on the same system, the application is known as a *local client*. If the application is installed on a system other than the DB2 server, the application is known as a *remote client*.



**Fig. 1–9** DB2 Universal Database - Runtime Client

The Runtime Client provides functions other than the ability to communicate with a DB2 server or DB2 Connect server machine. For example, you can do any of the following:

- Issue an interactive SQL statement on a remote client to access data on a DB2 server or DB2 Connect server.
- Run applications that were developed to comply with the Open Database Connectivity (ODBC) standard or OLE DB.
- Run Java applications that access and manipulate data in DB2 databases using Java Database Connectivity (JDBC) or SQLJ.

If you need to graphically administer and monitor a DB2 database server, then you should install the DB2 Administration Client. It includes all the graphical DB2 administration tools in addition to all of the functionality of the DB2 Runtime Client.

If you need to develop applications, then you should install the DB2 Application Development Client. This is a collection of developer's tools that are designed to meet the needs of database application developers. The DB2 Application Develop-

ment Client includes all of the graphical DB2 administration tools and the DB2 Runtime Client functionality.

The DB2 client product that you should install depends on your requirements and the operating system on the client machine. For example, if you have a database application developed for AIX, and you do not require the DB2 administration or application development tools, you should install the DB2 Runtime Client for AIX.

Some installations prefer having the DB2 Runtime Client reside remotely on another server. Remote workstations then need to access the DB2 Runtime Client code remotely, before getting access to DB2. Alternatively, DB2 can support a Java client, and with the use of the Type 4 JDBC driver, where the client code is accessed from the server. In this case, the DB2 Runtime Client code would not be required.

## DB2 Connect

The DB2 Connect product allows clients to access data stored on database servers that implement the Distributed Relational Database Architecture (DRDA). The target database server for a DB2 Connect installation is known as a *DRDA Application Server*.



**Note:** The most commonly accessed DRDA server is DB2 for OS/390.

DB2 Connect supports both the TCP/IP and APPC communication protocols. The protocol supported depends on the DRDA application server being connected to and the version of the host software being run. For instance, a DB2 Connect server acting as a DRDA Application Requester to a host DRDA Application Server can connect to DB2 for OS/390 at version 5.1 and higher. Any of the supported network protocols can be used for a DB2 client to establish a connection to the DB2 Connect server.

Some of the major capabilities provided by DB2 Connect include:

- Support for ODBC, OLE DB, CLI, JDBC, and SQLJ applications
- Distributed Join across all DB2 databases
- Distributed Join with heterogeneous databases (via Relational Connect)
- Connection Pooling
- zSeries Sysplex Exploitation for failover and load balancing

Some of the common uses of DB2 Connect are:

- Web-enabling DB2 OS/390 by providing browsers direct access
- Leveraging Microsoft applications written using ADO, ODBC, or OLE DB can transparently access DB2 on mainframe platforms
- Offloading mainframe development cycles

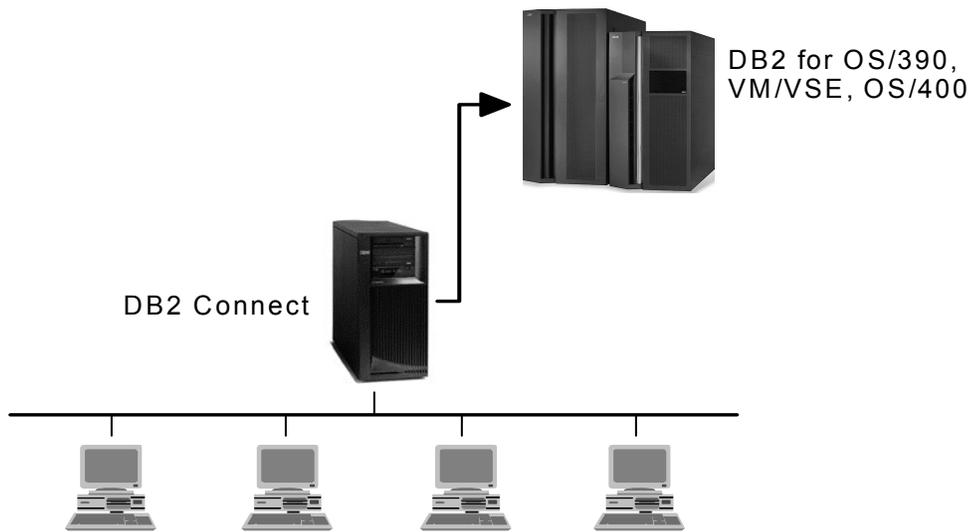
The database application must request the data from a DRDA Application Server through a DRDA Application Requester.



**Note:** The DB2 Connect product provides the *DRDA Application Requester* functionality.

The DRDA Application Server accessed using DB2 Connect could be any DB2 Server on OS/390, VM/VSE, or OS/400. If TCP/IP is the protocol of choice, then the following are prerequisites:

- DB2 for OS/390. Only version DB2 5.1 or higher supports TCP/IP in a DRDA environment.
- DB2 for OS/400. Only OS/400 version 4.2 or higher supports TCP/IP.
- DB2 Server for VSE/VM. Only version DB2 6.1 or higher for VM supports TCP/IP.



**Fig. 1–10** *DRDA application flow*

DB2 Connect is available as a server (Enterprise Edition) and a single-user package (Personal Edition). The DB2 Connect Enterprise Edition product provides the ability for multiple clients to access host data. A DB2 Connect server routes each database request from the DB2 clients to the appropriate DRDA Application Server (Figure 1–10). The remote client communicates with the DB2 Connect server using any of the supported communication protocols. DB2 Connect Personal Edition is available on Linux and the Windows platform. It provides access to host databases from the system where it is installed.

## DB2 Replication

DB2 Replication allows for data to be propagated from one location to another. It supports a wide variety of databases including DB2, Oracle, Microsoft, Sybase, Informix, IMS, Lotus Notes, and flat files. Replication is also a core technology that enables mobile users to keep their data synchronized with corporate data residing on a DB2 server.

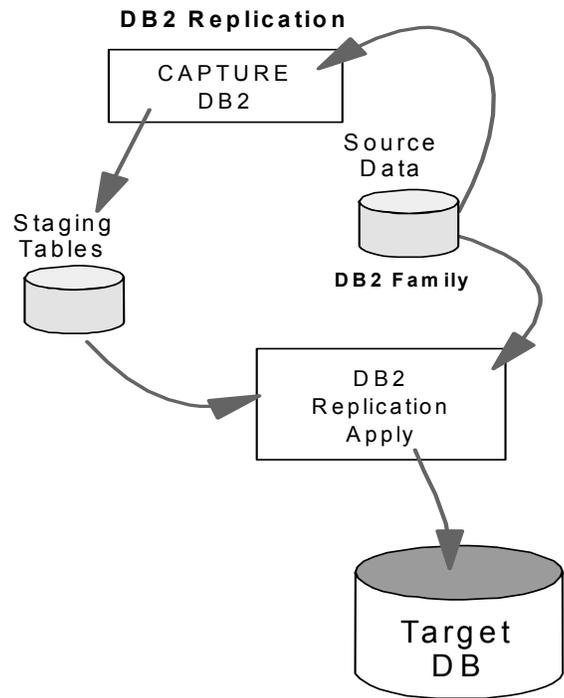


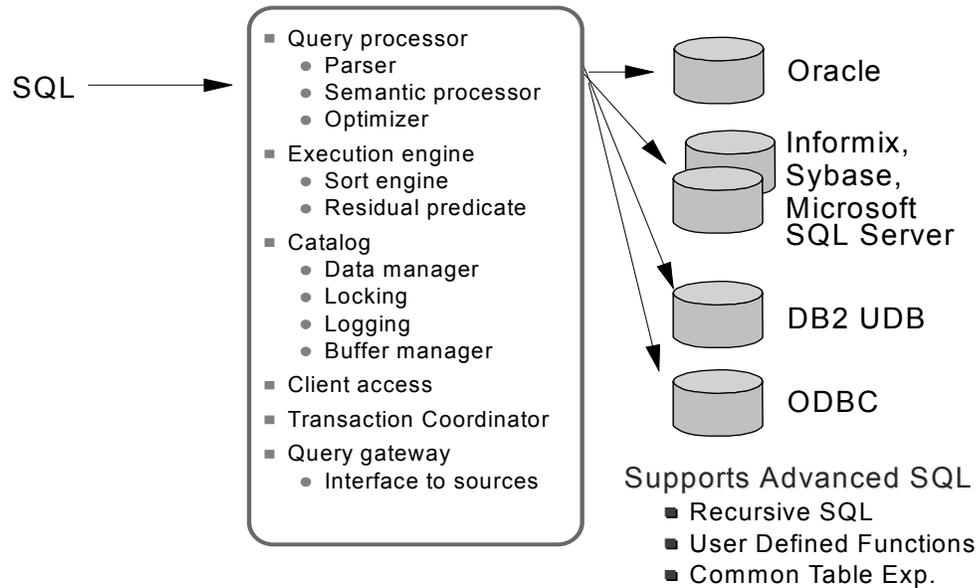
Fig. 1-11 DB2 replication environment

### DB2 Relational Connect

A federated system is distributed database environment that appears as a single virtual database. DB2 Relational Connect allows a single SQL statement to transparently access, join, or update data located across multiple data sources. DB2 Relational Connect enables the DB2 database to evaluate global statistics in order to optimize requests for remote data. The federated system also supports advanced object-relational SQL operations and can compensate for heterogeneous databases lacking support for specific functions.



**Note:** DB2 has built-in capability to access federated relational data from other DB2 UDB and Informix IDS database servers.



**Fig. 1–12** *Federated Database Environment*

In the life sciences industry, scientists need access to specialized data to support their research related to drug discovery. DB2 Life Sciences Data Connect is used to build a federated system that can access data useful to scientists, including: Excel spreadsheets, image data, flat files, and BLAST (Basic Local Alignment Search Tool).

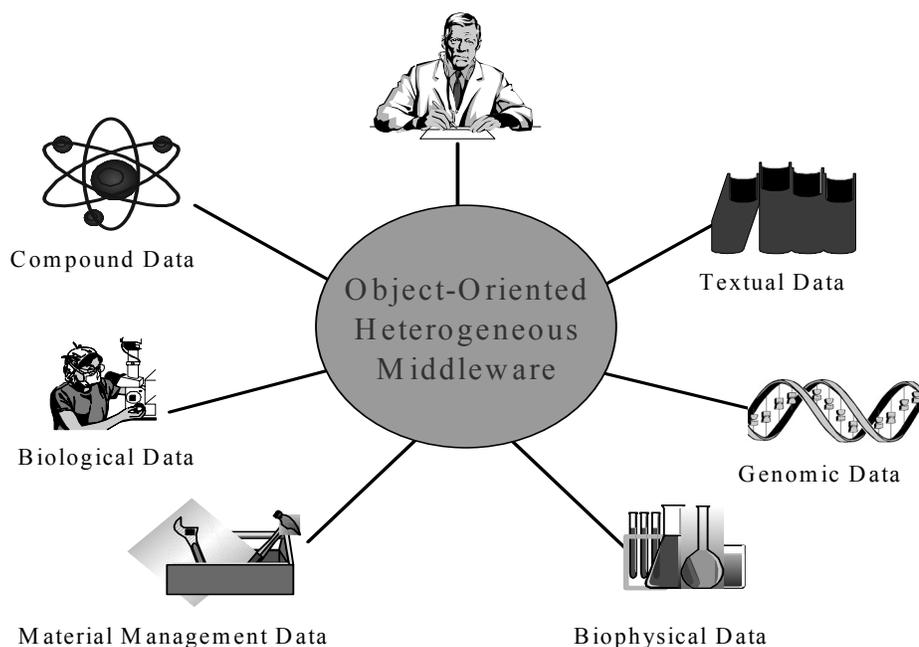


Fig. 1-13 DB2 Life Sciences Data Connect

## IBM WebSphere Application Server

IBM WebSphere Application Server is built on an open Java-based platform that enables applications to leverage existing application resources and access various databases, including DB2. Some of the capabilities of WebSphere Application Server include the use of Java servlets, Java Server Pages, and XML to quickly transform static Web sites into vital sources of dynamic Web content. Enterprise Java Beans (EJB) can also be used for implementing EJB components that incorporate business logic. WebSphere Application Server is packaged with DB2.

## DB2 Application Development

DB2 offers a rich application development environment that allows the developer to build databases supporting requirements from e-business and business intelligence applications. Many of these tools are integrated with the database; the major tools will be reviewed.

- **DB2 Universal Developer's Edition** - provides the tools for developers to build database applications, and includes: Personal Edition, Workgroup Server Edition, Enterprise Server Edition, DB2 Connect, Relational Extenders, all database clients, and WebSphere Application Server.
- **DB2 Development Center** - enables the creation of stored procedures and user-defined functions (UDFs). Also provides integration with WebSphere Studio Workbench and Microsoft Visual Studio.Net, Microsoft Visual Basic, and Microsoft's Visual Interdev.
- **DB2 Relational Extenders** - enables the SQL API to access unstructured data types including: text, image, audio, video, XML, and spatial data.
- **DB2 Data Warehouse Center** - provides the ability to build data marts/warehouses by automating the processes involved in managing, refreshing, moving, and transforming data, including the ability to define star schema model.
- **DB2 Data Warehouse Manager** - provides all the capabilities of the Data Warehouse Center, but includes support more data sources and includes the Information Catalog.
- **DB2 Data Links Manager** - DB2 offers supports for all popular programming languages and supports the latest Java-based application programming APIs, including: JDBC, SQLJ, ODBC, OLE DB, and CLI.

### DB2 Universal Developer's Edition

The DB2 development environment can be installed either on a DB2 server or on a DB2 client. The installation provides all of the necessary data access tools for developing database applications. There are two offerings:

- DB2 Personal Developer's Edition (PDE) - for Linux and Windows platforms.
- DB2 Universal Developer's Edition (UDE) - for all server platforms

The application development environment provided with both product packages allows application developers to write programs using the following methods:

- Embedded SQL
- Call Level Interface or CLI (compatible with the Microsoft ODBC standard)
- DB2 Application Programming Interfaces (APIs)
- DB2 data access through the World Wide Web
- Java applets or applications using JDBC or SQLJ.

The programming environment also includes the necessary programming libraries, header files, code samples, and precompilers for the supported programming languages. Several programming languages, including COBOL, FORTRAN, C, C++, and Java are supported by DB2.

An application developed using the Developer's Edition can be executed on any system with the same operating system that has the DB2 Run-time Client installed. To run the application on another operating system requires the application be rebuilt on the target operating system.

DB2 Personal Developer's Edition (PDE) includes:

- DB2 Personal Edition
- DB2 Connect Personal Edition
- DB2 Run-time, Administration, and Application Development Client
- DB2 Extenders
- WebSphere Studio, Entry Edition

DB2 Universal Developer's Edition (UDE) includes these components:

- DB2 Enterprise Server, Workgroup Server, and Personal Edition
- DB2 Connect Personal and Enterprise Edition
- DB2 Run-time, Administration, and Application Development Client (all platforms)
- DB2 Extenders
- WebSphere Studio, Professional Edition (Windows)
- Websphere Studio Websphere Application Server, Standard Edition

DB2 Universal Developer's Edition is supported on all platforms that support DB2 Enterprise Server Edition. It is intended for application development and testing only. The database server can be on a platform that is different from the platform on which the application is developed.

## DB2 Development Center

The Development Center is a graphical environment that supports the rapid development of DB2 SQL and Java stored procedures, SQL scalar and table user-defined functions (UDFs). This can be launched as a separate application from the DB2 program group or from the integrated development environments such as WebSphere Studio or Microsoft VisualStudio. The Development Center can be used to create stored procedures on local and remote DB2 servers and test, execute, and debug stored procedures.

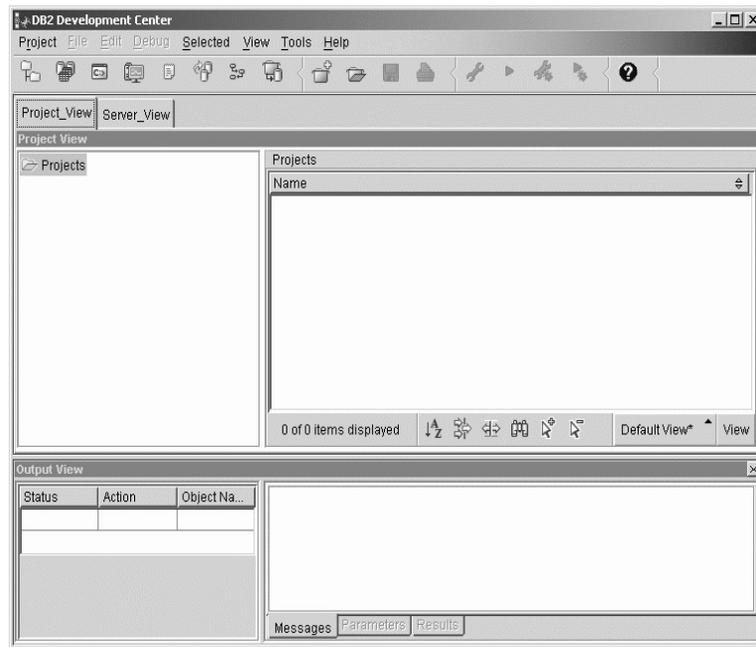


Fig. 1-14 DB2 Development Center



**Note:** The Development Center replaces the tool previously known as the Stored Procedure Builder and provides additional development enhancements such as the ability to build user-defined functions and structured data types.

## DB2 Relational Extenders

DB2 Relational Extenders offer the ability to manipulate data outside of conventional rows and columns to include the manipulation of these types of data: text, image, audio, video, and XML. The DB2 Relational Extenders encapsulates the

attributes, structure, and behavior of these unstructured data types and stores this information in DB2. From the developer's perspective, the DB2 Relational Extenders appear as seamless extensions to the database and enable the development of multimedia-based applications. The following DB2 Relational Extenders are provided by IBM:

- Net Search Extender
- Image Extender
- Audio Extender
- Video Extender
- Spatial Extender
- XML Extender



**Note:** The new Net Search Extender is a merge of the products previously known as Net.Search Extender and Text Information Extender.

SOLD	ONHAND	RATING	ARTIST	TITLE	COVER	VIDEO	MUSIC	SCRIPT
1212	64	G	Shirley	Inner Light				
39	60	R	Bill	Relics				
73	95	PG-13	Dana	All Good Things				
286	97	PG-13	Austin	Generations				

Mix traditional relational data with unstructured data

```

SELECT Artist, Music, Cover, Video,
FROM Video_Store
WHERE Rating = 'G' and
Video-length(video) ≤ 120 and
Text-Content-Search (SCRIPT, "Infinity")
ORDER BY Rating

```

Datatype Attribute  
 Specialized Search

Fig. 1-15 DB2 Relational Extenders

The purpose of the DB2 Relational Extenders is to provide for the management of unstructured data through the SQL API. By preserving the current investment in relational applications, new nontraditional applications can be introduced by leveraging existing skills. This open environment enables developers and independent software vendors to develop and introduce their own extenders as extensions to DB2.

The XML Extender is provided with DB2 and allows you to store eXtensible Markup Language (XML) documents as a new column datatype. You also have the ability to decompose and store XML in its component parts as columns in multiple tables. In either case, indexes can be defined over the element or attribute of an XML document for fast retrieval. Furthermore, text search and section search can be enabled on the XML column or its decomposed part via DB2 Net Search Extender. You can also formulate an XML document from existing DB2 tables for data interchange in business-to-business environments.

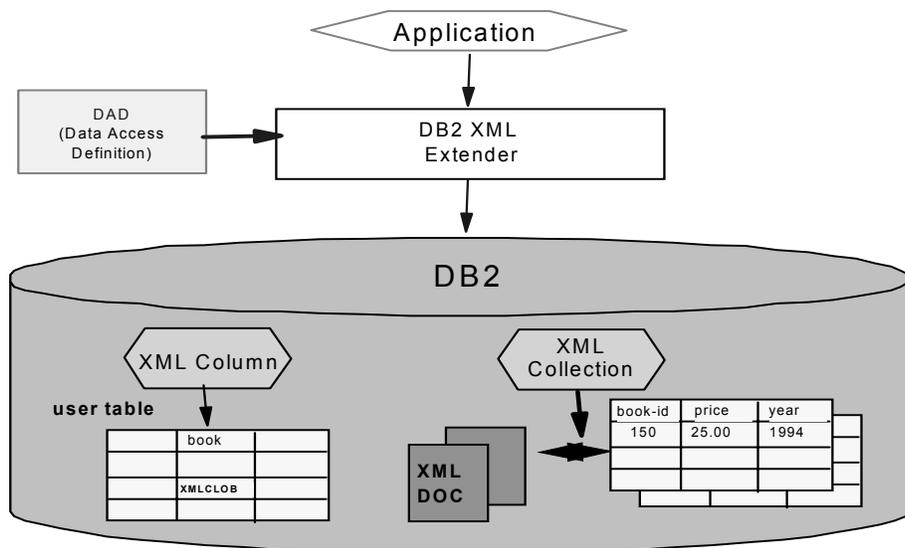


Fig. 1-16 XML Extender



**Note:** The new XML Extender now supports Web services with the Web Services Object Framework (WORF), a set of tools for implementing Web services with DB2. WebSphere MQ Series applications can also use the XML Extender.

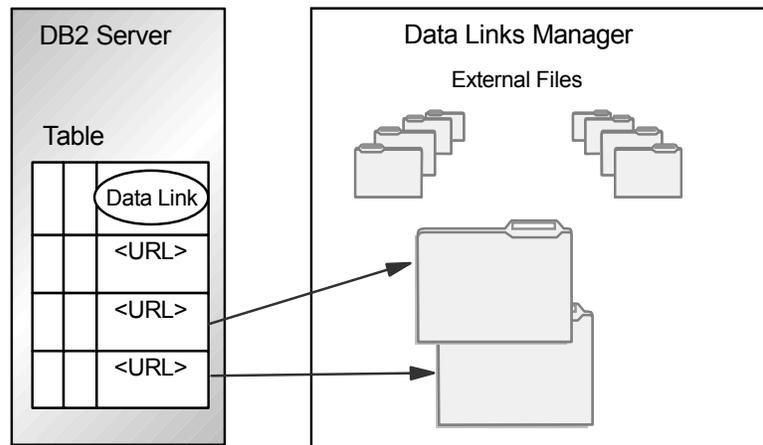
The DB2 Net Search Extender combines in-memory database technology with text search semantics for high-speed text search in DB2 databases. Searching with it can be particularly advantageous in Internet applications where performance is an important factor. Net Search Extender can add the power of fast full-text retrieval to JDBC and CLI applications. Its features let you store unstructured text documents of up to 2 gigabytes in databases. It offers application developers a fast, versatile, and intelligent method of searching through such documents.



**Note:** The new integrated DB2 Net Search Extender can be managed from the Control Center and includes built-in federated support to index and search text data stored in other DB2 or Informix Dynamic Server (IDS) databases.

## DB2 Data Links Manager

DB2 Data Links Manager can be used to build applications that need to combine the search capabilities of SQL with the advantages of working directly with files to manipulate raw data. A reference is stored for each external file, along with meta-data that describes the contents of each file.



**Fig. 1–17** DB2 Data Links Manager

DB2 Data Links Manager uses the DATALINK data type, which points to an external file, and the DB2 Data Links Manager components.

You use the DATALINK data type, just like any other data type, to define columns in tables. The DATALINK values encode the name of the Data Links Server containing the file and the file name in terms of a Uniform Resource Locator (URL).

Using DB2 Data Links Manager means that external files can be backed up with the database and SQL Data Control Language statements (GRANT and REVOKE) can be used to control permissions to those files.

Examples of applications that can use the DATALINK data type are:

- Medical applications, in which X-rays are stored on the file server and the attributes are stored in the database.
- Entertainment industry applications that perform asset management of video clips. The video clips are stored on a file server, but attributes about the clips are stored in a database. Access control is required for accessing the video clips

based on database privileges of accessing the meta information.

- World Wide Web applications to manage millions of files and allow access control based on database privileges.
- Financial applications, which require distributed capture of check images, and a central location for those images.
- CAD and CAM applications, where the engineering drawings are kept as files, and the attributes are stored in the database. Queries are run against the file attributes.

Even though the DATALINK column represents an object that is stored outside the database system, you can use SQL queries to search metadata to obtain the file name that corresponds to the query result. You can create indexes on videos, images, text (and so on), and store those attributes in tables along with the DATALINK column. With a central repository of files on the file server and DATALINK data types in a database, you can obtain answers to questions like “what do I have?” and “find what I am looking for.”

## DB2 Data Warehouse Center

This product brings together the tools to build, manage, govern, and access DB2 data warehouses. The DB2 Warehouse Center simplifies and speeds warehouse prototyping, development, and deployment. It helps satisfy user requirements for finding, accessing, and understanding information. It provides flexible tools and techniques for building, managing, and accessing the warehouse. And it meets the most common reporting needs for enterprises of any size.



**Note:** The Data Warehouse Center is integrated with DB2 UDB V8.1 and provides the basic extraction, transformation, and load capabilities to build data warehouses. This tool also provides a star schema builder and a process modeler for automating the steps of transforming data for end-users.

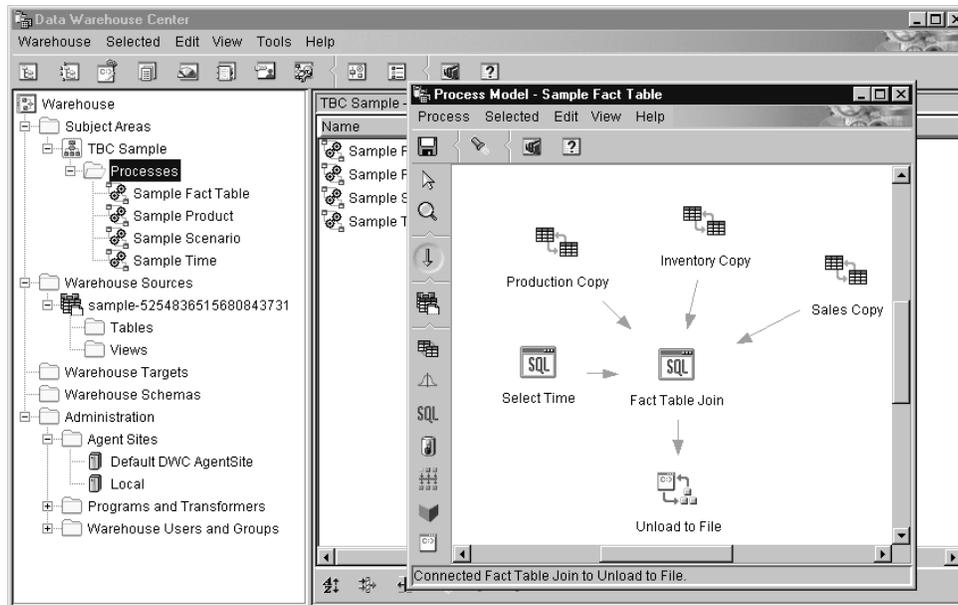


Fig. 1–18 Data Warehouse Center

## DB2 Warehouse Manager

The DB2 Warehouse Manager adds to basic warehouse and analytical functions available in DB2 Universal Database by providing:

- Additional warehouse scalability through warehouse agents co-located with the database. Warehouse agents manage the flow of data between warehouse sources and warehouse targets.
- Advanced transformations using Java stored procedures and user-defined functions including cleaning data, pivoting tables, and generating keys
- An integrated business information catalog to guide users to relevant information that they can use for decision making.
- Query report that satisfies the common reporting needs of most enterprises.
- An integrated business information catalog to guide users to relevant information that they can use for decision making.



**Note:** The Data Warehouse Manager extends the functions of the Data Warehouse Center by providing support for additional data sources, additional warehouse and statistical transformations, and an Integrated Information Catalog.

### Information Catalog Center

The Information Catalog Center is a tool for end-users to help easily find, understand and access available information in the corporation. This graphical tool uses a Business View Model, built by the Data Warehouse Administrator, that can be used by the business user to navigate through the data in an enterprise. It allows users to:

- Populate the catalog through metadata interchange with the Data Warehouse Center and other analytical and reporting tools
- Directly register shared information objects
- Navigate or search across the objects to find relevant information
- Display the metadata about the object
- Launch the tools used to render the information for the end-user



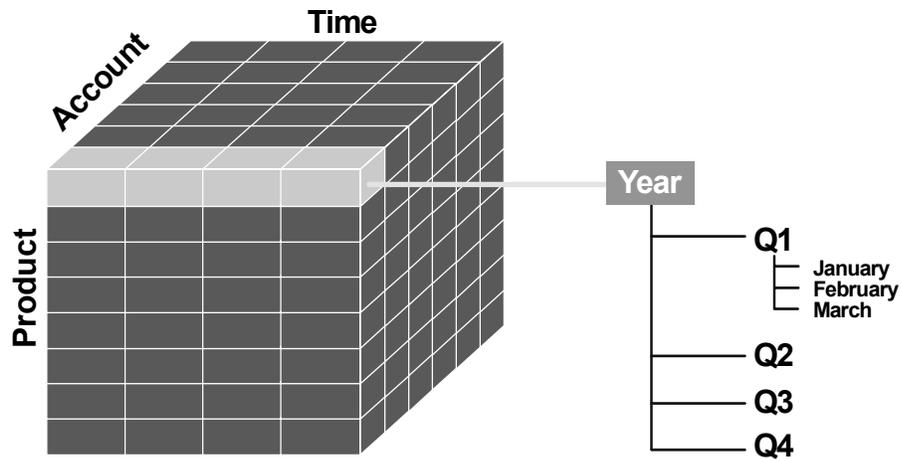
**Note:** The Information Catalog Center keeps metadata on data that end-users need to access. This tool automates the exchange of metadata with the Data Warehouse Center.

### DB2 OLAP Server

The DB2 OLAP Server is a scalable, industrial-strength Online Analytical Processing (OLAP) software that enables you to build sophisticated decision support, planning, and analysis applications for your enterprise. DB2 OLAP Server provides a fast path to turn your warehouse data into business insight. It delivers “speed of thought” query performance to a large set of online users. It is built for e-business with tools to help you quickly deploy Web-based analytical applications.



**Note:** DB2 OLAP Server provides DB2 integrated OLAP capabilities. Also, included to build OLAP solutions are: Integration Server to pull data from DB2, an OLAP spreadsheet plug-in for Excel and Lotus 1-2-3, and an Administration Manager.



Product Overview

Fig. 1-19 *Multidimensional cube*

Flexibility in storage is extended to the application level. Depending on individual application needs, the cube can be stored in either DB2 for added flexibility of SQL access or multidimensional storage for optimal performance.



details about each table is shown in the Contents Pane. A number of the existing tables in the SAMPLE database are displayed.

The main components of the Control Center are:

- Menu Bar - Used to access Control Center functions and online help.
- Tool Bar - Used to access the other administration tools.
- Objects Pane - This is shown on the left-hand side of the Control Center window. It contains all the objects that can be managed from the Control Center as well as their relationship to each other.
- Contents Pane - This is found on the right side of the Control Center window and contains the objects that belong or correspond to the object selected on the Objects Pane.
- Contents Pane Toolbar - These icons are used to tailor the view of the objects and information in the Contents pane. These functions can also be selected in the View menu.

Hover Help is also available in the Control Center, providing a short description for each icon on the tool bar as you move the mouse pointer over the icon.

## Replication Center

The Replication Center is a graphical tool that allows the administrator to set up and administer a data replication environment. Replication subscription definitions are stored in a central control server as scripts. The main functions in setting up a replication subscription can be performed with this tool, including:

- Registering replication sources
- Monitoring the replication process
- Operating the Capture and Apply programs
- Defining alerts

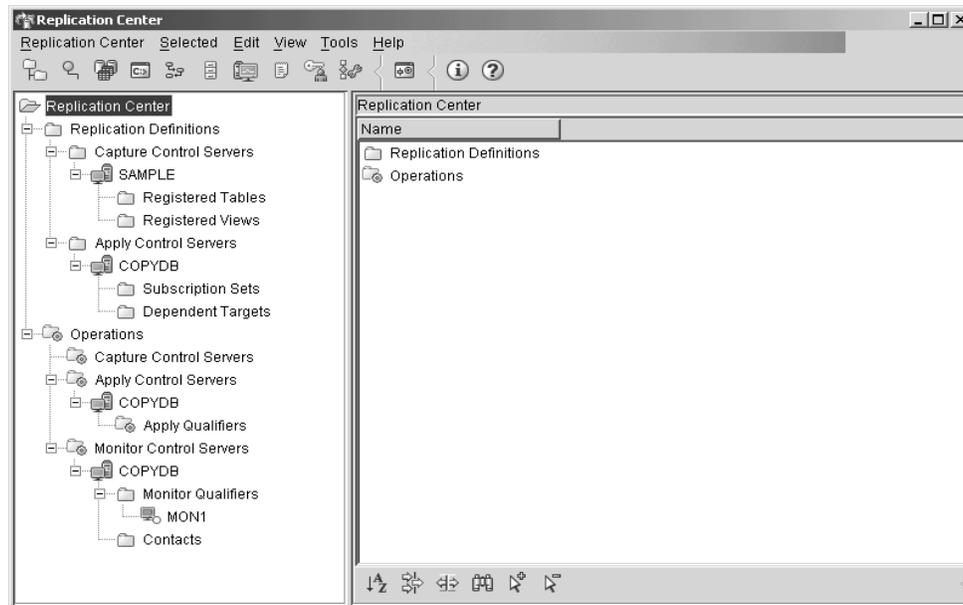


Fig. 1–21 DB2 Replication Center

## Other Tools Available from the Control Center

By using the Control Center Tool Bar, you can access other administration tools to help you manage and administer databases in your environment:

- Satellite Administration Center - Used to manage multiple DB2 Personal Edition or Workgroup Server installations.
- Data Warehouse Center - Used to build the data mart/warehouse.
- Command Center - This provides an interactive window that allows input of SQL statements or DB2 commands, the viewing of execution results, and explain information. The graphical command utility is the preferred method for text commands as it provides enormous flexibility and function; interacts with the Script Center.
- Task Center - Used to create, schedule, and manage scripts that can contain SQL statements, DB2 commands, or operating systems commands.
- Journal - This keeps a record of all script invocations, all DB2 messages, and the DB2 recovery history file for a database. It is used to show the results of a job and, contents of a script and also to enable or disable a job.
- License Center - Used to manage licenses and check how many connections are used.

- Development Center - This tool enables the creation and testing of DB2 stored procedures or user-defined functions.
- Tools Settings - Allows you to configure the DB2 graphical tools and some of their options.

The graphical tool set provided with DB2 is full functioned and very powerful, allowing you to administer and access your DB2 system from graphical interfaces.

The DB2 Web Tools allows the administrator to use an HTTP client to remotely execute:

- SQL statements against a DB2 database
- DB2 commands against a DB2 instance
- operating system commands against a DB2 server
- Stored Procedures through JDBC calls and DB2 administration server API

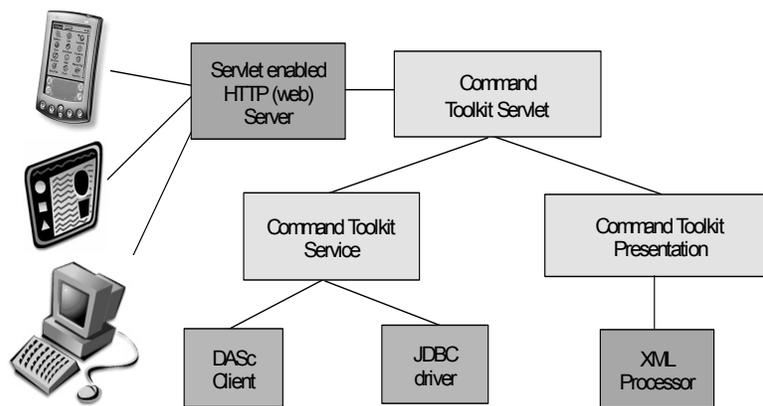
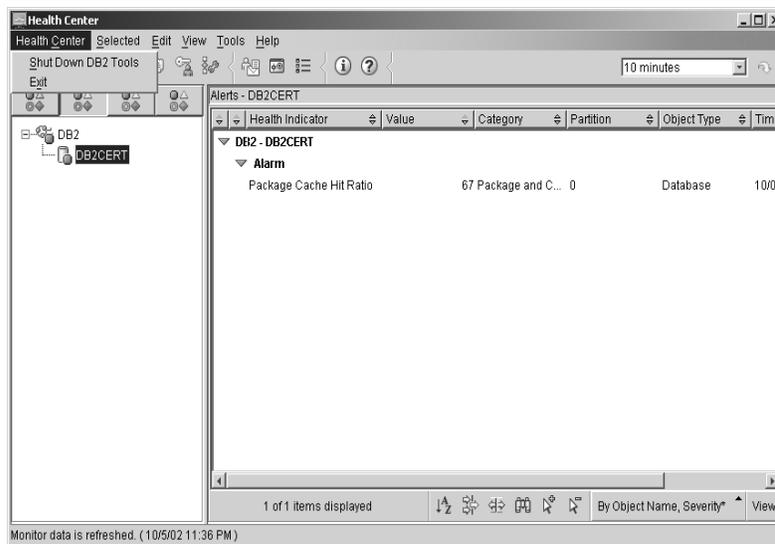


Fig. 1-22 DB2 Web Tools

## Health Center

IBM has an autonomic computing initiative that is focusing on creating an IT infrastructure that can self-manage, self-diagnose and self-heal itself. DB2 provides a variety of tools to enable the database to manage itself. Self-managing and resource tuning (SMART) database technology allows greater automation with respect to configuring, tuning, and managing database operations. The Health Center is a server-side tool that can monitor the health of a DB2 instance, even without user interaction. When defined thresholds are exceeded (i.e. log space not sufficient) alerts are activated. Alert notifications can be sent using e-mail or to a pager system, alternatively a preconfigured set of actions can be taken that has been defined in a script or task.



**Fig. 1–23** Health Center

The graphical user interface allows the administrator to select database objects and drill-down on its details, current alerts, and the recommended actions described to resolve the alert.

## Additional Administration tools

IBM has developed a number of system management tools to make administrators more productive in managing DB2. These tools include:

- DB2 Recovery Expert - users can obtain intelligent assistance when determining the most efficient technique for a recovery solution. Built-in self-managing and resource tuning (SMART) features provide intelligent analysis of a recovery sit-

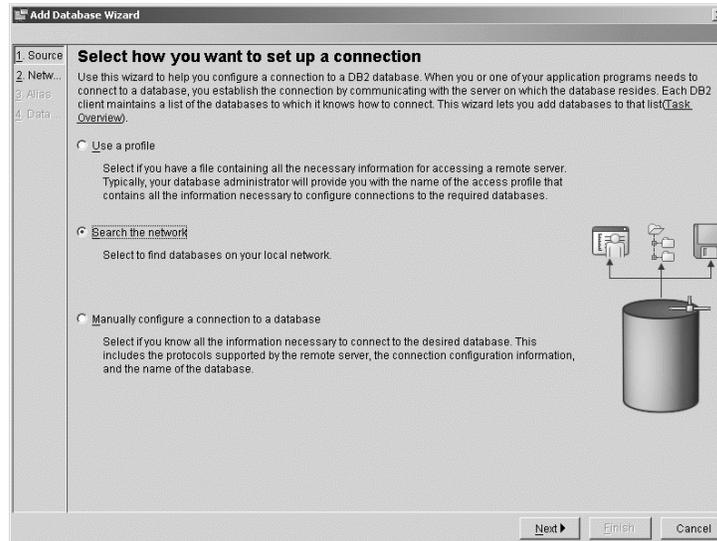
uation and automate the rebuilding of the database environment, as well as the ability to generate undo SQL.

- DB2 Performance Expert - recommends changes on DB2 performance-related information. Also provides recommendations for system tuning to gain optimum throughput.
- DB2 High Performance Unload - rapid unload table spaces of all types with various options.
- DB2 Table Editor - a tool that provides direct DB2 access to users creating, reviewing, or updating data.
- DB2 Web Query Tool - enables easy access to DB2 data.

All of the above tools run on multiple platforms.

## Advisors and Wizards

Wizards are tutors that help you create objects and perform specific database operations. Each Wizard has detailed information available to help the user. The DB2 Wizards are integrated into the administration tools and assist in a variety of administrative tasks. For example, the Add Database Wizard is used to set up communications between a DB2 client and a DB2 server, and is invoked from the Configuration Assistant.



**Fig. 1–24** Configuration Assistant - Add Database Wizard

Fig. 1–24 shows that there are a number of ways to add a remote database. You do not have to know the syntax of commands or even the location of the remote database server. One option searches the network, looking for valid DB2 servers.



**Note:** The Client Configuration Assistant has been renamed to the Configuration Assistant and significant enhancements have also been added.

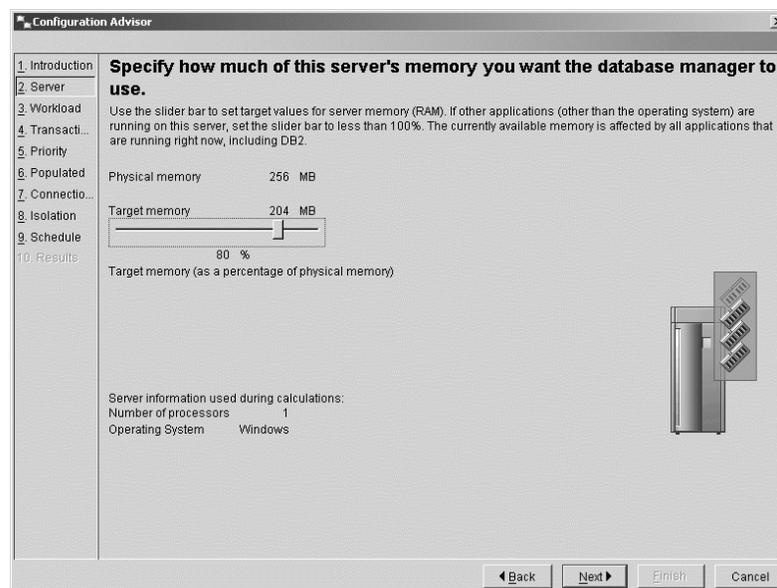
Invoked from the DB2 Control Center, the wizard extracts information from the system and asks questions about the database workload. It then runs a series of calculations designed to determine an appropriate set of values for the database and database manager configuration variables. You can choose whether to apply the changes immediately or to save them in a file that can be executed at a later time.

DB2 Wizards include:

- Create, Add, Backup, Restore Database
- Load
- Create Table
- Create Table Space
- Configure Database Logging
- Add Partitions

The Configuration Advisor (Fig. 1–25) can be used to get a database system up and running quickly. This can also be invoked from the Control Center, and gathers data about the hardware and application environment at your installation. The suggested configuration parameters can either be applied immediately or be used to create a script to be executed later.

This script also provides a list of the most commonly tuned parameters for later reference.



**Fig. 1–25** *Configuration Advisor*

The Design Advisor replaces the tool formerly known as the Create Index Wizard.

## The Command Line Processor (CLP)

The Command Line Processor (CLP) is a component common to all DB2 products. It is a text-based application that is used to execute SQL statements and DB2 commands. For example, you can create a database, catalog a database, and issue dynamic SQL statements from the CLP.



**Note:** The commands and statements issued through the DB2 CLP can also be issued through the DB2 Command Center, which is a preferred graphical interface.

```
db2 => list db directory

System Database Directory

Number of entries in the directory = 1

Database 1 entry:

Database alias           = DB2CERT
Database name           = DB2CERT
Local database directory = G:\DB2
Database release level  = a.00
Comment                 =
Directory entry type    = Indirect
Catalog node number     = 0
```

**Fig. 1–26** *Command Line Processor*

Fig. 1–26 shows a command and its output as executed from the Command Line Processor. The Command Line Processor can be used to issue interactive SQL statements or DB2 commands. The statements and commands can be placed in a file and executed in a batch environment or they can be entered in interactive mode.

The DB2 Command Line Processor (CLP) is provided with all DB2 Universal Database, DB2 Connect, and DB2 Developer's products and Clients.

All SQL statements issued from the Command Line Processor are dynamically prepared and executed on the database server. The output, or result, of the SQL query is displayed on the screen by default.

All of the DB2 commands are documented in the *DB2 UDB V8.1 Command Reference*.

## Visual Explain

Other graphical tools can be used for tuning or monitoring performance. Visual Explain is a graphical utility that provides a visual representation of the access plan that DB2 uses to execute an SQL statement.

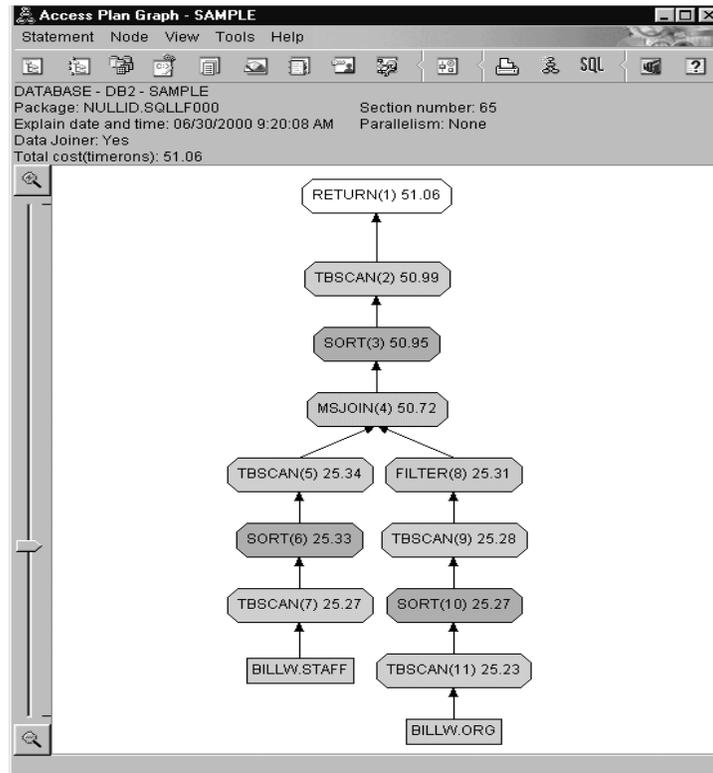


Fig. 1–27 Access Plan Graph in Visual Explain

Visual Explain can be invoked from the Control Center or from the Command Center. Fig. 1–27 shows the type of information that is displayed. This example shows that two tables are being accessed and an approximation of the cost of the query is also provided in the Visual Explain output. The estimated query costs represent the complexity and resource usage expected for a given SQL query. More details on the usage of Visual Explain are provided in “Monitoring and Tuning” on page 659.

## DB2 Query Patroller

DB2 Query Patroller controls, monitors query execution and works with queries to prioritize and schedule user queries based on user profiles and cost analysis performed on each query. Large queries can be put on hold and scheduled for a later time during off-peak hours. Queries with high priority (based on user profiles) are promoted to the top of the schedule. In addition, DB2 Query Patroller monitors resource utilization statistics to determine which CPUs are the least used. It can use this information to determine load distribution of the system, which can allow it to balance the number of users allowed to submit queries at any given time.

DB2 Query Patroller greatly improves the scalability of a data warehouse by allowing hundreds of users to safely submit queries on multi-terabyte class systems. Its components span the distributed environment to better manage and control all aspects of query submission.

DB2 Query Patroller acts as an agent on behalf of the end-user. It prioritizes and schedules queries so that query completion is more predictable and computer resources are efficiently utilized. DB2 Query Patroller obtains the query cost from the DB2 Optimizer and then schedules and dispatches those queries so that the load is balanced across the installation-specified nodes.

DB2 Query Patroller sets individual user and user class priorities as well as user query limits. This enables the data warehouse to deliver the needed results to its most important users as quickly as possible. If desired, an end-user can choose to receive notice of scheduled query completion through electronic mail.

DB2 Query Patroller consists of components running on the database server and end-users' desktops. DB2 Query Patroller is made up of several components, each having a specific task in providing query and resource management.

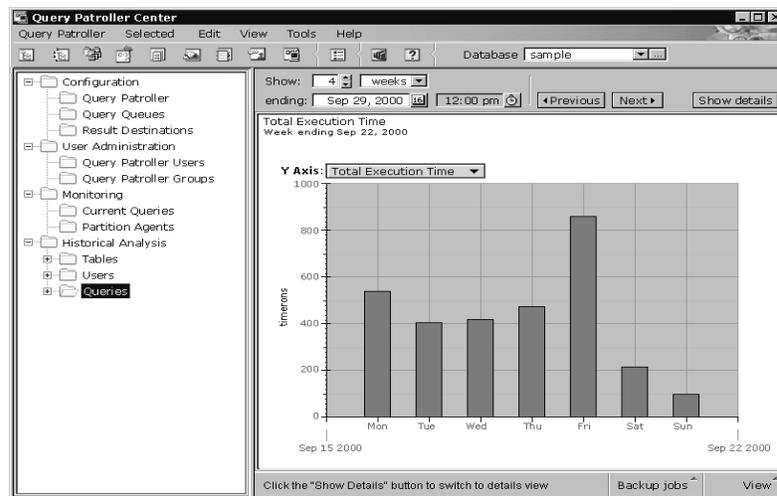


Fig. 1–28 DB2 Query Patroller Center

The Server is the core component of DB2 Query Patroller. It provides an environment for storing user profiles, setting up system parameters, maintaining job lists, and storing node information. The DB2 Query Patroller system administrator has an interface to the Server to perform these tasks. The Server component executes

on a node within the Database Management System (DBMS) called the Management Node.

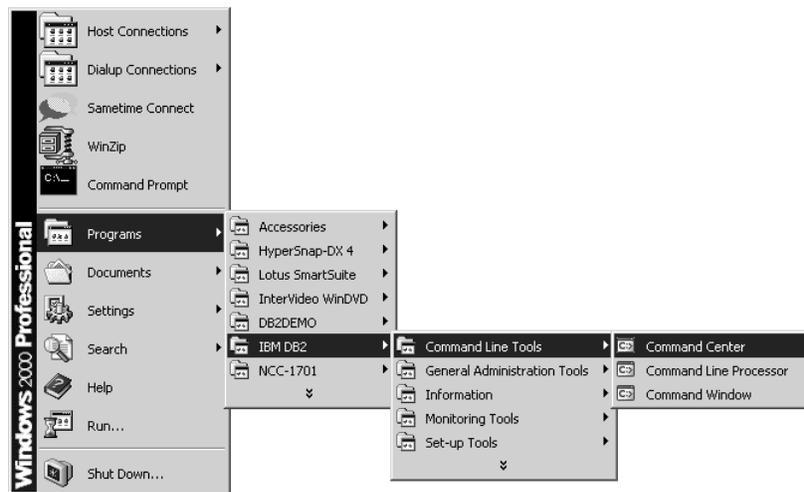
The Administrator component gives a DBA or system administrator the tools needed to manage the DB2 Query Patroller environment. This tool allows for the management and viewing of queries. The administrator provides menus to display job lists and history, user profiles, node information, and system parameters. It also provides for the display of utilization graphs.

### Database Monitoring

The *Snapshot Monitor* captures database information at specific intervals. The interval time and data represented in the performance graph can be configured. The Snapshot Monitor can help analyze performance problems, tune SQL statements, and identify exception conditions based on limits or thresholds.

The *Event Monitor* captures database activity events as defined by the event monitor definition. Event Monitor records are usually stored on disk and then analyzed after the data has been captured. The *Event Analyzer* graphical tool provided with DB2 can be used to analyze the captured data. More details on the usage of the DB2 monitors are provided in “Monitoring and Tuning” on page 659.

The DB2 folder is created on the desktop for environments such as Windows. This DB2 folder is typically used to invoke the graphical tools provided with DB2. Fig. 1–29 shows some of the components of the DB2 product as they appear in the DB2 folder. A component usually relates to an executable application or utility.



**Fig. 1–29** DB2 desktop folder (Windows)

The graphical tools integrated with DB2 support these functions:

- Create, alter, and drop databases, table spaces, tables, views, indexes, triggers, schemas, users, user groups, and aliases
- Load, import, export, and reorganize data and gather statistics
- Schedule jobs to run unattended
- Back up and restore databases
- Monitor and tune performance
- Tune queries using access path analysis
- Manage data replication

The following are also available in the DB2 desktop folder:

- Information Center - Provides the user with quick access to the DB2 product documentation. Information is available about common tasks, problem determination, DB2 online manuals and the sample programs provided with DB2.

- Configuration Assistant (CA) - This is a graphical tool that is used to configure access to remote databases. It can be invoked from the DB2 Desktop folder or from the command line with the DB2CCA command.
- Command Line Processor (CLP) - The CLP is a text-based program that allows you to enter DB2 commands or SQL statements.
- Command Window (Windows platform).
- Command Center - a graphical tool that allows you to enter one or more DB2 commands and statements, store the commands as scripts, view the explain information for a statement, and schedule scripts via the Script Center.
- Control Center - The central point of administration for DB2 Universal Database. This tool allows you to access other tools such as the Command Center and the Journal.
- Event Analyzer - This allows you to access information on database activities collected by event monitors.
- Event Monitor - This allows you to work with DB2 event monitors including creating, starting, stopping, and removing event monitors.
- First Steps - This allows you create the SAMPLE database, view sample contents, and access the on-line DB2 library.
- Start/Stop HTML Search Server - This enables (or disables) the searching of a topic in the online books.
- Development Center - assists with the creation of stored procedures and user-defined functions that run on DB2 servers
- SQL Assist - a graphical online tool that can be used to build Select, Insert, Update, and Delete statements

## Summary

This chapter discussed the DB2 Version 8.1 products for Linux, Windows, and UNIX. There are a number of offerings available:

- DB2 Enterprise Server Edition
- DB2 Workgroup Server Edition
- DB2 Personal Edition
- DB2 Everyplace Edition
- DB2 Connect Enterprise Edition
- DB2 Data Warehouse Manager
- DB2 Query Patroller

These products provide the flexibility to execute database applications running on pervasive devices up to multi-node clusters. DB2 provides support for the commonly used communication protocols and can access data from heterogenous database systems. DB2 replication is integrated in the database and can be used to transform data from other DB2 and non-DB2 database systems. DB2 Connect is used to access DB2 data from DRDA servers, typically DB2 running on the zSeries platform.

DB2 V8 has introduced a number of application development enhancements, including the Development Center, which can be integrated with popular tools such as WebSphere Studio and Microsoft VisualStudio. The DB2 Warehouse Manager and integrated Information Catalog provides a graphical interface to building data-marts and data warehouses. DB2 Query Patroller provides query and resource management for decision support systems.

Significant new features have been added to make DB2 easier to manage. SMART technology (Self-managing and Resource Tuning) has been integrated into a number of DB2 components, including: installation, configuration, utilities, problem determination and resolution, and availability. This is part of IBM's autonomic computing initiative and new features will continue to be added to make the DBA more productive.

This chapter also introduced some of the graphical and command line tools available in DB2. The Command Line Processor (CLP) is a text-based application that allows you to enter DB2 commands and SQL statements and is found in all DB2 products. From the desktop, an administrator can configure remote and local systems, administer instances and databases, and create database objects. In the remaining chapters, additional DB2 functions and tools will be examined in how they assist the end-user, application developer, and administrator.