

DB2 9 Fundamentals exam 730 prep, Part 1: DB2 planning

Skill Level: Introductory

Paul Zikopoulos (paulz_ibm@msn.com)
Database Specialist
IBM

27 Jul 2006

This tutorial introduces you to the basics of the DB2 9 products and tools, along with concepts that describe different types of data applications, data warehousing, and OLAP. This is the first in a [series of seven tutorials](#) to help you prepare for the DB2 9 for Linux, UNIX, and Windows Fundamentals exam 730.

Section 1. Before you start

About this series

Thinking about seeking certification on DB2 fundamentals (Exam 730)? If so, you've landed in the right spot. This [series of seven DB2 certification preparation tutorials](#) covers all the basics -- the topics you'll need to understand before you read the first exam question. Even if you're not planning to seek certification right away, this set of tutorials is a great place to start learning what's new in DB2 9.

About this tutorial

This tutorial introduces you to the basics of the DB2 9 products and tools, along with concepts that describe different types of data applications, data warehousing, and OLAP. It discusses how to use the Control Center, which is the central management tool for DB2 data servers. This tutorial also shows you how to use the Configuration Assistant, which lets you easily work with existing databases, add new ones, bind applications, set client configuration and registry parameters, and import and export configuration profiles.

This is the first in a series of seven tutorials designed to help you prepare for the

DB2 9 Family Fundamentals Certification (Exam 730). The material in this tutorial primarily covers the objectives in Section 1 of the test, "Planning." You can view these objectives at: <http://www-03.ibm.com/certify/tests/obj730.shtml>.

Objectives

After completing this tutorial, you should understand:

- The different versions of DB2, and the various DB2 products.
- The tools that are included with DB2.
- How to use the Control Center to manage systems, DB2 instances, databases, database objects, and more.
- How the Configuration Assistant lets you maintain a list of databases to which your applications can connect, manage, and administer.
- All of the standalone tools in the Control Center and the Configuration Assistant.
- What data warehousing is, and the DB2 products available to assist with data warehousing.

Prerequisites

The process of installing DB2 is not covered in this tutorial. If you haven't already done so, we strongly recommend that you download and install a copy of [DB2 Express - C](#). Installing DB2 will help you understand many of the concepts that are tested on the DB2 9 Family Fundamentals Certification exam. The installation process is documented in the Quick Beginnings books, which can be found at the [DB2 Technical Support](#) Web site under the Technical Information heading.

System requirements

You do not need a copy of DB2 to complete this tutorial. However, you will get more out of the tutorial if you download the free trial version of [IBM DB2 9](#) to work along with this tutorial.

Section 2. DB2 products

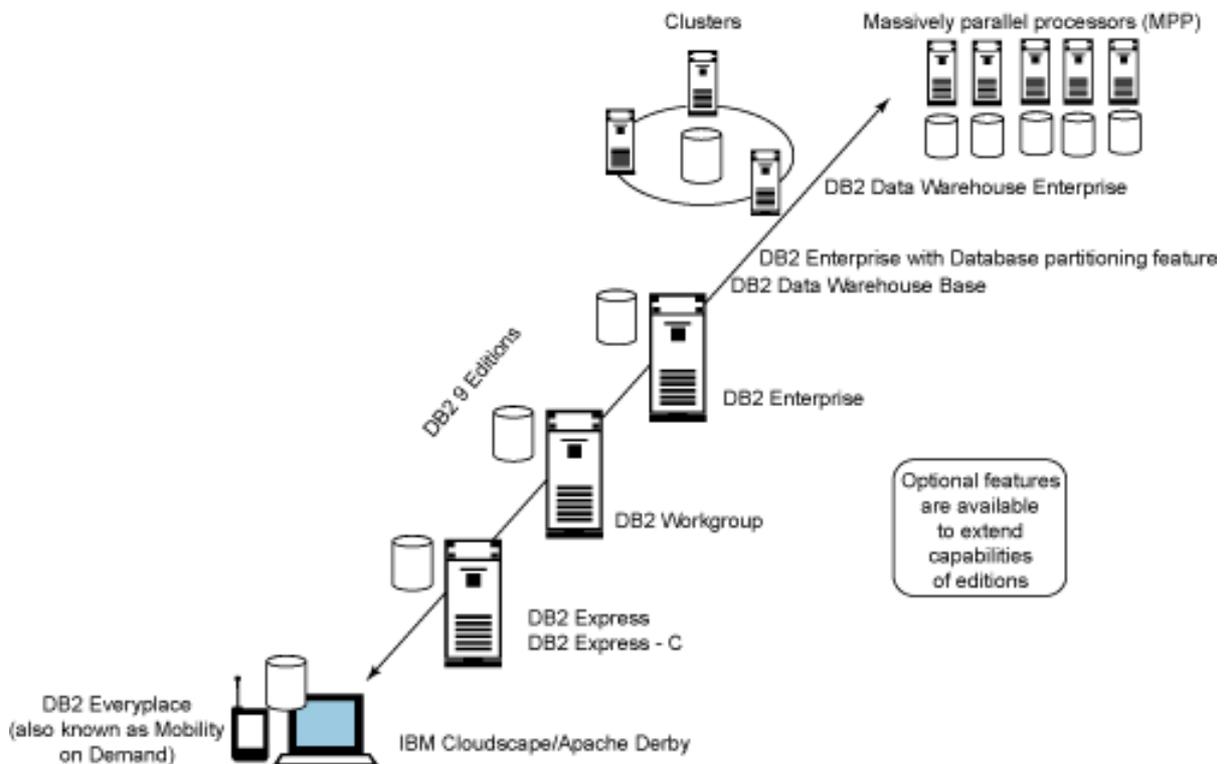
The different editions of DB2

DB2 9 delivers the right data management solutions for any business. No other database management system can match the advanced performance, availability, scalability, and manageability features found in DB2 9. However, there are different editions of DB2 available, each suited to a different part of the marketplace. On the Fundamentals exam you are expected to understand the different DB2 products and editions, so they are covered in this section.

All the available distributed editions of DB2 are shown in the figure below. This figure represents a progression: each edition displayed includes all the functions, features, and benefits of the editions to its right as you move up the stack, along with new features and functions. The code on the Linux, UNIX, and Windows (luw) platforms is about 90% common, with 10% of the code on each operating system reserved for tight integration into the underlying operating system. For example, using Huge Pages on AIX or the NTFS file system on Windows.

There are two other members of the DB2 family that are not shown in the figure below: DB2 for System i and DB2 for System z. While these databases share different code bases that are specifically tailored to their underlying operating systems and hardware architectures upon which they run, their SQL is 95% portable, truly making it a member of the DB2 family. For example, DB2 for System i is built into the i5/OS operating system. DB2 for z/OS leverages the hardware Coupling Facility on the System z servers and thus leverages a shared everything architecture, as opposed to DB2 luw, which uses a shared-nothing approach.

DB2 editions



Although it's outside the scope of this tutorial series to discuss the granular licensing of these editions, it's worth noting that some capabilities of DB2 9 are available for free in DB2 Enterprise. Where a capability is not included for free with DB2 Express

or DB2 Workgroup, you can purchase the function (in most cases) through an add-on Feature Pack.

For example, with DB2 Express 9 and DB2 Workgroup 9, you can add capabilities to your data server installations by buying one of the following feature packs:

Pure XML

Provides DB2 9's new XML data column type and indexes. DB2 9 comes with a hybrid engine that can handle SQL-based data, manipulated and stored relationally, and XML-based data that's manipulated and stored hierarchically.

High Availability

Provides online table reorganization, Tivoli System Automation for AIX and Linux, and the High Availability Disaster Recovery (HADR) function. Included free-of-charge in DB2 Enterprise.

Performance Optimization

Required for the use of Multidimensional Clustering (MDC) tables, Materialized Query Tables (MQTs), and query parallelism. Included free-of-charge in DB2 Enterprise.

Workload Management

Provides the Connection Concentrator, DB2 Query Patroller, and the DB2 Governor. The Connection Concentrator and DB2 Governor features are included free-of-charge in DB2 Enterprise.

DB2 Enterprise 9 comes with the following add-on features to extend the capabilities of this DB2 edition:

Pure XML

Provides DB2 9's new XML data column type and indexes. DB2 9 comes with a hybrid engine that can handle SQL-based data, manipulated and stored relationally, and XML-based data that's manipulated and stored hierarchically.

Advanced Access Control (LBAC)

For the provisioning of an extended security architecture that's based on role access to data.

Geodetic Data Management Feature

For the modeling of spatial and spherical data patterns used in various applications such as weather analysis, military defense, and applications that need to account for the curvature of the earth in their analysis.

Storage Optimization Feature

For row-level and backup/restore compression that can significantly increase the speed of operations, and minimize storage costs for you data.

Performance Optimization Feature

Provides the DB2 Performance Expert and DB2 Query Patroller products for use in a DB2 Enterprise server environment.

DB2 Everyplace

The true power of mobile computing lies not in the mobile device itself, but in its ability to tap into data from other sources. DB2 Everyplace brings the power of DB2 to mobile devices, leveraging their ability to synchronize data with other systems -- literally putting your enterprise data in the pockets of your mobile workforce and letting them update your enterprise data from remote locations.

DB2 Everyplace is more than just a mobile computing infrastructure. It's a complete environment that includes the tools you need to build, deploy, and support powerful e-business applications. DB2 Everyplace features a tiny "fingerprint" engine (about 200 KB) packed full of security features such as table encryption, and advanced indexing techniques that lead to high performance. It can comfortably run (with multithreaded support) on a wide variety of today's most commonly deployed handheld devices, such as: Palm OS, Microsoft Windows Mobile Edition, any Windows-based 32-bit operating system, Symbian, QNX Neutrino, Java 2 Platform Micro Edition (J2ME) devices like RIM's Blackberry pager, embedded Linux distributions (such as BlueCat Linux), and more.

If you need a relational engine and synchronization services, on a constrained device, you should use DB2 Everyplace. You should also consider this product for occasionally connected mobile users on laptops if their applications don't need features (such as triggers) that are not part of the DB2 Everyplace engine.

DB2 Everyplace was also shipped in DB2 8 as the Mobility-on-Demand feature. When you come across this feature in the DB2 8 or DB2 9 releases, you can assume the functions delivered by both products are identical. While the packaging changes between releases, DB2 Everyplace and DB2 Mobility-on-Demand deliver the same functions, features, and capabilities to your environment.

In DB2 9, Mobility on Demand is provided free with DB2 Enterprise. DB2 Express and DB2 Workgroup users need to purchase DB2 Everyplace Enterprise Edition to achieve this level of function.

DB2 Personal Edition

DB2 Personal Edition (DB2 Personal) is a single-user RDBMS that runs on low-cost commodity hardware desktops. DB2 Personal is available for Windows- and Linux-based workstations. DB2 Personal has all of the features of DB2 Express with one exception: remote clients cannot connect to databases that are running this edition of DB2. (However, workstations with the Control Center can connect to these databases to perform remote administration.) Because "DB2 is DB2 is DB2," applications that are developed for DB2 Personal will run on any other edition of DB2. For example, you can use DB2 Personal to develop DB2 applications before rolling them out into a production environment on DB2 Enterprise 9 for AIX.

DB2 Personal is useful both for PCs that are not connected to a network and for

those that are. In either case, it is useful for users who need a powerful data store, or who need to provide database storage facilities and be able to connect to remote DB2 servers.

Occasionally connected users may want to take advantage of DB2's built-in replication feature and the DB2 Control Server to set up a synchronized environment where mobile workers can keep in touch with their enterprise. Of course, this would only be suitable for users of laptops and certain workstations, such as those running point-of-sale (POS) applications.

DB2 Express - C

DB2 Express - C is not *really* considered an edition of the DB2 family, but it provides most of the capabilities of [DB2 Express](#). In January 2006 IBM announced this special free version of DB2 for Linux- and Windows-based operating systems. DB2 Express-C was designed for the partner and development communities, but as you get to know this version, you'll realize it has applicability almost anywhere. A defining characteristic of DB2 Express - C is that it doesn't have the limits that are typically associated with these types of offerings from other vendors. Where limits do exist, they are more than generous for the workloads designed for these systems.

For example, DB2 Express - C doesn't come with a database size limit and can address a 64-bit memory model. DB2 Express-C is perfect for developers and small and medium deployments, academic communities, and more. DB2 Express-C has all the resiliency and robustness of DB2 Express, but without some of the extended features with the fee-based DB2 Express edition. Features that are *not* included in DB2 Express-C include:

- Capability for features found in the DB2 Express Feature Packs - for example, HADR
- Replication Data Capture
- 24x7 IBM Passport Advantage support model

If you want to leverage any of these features in your environment, you need to at a minimum purchase DB2 Express.

DB2 Express Edition

DB2 Express Edition (DB2 Express) is a full-function, Web-enabled client/server RDBMS. DB2 Express is available for Windows- and Linux-based workstations. DB2 Express provides a low-cost, entry-level server that is intended primarily for small business and departmental computing. It has the same functions as [DB2 Workgroup](#), but is differentiated from DB2 Workgroup by the amount of memory and value units (which equate to the power of a server's processor cores) you can have on the server.

Additional features can be added to allow for extended capabilities, such as some found in DB2 Enterprise, without having to buy that edition. The Feature Packs available for DB2 Express 9 were outlined earlier in this tutorial.

DB2 Express can be licensed using a value unit determined by the processors running the application or a per Authorized User metric. Authorized users are a new concept to DB2 9, and represent users that are registered to access the services and data of a single data server in the environment. For example, if you had a user that needed to access two different DB2 Express 9 data servers and wanted to license this environment with authorized users, a single user would require two DB2 Express authorized user licenses (one for each server).

DB2 Express can play many roles in a business. It is a good fit for small businesses that need a full-fledged relational database store. They may not have the scalability requirements of some more mature or important applications, but they like knowing they have an enterprise quality database backing their application that can easily scale (without a change to the application) if they need it to. As noted, an application written for any edition of DB2 is transparently portable to another edition on any distributed platform.

DB2 Workgroup Edition

DB2 Workgroup Edition (DB2 Workgroup) is a full-function, Web-enabled client/server RDBMS. It is available on all supported flavors of UNIX, Linux, and Windows.

DB2 Workgroup provides a low-cost, entry-level server that is intended primarily for small business and departmental computing. Functionally, it supports all the same features as [DB2 Express](#). Additional features can be added to allow for extended capabilities such as those found in DB2 Enterprise, without having to purchase DB2 Enterprise. DB2 Workgroup can be licensed using the same options as DB2 Express.

In DB2 8, there were two types of Workgroup Edition: DB2 Workgroup Server Edition (DB2 WSE) and DB2 Workgroup Unlimited Edition (DB2 WSUE). DB2 WSE was only licensed by a named user license in addition to a base server license. DB2 WSUE was only licensed by a processor metric. In DB2 9, these editions merge into one edition -- DB2 Workgroup. The named user and server licenses have been replaced by a simplified Authorized User. The processor license still exists, albeit it by the conversion to Value Unit pricing per IBM pricing policies.

DB2 Workgroup can play many roles in a business. It is a good fit for small or medium-sized businesses (SMBs) that need a full-fledged relational database store that is scalable and available over a wide area network (WAN) or local area network (LAN). It is also useful for enterprise environments that need silo servers for lines of business, or for departments that need the ability to scale in the future. As previously noted, an application written for any edition of DB2 is transparently portable to another edition on any distributed platform.

DB2 Enterprise Edition

DB2 Enterprise Edition (DB2 Enterprise) is a full-function, Web-enabled client/server RDBMS. It is available on all supported flavors of Linux, UNIX, and Windows. DB2 Enterprise is meant for large and mid-sized departmental servers. DB2 Enterprise includes all the functions of the DB2 Express and DB2 Workgroup editions, and more. And, certain DB2 9 features are only available for this edition, such as the new DB2 9 Storage Optimization Feature.

DB2 Enterprise can be licensed using a value unit determined by the architecture of the processor running the application or a per Authorized User metric, just like DB2 Express and DB2 Workgroup. Authorized users are a new concept to DB2 9 (though this metric was available in DB2 8 Enterprise Server Edition), and represent users that are registered to access the services and data of a single data server in the environment. For example, if you had a user that needed to access two different DB2 Enterprise 9 data servers and wanted to license this environment with authorized users, a single user would require two DB2 Enterprise authorized user licenses (one for each server). Some features, such as the Database Partitioning Feature, are not available using the authorized user metric. DB2 Enterprise also officially supports sub-capacity licensing, like LPARs and dynamic LPARs.

DB2 Enterprise has the ability to partition data within a single server, across multiple database servers (all of which have to be running on the same operating system), or within a large SMP machine out of the box, thanks to its database partitioning feature (DPF).

You can purchase DPF as part of a DB2 Enterprise processor license, which also gets converted to Value Units. With DPF, the size of your database is only limited by the number of computers you have. DB2 Enterprise with DPF is meant for larger data warehouses, or for high-performance online transaction processing (OLTP) requirements. DB2 Enterprise with the DPF also allows multiple SMP machines to be clustered together under a single database image for very large-scale transaction volumes.

Data Enterprise Developer Edition

A special offering called the Data Enterprise Developer Edition (DEDE) is available for application developers. This edition offers several information management products that allow a single application developer to design, build, and prototype applications for deployment on any of the IBM Information Management client or server platforms. This comprehensive developer offering includes:

- DB2 Workgroup 9 and DB2 Enterprise 9
- IDS Enterprise Edition
- IBM Cloudscape/Apache Derby
- DB2 Connect Unlimited Edition

- And *all* the DB2 9 add-on features described earlier in this tutorial

This allows customers to build solutions that use the latest data server technologies with a reduced-price offering. The products found in DEDE are restricted to the development, evaluation, demonstration, and testing of your application programs.

DB2 8 had a free offering called DB2 Personal Developer's Edition that came with DB2 8 Personal Edition and DB2 8 Connect Personal Edition. This package was removed and replaced with DB2 Express - C in DB2 9.

DB2 clients

DB2 9 greatly simplifies the deployment of the infrastructure required to get your applications connecting to a DB2 database. DB2 9 provides the following clients:

DB2 9 Runtime Client

The best option if your only requirements are to enable applications to access DB2 9 data servers. They provide the APIs necessary to perform this task, but this client comes with no management tools.

DB2 9 Client

Includes all the functions found in the DB2 Runtime Client *plus* functions for client-server configuration, database administration, and application development through a set of rich graphical tools. The DB2 9 Client replaces the functions found in *both* the DB2 8 Application Development and DB2 8 Administration clients.

Java Common Client (JCC)

This 2 MB fully redistributable client provides JDBC and SQLJ applications access to DB2 data servers without installing and maintaining DB2 client code. If you are connecting to a DB2 for System i or DB2 for System z data server, you are still required to purchase the DB2 Connect product.

DB2 9 Client Lite

New in DB2 9, this client performs similar functions to the JCC client, but instead of supporting Java-based access to a DB2 data server it's used for CLI/ODBC applications. This client is especially well suited for ISVs that want to embed connectivity in their applications without redistributing and maintaining DB2 client code.

DB2 Extenders

The DB2 Extenders discussed in this section can take your database applications beyond traditional numeric and character data, and provide additional functions to the underlying data server.

XML Extender

DB2's XML Extender provides data types that let you store XML documents in DB2 databases, and adds functions that help you work with these XML documents while in a database.

You can store entire XML documents in DB2, or store them as external files managed by the database. This method is called *XML Columns*. You can also decompose an XML document into relational tables and then recombine that information to XML on the way out of the database. Basically, this means your DB2 database can strip the XML out of a document and just take the data, or take data and create an XML document from it. This method is known as *XML Collections*.

What about the pureXML Feature that's new in DB2 9?

You might be confused about the XML Extender and the pureXML add-on feature that's available in DB2 9 for all editions of this product. The DB2 XML Extender provides the XML capabilities that were part of the DB2 8 release. The pureXML feature enables DB2 servers to leverage the new hybrid storage engine that stores XML naturally in DB2 9. The performance, usability, flexibility, and overall XML experience of pureXML can't even be compared to the old XML Extender technology - however, the XML Extender is still shipped in DB2 9 free-of-charge. If you are planning to use XML in your data environment, it is strongly recommended you use the pureXML feature.

The pureXML feature lets you store XML in a parsed tree representation on disk, without having to store the XML in a large object or shred it to relational columns as you are forced to with the XML Extender. This can be very beneficial for applications that need to persist XML data.

With the XML Extender you need to use functions, and it doesn't support XQuery. If you're retrieving XML data, you can access just portions of the XML document without reading the entire document (if it was stored in a LOB), shredding it, and performing a join (if it were stored in relational tables), which are the only methods supported by the XML Extender.

Access to the data is a very natural experience when using the capabilities provided by the pureXML feature. For example, you can use SQL or XQuery to get to relational or XML data.

DB2 9 supports the shredding of XML data to relational in the same manner as the XML Extender, but it uses a different and far superior technology to do it. You may want to shred your XML to relational for any number of reasons, such as when the XML data is naturally tabular. To shred XML to relational using the DB2 XML Extender, you have to hand-generate Document Access Definition documents that map nodes to columns, and so on. With DB2 9, even without the pureXML feature, you can use the [DB2 Developer Workbench](#) to shred your data and automate the discovery of these mappings. The new mechanism in DB2 9 is also significantly faster than the XML Extender method.

DB2 Net Search Extender

This extender helps businesses that need fast performance when searching for

information in a database. High-performance in-memory searches are indispensable for e-commerce applications or any other application with high performance and scalability text-search demands. You are likely to see this used in Internet applications, where excellent search performance on large indexes and scalability of concurrent queries are needed. You also use this extender to search large XML documents. If you need a high-speed in-memory search, this is the extender for you. In DB2 8, the Text Information Extender was merged with the Net Search Extender. This extender is free in DB2 9 (in DB2 8 it was a chargeable feature).

DB2 Spatial Extender

This extender allows you to store, manage, and analyze *spatial data* -- information about the location of geographic features -- in DB2 alongside traditional data for text and numbers. With this capability, you can generate, analyze, and exploit spatial information about geographic features, such as the locations of office buildings or the size of a flood zone. The DB2 Spatial Extender extends the function of DB2 with a set of advanced spatial data types that represent geometries such as points, lines, and polygons. It also includes many functions and features that interoperate with those data types. These capabilities let you integrate spatial information with your business data, adding another element of intelligence to your database. This extender is free in DB2 9 (and has been since DB2 8.2).

DB2 Geodetic Extender

This extender allows you to enhance the type of applications you can build with the DB2 Spatial Extender. The DB2 9 Geodetic Extender lets you treat the earth as a globe and removes inaccuracies caused by operations like projections. Using the same spatial data types and functions provided in the DB2 Spatial Extender, you can use the DB2 Geodetic Extender to run seamless queries of data around the earth's poles and data that crosses the 180th meridian. You can maintain data that is referenced to a precise location on the surface of the earth.

DB2 Geodetic Extender is named for the discipline of geodesy, which is the study of the size and shape of the earth (or any body modeled by an ellipsoid, such as the sun or a celestial sphere). The DB2 Geodetic Extender is designed to handle objects defined on the earth's surface with a high degree of precision. The DB2 Geodetic Extender is only available for DB2 Enterprise 9.

DB2 Connect

A great deal of the data in many large organizations is managed by DB2 for i5/OS, DB2 for MVS/ESA, DB2 for z/OS, or DB2 for VSE and VM data servers. Applications that run on any of the supported DB2 distributed platforms can work with this data transparently, as if a local data server managed it. You can also use a wide range of off-the-shelf or custom-developed database applications with DB2 Connect and its associated tools. Quite simply, DB2 Connect provides connectivity to mainframe and midrange databases from Windows, Linux, and UNIX platforms.

There are a number of DB2 Connect editions available: Personal Edition, Enterprise

Edition, Application Server Edition, and two Unlimited Editions (one for i5/OS environments, and one for z/OS environments). The DB2 Connect products can be added on to an existing DB2 data server installation, or act as a stand-alone gateway. Either way, it's purchased separately (although some complimentary user licenses are provided in DB2 Enterprise). See [Resources](#) for more information on DB2 Connect.

DB2 Add-on tools

There are two kinds of tools for DB2: those that are free and those that are add-ons that can be purchased separately. The free tools come as part of a DB2 installation and can be launched from the Control Center, the Configuration Assistant, or on their own (you will learn about them in the next section of this tutorial).

A separate set of tools are available to help ease the database administrator's (DBA) task of managing and recovering data and making it accessible for distributed versions of DB2:

Tool	Description
DB2 Change Management Expert	Improves DBA productivity and reduces human error by automating and managing complex DB2 structural changes.
Data Archive Expert	Responds to legislative requirements like Sarbanes-Oxley by helping DBAs move seldom-used data to a less costly storage medium without additional programming.
DB2 High Performance Unload	Maximizes DBA productivity by reducing maintenance windows for data unloading and repartitioning.
DB2 Performance Expert	Makes DBAs more proactive in performance management to maximizes database performance.
DB2 Recovery Expert	Protects your data by providing quick and precise recovery capabilities.
DB2 Table Editor	Keeps business data current by letting end users easily and securely create, update, and delete data.
DB2 Test Database Generator	Quickly creates test data and helps avoid liabilities associated with data privacy laws by protecting sensitive production data used in test.
DB2 Web Query Tool	Broadens end user access to DB2 data using the Web and handheld devices.

Not all of these tools are available for all the DB2 9 editions. However, the licensing nuances are out of the scope of this tutorial.

Section 3. DB2 tools

Tools overview

The tools that are included with DB2 (hereafter called the *DB2 tools*, and not to be confused with the purchasable DB2 tools discussed in the [previous section](#)) provide a whole array of time-saving, error-reducing graphical interfaces into most of the DB2 features. With these tools, you can do the same tasks from a graphical user interface (GUI) that you can do from a command line or API. However, when you use the DB2 tools, you don't have to remember complex statements or commands, and you get additional assistance through online help and wizards -- so let's hear it for the DB2 tools!

The DB2 tools are part of the *DB2 Client*. When you install a DB2 server, you are actually installing all of the components of a DB2 Client as well (though most people don't realize it). The DB2 Client enables you to install the DB2 tools on any workstation, and lets you manage remote database servers. The DB2 Client also provides the required components to set up an application development.

The DB2 tools are really divided into two camps:

The Control Center (CC)

Is primarily used for administering DB2 servers. There are several other centers that are integrated and can be started from the Control Center.

The Configuration Assistant (CA)

Is used for setting up client/server communications and maintaining registry variables, though it can do more. We'll learn more about the CA in a bit.

Basic tool functions

There are about six basic features that you should be able to find in any DB2 tool (when applicable): Wizards, Generate DDL, Show SQL/Show Command, Show Related, Filter, and Help.

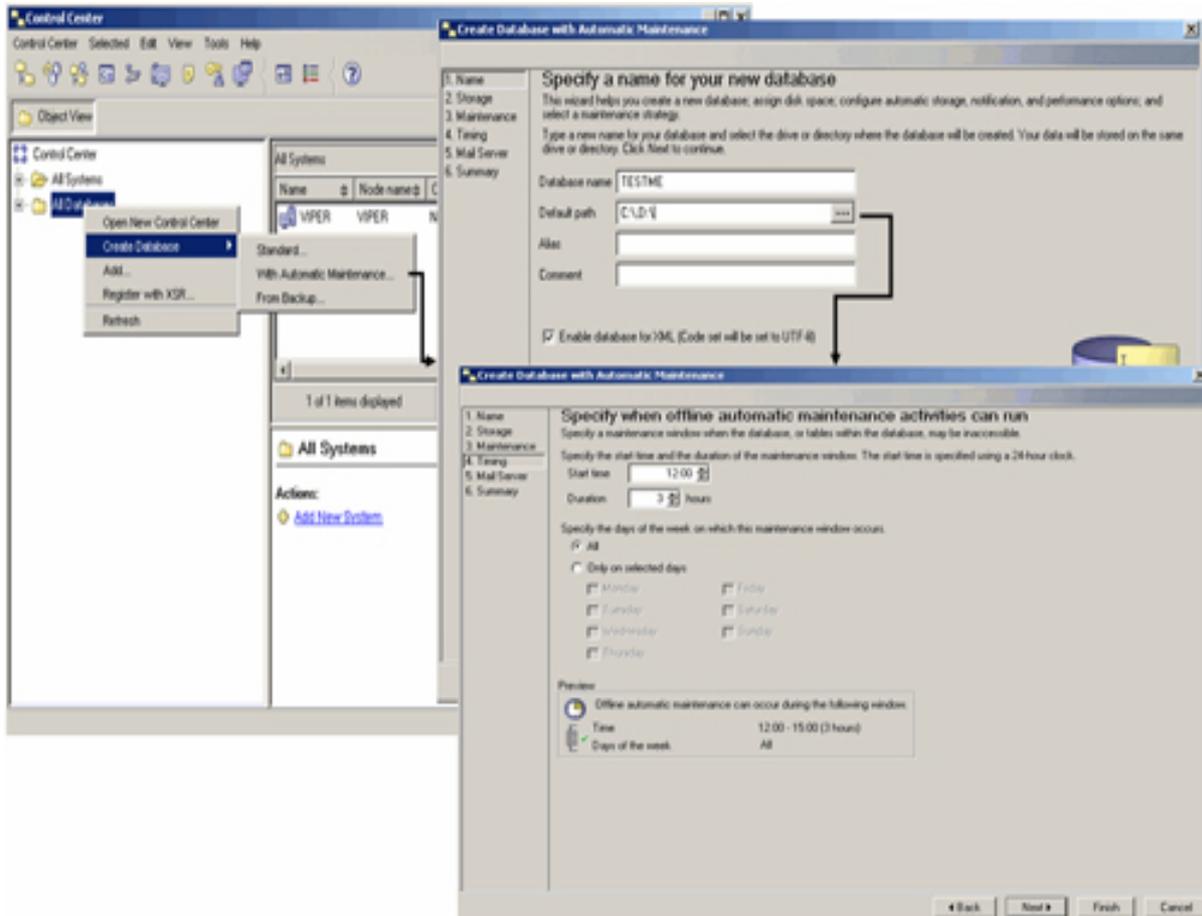
Wizards

Wizards can be very useful to both novice and expert DB2 users. Wizards help you complete specific tasks by taking you through each task one step at a time, and recommending settings where applicable. Wizards are available through both the Control Center and the Configuration Assistant.

There are wizards for adding a database to your system (cataloging it), creating a database, backing up and restoring a database, creating tables, creating tablespaces, configuring two-phase commits, configuring database logging, updating

your documentation, setting up a High Availability Disaster Recovery (HADR) pair, tuning your performance, and more. The following figure shows some panels of the Create Database wizard in DB2 9.

Creating a Database using a Wizard



If you were creating a database using this wizard, you could automate many of the post administration steps as well. For example, in the previous figure you can see that the TESTME database will be created with automatic maintenance. Also note the **Enable database for XML (Code set will be set to UTF-8)** checkbox. If you're leveraging the pureXML feature in DB2 9, you need to create your database in UTF-8 unicode format; this is another example of how the wizard can make you more productive. If you forgot to specify this option when creating a database from the command line processor, you would have to drop and recreate the database since this is a characteristic of a database that cannot be changed.

Advisors

There are special types of wizards that do more than just provide assistance in completing a task. Traditional wizards take you step-by-step through a task, simplifying the experience by asking important questions or generating the complex command syntax for the action you want to perform. When a wizard has more intelligence than just task completion and can offer advisory type functions, DB2 calls them *advisors*. They operate just like wizards, but have a lot of intelligence (some pretty complex algorithms) that churns out advice based on some input factors such as workload or statistics. Advisors help you with more complex

activities, such as tuning tasks, by gathering information and recommending options that you may not have considered. You can then accept or reject the advisor's advice. You can call advisors from the GUI, from APIs, and the command-line interface.

Advisors are part of the IBM autonomic computing effort, which aims to make software and hardware more SMART (self-managing and resource tuning)! Unlike some competitive offerings, the Advisors in DB2 are all included at no additional charge in *every* edition of DB2, including DB2 Express - C.

There are two main advisors in DB2 9: the Configuration Advisor and the Design Advisor.

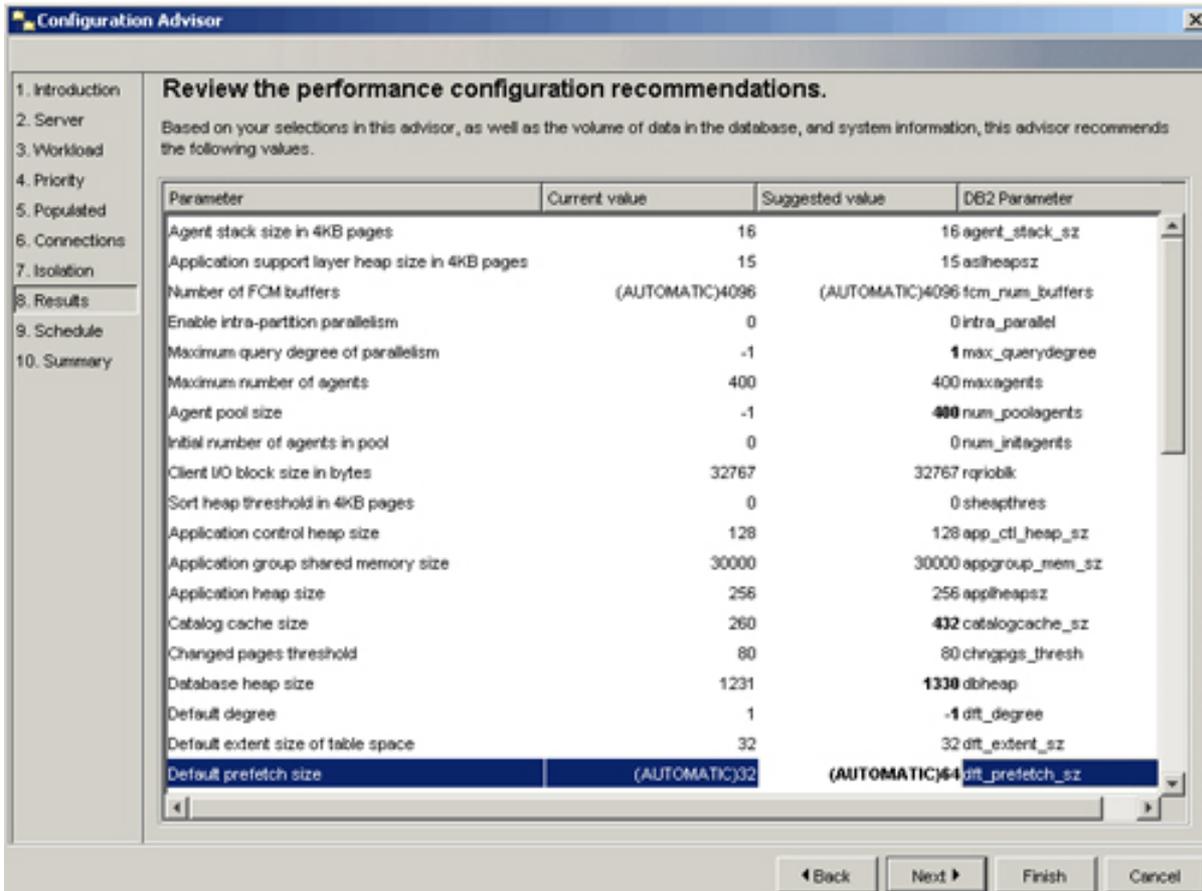
The DB2 Cube Views product also comes with an Optimization Advisor, but that topic is outside the scope of the DB2 Fundamentals Certification.

There is another Advisor that comes with DB2 called the DB2 Recommendation Advisor. This Advisor can only be accessed from the DB2 Health Center when DB2 surfaces an issue with its regular health check of your DB2 instances and their database (more on that in a bit).

The Configuration Advisor can be used to set instance and database-level configuration parameters for your DB2 environment. It asks you several high level questions that describe your environment (do you care more about the performance or availability of your database - or both equally, how many users will access a database concurrently, how much memory would you like to dedicate for DB2's use, and more). After converting the answers into input parameters that are passed to the underlying algorithms, DB2 SMARTly considers the answers you gave and makes several configuration recommendations based on your responses. The Configuration Advisor is especially well suited for OLTP workloads, but also works well with business intelligence-based workloads.

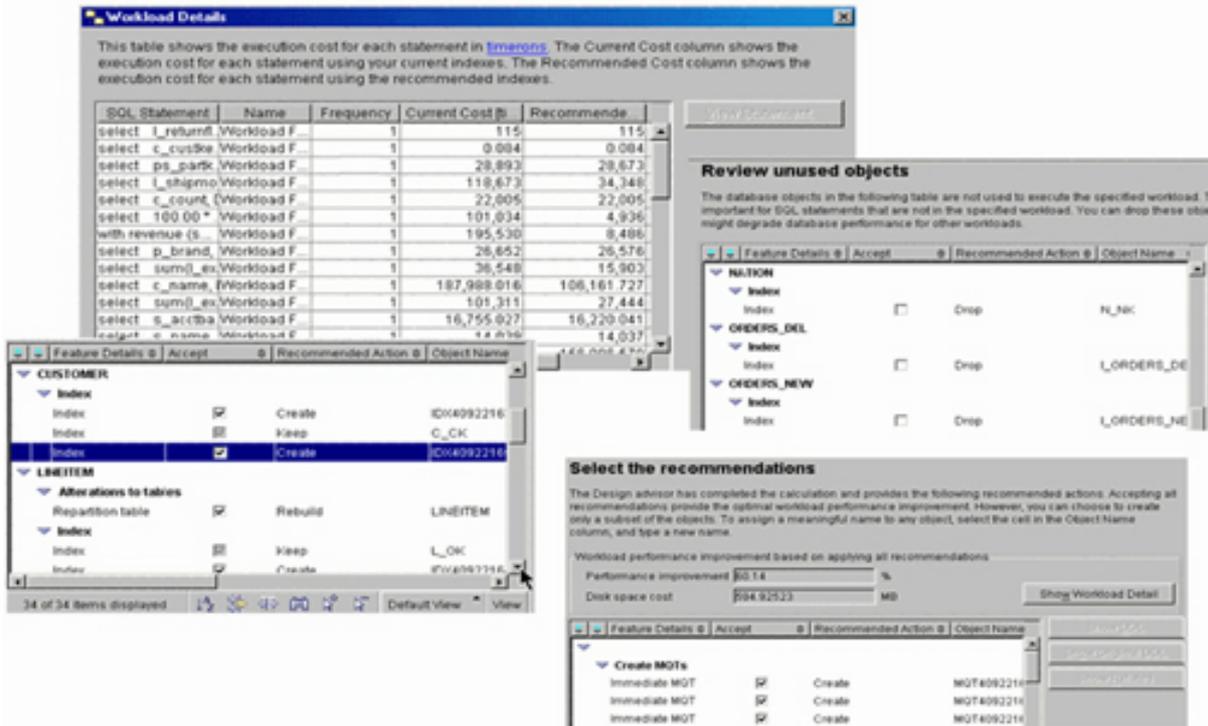
DB2 9 introduces a new feature for automated tuning for the shared database memory working set (also available free in all DB2 9 editions) called the Self Tuning Memory Manager (STMM). Using the Configuration Advisor with STMM is a great combination for an optimal, hands-off, dynamically tuning database system.

The Configuration Advisor works so well that in DB2 9 it's automatically started after you create a database (in some cases) using the Control Center. Even if you are an expert DBA, it is recommended that you use this tool. Think of the hours you can save by having DB2 provide you with what it thinks is an optimal configuration for your application. Then you can hand tune the performance to hit the expert level you'll undoubtedly attain after achieving your certification! An example of the Configuration Advisor is shown below.



The Design Advisor takes as input a workload that is either provided in a file, captured in the cache, in a DB2 Query Patroller repository, and more. Using the workload, the Design Advisor can suggest a change to the underlying database schema to attain optimal performance based on the submitted workload. The Design Advisor can suggest new (or changes to) indexes, MQTs, MDCs, and partitioning keys (used when you've installed the Database Partitioning Feature). It can also identify indexes that aren't being used for possible removal.

Keep in mind when you use this advisor, however, that the recommendations are only based on the submitted workload. It's an important point. The Design Advisor might tell you to drop an index or create an MDC table based on a query, but that could work against the performance of others queries. When using this tool, be sure you are profiling the most important parts of you application. An example of the Design Advisor is shown below.



The Design Advisor is different from wizards in that a wizard would help you create an index, but the advisor would actually suggest a specific index to create. Advisors truly let DBAs improve their productivity, and potentially their skills since it can be used as a learning tool, thereby reducing the effort and total cost of ownership of a DB2 solution.

Notebook

Another type of assistance tool, a notebook, differs from wizards because it doesn't step you through a particular process (such as creating a table). Notebooks simplify the task by reducing the time it takes to complete it. Essentially, notebooks are great for eliminating the need to memorize clunky syntax. Notebooks exist for such tasks as setting up event monitors, creating indexes, buffer pools, triggers, aliases, schemas, views, and more. The following figure shows the Create View notebook.

Using a notebook to create a view

Create View

VIPER - DB2 - DEMO

View schema: VIPER

View name: PZVIEW

SQL statement:

```
[(columns, ...)]  
AS [WITH (common_table_expression, ...)]  
SELECT <columns>  
FROM <tables>  
WHERE <search_conditions>
```

Check options:

None
 Cascaded
 Local

Comment:

Buttons: Undo, Clear, SQL Assist, OK, Cancel, Show SQL, Help

When taking the exam you should know about all the wizards, advisors, and notebooks, and how to use them. It is recommended that you go through the Control Center and the Configuration Assistant, exploring these helpers and performing the various tasks with their help. Right-click everywhere and explore with a test database: remember, practice makes perfect!

Generate DDL

The Generate DDL function lets you re-create, and optionally save in a script file, the Data Definition Language (DDL), authorization statements required to recreate the privileges on an object, the tablespace where the object resides, nodegroups, buffer pools, database statistics, and pretty much anything else that makes up the basis of your database (except the data).

By using the Generate DDL feature, you can save the DDL to create identically defined tables, databases, and indexes in another database -- using it as a cookie cutter, if you will. Administrators like to use this option to create a test environment that mimics the production environment. One nice thing about DB2, since you can manually update the statistics (something you should **never** do in a production environment), is that you can use this feature with the Generate DDL function to create a test database *without* having to load the data in the tables. When you click on the Generate DDL option, you are actually running the `db2look` DB2 system command.

If you want to move data into your new database objects to quickly set up a test database, you could use the traditional `LOAD` or `IMPORT` utilities, or the `db2move` command. This tool facilitates the movement of large numbers of tables between DB2 databases located on distributed workstations. `db2move` queries the system catalog tables for a particular database and compiles a list of all user tables. It then exports these tables in PC/IXF format.

Show SQL/Show Command

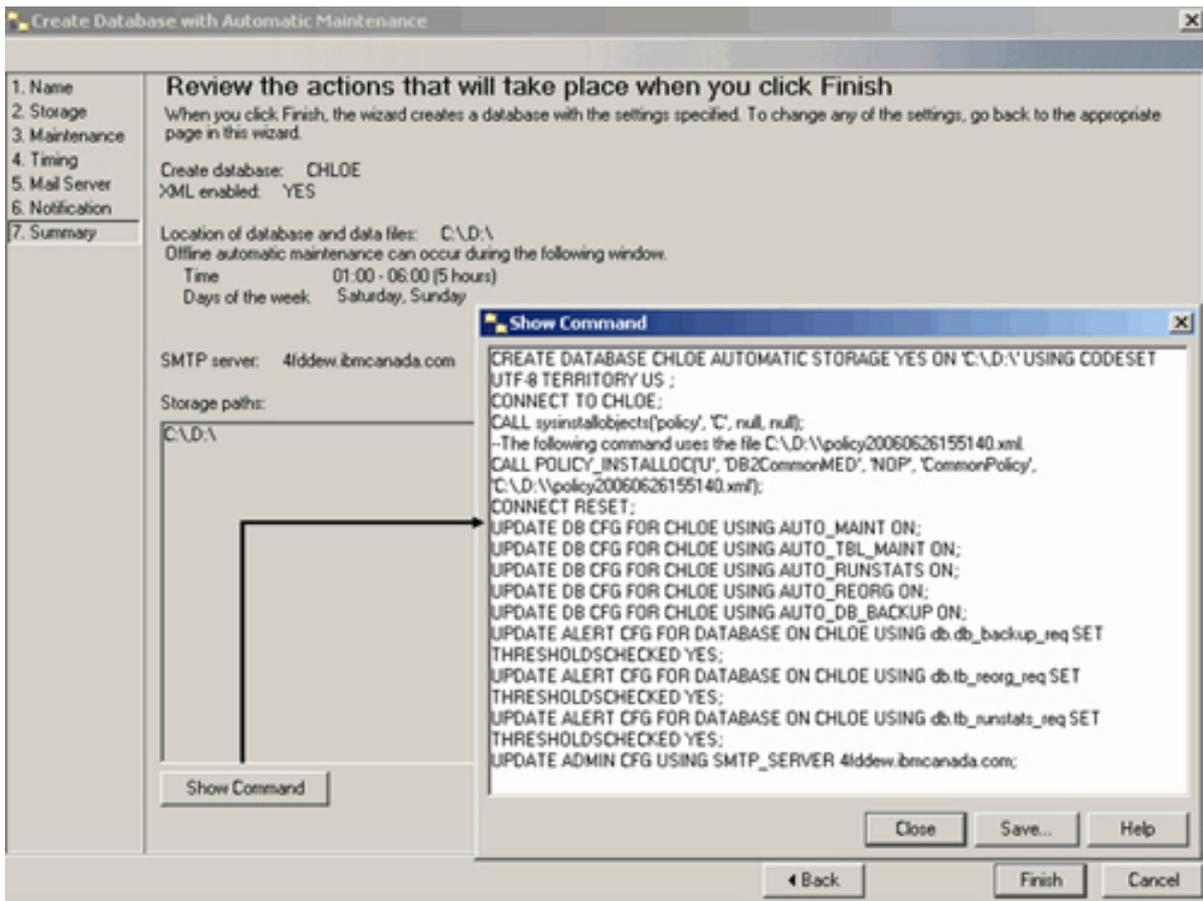
If a tool generates SQL statements or DB2 commands, then the Show SQL or Show Command button will be available on that tool's interface. Selecting this button will show the actual statement or command that DB2 will use to perform the task you've requested. You can save the information returned by this feature as a script for future reuse (so you don't have to retype it again), schedule it for execution later, or just use it to get a better idea of what's happening behind the interface. You can also use the copy and paste features of your operating system to work with the generated syntax in another application.

The following figure shows the `CREATE DATABASE` command that was generated by the Create Database Wizard (of course, if the wizard were generating SQL, the option would be to show the SQL generated for the task) for a database called CHLOE that:

- Will be used with the pureXML feature
- Has an automated maintenance plan whereby offline maintenance can be performed on Saturdays and Sundays between 1:00am and 5:00am
- Whose containers will be striped across the C: and D: drives using the DB2 automated storage management feature
- Will send e-mail notifications to DBAs through the `4fddew.ibmcanada.com` mail server to a pager

The Show Command option gives you the syntax for the task you are trying to do; this is a lot of hand written DDL you just saved yourself from writing.

The Show Command option

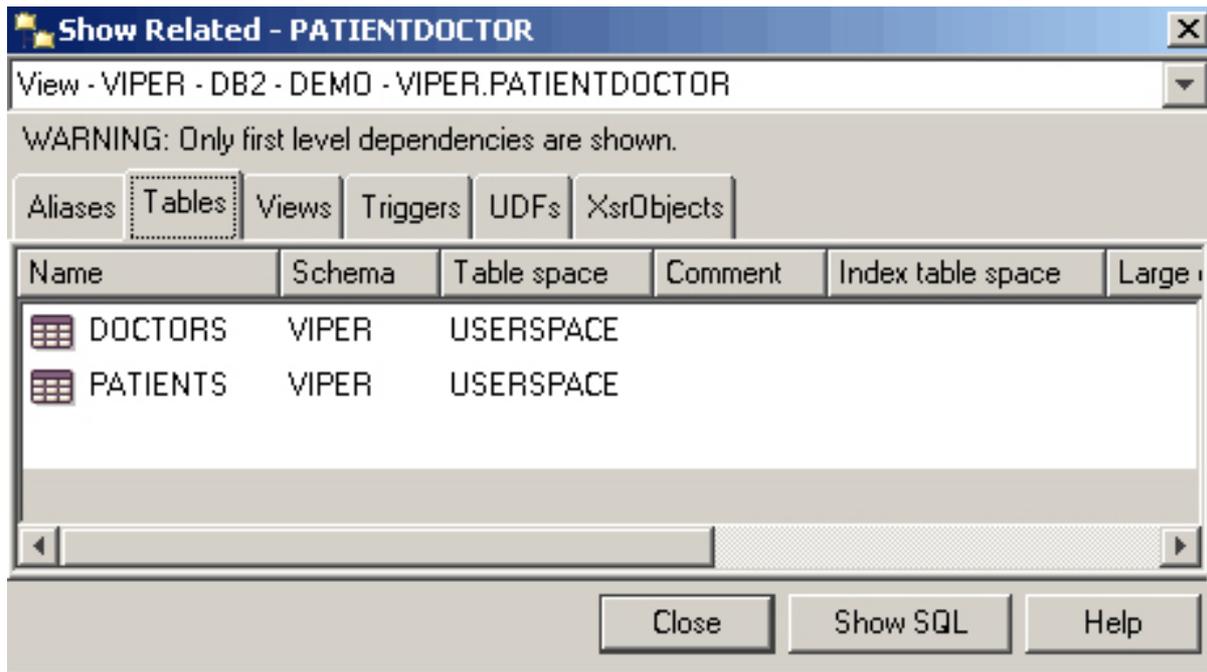


Show Related

The Show Related feature returns the immediate relationship between tables, indexes, views, aliases, triggers, tablespaces, user-defined functions (UDFs), and user-defined types (UDTs). For example, if you select a table and you choose to show the related views, you will only see the views that are based directly on that specific base table. You will not see views that are based on the related views, because those views were not created directly from the table.

By seeing a list of related objects, you can better understand the structure of a database, determine what objects already exist in a database and their relationships to one another, and much more. For example, if you want to drop a table with dependent views, the Show Related feature will identify which views will become inoperative as a result of dropping that object.

The following figure shows the results of using the Show Related feature on a view. As you can see, the VIPER.PATIENTDOCTOR view has dependencies on the VIPER.PATIENTS and VIPER.DOCTORS tables. Using this information, you should be able to tell that if either of these two tables were dropped, the VIPER.PATIENTDOCTOR view would become inoperable. The Show Related option shows you the relationships within or between database objects -- in this case, a view and its base table.



Filter

You can filter the information that is displayed in the contents pane of any DB2 tool. You can also filter information that is returned from a query (such as limiting the number of rows in a result set).

The tools let you save and name multiple filters and recall them at a later time. If you select the **View** button at the bottom right corner of the Control Center pane that displays the highlighted database objects, you'll see a pop-up dialog where you can create, save, and edit filters. Take a moment now to create a filter for all of the database objects that you create under your own user ID. In later sections of this tutorial, you can then use this filter to quickly and easily find the database objects you want to work with. You can imagine just how important these filters are, especially when working with supply chain management (SCM) or enterprise resource planning (ERP) applications like SAP, which have tens of thousands of tables.

Help

Extensive help information is provided with the DB2 tools using the Eclipse help engine. A Help button exists on most dialog boxes, as well as on the menu toolbar. These facilities provide you general help, and help on how to fill out the fields and perform tasks of a particular tool. From the help menus, you can also access a glossary and index of terms used in the dialog or reference information, along with the information provided in the product manuals.

The DB2 help is task-oriented, which should make it easier to locate the information required to do a particular task (for example, creating a database). DB2 also provides an update wizard to notify you that there are documentation updates awaiting your installation.

The DB2 processors: An introduction

The DB2 Command Line Processor (DB2 CLP), common to all DB2 products, is an application you can use to run DB2 commands, operating system commands, or SQL statements. This tool can be a somewhat cryptic method of invoking DB2 commands. However, the DB2 CLP can be a powerful tool because it extends its capability to store often-used sequences of commands or statements in batch files that can be run when necessary.

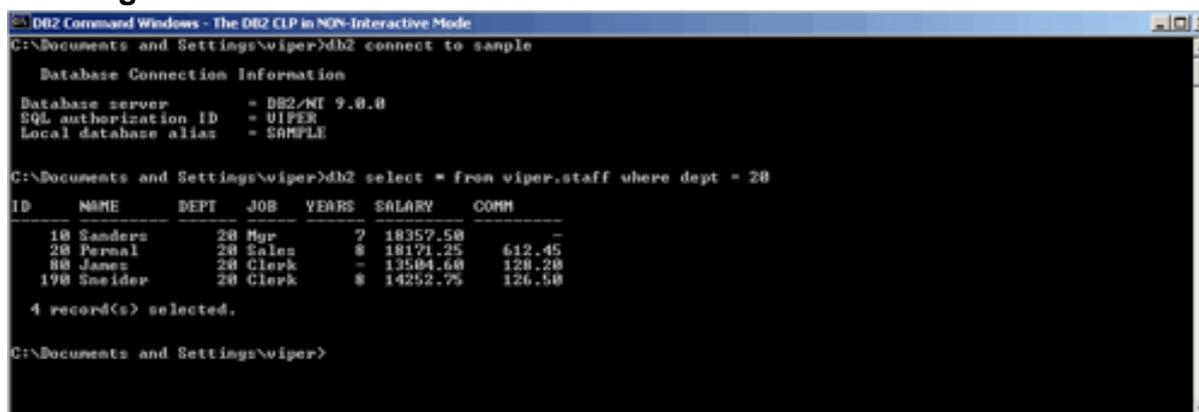
Some implementations of DB2 can use the operating system's native command line interface to enter DB2 commands; others cannot. For this reason, we'll refer to two different processors in DB2: the DB2 Command Line Processor (DB2 CLP) and the DB2 Command Window (DB2 CW). People are likely to call them by the same name since they share the same icon. In this tutorial we'll refer to the mode in which you do not have to prefix commands with the keyword `db2` as the DB2 CLP in *interactive* mode.

The DB2 CLP lets you enter DB2 commands interactively, without using the `db2` prefix to tell the operating system that you're planning to enter a DB2 command. However, if you want to enter an OS command, you have to prefix it with the exclamation mark, also called a bang key (!). For example, in the DB2 CLP, if you wanted to run the `dir` command, you would enter `!dir`.

For all operating systems other than Windows, the DB2 CW is built into the operating system's native CLP. In a Windows environment, you have to start a DB2 CW from a Windows command prompt by entering the `db2cmd` command or by selecting the appropriate option from the Start menu.

You can start the DB2 CLP from a DB2 CW by entering the `db2` command on its own. The following figure shows a command entered through the DB2 CW.

Entering a command with the DB2 CW



```
DB2 Command Windows - The DB2 CLP in NON-Interactive Mode
C:\Documents and Settings\vipier>db2 connect to sample

Database Connection Information
Database server      = DB2/NT 9.0.0
SQL authorization ID = UIPER
Local database alias = SAMPLE

C:\Documents and Settings\vipier>db2 select * from viper.staff where dept = 20
ID      NAME      DEPT  JOB    YEARS  SALARY  COMM
-----
10 Sanders 20 Mgr   7      18357.50 -
20 Permal 20 Sales 8      18171.25 612.45
80 James  20 Clerk -      13584.68 128.28
170 Snider 20 Clerk 8      14252.75 126.58

4 record(s) selected.

C:\Documents and Settings\vipier>
```

Notice that I had to enter the keyword `db2` to get this DB2 command to run. If I hadn't, the operating system would have thought this was an operating system command, and would return an error. If you are using the DB2 CLP, you don't need to do this, as shown in the figure below.

Entering a command with the DB2 CW in Interactive Mode

```

C:\Documents and Settings\vipier>db2
(c) Copyright IBM Corporation 1993,2002
Command Line Processor for DB2 ADCL 9.0.8

You can issue database manager commands and SQL statements from the command
prompt. For example:
  db2 -> connect to sample
  db2 -> bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
? CATALOG DATABASE for help on the CATALOG DATABASE command
? CATALOG for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 -> connect to sample

Database Connection Information
Database server = DB2/NT 9.0.8
SQL authorization ID = UIPER
Local database alias = SAMPLE

db2 -> select * from viper.staff where dept = 20
ID      NAME      DEPT  JOB    YEARS  SALARY  COMM
-----
10 Sanders 20 Mgr      7  18357.50  -
20 Pernal 20 Sales   8  18171.25  612.45
30 James 20 Clerk   -  13504.60  128.20
190 Snider 20 Clerk   8  14252.75  126.50

4 record(s) selected.

db2 ->

```

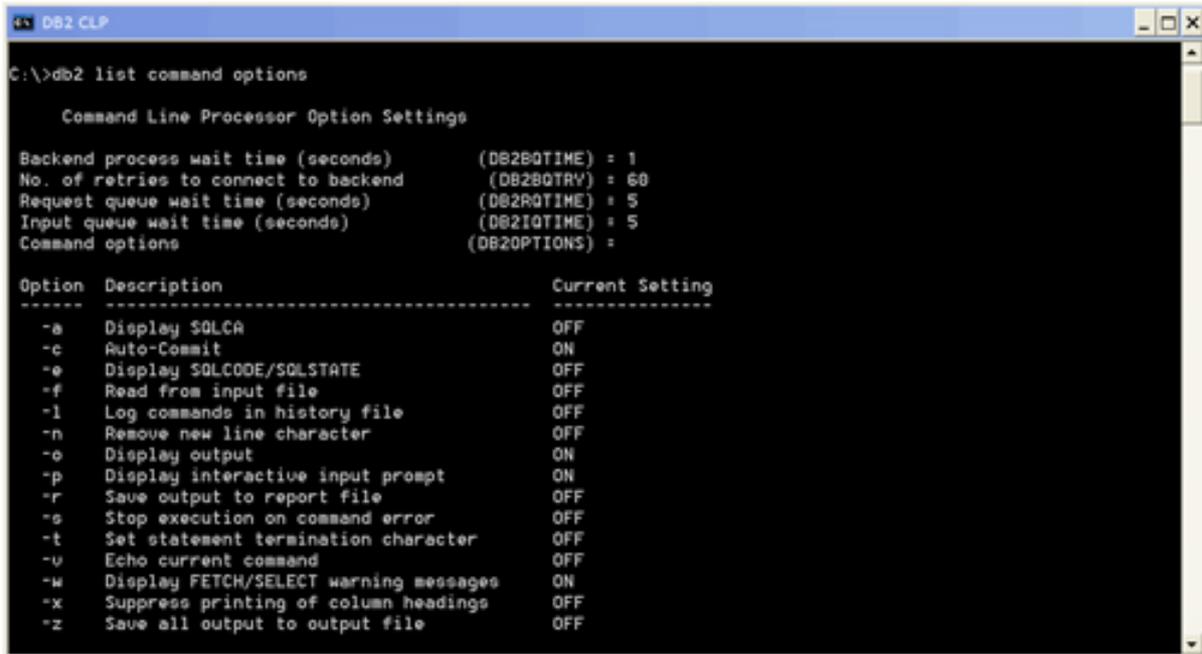
Using the DB2 processors

When using a DB2 processor, you can use command line options that alter the way the process, or a single statement or command entered from it, behaves. You can specify one or more processor options when you invoke a DB2 command. Some of the options that you can control are:

- The auto-commit of each statement that you can define using the `c` flag.
- An an input file that provides the DB2 commands and SQL statements which you can define using the `f` flag.
- The end-of-statement termination character (the default character is `;`), defined by the `t` flag.

You can get a list of all the valid options by entering `list command options` in a DB2 processor (don't forget when you would have to include the `db2` prefix to make this work). Run this command now, and you'll see there are over 15 different options, as shown below.

The various DB2 CLP options



There are two ways to change the options for a DB2 processor. You can set command options for a session by setting the `DB2OPTIONS` registry variable (which must be in uppercase), or by specifying command-line flags when you input a DB2 command. The latter method will override any settings made at the registry level. If you change the behavior for a single statement, that will override any settings in the session and the registry.

To turn an option on, prefix the corresponding option letter with a minus sign (-); for example, to turn the auto-commit feature on (which is the default), enter:

```
db2 -c command or statement...
```

To turn an option off, either surround the option letter with minus signs (-c-) or prefix it with a plus sign (+). Read the last two sentences again, because this can get confusing: a minus sign *before* a flag turns an option *on*, but a minus sign *before and after* a flag, or a *plus sign* before the flag, turns that option *off*. No, that isn't very intuitive (hey, I didn't write the code). Because this can be confusing, let's walk through an example with the auto-commit option.

Some command line options are on by default and some are off. The previous explanation (and the following example) describe the behavior and the effect of the command line options on options that are on by default. You would use the opposite logic if a command line option was off by default.

By default, the auto-commit feature is set to (-c). This option specifies if each statement is automatically committed or rolled back.

If a statement is successful, it and all successful statements that were issued before it with auto-commit set to off (+c or -c-) are committed. If, however, the statement fails, it and all successful statements that were issued before it with auto-commit set to off are rolled back. If auto-commit is set to off for the statement, you must

explicitly issue a commit or rollback command.

In the following figure, the value of the auto-commit feature on the command line was changed to illustrate this process.

Changing command line options at run time

```

C:\Documents and Settings\viper>db2 connect to sample

Database Connection Information
Database server      = DB2/NT 7.0.0
SQL authorization ID = VIPER
Local database alias = SAMPLE

C:\Documents and Settings\viper>db2 +c create table a (c1 int)
DB20000I The SQL command completed successfully.

C:\Documents and Settings\viper>db2 +c create table b (c1 int)
DB20000I The SQL command completed successfully.

C:\Documents and Settings\viper>db2 +c select * from a,b

C1          C1
-----
0 record(s) selected.

C:\Documents and Settings\viper>db2 rollback
DB20000I The SQL command completed successfully.

C:\Documents and Settings\viper>db2 select * from a,b
SQL0204N "VIPER.A" is an undefined name.  SQLSTATE=42704

C:\Documents and Settings\viper>_

```

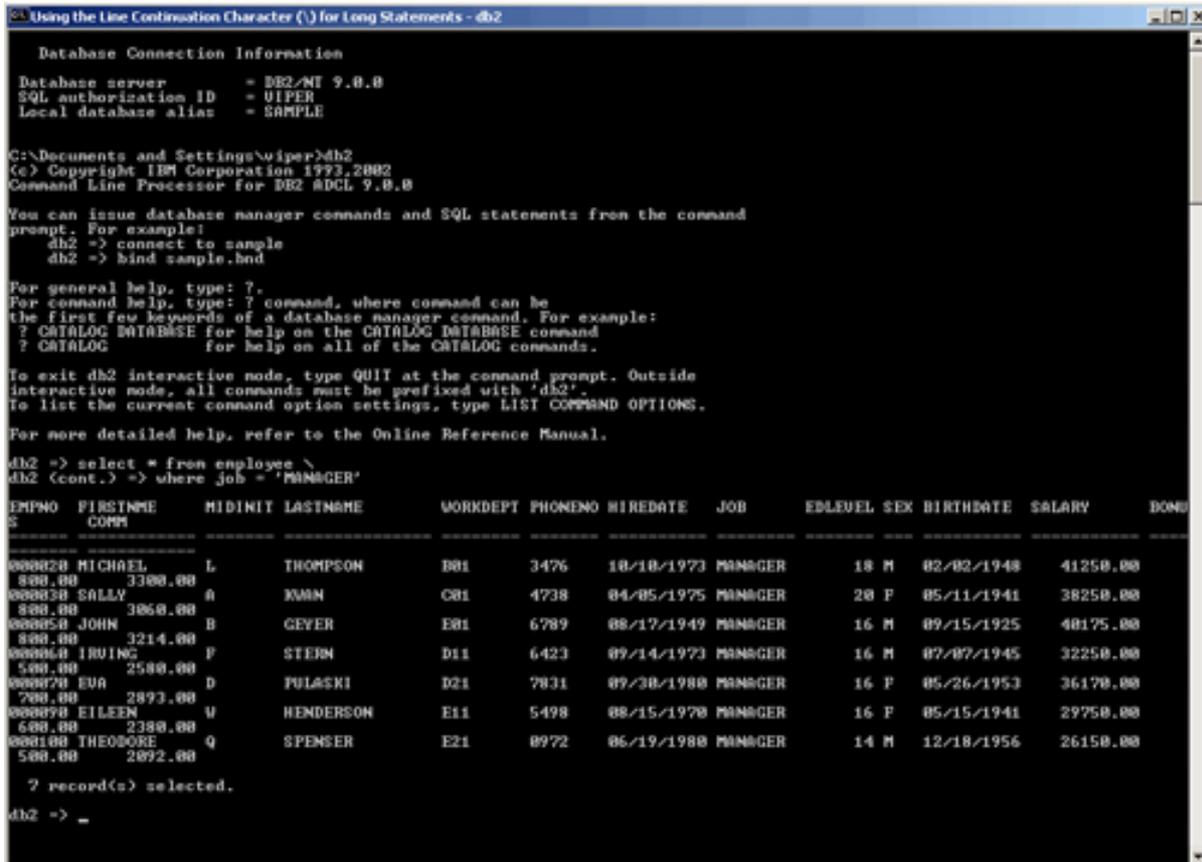
So what happened? Well, first I created a table called A, but did this while at the same time turning the default auto-commit option to off by using the `+c` option. (I could have surrounded this flag with a plus sign (`-c-`) and it would have done the same thing.) After creating table A (but not committing this action, remember), I created another table called B while turning off the auto-commit feature as well. Then I did a Cartesian join of both tables, again by dynamically turning off the DB2 CLP's auto-commit feature. Finally, I did a rollback and ran the same `SELECT` statement again and this time it failed.

If you look at the transaction, I never issued a commit. If the first `SELECT` hadn't included the `+c` option, it would have committed the creation of tables A and B and the tables (since the `SELECT` was successful) and this therefore successfully returns the same results as the first `SELECT` statement did.

Try the exact same sequence of commands, only this time use the `-c-` option. You should experience the same behavior. After that, try it without any options on the first `SELECT` statement and see if the second `SELECT` returns results or not.

Your operating system might have a maximum number of characters that it can read in any one statement (even when it wraps to the next line in your display). To work around this limitation when entering a long statement, you can use the line continuation character (`\`). When DB2 encounters the line continuation character, it reads the next line and concatenates the two lines during processing. You can use this character with both DB2 processors; however, be aware that DB2 has a limit of a 2 MB statement (that's a lot of command line typing). The figure below illustrates its use in the DB2 CLP.

Using the line continuation character with the DB2 CLP



If you're using the DB2 CW to enter commands, you may run into a problem with some of these special characters:

```
$ & * ( ) ; < > ? \ ' "
```

The operating system shell may misinterpret these characters. (Of course, this is not an issue in the DB2 CLP, since it is a separate application specifically designed for DB2 commands.)

You can circumvent system operators that you want to be interpreted by DB2 and not by the operating system by placing your entire statement or command within quotes, as follows:

```
db2 "select * from staff where dept > 10"
```

Try entering the previous command in a DB2 CW without the quotes. What happened? Look at the contents of the directory where you issued the command. I'll bet you'll find a file there called 10 that has an SQL error in it. Why? Well, DB2 interpreted your SQL statement

```
select * from staff where dept
```

and tried to place those contents in a file called 10. The > sign is an operating system instruction to pipe any output from the standard display to a specified file (in this case 10). The `select * from staff where dept` statement is of course an incomplete SQL statement; hence the error. The incorrect results are due to the operating system misinterpreting the special character.

You should experiment with both DB2 processors to get a feel for which is the best choice in any given circumstance. Take some time now to enter commands in both the DB2 CLP and the DB2 CW.

Section 4. The Control Center

Control Center overview

The Control Center (CC) is the central management tool for DB2 data servers. You can use the CC to manage systems, DB2 instances, databases, database objects, and much, much more. From the CC you can also open other centers and tools to help you optimize queries, schedule jobs, write and save scripts, create stored procedures and user-defined functions, work with DB2 commands, monitor the health of your DB2 system, and more. (Some of these functions are provided by tools launched from the CC.)

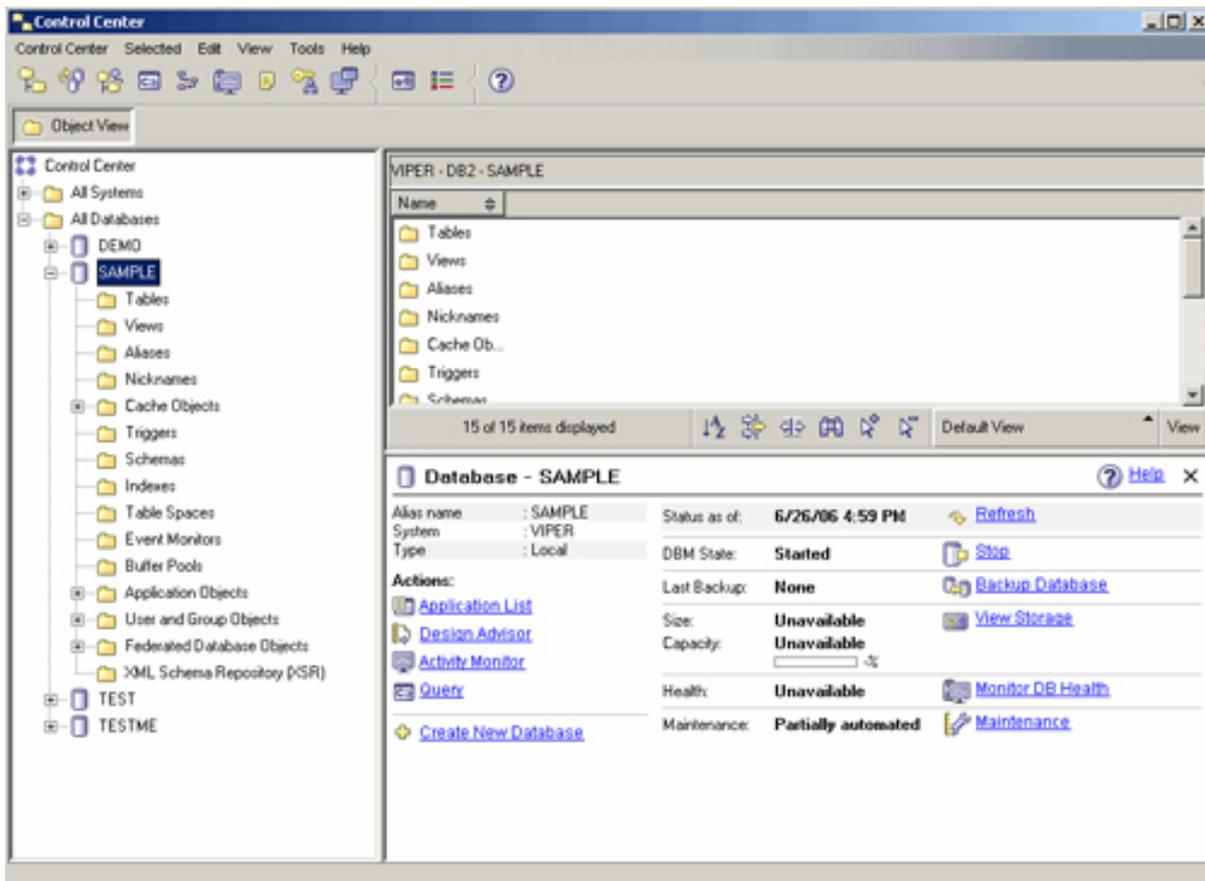
You can start the CC by entering the `db2cc` command from your operating system's command prompt, or locating the **Control Center** in your operating system's GUI-based interface.

Among other tasks, a DBA can use the CC to:

- Add DB2 systems, federated systems, DB2 for z/OS subsystems, IMS systems, and local and remote instances and databases to the object tree for management (some targets are limited with respect to the actions you can perform on them).
- Manage database objects. You can create, alter, and drop databases, tablespaces, tables, views, indexes, triggers, and schemas. You can also manage users.
- Manage data. You can load, import, export, reorganize, and collect statistics on your data.
- Perform preventive maintenance by backing up and restoring databases or tablespaces.
- Schedule jobs to run unattended. To schedule tasks through the CC, you must first create a TOOLS catalog database. If you did not create this database when you installed DB2 (it is an optional part of the installation process), you can do so from the **Tools Settings** option in the **Tools** action menu bar. If you haven't done so already, create it now.

- Configure and tune instances and databases.
- Manage database connections.
- Monitor and tune performance. You can run statistics, look at the execution path of a query, start event and snapshot monitoring, generate SQL or DDL of database objects or commands, and view relationships between DB2 objects.
- Troubleshoot.
- Manage data replication.
- Manage applications.
- Manage the health of your DB2 system.
- Launch other DB2 centers.
- AND MORE!!

The CC is shown below. You should be able to easily use this tool, as it is similar to many other interfaces in the marketplace; objects are on the left and details of those objects are on the right.

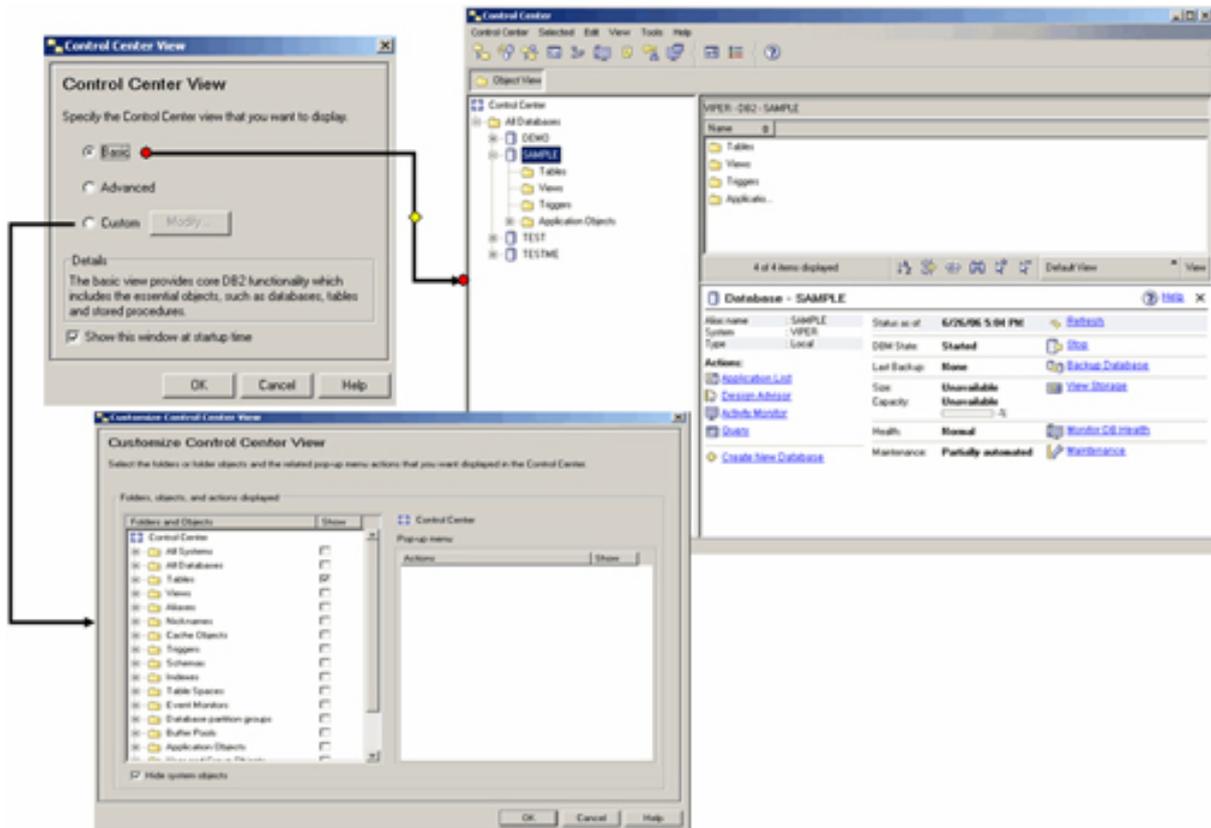


Notice the database dashboard that appears when you connect to a database. This provides DBAs with a one-stop location to quickly view the health of their DB2 databases. For example, in the previous figure you can see that no database backups exist for this database, and automated maintenance is only partially enabled. There is no information for the other categories that could either imply I've disabled them, or DB2 has not yet had time to return information on those objects.

To see everything that you can do with the CC, right-click on any object from the object tree. A pop-up menu shows all the functions you can perform on a selected object. For example, on the **Tables** folder, you can create a new table (you can also create a table from an `IMPORT` operation), create a filter for what is displayed in the contents pane, or refresh the view. The tasks that you can perform depend on the object that you select. I strongly recommend that you go through each folder and object and right-click your way to familiarity.

The CC lets you customize what is displayed in the navigation tree for a more simple view, or a customized view. The figure below gives an example of this process.

Customizing your DB2 Control Center



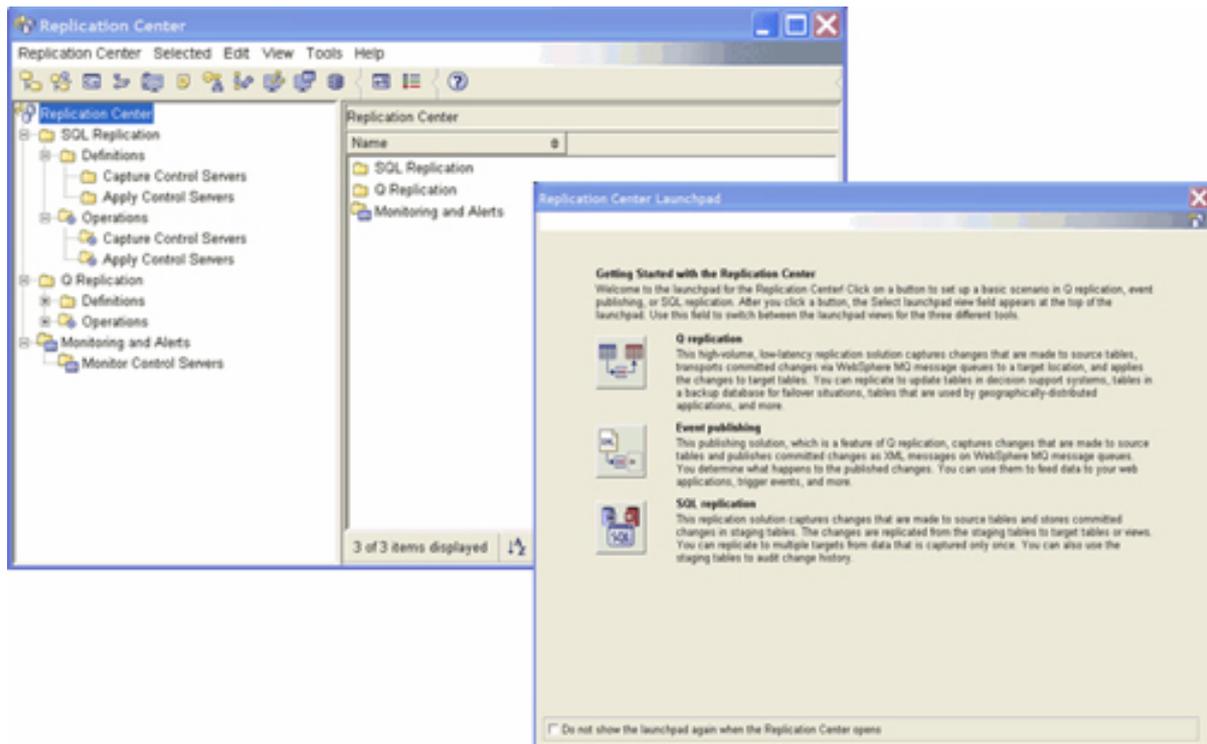
You can see in the previous figure that I chose a **Basic** view, which greatly simplifies the explorer tree. You can customize the folders and options that you can see when you right-click on an object. For example, you could customize the CC so it would only show the **Tables** folder, or even customize the actions that display when you right-click on this folder to only allow you to create a new table, not change an existing one.

The next several sections present a detailed description of the CC tools that can be started from the launchpad. Each tool that you launch includes the launchpad, so you can launch any DB2 tool from any other DB2 tool. This section covers the DB2 tools and centers that are used most often. Other sections in this tutorial will cover the remaining tools available from the launchpad that are not covered in this section.

The DB2 Replication Center

Use the DB2 Replication Center (DB2 RC) to administer replication between a DB2 data server and other relational databases (DB2 or non-DB2). From the DB2 RC, you can define replication environments, apply designated changes from one location to another, and synchronize data in two or more locations.

You can start the DB2 RC from the Start menu, from a DB2 tool's launchpad, or by entering the `db2rc` command at the command prompt. The following figure gives you an idea of what the RC looks like.



A replication-specific launchpad, shown above, is available to guide you through some of the basic replication functions. You can use the RC for all the kinds of replication that are supported in a DB2 environment. However, some of this function requires additional products. For example, Q-based replication is built on the WebSphere MQ Series product and is part of the WebSphere Information Integrator family of products.

Some of the key tasks you can perform with the RC include:

- Create replication control tables
- Register replication sources
- Create subscription sets
- Operate the Capture program
- Operate the Apply program
- Monitor the replication process
- Perform basic troubleshooting for replication.

The DB2 Satellite Administration Center

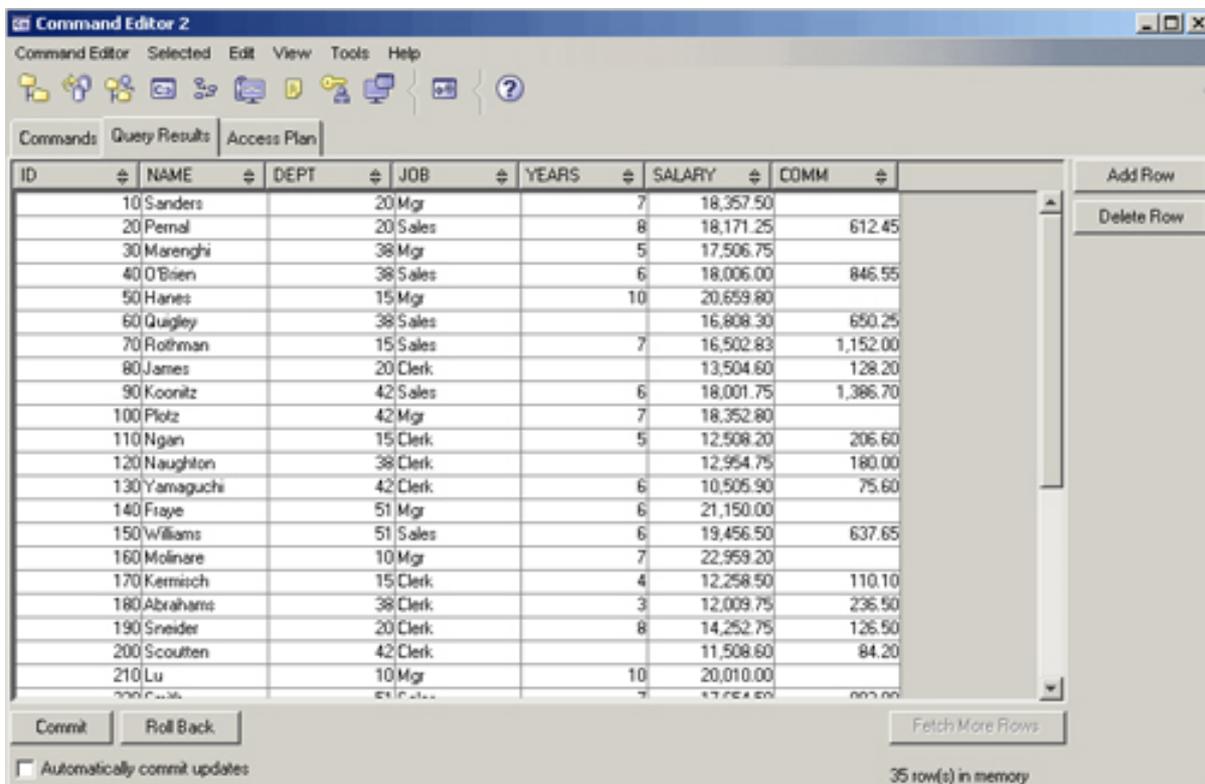
Use the Satellite Administration Center (DB2 SAC) to set up and administer a group of DB2 servers that perform the same business function. These servers, known as *satellites*, all run the same application and have the same DB2 configuration (database definition) to support that application. Using the DB2 SAC, you can have several DB2 data servers synchronize and maintain their configuration and data with

a master server. You can start the DB2 SAC from the launchpad in any DB2 tool.

The DB2 Command Editor

Use the DB2 Command Editor to build and execute DB2 commands and SQL statements, and to view a graphical representation of the access plan for an SQL statement.

You can start the Command Editor from the CC, from an operating system command prompt, or a Start menu in your operating system's graphical interface. The Command Editor is also available as a Web-based application that you can access from a Web browser, but its capabilities in this mode aren't as rich as when running it from a locally installed DB2 client. The **Query Results** page of the Command Editor is shown below.



Different tabs on the DB2 Command Editor provide different features:

Commands

Allows you to execute SQL statements or DB2 commands. (Entering DB2 commands in the Command Editor is like working in the interactive DB2 CLP mode: you don't need to use the `db2` prefix.) To run a command or statement that you enter, select the green **Play** button or press `Ctrl+Enter`. You can also enter operating system-specific commands from the DB2 CC by preceding the command with a bang (!) sign. For example, to list the contents of the current directory, enter `!dir`.

This tab also gives you the option to: retrieve from a history of commands you've run since starting the editor, specify an end of termination character for your session, and easily add database connections upon which you want to run

your statements.

Query Results

Shown in the previous figure, this tab lets you see the results of your query. You can also save the query's results or edit the contents of the table directly (assuming you have the privileges to do so).

Access Plan

Allows you to see the access plan for any explainable statement that you ran in this editor. DB2 will automatically generate the access plan when it compiles the SQL statement. You can use this information to tune your queries for better performance. If you specify more than one statement in a single operation, an access plan is created only for the first statement.

The DB2 Command Editor also comes with the SQL Assist tool, which is covered more in the next section. To invoke the SQL Assist tool, select the SQL Assist button on the **Commands** tab.

The DB2 Command Editor can also be run in a Web-based mode. In this mode, any Web browser, PDA, mobile, or other pervasive device that has access to the Internet can execute commands against a DB2 server. This helps DBAs continually keep in touch with their DB2 systems. In Web mode, the Command Editor does not have the Visual Explain or SQL Assist features.

The DB2 Task Center

Use the DB2 Task Center (DB2 TC) to run tasks, either immediately or according to a schedule, and to notify people about the status of completed tasks. You can start the DB2 TC from the Start menu in a Windows environment, from a DB2 tool's launchpad, or by entering the `db2tc` command from a command prompt.

A *task* is a script accompanied by associated failure or success conditions, schedules, and notifications. You can create a task within the DB2 TC, create a script within another tool and save it to the DB2 TC, import an existing script, or save the options from a DB2 dialog or wizard (such as the Load wizard) as a script. The script can contain DB2, SQL, or operating system commands.

For each task, you can:

- Schedule the task
- Specify success and failure conditions
- Specify actions that should be performed when this task completes successfully or when it fails
- Specify e-mail addresses (including pagers) that should be notified when this task completes successfully or when it fails.

You can also create a *grouping task*, which combines several tasks into a single logical unit of work. When a grouping task meets the success or failure conditions that you define, any follow-on tasks are run. For example, you could combine three backup scripts into a grouping task and then specify a reorganization as a follow-on task that will be executed if all of the backup scripts execute successfully. All of these features make the DB2 TC an indispensable resource for DBAs charged with managing a DB2 environment.

The DB2 Health Center

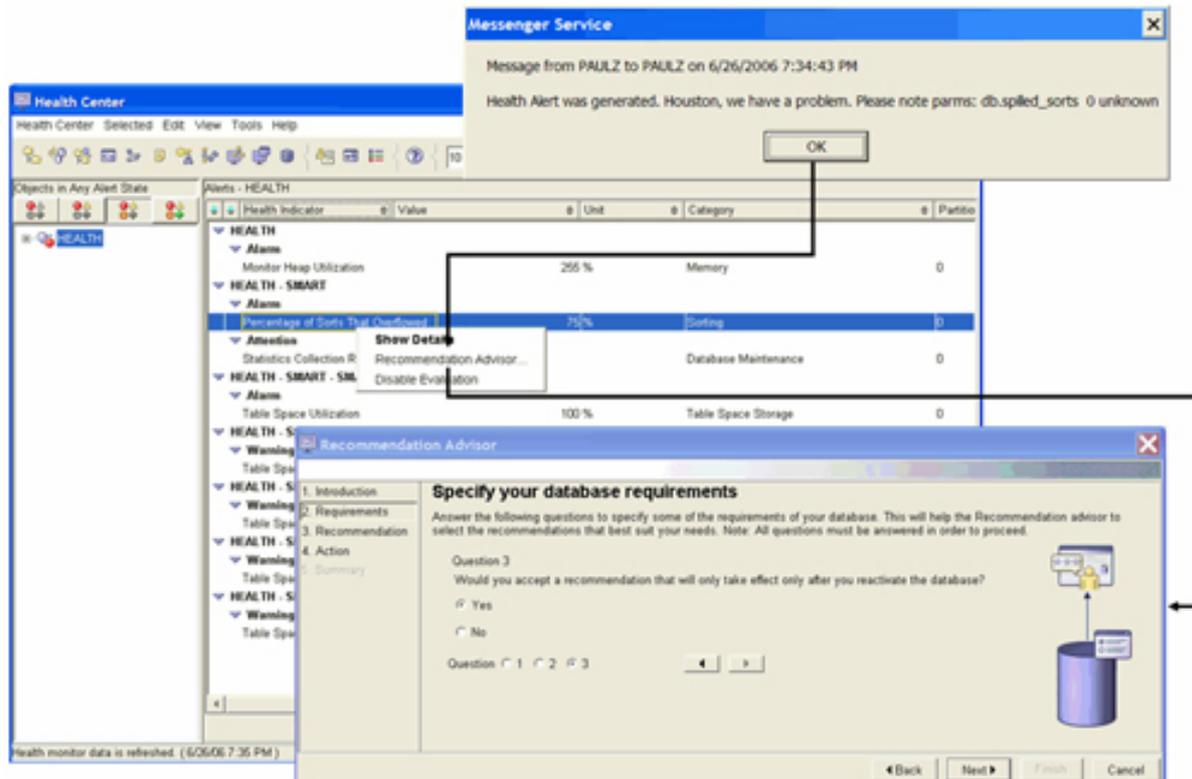
Use the DB2 Health Center (DB2 HC) to monitor the state of a DB2 environment and to make any necessary changes to it. You can start the DB2 HC from the Start menu in a Windows environment, from any DB2 tool's launchpad, or by entering the `db2hc` command at a command prompt.

When you use DB2, a monitor continuously keeps track of a set of health indicators. If the current value of a health indicator is outside the acceptable operating range defined by its warning and alarm thresholds, the health monitor generates a health alert. DB2 comes with a set of predefined threshold values, which you can customize. For example, you can customize the alarm and warning thresholds for the amount of memory allocated to a particular heap, or to notify you when sorts are spilling to disk, and so on.

Depending on the configuration of the DB2 instance, some or all of the following actions can occur when the health monitor generates an alert:

- An entry is written in the administration notification log, which you can read from the Journal.
- The DB2HC status beacon appears in the lower right corner of a DB2 tool's window.
- A script or task is executed.
- An e-mail or pager message is sent to the contacts that you specify for this purpose.

The figure below depicts the DB2 HC and a suggested resolution strategy in response to a condition it has detected.



You can see that an instance of DB2 running on Windows is using the integrated messaging service in Windows to notify a DBA that there is a problem with sorting.

There are a number of key tasks that you can perform with the DB2 HC. For instance, you can:

- View the status of the DB2 environment. Beside each object in the navigation tree, an icon indicates the status for that object (or for any objects contained by that object). For example, a green diamond icon beside an instance means that the instance and the databases contained in the instance do not have any alerts.
- View alerts for an instance or database. When you select an object in the navigation tree, the alerts for that object are shown in the pane to the right.
- View detailed information about an alert and recommended actions. When you double-click an alert, a notebook appears. One page shows the details for the alert; another shows any recommended actions, and so on. (The second page of this Advisor is shown in the previous figure.)
- Configure health monitor settings for a specific object, or the default settings for an object type or for all the objects within an instance.
- Select which contacts will be notified of alerts by e-mail or pager

message.

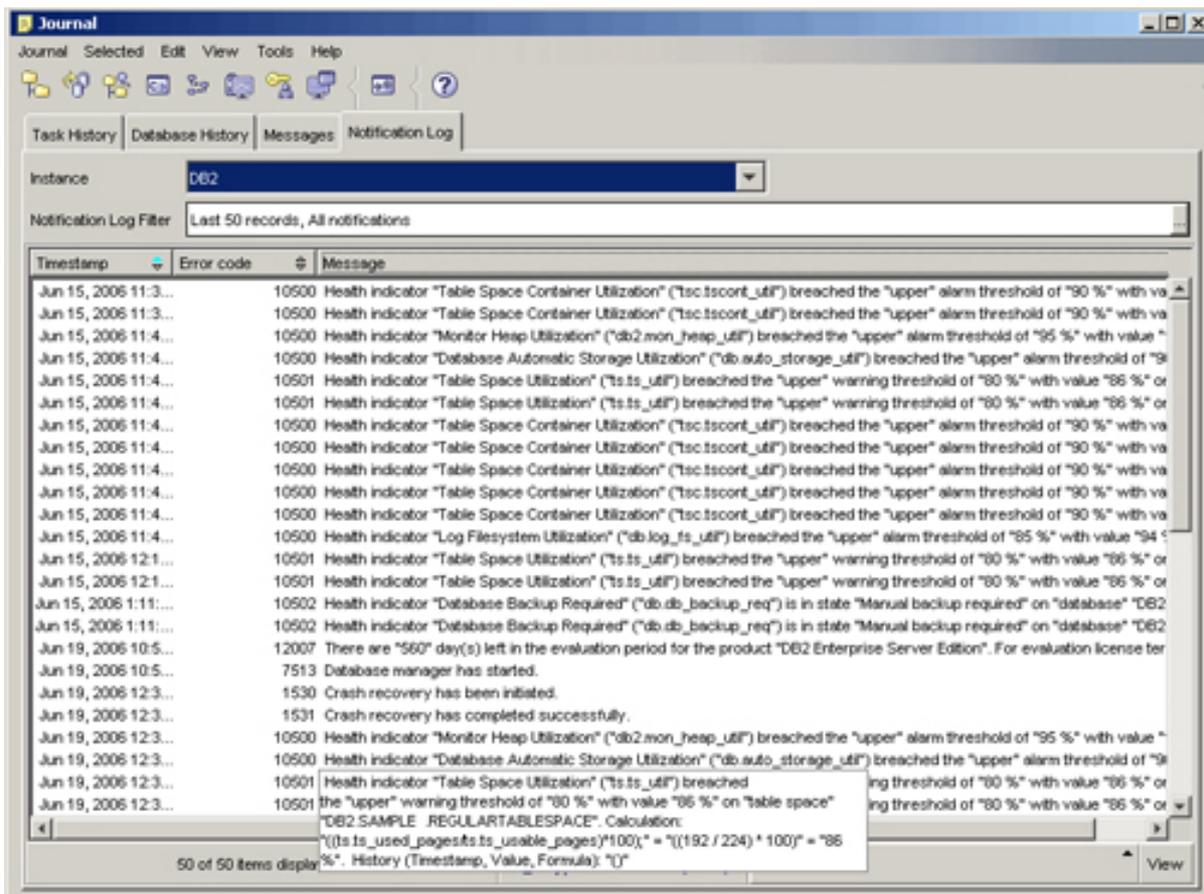
- Review the history of alerts for an instance or database.

Some of the actions you can perform with the DB2 HC are shown in the previous figure. When you open the Health Center you can see there are several other issues that DB2 has identified as unhealthy. A DBA can select a problem and ask DB2 to recommend the best method to resolve the problem. The Recommendation Advisor will ask a number of questions and, based on your responses, suggest the best course of action to resolve the issue-at-hand.

The DB2 Journal

The DB2 Journal displays historical information about tasks, database actions and operations, Control Center actions, messages, surfaced health alerts, and more. You can start the DB2 Journal from the Start menu in a Windows environment, or from the launchpad of any DB2 tool.

The figure below shows the DB2 Journal, with some information from past events displayed.



There are four tabs in this tool, each providing a DBA with valuable information:

Task History

Shows the results of tasks that were previously executed. You can use this information to estimate how long future tasks will run. This page contains one row for each execution of a task.

For each completed execution of a task, you can perform the following actions:

- View the execution results
- View the task that was executed
- Edit the task that was executed
- View the task execution statistics
- Remove the task execution object from the Journal

Database History

Shows information from the recovery history file. This file is updated when various operations are performed, including: backup, restore, roll forward, load, and reorganization, and more. This information could be useful if you need to restore a database or table space.

Messages

Shows messages that were previously issued from the Control Center and other GUI tools.

Notification Log

Shows information from the administration notification log.

The DB2 License Center

The DB2 License Center displays the status of your DB2 license and usage information for the DB2 products installed on your system. It also enables you to configure your system for proper license monitoring. You can use the DB2 License Center to add new licenses, set authorized user policies, upgrade a try-and-buy license to a production license, and more. You can also control DB2 licenses through the command line using the `db2licm` command.

The DB2 Information Center

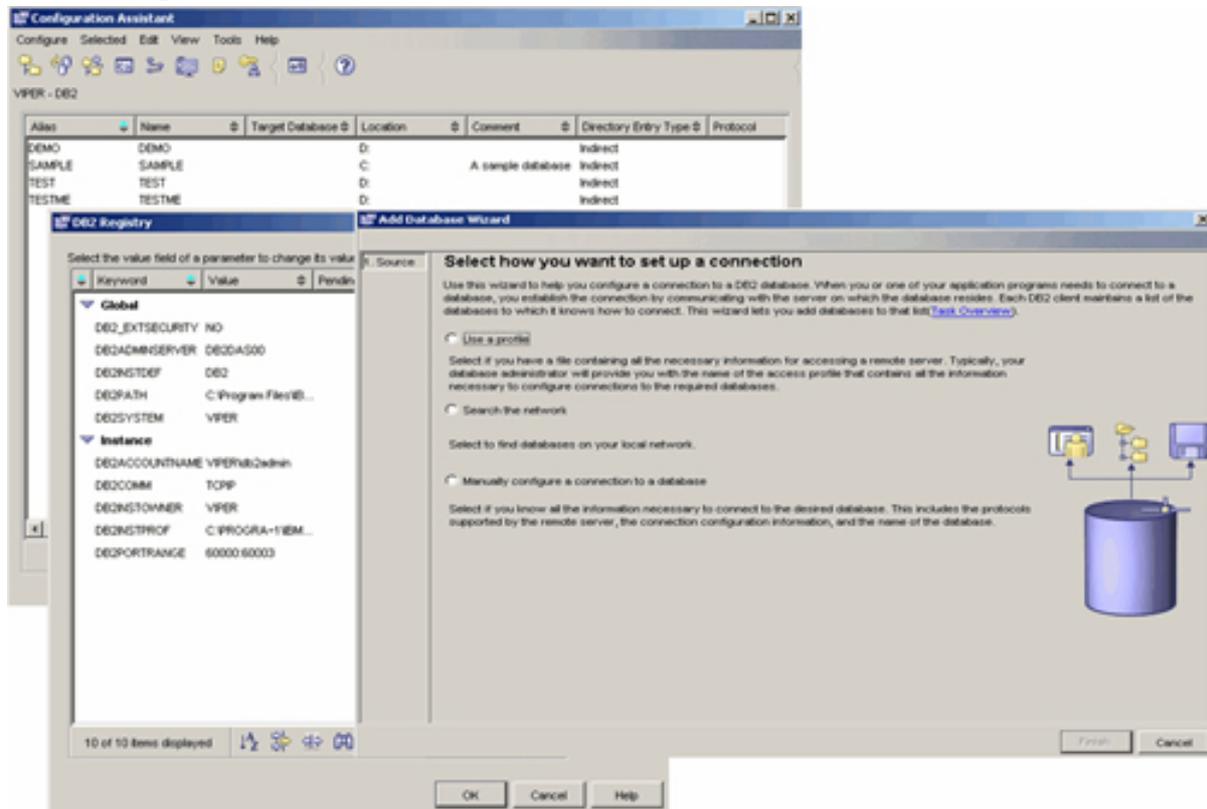
Use the Information Center to find information about tasks, reference material, troubleshooting, sample programs, and related Web sites. This center will prove to be your one-stop shop for DB2 information. You can start the Information Center from the Control Center, from the Start menu in a Windows environment, or by entering the `db2ic` command.

Section 5. The Configuration Assistant

Overview

The Configuration Assistant (CA) lets you maintain a list of databases to which your applications can connect, manage, and administer. It is mostly used for client configurations. You can start the CA by entering the `db2ca` command on a command prompt, or, in a Windows environment, from the Start menu. Some of the functions available in the CA are shown in the following figure.

The Configuration Assistant



Each database that you wish to access from the CA must be cataloged at a DB2 client before you can work with it. You can use the CA to configure and maintain database connections that you or your applications will be using. The Add Database wizard (shown above) will help you automate the process of cataloging nodes and databases, while shielding you from the inherent complexities of these tasks.

From the CA, you can work with existing databases, add new ones, bind applications, set client configuration and registry parameters (also shown above), and import and export configuration profiles. The CA's graphical interface makes these complex tasks easier through:

- Wizards that help you perform certain tasks.

- Dynamic fields that are activated based on your input choices.
- Hints that help you make configuration decisions.
- The Discovery feature, which can retrieve information that is known about databases that reside on your network.

As you can see in the figure above, the CA displays a list of the databases to which your applications can connect from the workstation where it was started. Each database is identified first by its database alias, then by its name. You can use the Change Database Wizard to alter the information associated with databases in this list. The CA also has an Advanced view, which uses a notebook to organize connection information by the following objects:

- Systems
- Instance nodes
- Databases
- Database Connection Services (DCS) for System i and System z databases
- Data sources

You can use the CA to catalog databases and data sources (like CLI and ODBC parameters), configure your instances, and import and export client profiles.

Take the time now to add a database that you've created to your connection list. If you have the resources, try to add a database that resides on a separate server using the discovery feature. Now go through each of the other features in the CA to understand the framework in which the CA presents them. Especially focus on profiles; any DBA worth their salt will want to understand profiles because they allow you to cookie-cut a DB2 client installation for mass deployment (including the database connections and the configuration settings).

Section 6. Other DB2 tools

There are a multitude of other DB2 tools you can use to make your work easier. These tools are delivered free of charge as standalone tools in the Control Center or the Configuration Assistant. Don't confuse these tools with the add-on DB2 tools that are available separately, which you learned about in [DB2 Add-on tools](#).

Visual Explain

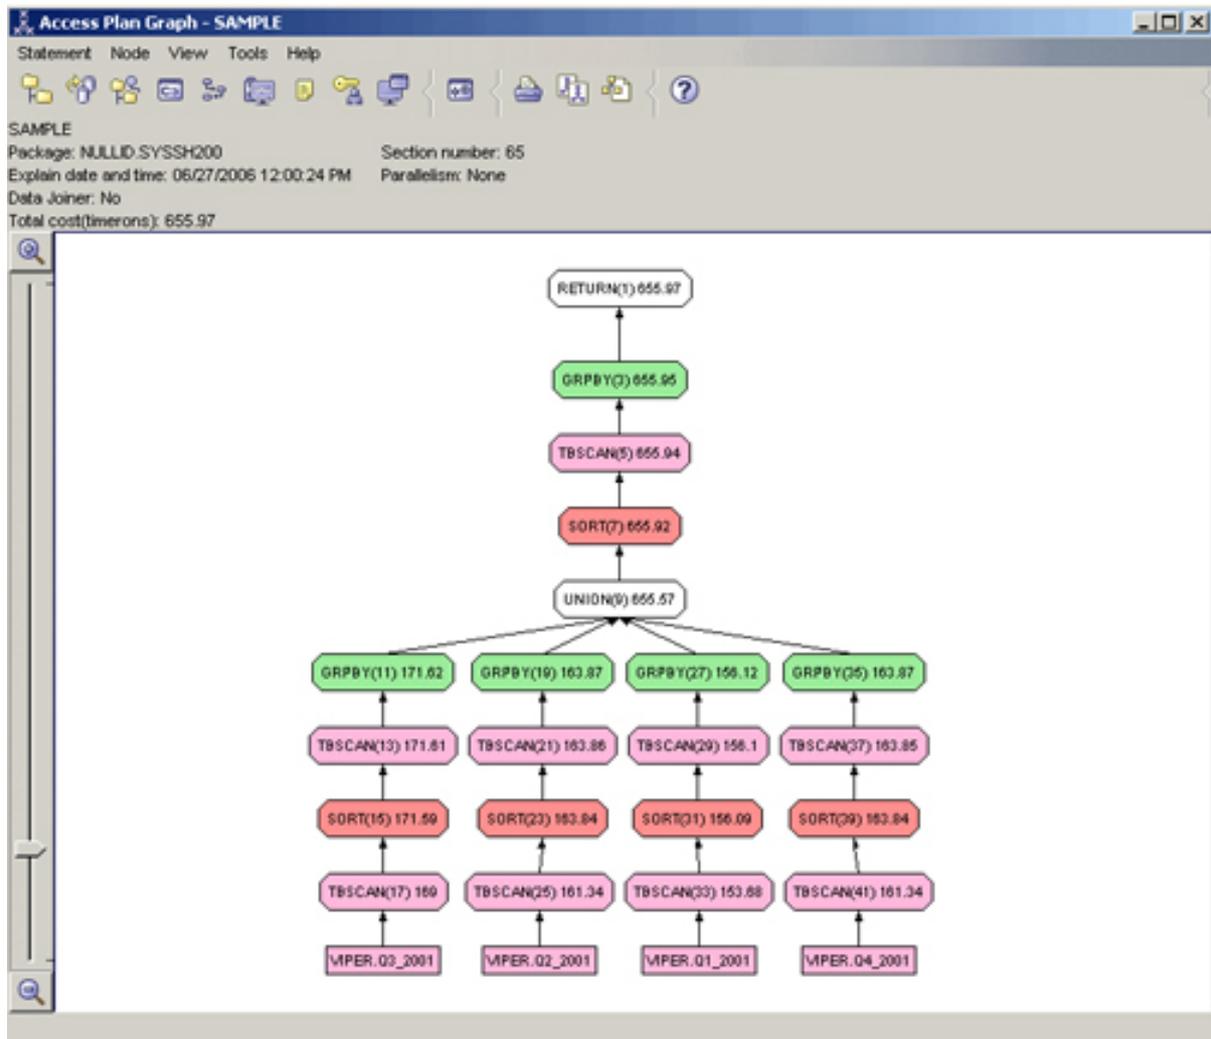
Visual Explain lets you view the access plan for an explained SQL statement as a graph. You can use the information available from the graph to tune your SQL query for better performance. Visual Explain also lets you dynamically explain an SQL statement and view the resulting access plan graph. Visual Explain is available as a stand alone tool from the Control Center, or through interfaces associated with the Command Editor and the Developer Workbench.

The DB2 optimizer chooses an access plan and Visual Explain displays this plan, in which tables and indexes -- and operations on them -- are represented as nodes, and the flow of data is represented by the links between nodes. To get more information on any step of the query plan, double-click on that object in the explain output.

The best part of Visual Explain is that you don't even have to run the query to get the information you are looking for. For example, let's say you suspect a query you have is inefficiently written; using Visual Explain, you can graphically look at the cost of the query without actually running it.

You can get the graphical access plan for a query without running it by entering it in the Control Center. From the Control Center tree view, select the database you want to work with, right-click, and select **Explain SQL**. Enter the SQL statement that you want to explain and select **OK**. An example of a visually explained query is shown below.

Visual representation of a query's access plan using Visual Explain



The Snapshot and Event Monitors

There are two utility monitors provided in DB2 to help you better understand your system and the impact of operations upon it.

The *Snapshot Monitor* captures database information at specific points in time. You determine the time interval between these points and the data that will be captured. The Snapshot Monitor can help analyze performance problems, tune SQL statements, and identify exception conditions based on limits or thresholds. In DB2, you have the ability to retrieve snapshot information into a DB2 table using an SQL UDF or programmatically using a C API.

The *Event Monitor* lets you analyze resource usage by recording the state of the database at the time that specific events occur. For example, you can use the Event Monitor when you need to know how long a transaction has taken to complete, or what percentage of available CPU resources an SQL statement has used.

Tool Settings

The Tools Settings notebook allows you to customize the DB2 graphical tools and some of their options. You can use this notebook to:

- Set general property settings, such as the termination character, the automatic startup of the DB2 Tools, hover help and infopop characteristics, the maximum number of rows in a result set, and more
- Change the fonts for menus and text
- Set DB2 for z/OS Control Center properties
- Configure Health Center notification preference
- Set documentation preferences
- Set up a default DB2 scheduling database used for maintaining scheduled tasks

DB2 Governor

The DB2 Governor can monitor the behavior of applications that run against a database and can change certain behaviors, depending on the rules that you specify in the Governor's configuration file.

A Governor instance consists of a configuration file and one or more daemons. Each instance of the Governor that you start is specific to an instance of the database manager. By default, when you start the Governor, a Governor daemon starts on the database. Each Governor daemon collects information about the applications that run against the database. It then checks this information against the rules that you specify in the Governor configuration file for this database.

The Governor manages application transactions as specified by the rules in the configuration file. For example, applying a rule might indicate if an application is using too much of a particular resource. Perhaps after a query has run for such amount of time, or it's consuming a certain amount of CPU cycles, the rule would specify the action to take (such as changing the priority of the application or forcing it to disconnect from the database).

The DB2 Governor is free with DB2 Enterprise, and can be added to a DB2 Express or DB2 Workgroup installation by purchasing the DB2 Workload Management Feature.

DB2 Developer Workbench

DB2 9 has a brand new integrated development environment (IDE) called the DB2 Developer Workbench (DB2 DWB), specifically designed for power DBAs and developers who are required to build business logic for their databases. The DB2 9 DWB replaces the DB2 8 Development Center (DB2 DC), which was a successor to the DB2 7 Stored Procedure Builder (DB2 SPB). The DB2 DWB is a completely

different tool that's built on the Eclipse foundation, so the interface is identical to other DB2 products that use this architecture, such as DB2 SQL Warehousing Tool, Rational Application Developer, Rational Data Architect, and more.

The DB2 DWB is a rapid iterative development environment that can be used for building stored procedures (SPs), user-defined functions (UDFs), and structured data types. It is also useful for creating XQuery statements, registering XML schemas, annotating XML Schema documents for shredding, SQLJ development, and much more.

The DB2 DWB has been completely rewritten on the Eclipse platform, whereas the DB2 DC used to be based on Java-Swing classes. The new architecture allows for enhanced scalability and higher productivity of this tool. Unlike the DB2 DC, the DB2 DWB is a separately installable product (and it's free). You have to download it and install it separately with a DB2 Client. The DB2 DWB is kept as a separate installation in DB2 9 so that it can be updated more frequently than the data server itself.

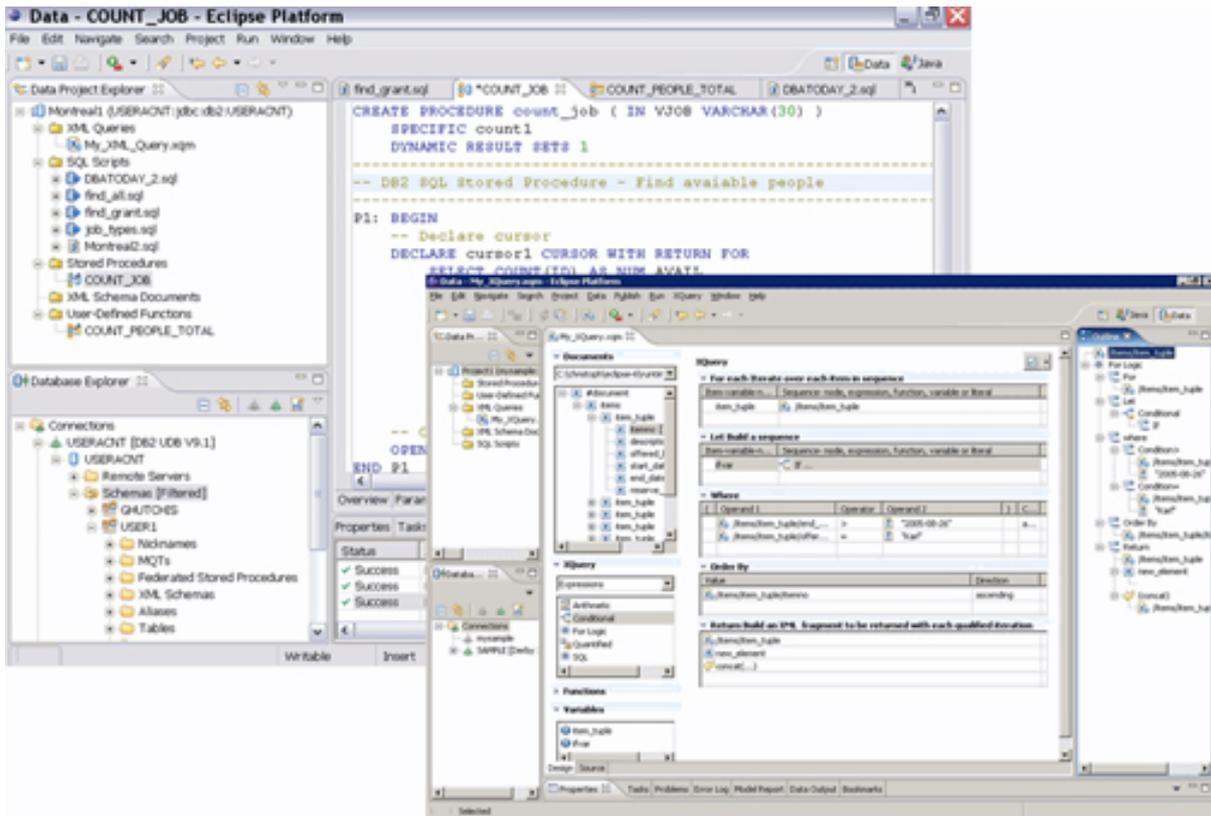
In addition to the features offered in the DB2 8 DC (mainly the building of stored procedures using Java and SQL/PL, and user defined functions and types), the DB2 DWB includes support for:

- A Data Development Project where you can logically group the development of routines, SQL statements, scripts, and XQuery statements.
- Tools to manage the life cycle of the objects in a Data Development Project. For example, routines can be copied or dragged from the Database Explorer view to a Data Development Project in the Data Project Explorer view. (You can also copy and paste or drag-and-drop routines between projects.)
- A content assistant that provides context sensitive SQL statement completion within the various editors for UDF, stored procedure, and SQL statement development. For example, when you enter a schema name followed by a period (.), the DB2 DWB will automatically enumerate the objects within the schema for easy selection.
- The ability to import routines into a Data Development Project. Using this feature, you can create a stored procedure or user-defined function by importing source files containing the routine's code or by selecting one or more existing routines from a project within the workspace.
- The deployment of routines to different data servers that aren't on the same platform. You can use the Deploy Wizard to deploy your routines to a DB2 data server from either the Database Explorer or Data Project Explorer views. For example, you can create a routine on a DB2 for

Windows server and then deploy it to a DB2 for z/OS platform. You can even move routines from an Apache Derby/IBM Cloudscape data server to a DB2 for Linux, UNIX, Windows, or z/OS data server. You can also specify options for the deployment, such as how to handle duplicate routine names, error handling, and so on.

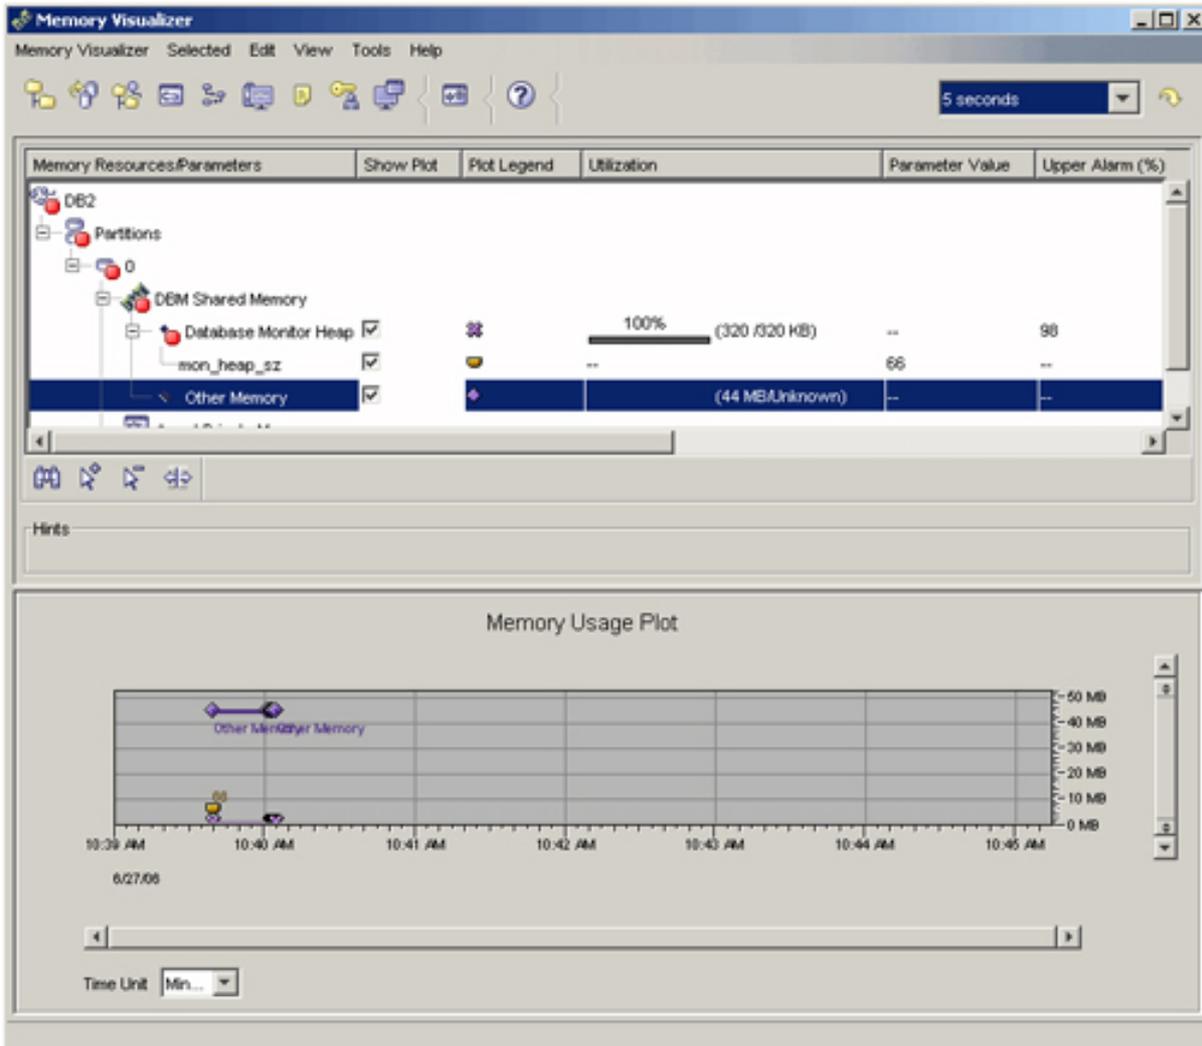
- A comparison utility that can compare the source code and attributes between two routines, and optionally synchronize the code contained within a Data Development Project.
- Performance Tuning. Users can launch the Visual Explain utility for DB2 z/OS and DB2 for Linux, UNIX, and Windows within a routine's body or a stand-alone SQL statement. (Visual Explain is launched as a separate application program.)
- For the development of SQLJ applications with features to generate an SQLJ template file with a wizard, automatic translation and compilation, customization with a wizard, printing of the profile file, editing of your SQLJ applications with code assist and templates, and the ability to debug SQLJ files.
- A complete set of SQL development tools that include an SQL Query Wizard, SQL Query Builder, SQL Script Editor, SQL Query Model, and an SQL DML parser. Using these tools, you can create SQL statements using the SQL Query Builder, which gives users graphical assistance to build their queries.
- Migration of DB2 8 Development Center projects.
- Teaming Support that lets you share the objects you build in the DB2 DWB with teaming software such as CVS or IBM Rational ClearCase. Once you've shared a project, you can manage all the changes, update history, and synchronization with the repository using menu actions from your project.
- The editing and browsing of XML and relational data.
- The creation, registration, and editing of XML Schema documents within the DB2 XML Schema Repository (.XSR).
- The building of XQuery statements with the XQuery Builder.

An example of the DB2 DWB is shown below.



Memory Visualizer

The Memory Visualizer helps DBAs understand their systems' memory allocations on a per-instance basis. You can start this tool by right-clicking the instance you want to work with in the Control Center and selecting **View Memory Usage**. You can also select the tool directly from the Start menu in a Windows environment. The following figure shows the Memory Visualizer monitoring an instance called DB2.



The top pane of the Memory Visualizer shows memory components organized in a navigation tree. Historical values, alarm, and warning thresholds are shown to the right of each component. The lower pane shows a memory usage plot.

The high-level memory components that this tool monitors include: database manager shared memory, database global memory, application global memory, agent/application shared memory, and agent private memory. Each high-level component is divided into lower-level components that determine how the memory is allocated and deallocated. For example, memory is allocated and deallocated when the database manager starts, when a database is activated, and when an application connects to a database.

Some of the key tasks you can perform with the Memory Visualizer include:

- Viewing overall memory usage.
- Specifying which memory information to display and which information to hide for a DB2 instance and its databases.
- Updating the configuration parameters for an individual memory component to prevent it from using too much or too little memory.

- Saving the memory allocation data.
- Loading memory allocation data from a file into a Memory Visualizer window.

Storage Management

DB2 has a Storage Management tool that helps DBAs understand their storage requirements and potential considerations. DB2 9 includes automated storage management where DB2 can automatically grow containers that require more space, and even manage the tables spaces (this feature is available with DMS FILE-based table spaces only). If you're using the automated storage management capabilities in DB2 9, the Storage Management tool may not be as useful if you choose a manual-based storage mechanism for your data server.

You can use the Storage Management tool to schedule or immediately capture a snapshot of currently allocated storage resources. Start the Storage Management tool by selecting a database in the Control Center, right-clicking, and selecting **View Storage**.

The Storage Management tool also allows you to specify warning and alarm thresholds (there are defaults provided out of the box) for the following storage criteria:

Space usage

Measures the amount of disk space used in a database, database partition group, or even at the tablespace level.

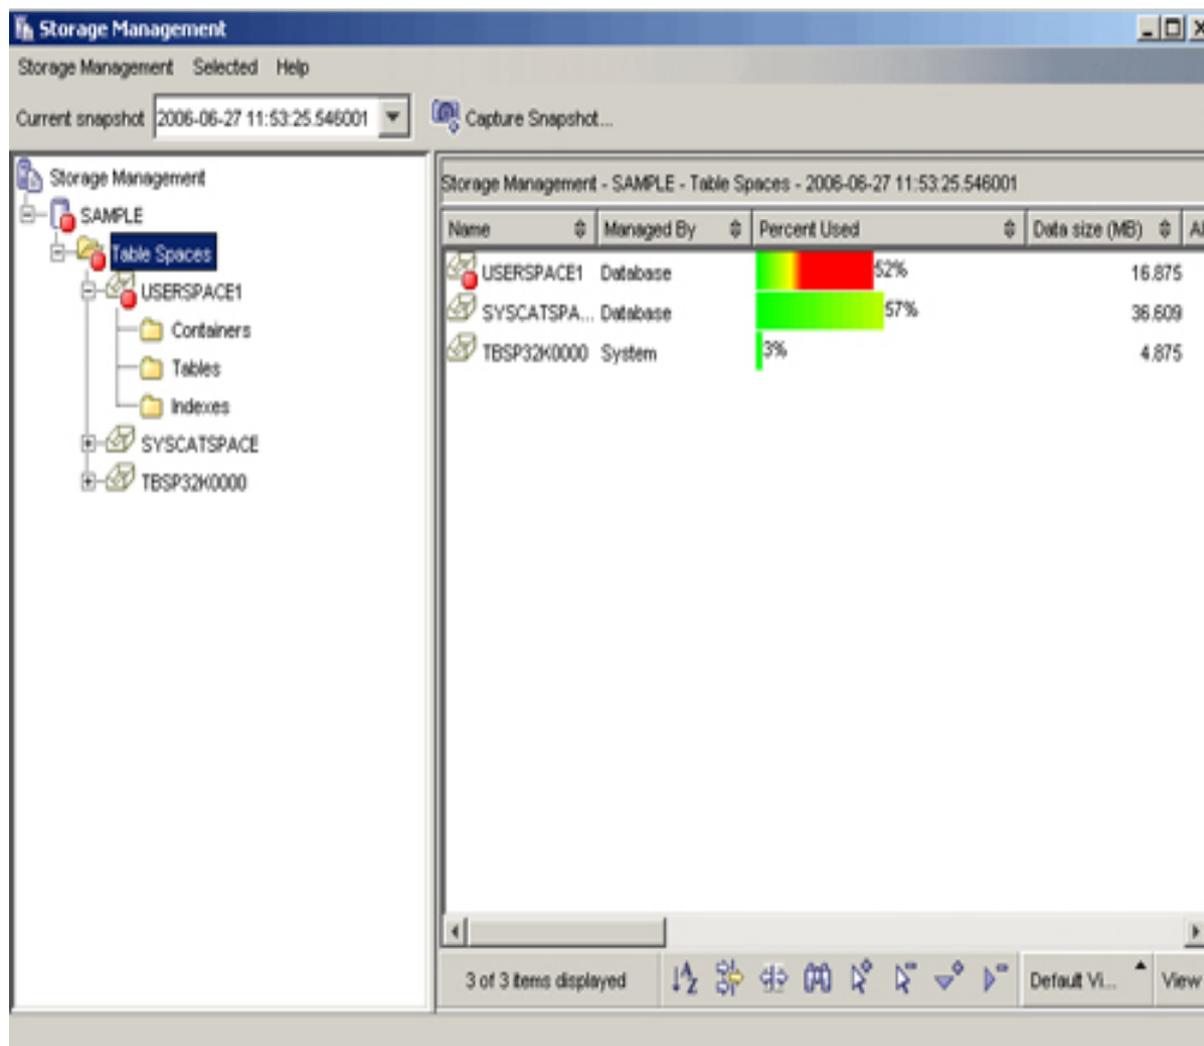
Data skew

Measures the balance of data between database partitions or tables.

Cluster ratio

Measures the quality of coverage an index has on one table.

An example of the Storage Management tool in DB2 9 is shown below.



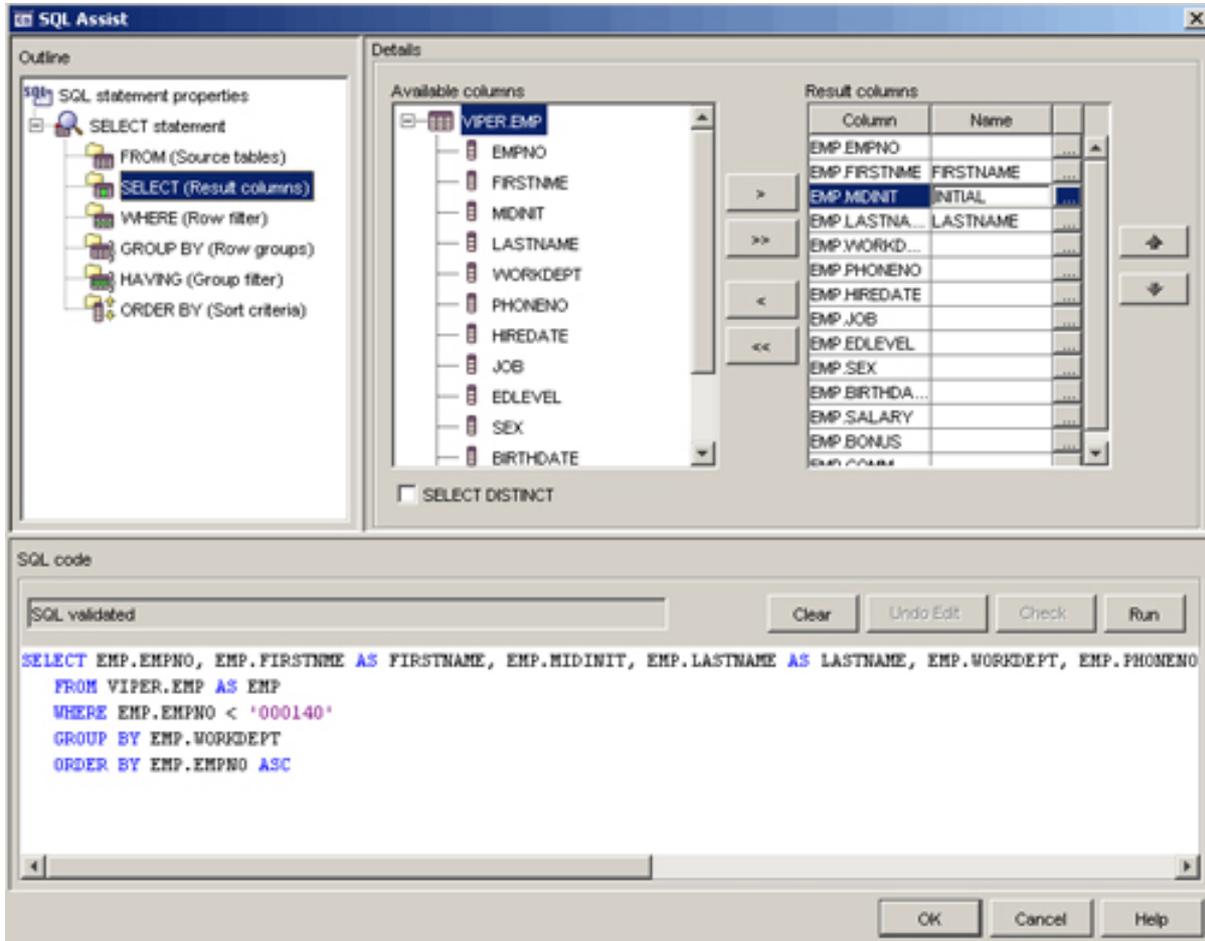
In-doubt Transaction Monitor

The In-doubt Transaction Monitor helps DBAs work with global transactions that are in an in-doubt state. For example, a broken communication line could leave a transaction prepared but not yet committed or rolled back. In most cases, you won't have to use this tool and can simply wait for the Transaction Manager to resynchronize; but if you cannot wait, you can use the In-doubt Transaction Manager. Only experienced DBAs should try this -- this tool is for people who *really* know what they are doing. You can start this tool by either selecting it from the Start menu, or entering the `db2indbt` command from a command prompt.

SQL Assist

SQL Assist is a graphical tool you can use to build SELECT, INSERT, UPDATE, and DELETE statements. The tool uses a notebook to help you organize the information you need to create an SQL statement. SQL Assist is available from multiple tools and wizards.

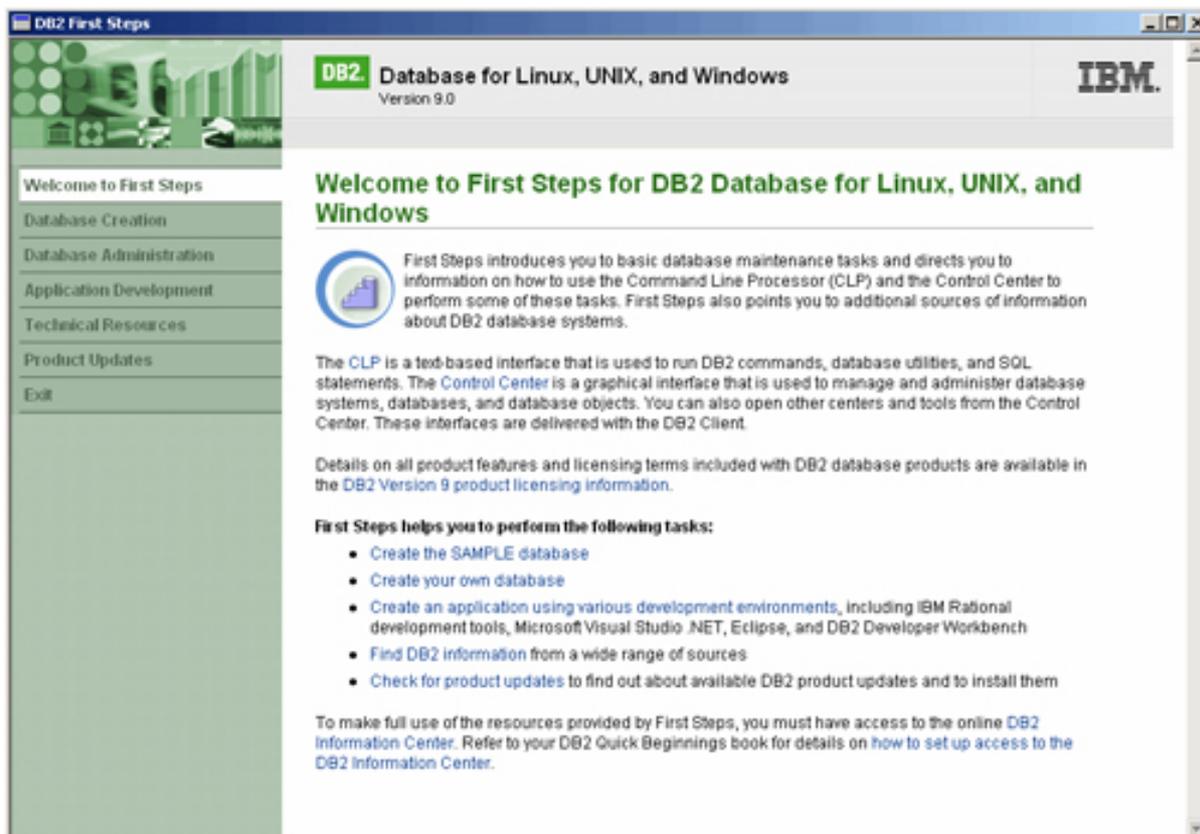
SQL Assist makes it really easy for novice users to point and click their way through query creation and build more complex SQL statements. The next time you query your database with the Command Editor, or build a view using the Control Center, invoke the SQL Assist tool to get a feel for how much it can help you. The following figure shows the SQL Assist tool.



First Steps

First Steps is a graphical tool that helps get you started with DB2. First Steps has several options; all are available by clicking on the icon next to the desired action.

With First Steps, you can create sample databases, launch the DB2 tools for management and connectivity, create new databases, check for product updates, view the product library, access application development resources, find more DB2 resources on the Web, and more. If you haven't already done so, make sure you create the SAMPLE database using this tool to help reinforce what you will learn in this and the follow-on tutorials. The First Steps tool in DB2 9 is shown below.

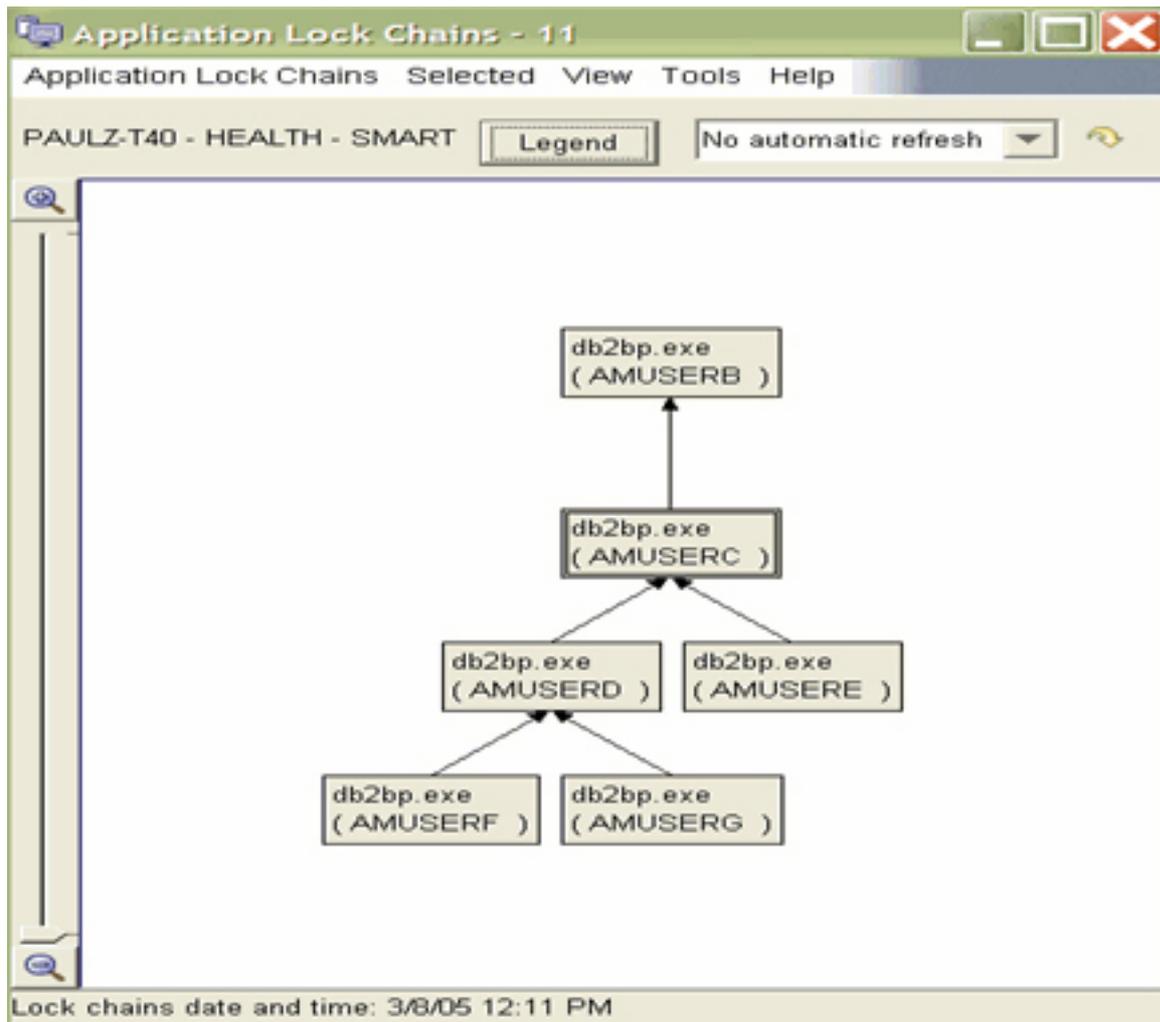


The Activity Monitor

The Activity Monitor can help you monitor application performance, application concurrency, resource consumption, and SQL statement usage. It can assist you in diagnosing database performance problems such as lock-waiting situations, and the tuning of queries for optimal usage of the database resources. The Activity Monitor also comes with a number of reports that DB2 automatically generates for you.

In the following figure you can see the different application lock chains associated with their respective authorization IDs. AMUSERB's application is causing blocking on AMUSERC's application, which in turn is blocking AMUSERD and AMUSERE. AMUSERD, in turn, is blocking two other applications.

The Activity Monitor



If you right-click any of the applications connected to the database, you can find out more details about the lock, show the SQL statement that is causing the lock, or force the lock.

Section 7. Data warehousing

What is data warehousing?

There are two main types of database applications: online transactional processing (OLTP), and data warehousing, which includes reporting, online analytical processing (OLAP), and data mining applications. DB2 9 is exceptional at both. What differentiates an OLTP system from a Business Intelligence (BI) data warehousing system? The typical queries that are accessing the data.

An OLTP system is typical of a Web order system, where you perform transactions

over the Web (such as ordering a product). These applications are characterized by granular single-row look-ups with logic that likely updates a small number of records. In contrast, BI-type queries perform large table scans as they try to find data patterns in vast amounts of data. If you've ever been asked to summarize all of the sales in the West region, that's a warehousing query.

Quite simply, think OLTP, think short-and-sweet. Think BI, think looking for needles in a haystack or aggregating a lot of data for reporting. Of course there's more to it, but you get the point.

Systems that contain *operational data* -- the data that runs the daily transactions of a business -- are OLTP systems. However, these systems contain information that business analysts can use to better understand how the business is operating. For example, they can see what products were sold in which regions at which time of year. This helps identify anomalies or to project future sales.

However, there are several problems if analysts access operational (OLTP) data directly for reporting and other BI activities:

- They might not have the expertise to query the operational database. In general, the programmers who have the expertise to query the operational database have a full-time job maintaining the database and its applications.
- Performance is critical for many operational databases, such as a database used to process banking transactions. These systems can't handle users making ad-hoc queries on operational data stores. Consider, for example, the time you take to pay bills online. When you select **OK**, it usually takes only a few seconds to process a payment. Now, consider a bank analyst trying to figure out how to make more money from an existing customer base. The analyst runs a query that is so complex that your banking transaction now takes about 30 seconds to complete! Obviously that performance time is not acceptable (and neither are the new charges that the analyst is dreaming up). For this reason, operational data stores and reporting data stores (including OLAP databases) are generally separated.

Over the last few years, though, reporting data stores have tended to become pseudo-operational and current. Such stores are called *operation data stores* (ODSs) or even *active data* warehouses. Consider the telecommunications industry, for example. ODSs are popular with these companies as they try to identify fraudulent charges as soon as possible. DB2 is one of the few databases that is well suited for both operational and reporting workloads.

- Operational data is not generally in the best format for use by business analysts. Sales data that is summarized by product, region, and season is much more useful to analysts than raw transaction data.

Data warehousing solves these problems. In data warehousing, you create stores of *informational data* -- data that is extracted from operational data and then transformed and cleansed for end-user decision making. For example, a data warehousing tool might copy all the sales data from an operational database, perform calculations to summarize the data, and write the summarized data to a database that is separate from the operational data. End users can query the separate database (the warehouse) without affecting the OLTP databases.

DB2 products for data warehousing

DB2 is built for business intelligence. Not only is DB2 the world's most scalable database, it also has a robust set of business intelligence functions. DB2 has two Data Warehouse Editions (DB2 DWE) that provide steep discounts on a complete set of data warehousing tools. Some of these tools are available for purchase outside of the DB2 DWE packages, and some are not. DB2 DWE comes in the Base and Enterprise Editions.

The DB2 DWE builds on and extends the powerful DB2 Enterprise 9 product to provide improved performance and usability of data warehousing, and analytics features, for real-time insight and decision making to achieve information on demand. DB2 DWE features include tools for warehouse management, analytic application development, OLAP, data mining, and very large database (VLDB) query and resource management.

The products included in DB2 DW Enterprise Edition are:

DB2 Enterprise

DB2 Data Partitioning Feature

This feature lets you partition a database within a single server or across a cluster of servers. DB2 Data Partitioning provides businesses with the scalability to support very large databases, common in a data warehousing environment, as well as complex workloads and increased parallelism for administration tasks. This feature is available for sale outside of the DB2 DWE product.

DB2 Data Warehouse Edition Design Studio

Some of the Eclipse-based Rational Data Architect (RDA) modeling functions are included and extended to form the DWE Design Studio, a development environment for BI solutions. The Design Studio integrates the following tasks in a unified graphical environment: physical data modeling (RDA), DB2 SQL-based warehouse construction, OLAP cube modeling, and data mining modeling.

This design tool lets designers connect to source and target databases, reverse-engineer physical data models, build DB2 SQL-based data flows and mining flows, set up OLAP cubes, and prepare applications for deployment to

runtime systems. Because it's based on the Eclipse framework, the tool looks identical to the [DB2 Developer Workbench](#).

SQL Warehousing Tool

To support end-to-end business intelligence solutions, DB2 DWE provides an infrastructure for warehouse building and maintenance, including tools for application design, deployment, execution, and administration.

The SQL Warehousing Tool solves data integration problems in a DB2 data warehouse environment. Users can model logical flows of higher-level operations, which generate units of code that are organized inside execution plans. The tool provides a metadata system and an IDE to create, edit, and manage these flows, and a code generation system that understands the source graph and translates it into optimized SQL code for execution. When the development of the flows is complete, the second part of the system comes into play: the packaging of the generated code and associated artifacts into a data warehouse application that can be deployed into various target runtime systems.

Data warehouse application developers and architects use the Design Studio to create:

- Physical data models for source and target databases and staging tables.
- Data flows that represent the movement of data from sources through transformation steps and into target databases. For example:
 - SQL data flows that use the SQL-processing power of DB2 for warehouse building operations that work on data from relational tables, flat files, and so on
 - Mining flows that integrate key data mining operations into an SQL-based model
 - SQL data flows that incorporate optional DataStage ETL jobs as subflows
- Control flows that order a set of related data flows and define processing rules for the execution of those data flows.
- Deployment packages that contain all of the files and parameters that form a complete application that is ready for installation and execution.

Control flow processes can be executed and managed, with DB2 Data Warehouse Edition serving as the runtime engine and IBM WebSphere® Application Server software providing control and scheduling capabilities. WebSphere Application Server provides the run-time deployment environment for data warehouse applications. Administrators access WebSphere by running the browser-based Data Warehouse Edition Administration Console from a Web client computer.

DB2 Data Warehouse Edition Administration Console

The DB2 DWE Administration Console is a Web-based application for managing and monitoring BI applications. Installed with the WebSphere Application Server, the DWE Administration Console uses Web clients to access and deploy data warehouse applications modeled and designed in DWE. The console makes use of the WebSphere security model, allowing users with administrator, manager, and operator roles to perform a variety of administrative tasks from a single, unified set of console pages. The console supports five categories of functions:

- **DWE Common:** Create data sources and enable databases for OLAP and mining.
- **SQL Warehousing:** Deploy, schedule, and monitor data warehousing applications that were created in the Design Studio. View statistics and logs associated with processes and troubleshoot runtime failures.
- **OLAP:** Import and export cube models, use the OLAP optimizer, and display metadata content for the cube model, such as its tables, joins, measures, and attributes.
- **Mining:** View, export, update, and delete models in the mining database. You can also import mining models into the database and load models into the cache. The mining visualization tool provides graphic representations of the results of the mining model.
- **Alphablox:** Launch the native Alphablox Administration tool.

DB2 Data Warehouse Edition OLAP Acceleration

Formerly known as DB2 Cube Views. With the OLAP capabilities included in the DWE Design Studio and Administration Console, users can create, manipulate, optimize, deploy, import, or export cube models, cubes, and other metadata objects to be used in OLAP analysis. The DWE Design Studio provides easy-to-use wizards and windows to help users work with OLAP metadata.

DB2 Data Warehouse Edition Data Mining

By using the DB2 DWE data mining features, you can discover hidden relationships in your data without exporting data to a special data mining computer or resorting to small samples of data. DB2 DWE delivers DB2 Extensions for the following modeling operations: market basket analysis, market segmentation, profiling, and more.

DB2 Data Warehouse Edition Data Visualization Feature

To visualize the mining models you create.

DB2 Alphablox analytics

DB2 Alphablox provides the ability to rapidly create custom Web-based

applications that fit into the corporate infrastructure and reach a wide range of users, inside and outside the corporate firewall. Applications built with DB2 Alphablox run in standard Web browsers, allowing real-time, highly customizable multidimensional analysis from a client computer. DB2 Alphablox is tightly integrated with DWE Cube Views, which provides common metadata and database optimization for Alphablox multidimensional analysis.

DB2 Data Warehouse Edition Query Workload Management Feature

DB2 Query Patroller is the DWE query workload management feature. This product is also available outside of the DB2 DWE package. DB2 Query Patroller is a powerful query management system that you can use to proactively and dynamically control the flow of queries against your DB2 database in the following key ways:

- Define separate query classes for queries of different sizes to better share system resources among queries, and to prevent smaller queries from getting stuck behind larger ones.
- Give queries submitted by certain users high priority so that these queries run sooner.
- Automatically put large queries on hold so they can be canceled or scheduled to run during off-peak hours.
- Track and cancel runaway queries.

The features of DB2 Query Patroller allow you to regulate your database's query workload so that small queries and high-priority queries can run promptly and your system resources are used efficiently.

Information about completed queries can also be collected and analyzed to determine trends across queries, heavy users, and frequently used tables and indexes. Administrators can use DB2 Query Patroller to:

- Set resource usage policies at the system and user levels.
- Actively monitor and manage system usage by canceling or rescheduling queries that could impact database performance.
- Generate reports that assist in identifying trends in database usage, such as which objects are being accessed, and which individuals or groups of users are the biggest contributors to the workload.
- Query submitters can use DB2 Query Patroller to: monitor the queries they have submitted; store query results for future retrieval and reuse, effectively eliminating the need for repetitive query submission; and set a variety of preferences to customize their query submissions, such as whether to receive e-mail notification when a query completes.

Section 8. Summary

Summary

In this tutorial you've learned about the different versions of DB2, and the various DB2 products. We discussed the tools that are included with DB2. You also learned about the Control Center, the Configuration Assistant, and all of the standalone tools in both the Control Center and Configuration Assistant. We discussed data warehousing, and the associated DB2 products.

To keep an eye on this series, bookmark the series page, [DB2 9 Fundamentals exam 730 prep tutorial series](#).

Resources

Learn

- Check out the other parts of the [DB2 9 Fundamentals exam 730 prep tutorial series](#).
- [Certification exam](#) site. Click the exam number to see more information about Exams 730 and 731.
- [DB2 9 overview](#). Find information about the new data server that includes patented pureXML technology.
- [DB2 XML evaluation guide: A step-by-step introduction to the XML storage and query capabilities of DB2 9](#)
- Learn more about DB2 9 planning from the [DB2 9 Information Center](#).
- Visit the [developerWorks DBA Central zone](#) to read articles and tutorials and connect to other resources to expand your database administration skills.
- Check out the [developerWorks DB2 basics](#) series, a group of articles geared toward beginning users.

Get products and technologies

- A [trial version of DB2 9](#) is available for free download.
- Download [DB2 Express-C](#), a no-charge version of DB2 Express Edition for the community that offers the same core data features as DB2 Express Edition and provides a solid base to build and deploy applications.

About the author

Paul Zikopoulos

Paul C Zikopoulos, BA, MBA, is an award-winning writer and speaker with the IBM Database Global Sales Support team. He has more than nine years of experience with DB2 UDB and has written numerous magazine articles and books about it. Paul has co-authored the books: *DB2 - The Complete Reference*, *DB2 Fundamentals Certification for Dummies*, *DB2 for Dummies*, *A DBA's Guide to Databases on Linux*, and *DB2 Version 8: The Official Guide*. Paul is a DB2 Certified Advanced Technical Expert (DRDA and Cluster/EEE) and a DB2 Certified Solutions Expert (Business Intelligence and Database Administration). You can reach him at: paulz_ibm@msn.com.