

Files on the run

DATABASES

FOR THE MODERN

WORLD

Keeping your information under the bed in a shoe box is not the best way to store away any sort of data - unless you happen to be using shoes as an *aide memoir* for something. The value of your information only make itself apparent once you have found what you are looking for. To do this effectively, you need to be very strict with the rules you set yourself on how you should file away your knowledge later retrieval. Problem is I am appalling lax at these things, maybe you are too.

Luckily, computers are very strict about what they will do, give them a rule and they are quite happy to continue following it until the electricity runs out. But it is still down to use to set out those rules and that is why we have **Database Management Systems**. The more complex your data, and the way in which you want to handle it, the more complex a **DBMS** you will need.

There are three main parts to how you would assess the value of a **DBMS**:-

- Data acquisition
- Data Storage
- Data retrieval

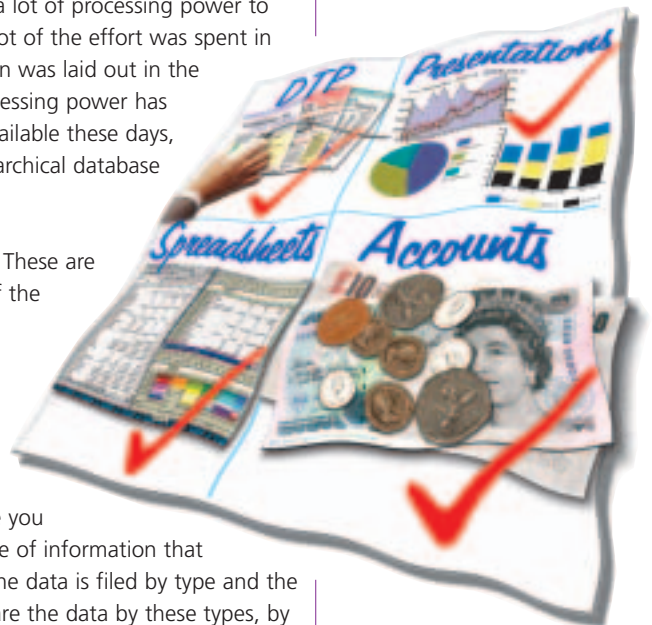
You will also need to be aware of the differing types of database available

Hierarchical databases Excluding flatplan type databases, which you might find yourself creating on a spreadsheet, hierarchical databases have the longest legacy in computer databases. It is this type of database which would first come to mind. All of the items of data in this type of database have a single link leading from a parent. Locating a piece of information in this database requires you to know something about why it is in the database, to know who its parent link is. Imagine, if you will, a database of published works. To find any given work you need to know under what criteria you will be able to find it

filed away. This could be by title, then by year of publication, then by type of publication, name of publication, name of the piece of work. If we have the answers to all of these questions we get progressively nearer to finding what we are looking for. Since the links to these types of criteria need to be set in advance, and can't be changed, they are very important to get right from the outset of building the database. Should you set a database up like this but then find that you hardly ever know when a piece was published, your database will prove worthless to you because you won't even be able to get started with the searching. This type of ridged structure doesn't take a lot of processing power to get results, because a lot of the effort was spent in the way the information was laid out in the first place. Luckily, processing power has become much more available these days, so the days of the hierarchical database must be numbered.

Relational databases These are the current mainstay of the database business. The majority of all of the databases we will mention will be relational. In relational databases, gone are the hard links that take you relentlessly to that piece of information that you seek. In this case the data is filed by type and the DBMS is able to compare the data by these types, by predefined characteristics. This then means that the DBMS can make some comparisons and create new data from it. With the hierarchical database, you would only ever get out what you put in, now you can have more. A hierarchical database might be able to tell you how many widgets I bought this week, but with the help of a relational database and buying details about John and Steve, it can tell you how

Power doesn't come from information, it comes from the ability to use that information effectively. Colin Murphy looks at what is available for the Linux user.



JDBC - Connecting Java and Databases

Java Database Connectivity is a standard SQL database access interface, providing uniform access to a wide range of relational databases. It also provides a common base on which higher level tools and interfaces can be built. This comes with an "ODBC Bridge". The Bridge is a library which implements JDBC in terms of the ODBC standard C API.

many were sold in total. If you keep information on how many were purchased at a time and when, you could work on pack quantities and stock control. On the days when none are sold, you know you can close the shop, especially since its a Sunday. Go on, fill your life with conjecture. There is still the need to set up the conditions and the relationships between the data, but this can be amended afterwards to a much greater extent.

Network databases Here we almost have a melding of the two previous designs. Best described as a data organisation, links are made between data that is 'sees' as having a valuable relation to, and this is not limited to just the one link. Think of them more like web pages, they usually have more than one link, so the World Wide Web is nothing more than a network database, it's just that it has very little predefined structure.

Some history

The first Relational Database is widely regarded as the the Multics Relational Data Store, and dates back to 1976. This database management system is widely regarded as the first relational DBMS ever offered by a major computer manufacturer. Things have come on a long way since.

Here we have a list of some of the database application that you can find working with Linux.

MySQL - <http://www.mysql.com/>

The MySQL database server is the world's most widely used open source database. It boasts speed and flexibility and can be customised to a great degree. Extensive reuse of code within the software and the concentration given to providing a core system have allowed the developers to perfect a rich set of functions and features, giving you a database management system which is compact, stable and easy to set up and administer. The unique separation of the core server from the table handler makes it possible to run MySQL under strict transaction control or with ultrafast transactionless disk access, providing you with the best options for the majority of all database applications.

MySQL AB, the company behind MySQL, offer a full range of professional services, again, showing the flexibility and power of open source projects in a commercial world.

GNOME-DB - <http://www.gnome-db.org/>

The GNOME-DB project aims to provide a free unified data access architecture to the GNOME project. It makes use of its own database front-end for administrators and **libgda** - a data abstraction layer. It can manage data stored in databases or XML files and it can be used by non-GNOME applications. Although focused on the GNOME desktop environment, care has been taken to cleanly separate the data access framework (libgda) from the user interface.

Ksql - <http://ksql.sourceforge.net/>

KSql, previously known as KMySQL, this is a KDE database client. It was originally MySQL specific, but uses now plugins to access databases like miniSQL and PostgreSQL. With the latest version you can view queries result in multiple tabular views, print them or export then in HTML, edit your queries in a comfortable edit box with history, create nice forms with a WYSIWYG editor and save usual queries and recall them with a double-click.

Mimer SQL - <http://developer.mimer.com/>

Mimer SQL 9 is a high performance, easy-to-use Relational Database Management System (RDBMS). Mimer SQL offers scalable performance, including multi-processor support, and with its availability on all major platforms is ideally suited for open environments where interoperability is important. Recent additions include the new New Mimer JDBC Driver, providing Java support for: array fetches, batch operations, CallableStatements, DataSource, distributed transactions (XA), LargeObjects (LOBs), scrollable ResultSets, Set FetchSize and Unicode. Development is also underway for Mimer SQL Embedded a small footprint DBMS especially for use on hand held devices, mobile phones and other small appliances.

MySQL Data Manager - <http://www.edatanew.com/>

Fully featured web client for MySQL allows users to manage database records and structure via user friendly graphic interface, which allows you to remotely manage MySQL databases and user access over the Internet using web browser based interface. MySQL Data Manager is an all-inclusive web based



The MySQL Data Manager front end

MySQL front end with a powerful interface for management, development and support of databases on the web.

PhpMyAdmin - <http://www.phpmyadmin.net/>

phpMyAdmin is a tool written in PHP3 intended to handle the administration of MySQL over the WWW. Currently it can: - create and drop databases - create, copy, drop and alter tables - delete, edit and add fields - execute any SQL-statement - manage keys on fields - create and read dumps of tables - export/import CSV data.


Tora - Toolkit For Oracle -

<http://www.globecom.se/tora/>

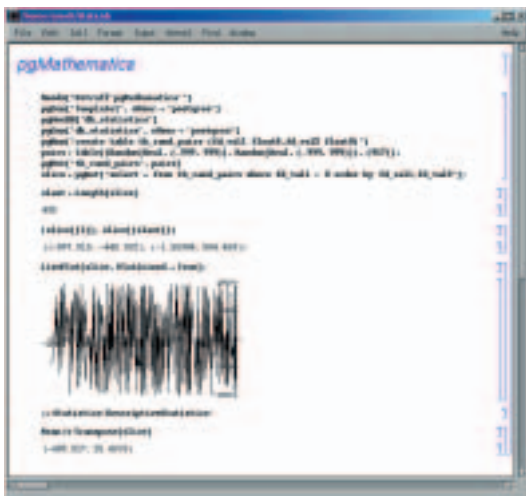
Tool for Oracle (Tora) is a tool for DBAs and database software developers. It currently features a schema browser, SQL worksheet, PL/SQL editor & debugger, storage manager, rollback segment monitor, instance manager, security manager, SQL output viewer, schema comparison and extraction and SQL templates.

LEAP RDBMS 1.2.6 - <http://leap.sourceforge.net/>

LEAP is an RDBMS (Relational Database Management System). It is used as an educational tool around the world to help students, and assist researchers and teachers as they study and teach databases.



Tora - Powerful statistical analysis of the performance of your database.



pgMathematica combines the power of a PostgreSQL database with the Mathematica engine

PHP - <http://www.php.net/>

PHP is an HTML-embedded scripting language. With special sets of tools you can develop web pages that are able to access databases directly. Much of the PHP syntax is borrowed from C, Java and Perl with a couple of unique PHP-specific features thrown in. With PHP you can create web pages which are dynamic, showing details based on client input with results from your database.

PHP Builder - <http://www.phpbuilder.com/>

This is probably the most valuable resource for anyone that is hoping to develop with PHP. It features Articles, chat rooms, code libraries, documentation and lots more.

PgMathematica 2002.0 -

<http://www.petroff.ch/pgmathematica/>

Mathematica <http://www.wolfram.com/products/mathematica/> a technical computing tool can be used with a PostgreSQL database.



Oracle9i -

<http://www.oracle.com/>

Another one of the main players in the database market. Increasing their Linux support with a concerted effort with Linux cluster solutions for the enterprise market with its 'Unbreakable Linux' product.

SuSE Linux Database server -

http://www.suse.co.uk/uk/products/suse_business/database_server/index.html

SuSE Linux Enterprise Server is the cross-architecture operating system solution for company-wide IT infrastructures. SuSE Linux Database Server couples the advantages of SuSE Linux - speed, security, cost efficiency, and high quality - with a professional and proven database system. Based on an optimised operating system basis, SuSE Linux Database Server offers security and stability for DB2 Universal Database and all relevant server services.

Red Hat database -

<http://www.redhat.com/software/database/>

Red Hat also offers a boxed set solution featuring Red Hat Linux 7.1 and PostgreSQL 7.1.2. The value comes from a supplied with an understanding for the Linux market. In addition to the software, the product includes Red Hat Installer, extra documentation and support packages, designed for seamless integration. This reduces the amount of time spent on installation, allowing businesses to get up and running as quickly as possible. With business growth, such as Web and e-business deployments, in mind, this package will allow for a quick implementation by businesses with little database expertise.



If you are a SuSE user then you have a database solution ready, out of the box.