


IBM DB2 Universal Database

Version 7.2



An extract from the Bloor Research Database 2001 report

 **Bloor**
Research



Fast facts

DB2 Universal Database (UDB) is not a single product. This is true both in terms of the database itself and in terms of the options that are available. While the latter may be commonplace, it is less common for there to be multiple versions of what is ostensibly the same product. DB2 UDB aims to provide common functionality across a wide variety of platforms ranging from mainframes, through Unix and Linux, to AS/400 systems and PCs, plus, with DB2 Everyplace, hand-held and other mobile devices. Rather than attempt to shoehorn a single product into all of these environments, IBM instead maintains different code sources that are optimised for their respective targets but offer common functionality (cut down in the case of DB2 Everyplace) and the same look and feel.

Key findings

In the opinion of Bloor Research the following represent the key facts of which prospective users should be aware:

- The parallel facilities in DB2 UDB are at least as good as, and probably better than, any competitive product on the market. This will make the product suitable both for very large data warehousing environments and OLTP/e-Commerce environments.
- The DB2 Everyplace version provides support for mobile users at the PDA and Smartphone level.
- The relationship between DB2 UDB and external data sources is excellent, and is provided by a number of products, including Data Links Manager, DB2 DataJoiner and DB2 Relational Connect.
- The object extensions in DB2 UDB have been significantly improved since their first release, to the extent that they are superior to those of other relational vendors.
- We are pleased to see that XML may be stored in parsed format, thereby eliminating the need to parse data whenever a query is run. We also like the ability to map XML across multiple columns.
- Management facilities in DB2 UDB are very good and we especially like the range of features for managing data warehouse implementations.

The bottom line

Of all the vendors that have embraced the concept of the universal database it is arguably DB2 that is the most comprehensive. In particular, IBM appears to be the most firmly committed to offering an all-embracing product set. The major exceptions to this are that it provides a separate content management product and that it does not provide a metadata repository for developers. On the other hand, its data warehousing facilities are extensive and we like the close integration introduced with WebSphere and MQSeries in this release.



It is also worth noting a major difference in approach that differentiates IBM from Oracle. IBM takes the view that DB2 should support the ability to integrate and federate information rather than centralise it. This is in direct contrast with Oracle, which is pushing centralisation. Which approach is best is a matter of opinion (although we incline towards distributed environments and believe that customers do too) and will depend on circumstances. So, the real issue becomes a philosophical one: do you want to centralise or not? In such situations it is worth considering what will happen if you make the wrong decision. In this case the answer is clear cut: if you find that you have chosen unwisely, then reversing that decision by centralising with IBM will be much easier than if you have to try federating with Oracle.



Product Overview

Introduction

IBM sees DB2 UDB as the heart of a set of products that together offer comprehensive data management capability. It is this concept of data management, as opposed to database management, that is at the centre of IBM's marketing of DB2 and that is driving the increasing emphasis on DB2 as a brand.

The trend over the last few releases, and IBM is not alone in this, has been to encompass more and more facilities within this overall context of data management. However, it is worth commenting that IBM does not pretend that it can provide a comprehensive data management solution on its own. It has therefore taken on a number of third-party products, which it is re-selling, to complement its own facilities. This is particularly true within data warehousing, where IBM licences software from vendors such as Hyperion and ETI. This is a notable contrast from its major competitors, which appear to believe that they can do everything themselves.

IBM web address: www.ibm.com

DB2 web address: www.ibm.com/software/data/db2

Product availability

DB2 products basically break down into five groups:

1. Clients - these may run under OS/2, AIX, Solaris, HP-UX, Linux, all flavours of Windows, SGI and PTX (formerly Dynix/PTX).
2. Personal and Workgroup Editions - Personal Editions of the product are OS/2, Linux or Microsoft flavoured, while Workgroup Editions run on the same platforms as the Personal Edition plus AIX, Solaris and HP-UX. The Satellite edition is to be discontinued as a separate product with its facilities (particularly hands-off administration) being merged back into DB2. There is also a small footprint product, DB2 Everyplace that is used for hand-held devices such as PDAs and Smartphones.
3. Enterprise Edition - this runs under OS/2, Windows NT/2000/XP, AIX, Solaris, HP-UX, Linux and MUMA-Q (PTX).
4. Enterprise-Extended Edition - this runs under Windows NT, AIX, HP-UX, Linux, NUMA-Q (PTX) and Solaris.
5. Hosts - these can be accessed from distributed Windows, OS/2 or Unix platforms via the DB2 Connect product and include DB2 UDB for OS/390 and z/OS, DB2 Server for VSE & VM and DB2 UDB for AS/400 and iSeries.



It should be noted that IBM uses different numbering schemes (and code bases) for DB2 on its host platforms (AS/400 and iSeries, OS/390 and zSeries and so on), as opposed to the distributed platforms on which the various editions described above run. This is because IBM's strategy is to exploit the features of its own platforms while being broadly portable on Unix, Linux and Windows platforms (yet reserving roughly 10% of the code to exploit unique features of any particular operating system). Nevertheless these share essentially the same features.

In addition, it is important to note that the differences between the different editions of DB2 are primarily scale differences. The different editions share the same functional capabilities.

It is also worth noting that, unlike many of its competitors, IBM has opted to support all the major Linux distributors (Caldera, TurboLinux, Red Hat and SuSE) rather than just a single vendor.

As one might expect there are many options that are in one version of the product but not another. In addition, even in the case of the Enterprise-Extended Edition, there are optional extras that you may or may not choose to deploy. In general we will not distinguish between features in different versions of the product, and nor will we differentiate between built-in features and those that are optional. However, we will not be discussing products that are complementary rather than optional. For example, the latest release has extended facilities for integrating DB2 with MQSeries. We will discuss the features that support integration but not MQSeries itself. The same applies to VisualAge.



Transaction Processing

The truth is that all the leading relational databases have a long history of supporting transactional applications, as well as the query processes that complement these environments. For this reason there is little purpose in discussing those features that are common across different vendor's products and we shall therefore focus on those elements of DB2 UDB that distinguish it from its competitors. In a few areas, these differences are significant enough to warrant a sub-section of their own but in most cases there are only specific features to be considered. We will begin with the latter.

- Storage engine – Over the last several releases IBM has increased capacity extensions in a number of areas, most notably to include support for multiple page sizes (4, 8, 16 and 32K options) and multiple buffer pools. The optimiser is aware of both of these features. In addition, multiple TCP/IP listeners were introduced in the 7.1 release so that more concurrent connections can be handled. In the latest version there is support for Windows 2000 Address Windowing Extensions, which allow you to allocate buffer pools of to up 64 Gb as well as 64-bit support for Windows, Linux, HP-UX and so on.

There has also been a significant emphasis on increased resilience. Filtered log recovery and a flushing log option were both introduced in the previous version but now the company has added on-demand log archiving and dual logging (that is, mirroring of active log files). Further enhancements include extensions to incremental back-up, the ability to back-up from a split image and parallel forward recovery facilities.

It is important to note is that DB2 does not include its own Java Virtual Machine (JVM). This contrasts with Oracle's stance. Which approach will provide better performance will depend upon the particular implementation.

- Partitioning. In addition to the conventional requirement of partitioning tables across storage devices in order to reduce the I/O bottleneck of large queries, partitioning has a secondary, and perhaps more important, role within the field of parallelism. This is to allow the processing of a query to be divided among several partitions in order to, ideally, divide the processing time of a query by a factor of the number of the data partitions.

Partitioning will be particularly important in federated database and similar environments where the quality of the partitioning will significantly impact on parallel performance. Unfortunately, DB2 UDB is somewhat less flexible than some of its competitors in this regard, in that it only offers hash partitioning. That said, the implementation of the hash partitioning is particularly good. The granularity of the hash partitioning, which can use data from one or more columns, is adjustable and there is also a



mechanism for defining node-groups over which a given table or set of tables may be partitioned.

- Indexes – DB2 is limited to the use of Btrees for conventional indexing (though other indexing options are available for particular objects such as spatial datatypes). There is an index wizard, which is designed to reduce the complexity of performance analysis and tuning, by making it easy for the DBA to find the ‘right’ set of indexes. There is also a capacity management feature that provides an estimate of table and index requirements when you are adding a given number of rows. Online index re-organisation with automatic space reclamation is provided. In the latest release you can specify columns with a length of more than 255 bytes (this was the previous limit) as part of an index key. There is no limit to the number of levels in an index. One useful feature, added in version 7.1, is the ability to define an index against a generated column.
- Optimiser - DB2’s cost based optimiser is particularly rich in terms of the statistical information that it uses. We particularly like its SQL re-write capability. This is the facility whereby the optimiser will recognise poorly written SQL and automatically re-write it so as to improve performance. Also noteworthy are the global optimisation features of DB2 DataJoiner (see below).
- Datatypes – apart from object-relational and XML support (both of which are discussed below) the most interesting facet of DB2 in this respect is its support for user defined datatypes, which are called distinct types by IBM, and are derived from base types as supplied with the dbms. Effectively, new datatypes inherit their attributes from the base type and the database provides casting so that the value of the new type can be converted to that of the base format and back again. For example, you might derive Francs from a money datatype. Similarly, Yen might be so derived. For obvious reasons one would not want to allow the addition of Francs to Yen and this is prevented by a facility called strong typing. Casting has been extended to Unicode databases in the latest release.

User defined functions are also supported. However, one of the problems with these is that users may not be entirely happy with the safety of third party (or even in-house) user defined functions or, for that matter, data extenders. So IBM has implemented a facility called fencing that allows you to deploy these either within the DB2 address space for maximum performance, or in their own spaces, where safety is paramount.

- Multimedia – while data extenders are discussed below, you may still want to store binary objects directly. As far as LOBs are concerned, IBM supports three types, which can be stored either internally or externally, as required. These are BLOBs, CLOBs and DBCLOBs, which are binary, character and double byte LOBs respectively. The facilities for manipulating these are extensive, with support for locator variables, file referencing and metadata held separately from the LOBs themselves. However, while recovery can be turned off for LOBs there are no special buffering or caching facilities for



them. Bufferpools are not used and there is nothing specific to replace their functionality. However, it is possible to define LOBs in an SMS or DMS file tablespace, in which case all access to the LOBs is through the file system. Since most (if not all) of today's file systems support caching, this can be used to control caching for LOBs.

Database operations

DB2 UDB supports all the standard constraints, stored procedures and trigger facilities that one would expect from a leading database product. However, built-in support for alerts is limited to system management events.

A Java stored procedure builder was introduced in version 6.1. This is a graphical tool for building, editing, installing, testing and debugging both SQL and Java stored procedures. It also includes what IBM refers to as 'SmartGuides' for any code that is required, as well as a feature that provides relevant cost estimates. This tool is common across the entire DB2 family and provides a one-step build to any DB2 platform. In addition, it is integrated with Microsoft Visual Studio, Visual C++ and Visual Basic as well as IBM's VisualAge for Java.

While on this subject it is worth commenting that versions 6.1 and 7.1 both had substantial additional facilities to support application development. These included enhancements to support Microsoft Windows environments, and new facilities for dynamic SQL. Some of the more useful features added in version 7.1 included support for nested stored procedures, the ability to declare temporary tables and the ability to define a column when its values are pre-computed automatically using an expression. Some of these features were specifically intended to help migration from other platforms and, in version 7.2 there is a free migration toolkit, which builds on these facilities, to assist companies migrating from Sybase environments. IBM has also added functionality that equates more closely to that of Oracle environments so that it encourages those users to migrate also. There is a specific promotional offer available for PeopleSoft users to migrate to DB2.

Finally, it is worth mentioning that there is support for moving compiled stored procedures in the latest release. This means that you do not have to have a relevant compiler on the target system.

Data Extenders

There are a number of data extenders provided by IBM, including text, images, video, and audio data. More recently the company has included spatial, XML and Net Search extenders, which are discussed below. A variety of other extenders, including currency, GIS, encryption and optical products, are available from third-party vendors.

Unlike most other products on the market, the text extender is not based on third party work but on the Stairs product (originally designed for MVS). This is highly versatile and provides virtually every facility one might want, based on a linguistic index that filters out insignificant words and reduces all words (in any of 18 languages) to their stem form. It is also independent of punctuation and supports searching for synonyms. One nice feature is the ability to rank



documents as to how well they fit a particular search pattern. However, the text extender does not support searching against a theme or by gist.

In the latest release the text extender has been enhanced so that you can now search documents using an SQL query and, more specifically, you can search against columns that contain either HTML or XML text.

The image extender supports various formats and can convert between them, automatically extracts thumbnails for browsing purposes, and offers full content based comparisons. The video extender also supports various formats and can automatically segment a video into shots by detecting scene changes, with extraction of a representative scene from each such shot. The audio extender once again supports multiple formats, and can detect the characteristics of audio data such as playing time, and the number and names of audio tracks. All of these data extenders have full security and recovery capability.

DB2 Data Links Manager

Data Links Manager arose out of a requirement to be able to extend the management capabilities of the database to external file systems. Legacy applications based on these often have performance advantages and API characteristics that mean that their users do not want to convert the relevant data into a tabular format. However, they would like the ability to use the database's backup and recovery, security, referential integrity and other features. Data Links Manager was developed to allow these external files to be managed from within the DB2 database.

This technology is implemented by means of a software layer on the relevant system, together with a new DataLinks file-link datatype. This is illustrated in Figure 1 where DB2 controls images that are stored in the picture column while the DB2 File-Link Filter intercepts any file-open, file-rename, or file-delete commands to check for authorisation.

The Data Links Manager has a useful corollary. One of the issues about data extenders is how and where to manage the objects in question. In particular, should the indexes to these objects reside in the database or externally? Most vendors have concluded that indexes, at least, should be in the dbms. This allows them to be under the management control of the database and to use all of its facilities for backup, recovery and so on. However, IBM has not been able to take this course because existing products (like Stairs) would have required a major re-write. It was therefore easier and faster not to enforce internal indexes. Instead, IBM now offers the choice of whether to implement these within the database or not. For those who do not, Data Links Manager offers a convenient way to manage these indexes.

We are very impressed with this concept and believe it lays the foundation for a technology that will become increasingly important. We are also pleased that IBM has extended support beyond AIX to include both Windows and Sun Solaris platforms.

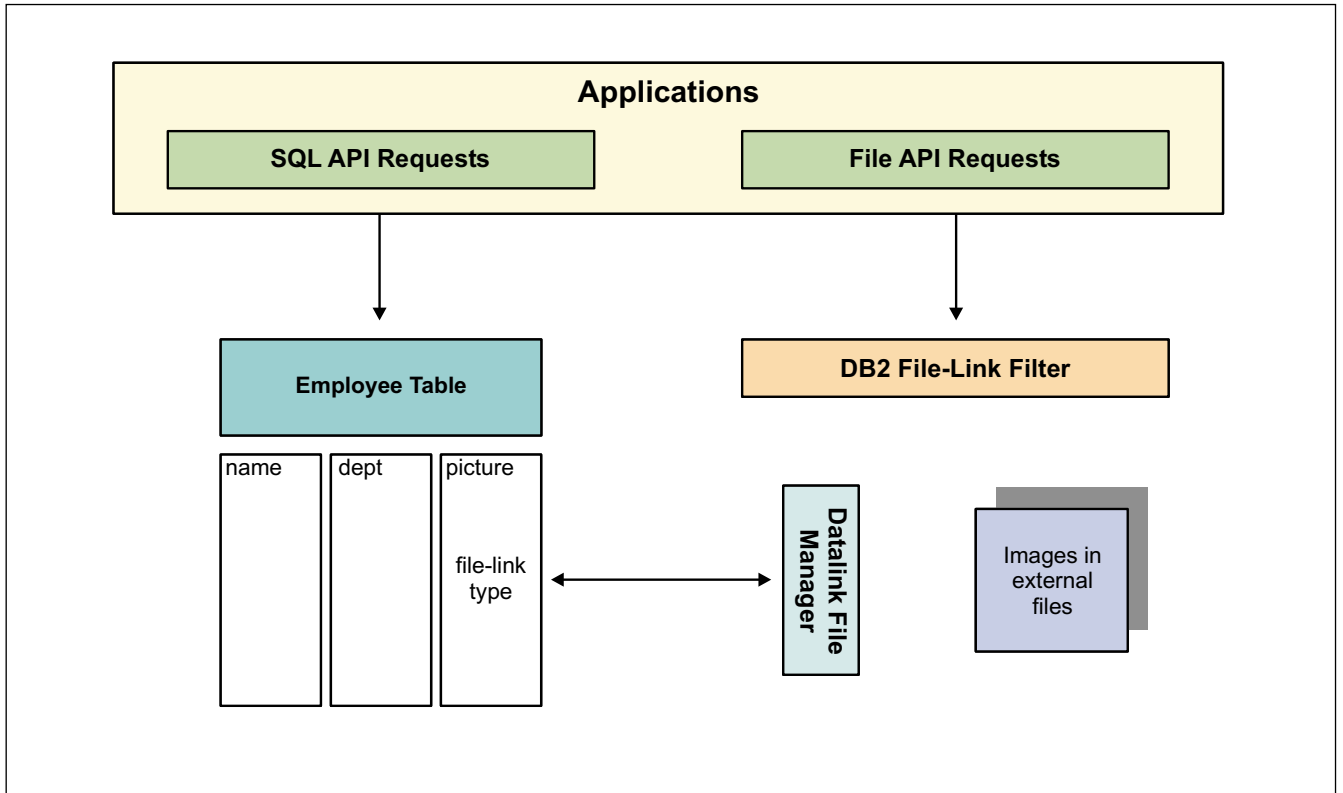


Fig. 1: DB2 Data Links Manager

Object-relational facilities

IBM uses abstract datatypes (which it calls structured types) to allow complex objects to be modelled and mapped into DB2. However, prior to version 7.1 you either built (structured) type tables or you didn't. With 7.1, that changed so that you can now create tables with structured type columns. Moreover, structured types can be embedded within structured types (so DB2 now supports nested tables). In addition, whereas previously you were limited to using functions with structured types, you can now define true methods that support encapsulation.

Typed tables are organised into hierarchies. Under the covers (i.e. where only the DBA needs to know about it), table hierarchies are stored using a method by which each hierarchy corresponds to one physical implementation table, where this table contains the union of the columns required to store rows of any subtable in the hierarchy. This is a relatively new technology and there are several theoretical approaches to storing this data, each of which has both advantages and disadvantages. In IBM's view, the approach adopted should provide the best overall performance with the fewest problems and restrictions. However, even the company's own tests suggest that this method will not be optimal in all circumstances and only time and experience will show whether this is indeed the best option.

Object ids (OIDs) are supported. However, these are user generated rather than system generated. Moreover, OIDs are only guaranteed to be unique for



the typed table hierarchy in which they exist. This means that you cannot use OIDs for clustering purposes.

Another major feature of object orientation is inheritance. This is implemented as single (not multiple) inheritance amongst the tables in a table hierarchy. However, IBM also feels, rightly, that inheritance should be applied to the database operations that apply to those tables. That is, that constraints and triggers should be inheritable. This was released with limited functionality in version 6.1 but has since been extended to include referential integrity constraints and triggers.

The other major OO feature that is included in DB2 UDB is support for object views and object view hierarchies. These provide the equivalent flexibility and functionality for typed tables as views do for conventional relational tables. Indeed, you can also define object views against existing relational tables, which can be useful for migration purposes.

As far as methods are concerned, one point to note is that IBM's literature claims that methods can be written in C, C++, Java, OLE (Basic) or any other language conforming to C calling conventions. What it doesn't say is that the support for C++ is very restricted. In fact, it is limited to the subset of C++ that is C itself.

Distributed Processing

In terms of conventional distributed processing, IBM supports all the different distributed topologies (master-slave, peer-to-peer and so forth) that one would expect, as well as an XA compliant two-phase commit.

Traditionally, of course, IBM's approach has always been hierarchical, and this is reflected in its DB2 Connect product, which is an implementation of IBM's distributed database connectivity strategy known as DRDA (Distributed Relational Database Architecture). This provides the aforementioned two-phase commit, static and dynamic SQL support, ODBC and JDBC access, dynamic load balancing and connection pooling. This last is particularly interesting as it allows the prioritisation of SQL requests submitted by DB2 Connect users. DB2 Connect integrates with both WebSphere and Net.Data (see below) so that host data can be published on a Web server and accessed via a browser. DB2 Connect also works with a variety of TP monitors, including IBM's TXSeries and CICS products, as well as Encina and MTS. In the latest version of DB2 Connect the product has been updated to allow data to be shared (across memory) between OS/390 and Linux applications that have been installed on the same system.

Distributed queries are also supported, with heterogeneous environments that comprise DB2 family databases, Microsoft (or other) OLE DB data sources, and any other data sources that are accessible via a TABLE function appearing as a single database. This facility makes use of the distributed query optimisation facilities in DataJoiner (see below). It is also worth noting the facilities of DB2 Relational Connect (which is effectively a read-only version of DataJoiner) that provides query capabilities against Sybase, Oracle, SQL Server and Informix databases.



Data Replication Services

DB2's replication facilities are provided through what used to be called the DataPropagator product (it still is on OS/390 and zSeries). This supports both a master/slave approach, where there is a single centrally managed master copy of the database feeding updates to replicas, and an update-anywhere architecture where updates are accepted at any copy and the results need to be synchronised throughout. On-demand replication is also provided. Further, as one would expect, replicas may be whole or partial copies, and cascade replication (which IBM calls fan-out distribution) is supported.

Replication in heterogeneous environments is supported, in conjunction with DataJoiner (see below), with bi-directional replication provision for non-DB2 sources such as Oracle, Sybase, Informix, SQL Server, IMS, and VSAM. In addition, data may be replicated to (but not from) Lotus Notes and any ODBC or JDBC compliant data store. The extent of the bi-directional replication here is a major competitive advantage although the conflict detection and resolution facilities provided are limited to DB2 environments.

The facilities provided use trigger-based replication for non-DB2 sources and log-based replication for DB2. The log-based approach has a number of advantages over triggers (particularly in performance terms) and, moreover, we particularly like IBM's implementation, which is very flexible. Nevertheless, triggers are simpler to implement so it would be nice to have both these alternatives as options.



m-Commerce

Apart from replication the main features of DB2 that support mobile commerce are its database options and, potentially, its spatial extender.

DB2 Everyplace

Currently, IBM has two DB2 products addressing the mobile environment: the Satellite Edition and DB2 Everyplace. The Satellite Edition was designed to support hands-off, remotely administered laptops and servers. However, this functionality is being gradually implemented within the Personal Edition, so that the Satellite Edition will be phased out.

DB2 Everyplace, on the other hand, is a small footprint (about 100K) version of DB2 that runs on PalmOS, Windows CE, EPOC, embedded Linux and Neutrino devices. It includes a synchronisation server that supports bi-directional replication with DB2 systems on other platforms.

Spatial Extender

IBM's spatial extender was originally developed jointly by ESRI and IBM using DataJoiner (see below). However, in version 7.1 this was integrated into the database product in a further joint development between the two companies. The product supports concepts such as the distances between two points (addresses) and whether a point is inside or outside a defined area. This has obvious potential for location-based services that are likely to be important in m-Commerce environments, although it clearly has application in other environments (notably business intelligence) as well.



e-Commerce

DB2 provides a number of facilities that will be useful within an e-Commerce context. Some of these, such as the closer integration with MQSeries, have wider applicability but are most usefully treated here.

Net.Data

Net.Data is an application development and deployment environment that is intended for non-programmers, which makes the connection between Web applications and the data sources they need to exploit. In particular, while it can be used for developing new Web applications, it is especially aimed at Web-enabling existing data.

Briefly, it supports native access not only to DB2 but also to Oracle, Sybase, IMS and other sources supported by DataJoiner, as well as more generic ODBC and JDBC connectivity. In development terms it supports Java applets, applications and servlets (which operate in conjunction with WebSphere) as well as Perl, C and C++.

Notable features include a set of high-performance index search functions for full text searching, support for MIDI file types and Unicode support. In addition, the 7.1 release added support for XML datatypes, a macro switching ability that allows you to change between HTML and XML output, and the ability to generate XML documents from DB2 tables. You can also use Net.Data to attach files to e-mail messages and uploading browser files.

XML Support

DB2 provides both an XML datatype and an XML extender. Basically you have two choices: either you can store an entire document in a single column (using the datatype) or you can decompose the document across multiple columns (using the extender). In either case, you can define indexes against the relevant column or columns. However, the optimiser only understands the datatype approach, so use of the text extender's search capabilities (which understands XML) will be preferred when the XML extender is used.

When using the XML extender the XML is parsed into a DTD and then mapped into a DAD (data access definition), which allows you to define the columns to which the XML is mapped. The fact that the document is pre-parsed is particularly important. It means that you do not have to parse the XML as a part of any query process, which could cause a substantial overhead.

As we have noted above you can use Net.Data in conjunction with the XML extender to generate XML documents from DB2 data. One weakness of the product however, is that it is still DTD-based and it does not yet support XML schemas.

Net Search Extender

This is an optional add-on that provides in-memory (and therefore very fast) text searching for e-Commerce and Web self-service applications. It supports word, phrase, fuzzy and wildcard searches and is designed for rapid



searching without needing to lock tables. This extender is primarily designed for such things as catalogue searching.

WebSphere and MQSeries

A notable feature of the latest release is the much closer integration with both WebSphere and MQSeries.

As far as WebSphere goes, we have already noted the integration with Net.Data. Perhaps more significantly, IBM's support for Web Services includes SOAP (simple object access protocol). This is dependent upon WebSphere since this combines stored procedures, XML and EJBs through WebSphere. While on this topic it is worth noting that UDDI (Universal Description, Discovery and Integration) is implemented by using DB2 as the repository for UDDI.

In version 7.2 you can now access data from an MQSeries message queue as if it was a DB2 database view. In other words, SQL can address a message queue as if it were a table and there is a wizard available to help to create the relevant table function and view. As an aside, there is also a comparable facility and wizard to support Microsoft OLE DB and Data Transaction Services.

It is worth noting that the advantages of the link with MQSeries are not just one way. For example, new data sources become available to applications. In particular, this means that a message queue may become a data source in a data warehousing environment while, on the other side of the coin, a message queue may be used to broadcast information to a wide audience.



Business Intelligence

Business intelligence is a major market for IBM and DB2 is its core technology. At a fundamental level, DB2 is eminently suited to act as the basis for a data warehouse, because of its performance and scalability. In addition, it also has special features to support both very large databases and tables, the latter being common in environments that adopt a ROLAP approach to information processing. For example, there is support for automatic summary tables (including incremental refresh with joins and aggregates), extended SQL functionality to handle Cube/Rollup/Grouping sets, as well as a variety of other functions. Another major feature is the particular facilities for optimising star joins.

However, IBM's vision of business intelligence goes much further than simply providing the database technology on which to build a data warehouse. First, it encompasses a two-tier model that includes a central data warehouse integrated with local data marts. In general we agree with such a view, though it can be a bit prescriptive for smaller organisations or for those wishing to undertake a small-scale pilot. In addition, IBM aims to provide all the tools necessary to build a complete business intelligence solution. These encompass a variety of products and functionality and include:

1. Data extraction via ETI*Extract.
2. Data transformation and cleansing via ETI*Extract, Vality Technology's Integrity, the Trillium Software System and certain features of IBM's own warehouse management facilities.
3. Warehouse management via Data Warehouse Center and DB2 Warehouse Manager (which includes Query Patroller, the Information Catalog Manager and QMF for Windows). These products have replaced what used to be Visual Warehouse.
4. Access to external data sources via DB2 DataJoiner (see below).
5. Warehouse Manager Connectors – these are extensions to the concept of Relational Connects (see below), specifically designed for business intelligence environments. Specifically, there is support for accessing and importing data from i2 TradeMatrix and SAP R/3 environments (no doubt others will follow). In addition there is a Warehouse Manager Connector for the Web that allows you to extract information from a WebSphere Site Analyzer database or webmart.
6. Data mart support via DB2 OLAP Server. Built-in facilities include CUBE and ROLLUP, statistical analysis functions, and aggregation features such as rolling sum and rolling average. DB2 OLAP Server is discussed further below.



7. DB2 Intelligent Miner for data mining. A new extension to this product is known as DB2 Intelligent Miner Scoring Services. The intention behind this product is to enable users to incorporate mining analytics into business intelligence, e-Commerce and OLTP applications. It is worth noting that DB2 Intelligent Miner is used as the basis for the recently released WebSphere Commerce Analyzer Advanced Edition, which provides e-Commerce mining capabilities.
8. Query processing facilities via partnerships with Business Objects, Cognos, Brio et al, plus the company's own Query Management Facility (QMF).
9. Life Sciences Data Connect – this is a recent introduction that allows you to work with a virtual database, whose underlying sources may be in diverse locations (for example, the human genome project). The techniques used here allow you, effectively, to wrap almost any data source into something that can be manipulated by SQL.

Of these facilities, the third-party products are not discussed here. Many of these are the subject of independent reviews from Bloor Research, as is DB2 Intelligent Miner. Of the remaining products, DataJoiner is not confined to the Business Intelligence arena and is considered below. However, one point is worth noting. ETI now positions its products as a data movement vendor and not just for ETL in a data warehousing environment. There are many situations in which you want to extract and transform data in a B2B or enterprise application integration environment. Arguably, therefore, ETI should have been considered under the e-Commerce heading above.

Data Warehouse Center

The Data Warehouse Center integrates closely with the ETL, OLAP and end-user products discussed above, to provide a graphical management environment to help automate, and monitor the progress of, the complex tasks involved in data warehouse creation and management. The tasks it performs include:

- Extraction of raw data from heterogeneous operational systems.
- Cleaning the extracted data to remove duplicates, filling in missing information, correcting detectable errors, etc. This facility has been expanded in the latest release through support for Trillium Software's name and address cleansing capabilities.
- Transformation of data types and the construction of derived fields to conform with the warehouse architecture.
- Summarisation and aggregation of data as is required to fit the warehouse data model. In this respect the latest version of DB2 supports integration with Computer Associates' ERwin.
- Movement and scheduling of transient data from source to target platforms.



- Managing the loading of new data into the warehouse.
- Integration of the metadata between the tools of the different vendors to provide a common look and feel for data warehouse administrators.

Design components also help the administrator to define appropriate multidimensional or star schema architectures, and implement the resulting designs.

Through DB2 DataJoiner, the Warehouse Center can both extract from and load to non-DB2 databases, so that a very wide choice of target environments is available. The flow of data can also be integrated with DB2's replication services (see above) to co-ordinate the level of currency required of the data warehouse contents.

DB2 Warehouse Manager

While Warehouse Center is focused at creation and management of a data warehouse, the Warehouse Manager (which is an optional add-on) is more aimed at maximising its scalability and accessibility. It includes three sub-projects:

- DB2 Query Patroller is a tool for the tracking and management of queries. While on the subject of queries, one particular feature of ODBC queries is worth noting. If you have repetitive dynamic queries than DB2 treats them as if they were static, which should provide better performance.
- Information Catalog Manager – basically this allows users to browse, understand and access metadata that is created and stored by the Warehouse Center. It is noteworthy that IBM supports the Common Warehouse Metamodel (CWM).
- QMF for Windows – this Query Management Facility is a multi-purpose query tool that can be used as easily in an operational environment as one dedicated to business intelligence. In the most recent release its facilities have been integrated with those of Query Patroller.

DB2 OLAP Server

DB2 OLAP Server was originally released at the beginning of 1998. This resulted from a joint development with Arbor Software (now Hyperion), the leading independent provider of multi-dimensional databases. The product consisted of the entire Essbase (Hyperion's product) multidimensional server functionality based on a physical structure implemented on a DB2 UDB core. This has the advantage that all Essbase-compliant tools and applications can run against DB2 OLAP Server without change, while allowing the use of IBM's management tools for administering and building the environment.

Other improvements over using a native Essbase model are the rapid data loading, the more compact structure, and the ability to manage larger data models from within an IBM environment. Nevertheless there is a translation issue that can result in degradation in user response times. For environments in which this is critical, IBM is now reselling the Essbase product set in its entirety, for those who prefer an mddb at the back-end. It is also worth



noting that in the latest release you can now choose your data store (ROLAP or OLAP) on an application-by-application basis, rather than all at once when you first install the DB2 OLAP Server.



Tools for the Administrator

Finally, another new feature of version 6.1 is the introduction of support for the Light Directory Access Protocol (LDAP). This can be used to create a repository that stores location and other information about system resources, including DB2 database servers.

DataJoiner

While DB2 UDB provides a number of features that facilitate interoperability (a number of which have already been discussed) the main focus is on DataJoiner.

DB2 DataJoiner has been mentioned several times in the course of this review. It is used within the Business Intelligence arena (amongst other things to provide distributed query optimisation), it is integrated with DB2's replication services to provide both synchronous and asynchronous access to data, and it is used in conjunction with DB2 Connect. DataJoiner does exactly what its name suggests: allows you to join data from diverse sources. These may be heterogeneous databases (whether DB2 or not, relational or not, local or not) or they may be data from diverse data extenders. In the case of the latter, this means that you can mix relational data with multimedia data in the same table when stored locally in DB2 DataJoiner.

However, DataJoiner is not as simple as this. It exhibits the characteristics of a gateway to other data sources (including native access to external databases such as Oracle, SQL Server, Informix and Sybase through the DB2 Relational Connect option), a distributed database and replication manager, a global optimiser and a dbms. We know of no other middleware product with this breadth of functionality.

Apart from its gateway capabilities, DataJoiner provides a global catalogue of metadata (simplifying the use of ad hoc query tools). Moreover, it adds the ability to construct views of data that consist of joins across these data sources, and will impose a global two-phase commit. Also at the global level, DB2 DataJoiner has an optimisation component. That is, it aims to optimise SQL according to whichever data source it is addressing. This is unusual. Most vendors tend to give their own products an edge when it comes to optimisation. IBM appears to be taking a genuinely dispassionate approach in this respect, which we commend.

In addition, DataJoiner includes a 'functional compensation' function. Basically, this is the idea that if a non-DB2 data source does not support a given SQL capability, then DataJoiner can compensate for this lack so that an application will see that data source as if it had that capability.

DataJoiner is built on its own DB2 database. This provides a number of further benefits. In particular, some types of join operations are made more practicable through the ability to store temporary tables. For example, consider the need to perform a Cartesian join (the full product of joining every



row in one table with every row in another) across two tables with 1,000 rows each, resulting in an answer set with one million rows. Joining the two tables at the source database would result in all of those rows being shipped for processing – but DB2 DataJoiner permits the two sets of 1,000 rows to be stored as temporary tables within its own database, where the join is performed. This offers a potential saving of 99.8% in traffic between the different components of the system!

Resilience

As far as 24 x 7 features are concerned, DB2 UDB can carry out all administration functions in flight, including index reorganisation in this release. These functions can also happen in parallel, on discrete parts of the database, with the exception of updates to optimisation statistics.

As far as failover and failback are concerned, DB2 UDB relies on (or exploits, depending on how you look at it) operating system services to support these functions and IBM is not implementing anything in the database itself, other than support for the relevant clusters or SMP/MPP systems. Thus user connections can be failed over and back, but there is no support for maintenance of either in-flight queries or updates. This is par for the course in so far as relational vendors are concerned.

Usability

IBM, particularly at the mainframe level, has historically had a reputation as a vendor whose products need serious management capability and it is in this environment that the management tools provided are richest. Thus this section focuses particularly on the OS/390 and zSeries versions of DB2, rather than the Unix, Linux and Windows versions of the product that we have mostly concentrated on in this review. There are a number of significant products worth mentioning:

- DB2 Visual Explain – allows you to examine and tune the functioning of the DB2 optimiser and DB2 DataJoiner;
- DB2 Installer – provides a graphical interface for customising database jobs on a workstation. In addition to installation this also provides migration and update capabilities.
- DB2 Estimator – is a stand-alone tool for estimating the performance of DB2 applications.
- DB2 Performance Monitor – provides analysis, control and tuning capabilities for DB2 applications, with a real-time online monitor and an explain feature.

In addition to these products there is also a DB2 Buffer Pool Tool and DB2 Administration Tool, plus a variety of utilities, including Copy, Load, Re-organisation, Rebuild, RunStats and recovery utilities. DB2 also supports SNMP for linking to Tivoli and other systems management products.



Summary

We are very impressed with the whole DB2 UDB package. In particular we like the 'data trio' of Data Links Manager, DataJoiner (and its subset DB2 Relational Connect), and the replication services that used to be known as DataPropagator. As IBM has recognised through its message of data management, rather than database management, database products can no longer be considered as isolated entities. In terms of pure database features there may be little to choose between the leading vendors in the market; it is rather in a product's relationship to a wider milieu that product selection will be increasingly focused and we believe that IBM is well placed in this regard.



Suite 6, Challenge House, Sherwood Drive,
Bletchley, Milton Keynes, MK3 6DP, United Kingdom

Tel: +44 (0) 1908 625100 - Fax: +44 (0) 1908 625124
Web: www.bloor-research.com - email: info@bloor-research.com