

# **DB2 Universal Database and Total Cost of Ownership (TCO)**

**White Paper**

**October 2000**



Richard Yevich  
Principal  
Yevich, Lawson and Associates, Inc.  
1-888-246-5049  
[info@YLAssoc.com](mailto:info@YLAssoc.com)  
[www.ylassoc.com](http://www.ylassoc.com)

Trademarks included in this document are the property of their respective owners.

---

## **Table of Contents**

DB2 Universal Database and Total Cost of Ownership (TCO)	3
Scalability / Performance	3
People Considerations	3
Technology Factors	4
Business Factors	6
Manageability	7
People Considerations	7
Technology Factors	7
Business Factors	8
Availability / Reliability	8
People Considerations	8
Technology Factors	9
Business Factors	9
Accessibility	9
People Considerations	9
Technology Factors	10
Business Factors	10
Application Development	10
People Considerations	11
Technology Factors	11
Business Factors	11
Applications	11
People Considerations	12
Technology Factors	12
Business Factors	13
Conclusion	13

---

## **DB2 Universal Database and Total Cost of Ownership (TCO)**

A growing concern we have seen from companies is the focus on the total cost of ownership (TCO) when acquiring and managing technology. With respect to software, there is no universally accepted way to compare the TCO between two products. Few customers will totally rely and accept vendor claims on which software is superior with respect to cost of ownership. Currently, at the home pages of both Microsoft and Oracle each claim their database software is superior to the other and both refer to third-party consultants to support their claims of leadership with respect to TCO.

Given this environment, the scope of this paper examines some of the advantages of DB2 UDB that have been acknowledged by the customer in production environments. This paper does not attempt to provide the definitive financial justification of using DB2, instead, it is a tactical guide to address questions we have received from customers and to share our experience. The experiences we have collected focus on the people, technology, and business cost variables customers often ask us to comment on. We have grouped the factors that impact TCO under the following headings: Scalability / Performance, Manageability, Availability / Reliability, Accessibility, Application Development, Applications.

### ***Scalability / Performance***

There is a definite cost associated when selecting technology that does not perform nor scale. Some of the issues that customers have asked us to address are:

### **People Considerations**

- Database tuning requirements

The industry has long recognized that IBM delivers the world's best optimization technology. This benefit allows DBAs to be much more productive and not spend time tuning each and every SQL statement as Oracle DBAs typically must do. In five major customer engagements, we made the following observation: In DB2 applications, only 7% of the medium or complex queries require further tuning as compared with 62% for Oracle. Tuning typically requires approximately 2 man-hours per query. In a medium sized application of 1000 queries, the savings with DB2 would be 1060 hours or over one-half a man year. In fact, we found that up to 6 hours may be required to optimize complex queries, but a lower number was chosen anticipating that the incumbent Oracle DBAs would scoff at the results. A major airline shared this following anecdote with us: approximately 60% of their developer's time is spent writing and optimizing SQL - while the other 40% is spent doing real Oracle DBA work.

In an e-business environment, Adrian Challinor, Research & Development Director e-zData.net

---

(www.e-zdata.net) shares his performance experiences:

"After lots of tweaking, special coding, etc., we managed to get good performance out of Oracle. We had to do a lot of work, special coding, use of large memory arrays, etc. to get decent performance with our applications running on Oracle. We ported to DB2 recently. With no special code, no use of arrays, DB2 straight out of the box runs faster than our Oracle 8i version during tests. This saved approximately 2 man-months work. We now develop for DB2 and port to Oracle! I fully concur with your comments on Oracle (and this from a die-hard Oracle techie)."

In an Oracle environment (mostly UNIX databases) where DBAs "manually write an access path" these applications will not only take longer to write, they will not benefit from improvements to the environment such as faster CPUs or improved optimizer cost algorithms in subsequent database releases. To leverage the new environment, applications will be re-coded, which would be another labor cost.

Oracle has not published any recent query benchmarks, including TPC-H, TPC-R, or SAP's BW benchmark. One has to examine the retired TPC-D benchmark to understand how Oracle tunes its database for large query environments. Some of the techniques Oracle employed for query benchmarks include:

- 18,000 SQL statements to define 5 tables (9x more than DB2 for a similar benchmark -100GB)
- 31 pages to define indexes (31x more than DB2 for a similar benchmark - 100GB)

We expect Oracle will continue to require excessive tuning and employ benchmark techniques that would never be accepted in a customer's production environment. In fact, third-party tools are available that will re-write Oracle queries, some of these tools claim up to 400% performance improvement, these tools are not necessary in a DB2 UDB environment. DB2 UDB will continue to prove to be much easier to tune, and provide greater performance and scalability.

## **Technology Factors**

- Database performance and scalability for large data warehouses
- Scalability Capabilities

While on paper, Oracle, Teradata, Informix claim that their databases can perform and scale, typically for very large data warehouses, most of the largest implementations are using DB2 UDB or Teradata today. Oracle customers historically have 'broken' up the data warehouse so that very large databases are not required. It is very rare to find a real TB+ Oracle databases in production.

---

In a recent Bloor Research report, "IBM and Oracle in Marketing Battle of Words (December, 1999)", the following assessment was made: The data warehouse market is a little more complex, although IBM does have some justification for complaint here. In terms of the very large data warehouse sites, IBM is certainly ahead of Oracle. Indeed IBM's own ordering of DB2, then Teradata, then Oracle is probably correct at the high end.

DB2 UDB also has the distinction of being able to scale all functions. Oracle8i claims, "Parallel Everything" but Oracle8i Parallel Server has the following restrictions (documented in the Oracle manuals):

- A transaction can only contain one Update/Insert/Delete statement and it must be the last statement (parallel query must be executed before)
- Triggers are not supported for parallel DML, which also means that any tables being replicated cannot support parallel DML
- Parallel DML cannot occur in the presence of certain constraints: self-referential integrity, delete cascade, and deferred integrity
- Parallel update/insert/delete cannot occur on tables with object columns, and/or LOB columns

Finally, the main reason why Oracle is not used for very large databases, is that its shared disk architecture is not suitable for scalability in the terabyte range. In a recent evaluation of parallel database architectures (Data Decisions, eWEEK, May 22, 2000), the following assessment was made:

"Choosing the right database clustering strategy is a critical, far-reaching decision for e-businesses, and evidence is mounting that the "shared-nothing" database cluster design used by Microsoft Corp. and IBM is the best way to meet the challenges inherent in high-performance e-commerce."

The eWEEK evaluation also interviewed Hal Berenson, who was Digital's head of Rdb development when clustering was being added to the database.

"The simplest way to [add clustering] is to not change much in the database system. Use the same buffering system, the same I/O mechanism, and take the same locking construct in the database. But extend that, so no matter what node you're running on ... you make sure multiple instances don't step on each other." ... "What we found with Rdb is we never got [traffic levels] so low you would get linear scaling in the cluster," ... "You could get some level of scaling up to three or four nodes, but after that, things would flatten out."

Oracle Parallel Server adopted the same clustering architecture as Rdb, while this method proved to enable faster delivery of a parallel database to the market, it does not exploit today's environment of commodity processors and disk. In fact, Microsoft later hired

---

Berenson to help revamp its SQL Server 7.0 release, and he lobbied for a shared-nothing design.

## **Business Factors**

- Disclosure
- Initial Software Costs
- Service and Support Costs

It is no secret that Oracle's Terms and Conditions does not allow for open disclosure of its performance by anybody who has purchased its database. This has created a restricted environment where only Oracle-approved benchmarks or performance numbers are available for comparison. Unfortunately, customers attempting to compare databases based on performance are have been limited on what data can be compared. DB2 UDB has no such restrictions and does not prevent neither customers nor consultants in openly publishing performance experiences with their database.

With respect to initial acquisition cost of the software, DB2 UDB compares favorably against the other databases and even greater price differences exist between DB2 UDB EE vs. Oracle's Enterprise Edition. (Check the Oracle Home Page for prices).

In a parallel environment, the customer must buy Oracle Parallel Server and the Partitioning Option to get the capabilities packaged in DB2 UDB EEE. We fully expect that the trend of charging the customer by feature enhancement will continue with Oracle.

Finally, in an effort to cut its costs, Oracle Corporation has recently announced this year that it will use the Internet to provide self-service technical support and will no longer allow customers to phone Oracle directly for technical support. Attendees at this year's International Oracle Users Group expressed strong concerns and that they don't want their phone privileges taken away. One user comments:

"I'm reluctant to use (the online system)," said Jerry Webb, Oracle support project manager for the computer operations division at Deere & Co. in Moline, Ill. "With the phone, at least I get to talk to somebody, even after being put on hold." (Oracle Users Wary of Its Web-Based Support Plan, ComputerWorld 5/15/2000).

It should be noted, that even with a lower level of support, Oracle charges one of the highest support and maintenance costs in the industry. Annual Basic support (Silver) costs starts at 22% of list price. Many customers who use both DB2 and Oracle comment that the level of service and support Oracle provides does not merit the higher costs.

---

## **Manageability**

The ease in which a customer can incorporate the managing DB2 in a production environment will often have a direct impact on the length of the sales cycle. Some of the issues that customers have asked us to address are:

### **People Considerations**

- DBA skill required

The customer may present the lack of available DB2 skills in the marketplace as a sales objection. In fact, the customer may argue the because of the scarcity of DB2 skills, it would be more expensive to implement and support a DB2 environment. However, in every customer situation we have encountered, there are always more DBAs required (Oracle, Sybase, Informix, Microsoft) to manage the third-party database environment, than DB2. Also, because DB2 strongly adheres to SQL standards, it is very easy for an existing DBA trained in one environment to quickly become proficient in supporting DB2. There is an IBM Education course titled, *CF281--Fast Path to DB2 UDB for Experienced DBAs*, that was designed specifically to get Oracle/Informix/Sybase/etc DBAs to quickly be productive in a DB2 environment. Additional educational support includes: International DB2 Users Group, DB2 Technical Conferences, IBM Learning Services (demand rose up 40% last year, over 54,000 students took courses, IBM Partner Education (over 7500+ partners educated), and the DB2 Scholars Program (over 3000 universities worldwide involved), and independent education programs such as the international DB2 Symposium held worldwide every month.

If the customer already has DB2 installed on the OS/390 platforms, these skills can be leveraged and lower the overall support costs required for introducing a distributed computing environment. We have seen many customers able to leverage their DB2 skills from other environments on the distributed platform (Intel and Unix). In a recent large DB2 UDB installation, the customer acknowledged there would be a cost savings of up to \$250,000 over a three-year period.

### **Technology Factors**

- Utilities Provided
- Utility Speed

DB2 UDB provides a rich set of utilities with its Control Center. Many of these utilities are not offered by IBM's database competitors. One competitive advantage IBM offers is the REORG utility. This utility addresses the problem of data fragmentation, over time, data becomes de-clustered, performance suffers, and reorganization of data enables faster access to data.

---

Third-party tools for reorg are available and recommended for Oracle installations. One reference customer for these third-party reorg tools (a pharmaceutical research company, managing multiple Oracle instances (20 instances running mostly on Sun) with no more than 10GB of daa per instance) gave the following quote:

"We were performing tablespace reorgs manually each week, and complete database reorgs each quarter, to minimize fragmentation," he said. "It was very repetitive and time-consuming. Because the DBAs here are responsible for a lot of other tasks, the reorg time was becoming too demanding."

With respect to utility speed, DB2 is much faster than Oracle on load performance and index creation. DB2 UDB has both of these utilities SMP-enabled, while in the Oracle environment, the DBA must split the input file into multiple files to get any parallelism for load jobs. This procedure is often complex and time consuming. Also, Oracle's index creation runs serially, so there is another significant productivity gain we have observed using DB2 utilities over Oracle.

## **Business Factors**

- Cost of utilities

DB2 UDB includes the Control Center and various utilities at no additional charge. Oracle's Enterprise Manager provides some management capabilities, but additional chargeable options are necessary to provide complete support. Many Oracle customers have complained about Oracle's pricing in this area: Some 88% of respondents are unhappy paying for enterprise manager and its options - the tools that manage an Oracle database (Oracle Users Go on the Warpath, Reference Computing 06/05/1999 pg.1)

## ***Availability / Reliability***

Customers have often asked us what are key strengths are in availability. Customers have generally acknowledged that IBM provides very reliable software.

## **People Considerations**

- Customer satisfaction

When the end user has an environment where the system is unavailable, either because the database had to be taken off-line to execute utilities, or the lack of reliability of the database continues to bring the system down, this can create an opportunity to sell DB2. Customers generally have had good experiences with IBM with respect to availability and reliability.



---

## Technology Factors

- Takeover time

Customers often ask us how IBM can make a system 100% available. In this scenario, focusing on just the database would not be sufficient; the entire environment must be considered. So no matter what vendors may say about their replication or takeover features, the system will never be 100% available until you have an environment where each transaction is duplicated and sent to two different independent databases.

However, with respect to takeover times, we often get asked why DB2 UDB does not support HACMP Mode 3, a shared disk implementation. We have observed that DB2 recovery from a failure was competitive with and generally faster than Oracle.

## Business Factors

- Cost of Unavailability

There is a definite cost that can be associated with the database not being available because the DBA is manually unloading and loading the data instead of running a reorganization or redistribute utility. These costs will increase as the size of the database grows, the costs associated with the unavailability of the database is often best determined by asking the customer to put a value on downtime.

## *Accessibility*

There is a significant cost associated with the inability to access and integrate data from different platforms and heterogeneous databases.

## People Considerations

- Support staff requirements

It is not uncommon to for a customer to dedicate six people to oversee a project to integrate data from one database to another database. These can be quite complex projects, especially when heterogeneous databases are involved. The biggest costs involved will be the integration effort required, quite often that integration effort requires accessing data from DB2 in the OS/390 environment.

With respect to accessing data from heterogeneous platforms, clearly there is an advantage of being able to leverage one support staff knowledgeable about DB2, rather than two different support groups, each only knowing how their database operates.

---

## Technology Factors

- Heterogeneous Environment
- Performance

In a data-warehousing environment, we have seen many businesses write their own extract, transformation, and load programs (ETL programs) when trying to get data out of DB2 OS/390 into a database on the distributed platform. DB2's replication and DB2 Connect can replace several of these user-written extraction, transformation and load program.

With respect to performance, both DB2 Connect and DB2 DataJoiner have out performed third-party gateways and the middleware provided by the other database companies. Today, DB2 Connect is the market leader today for accessing data from DB2 OS/390. DB2 DataJoiner's focus on global optimization also makes it the leader for performance and integration of heterogeneous databases. IBM has continued to improve its offering and has introduced DB2 DB2 Relational Connect to provide heterogeneous access to Oracle databases for customers needing to integrate or migrate from Oracle.

## Business Factors

- Initial Software Costs

DB2 Connect is software that is installed on the workstation. Compared to mainframe-based gateways, such as Oracle Transparent Gateways, DB2 Connect is much more economical to install. As the number of DB2 OS/390 systems needed to be accessed increase, the benefits of using DB2 Connect become greater.

The Internet has driven the database vendors to reevaluate their pricing structures. DB2 UDB uses a very simple processor-based pricing model, where customers can benefit by paying a flat rate for each processor regardless of the number of web users using the system. Oracle uses a per-MIPS pricing model called Universal Power Unit. Oracle customers have reacted very unfavorably to this because they see huge costs increases. We have seen many customers ask us for alternatives to Oracle because of this new pricing model.

## *Application Development*

Customers who fully exploit the object-relational capabilities in DB2 UDB can save significant effort and costs in implementing complex applications. Some of the customer considerations we have seen:

## People Considerations

- Heterogeneous Databases

---

DB2 provides object-relational capabilities that can also scale into the parallel environment and lacks the Oracle has restrictions has working with objects (see Oracle8i Parallel Server restrictions mentioned earlier in this paper)

## **Technology Factors**

- Openness
- Performance

Every database vendor has introduced a proprietary stored procedure language, except IBM whose stored procedure language is based on industry standards. Applications written for one database using stored procedures will not work with a different database. DB2 UDB V7 and DB2 UDB V6 on OS/390 introduce SQL/PL, which implements SQL-based stored procedures based on industry standards (SQL99).

Recently, we have received a few questions on how to position Oracle's Internet File System (iFS). The iFS actually stores external data inside the database, this creates a performance overhead when an application (i.e. CAD/CAM requiring an engineering drawing) needs data. The iFS does not allow the application to access the data directly, the files must be released from the database first, before the application can access the data. With IBM's DB2 DataLinks, the benefits of accessibility, backup and restore are there, but without the performance overhead.

(Note, that Oracle's iFS product was over a year late in delivery from its expected delivery date)

## **Business Factors**

- Initial software costs

We have noticed that even Oracle cannot ignore DB2 UDB's growing popularity. DB2 UDB has always included its rich DB2 Relational Extenders at no additional charge to support multimedia data types. Oracle's InterMedia, which used to carry a significant charge, is now included with its database.

## **Applications**

The ease in which new business applications can be supported can either speed up or slow down the customer's implementation schedule. Here are some customer considerations we have seen:

## **People Considerations**

- Leverage existing skills

---

We often get asked if non-Microsoft databases can be integrated into the customer's existing Microsoft programmers. This community should not to be ignored, they can often influence a database decision, by stressing the integration requirements to their Microsoft tools. Often there is a lack of awareness in the Microsoft community of what other databases can provide.

DB2 was first database to enable itself on Windows 2000 (DB2 UDB V6.1 and V7.1) and the DB2 Stored Procedure Builder excellent integration with Microsoft's Visual Studio. Naturally, Microsoft's OLE DB is supported, DB2 V7.1 provides both consumer and provider support. While Oracle does support much of the Microsoft development environment, we have observed that Microsoft has often had to build support for its technology to integrate with Oracle (i.e. ODBC, OLE DB).

## Technology Factors

- Transactions vs. Queries

DB2 UDB's unique code base allows scalability in both transaction and decision support environments. DB2 UDB has published the best OLTP (TPC-C) benchmark and the best query (TPC-H) benchmark to date. DB2 UDB's newer code base has been architected to handle both kinds of workloads. In a recent InfoWorld ("Garden.com sees flowers blooming on Web" April 3, 2000) interview with the CTO from Garden.com, he quotes the following regarding Oracle technology:

"In the last four years, the biggest thing I've spent my time on is performance. Everything comes from performance. If you have performance wrong, you're never going to be able to get your architecture optimized. If you get the performance right, the sky's the limit. I've spent more time on caching in the last four years than anything else; that's the only way you get performance [with Oracle]. And the reason is that the databases are so damn slow. I don't understand why everyone loves Oracle....The level of work I have to put in just to get around the performance is unbelievable.

When the CTO was asked how he gets around the Oracle performance problem, he quotes:

"I have 300,000 cache files on every single one of my Web servers, all storing this data locally. I have to tell people, "Don't hit the database. Put everything in it, but don't hit it." It's really unbelievable that we had to do that and it's all because the database is based around an architecture that's 20 years old."

---

## Business Factors

- Initial Software Costs
- Third-party Vendor Support

In keeping competitive with Microsoft, DB2 V7.1 packaging includes the ability to build datamarts / data warehouses, and on OLAP starter kit (supports up to 3 users). Alternatively, this could easily cost up to \$100,000 or more in an Oracle environment.

DB2 has come a long way in being accepted by the ISV community. SAP, Siebel, i2, Arriba, and PeopleSoft are just a few of the vendors that have selected DB2 as their strategic development platform. In fact, a growing number of application packages do not support some of the older legacy databases such as Sybase, Informix, and Teradata. To see what applications support DB2 UDB, check [www.softare.ibm.com/data/partners/solutions](http://www.softare.ibm.com/data/partners/solutions), or ask us for documentation.

## Conclusion

The demand for e-business and data warehousing applications will continue to benefit the major database companies and fuel a need for DBAs to support these applications. This paper has attempted to summarize what customers and independent analysts have experienced in using DB2 UDB in a production environment. We have also shared our first-hand experience in working with IBM's data management technology. We expect that DB2 UDB will continue to set a benchmark with respect to TCO as future enhancements are made.

### Reference Web Sites:

DB2 UDB Home Page - <http://www.ibm.com/data>

Oracle Home Page - [www.oracle.com](http://www.oracle.com)

TPC Home Page - <http://www.tpc.org>

---

Yevich, Lawson and Associates, Inc. is an international consulting firm specializing in all areas regarding database architecture performance. Its consultants are some of the most notable in the industry in terms of experience and depth of knowledge. YL&A offers consulting services and education in the areas of database design, large data warehouses, SQL coding, application design, and performance audits and reviews. In addition, it also has a development staff who can offer a variety of services in the areas of project management and product development, especially in the area of customized applications for the web, e-Business, and information portals.