Informix Software, Inc.

Competitive Analysis of Extended - Relational Database Management Systems

Thomas X Yates
Product Marketing

Version 1.5
Nov 11, 1997

Table of Contents

1. Introduction

1.1 Purpose of This Document

This document provides a detailed competitive analysis of the extended-relational database offerings from the four major RDBMS vendors—Informix, Oracle, IBM, and Sybase. The document contains a matrix of extended-relational functionality and whether or not the functionality was implemented by each vendor. The chart provides a high-level analysis of the product functionality. Elsewhere in the document, Each feature in the matrix is rigorously defined, the business value identified, and the implementation details for each vendor are discussed.

1.2 About This Document

This information contained in this document is the result of extensive competitive research. The following list contains some, but not all, of the sources of information for this document: Oracle 8 product documentation; Oracle white papers; Oracle marketing materials; IBM DB2 Universal Server v5.0 product documentation; IBM white papers; IBM marketing materials; Sybase Adaptive Server 11.5 and SQL Server 11.5 product documentation; Sybase white papers; Sybase marketing literature; industry analyst reports from Gartner Group, Forrester Research, Hurwitz Consulting, The Bloor Group, IDC, Ovum, and others; computer industry periodicals including DBMS Magazine, Database Programming and Design, Object Magazine, Computer World, Information Week, Info World, and others. The analysis of Informix Dynamic Server Universal Data option was conducted based on functionality available today or due to be release in or before v9.14.

It is believed that the information contained herein is 100% accurate as of the date of publishing of this paper. Every effort was made to ensure that the information is correct. Refer to Section 1.3, Providing Feedback on This Document if you have any questions or comments.

1.3 Providing Feedback on This Document

Email is the preferred mechanism for feedback. Please send comments to Andrew Hawthorn (hawthorn@informix.com) and carbon copy Steve Lambright (slam@informix.com), Tom Yates (tyates@informix.com), and Chris Okon (cmokon@informix.com). To provided telephone feedback, please call Andrew Hawthorn (510.628.3852) or one of the individuals listed above.

The most important aspect of your feedback is how the information helped you with a customer or prospect, the audience who was most influenced by the document, the particular issues they found compelling, and who the competition was.

If you see anything that you believe is incorrect, please provide that feedback as well. Since this document is the product of extensive research and the information contained is believed to be accurate based on public information, please provide a verifiable reference for your input in the form of a URL, periodical, analyst/research report, product documentation, or other equivalent source.

This document will be maintained as part of an ongoing effort to supply Informix employees with the critical information they need to successfully compete in the market. Specific suggestions on how to improve the document are always appreciated.

2. Competitive Analysis Matrix

	Informix	Oracle	IBM	Sybase
Extended Relational Architecture				
Single Server	Yes	No	Yes	No
Extendible Query Language	Yes	No	Yes	No
Extendible Optimizer	Yes	No	No	No
Extendible Indexes	Yes	No	No	No
SQL3 Enhanced	Yes	No	Yes	No
External Data Tables	Yes	No	Yes	No
Upgrade w/o Migration	Yes	No	No	No
Certified Data Extensions	Yes	No	No	No
Data Types				
SQL3 Types				
List	Yes	No	No	No
Set	Yes	No	No	No
Multi-Set	Yes	No	No	No
Specific Complex Types				
Text	Yes	Yes	Yes	No
Spatial	Yes	Yes	No	Yes
Time Series	Yes	Yes	No	No
Image	Yes	No	Yes	No
Video	Yes	Yes	Yes	No
SQL Type Extensibility				
Distinct Types	Yes	No	Yes	No
Row Types	Yes	Yes	No	No
Complex Type Extensibility				
New Base Types	Yes	No	Yes	No
Developer's Kit	Yes	No	No	No
Data Modeling				
O-O Features				
Data Type Inheritance	Yes	No	No	No
Table Inheritance	Yes	No	No	No
Polymorphism	Yes	No	Yes	No
Server Routines				
Stored Procedures	Yes	Yes	Yes	Yes
C/C++ Routines	Yes	No	Yes	No
Java Routines	Yes	No	Yes	Yes
Extended Indexing				
User Defined Indexes	Yes	No	No	No
Index UDT's	Yes	No	Yes	No
Support for New Index Types	Yes	No	No	No
Extend Index Types	Yes	No	No	No

	Informix	Oracle	IBM	Sybase
O-O Application Development				
C/C++ Object Interface	Yes	No	No	No
JDBC Interface	Yes	Yes	Yes	Yes
Java Object Interface	Yes	No	No	No
DBMS Features				
Query Optimization	Yes	No	Yes	No
Referential Integrity	Yes	No	Yes	No
Triggers	Yes	No	Yes	No
Data Integrity	Yes	No	Yes	No
Partition Data	Yes	No	No	No
Parallel DML	Yes	No	Yes	No
Random Access LOBs	Yes	Yes	Yes	Yes
Incremental Transaction Logging for LOBs	Yes	No	No	No
Extendible Storage Manager				
User Defined Primary Access	Yes	No	Yes	No
Parallelism (for above)	Yes	No	No	No
External Data Sources	Yes	No	Yes	No
Predicate Push	Yes	No	No	No
Manageability				
Single Admin Tool	Yes	Yes	Yes	No
Centralized Backup	Yes	No	Yes	No
Centralized Restore	Yes	No	Yes	No
Parallel Backup for UDT's	Yes	No	Yes	No
Parallel Restore for UDT's	Yes	No	Yes	No
Incremental Backup for LOBs	Yes	No	No	No
Incremental Restore for LOBs	Yes	No	No	No
Interoperability				
Legacy Data Access (gateways)	Yes	Yes	Yes	Yes
XA Interface	Yes	No	No	No

3. Extended Analysis by Feature

3.1 Extended Relational Architecture

The manner in which vendors have extended the capability of their RDBMS products. The term "extended relational" is used in favor of "object-relational" since the O-R paradigm is not required to extend the capabilities of an RDBMS.

Extended-Relational architectures provide a mechanism for IT organizations to expand the types of data managed by the DBMS while simultaneously preserving existing investment in RDBMS technology. The degree to which RDBMS can be extended, how much investment is preserved, the competitive advantage to be gained by upgrading to an extended-relational database, and the cost associated is dependent on each vendors implementation.

INFORMIX: INFORMIX-Dynamic Server is an O-RDBMS.

ORACLE: Oracle 8 has some object-based capabilities and Oracle marketing occasionally refers to it as an O-RDBMS, but its implementation of O-R functionality is incomplete in comparison to Informix and IBM.

IBM: IBM DB2 Universal Server is on O-RDBMS based on IBM DB2 Common Server. The functionality is quite similar to that of Informix—the differences are covered elsewhere in this document.

SYBASE: Sybase Adaptive Server is not an O-RDBMS. Sybase Adaptive Server is a middleware engine which uses separate data stores for new data types and has no support for any object-oriented capabilities.

3.2 Single Server

All data types and server routines are managed (read, write, and execute) by a one and only one database management system. IT users must be able to create and manage new base types which are not created from pre-existing SQL types.

This is the preferred mechanism for integrating extended data types into the RDBMS. The perceived benefits of the single server architecture are numerous, including: greater query efficiency; higher levels of data integrity and server availability; easier data management.

INFORMIX: INFORMIX-Dynamic Server is a single-server system as defined.

ORACLE: Oracle 8 is not a single server architecture because IT users can not create new base data types and add them to the server. Only Oracle can create and add new base data types. Oracle markets Data Cartridges as a mechanism for adding new data types to the server, but it is important to realize that these are not base data types, they are roughly equivalent to SQL3 row types. Data Cartridges are implemented by defining a PL/SQL structure and related procedures. As a reaction to Informix's and IBM's superior capabilities, Oracle has committed to enhancing Data Cartridges to support new base types by the

end of 1998—which means the functionality may be available in late 1999, but most like users will have to wait until the new millennium.

IBM: IBM DB2 Universal Server is a single-server system as defined.

SYBASE: Sybase Adaptive Server is not a single server architecture because IT users can not create new base data types and add them to the server. In fact, Sybase has eschewed this architecture and promotes that multiple independent servers, each focused on a specific data type is a superior solution. For once, Oracle and Informix (not to mention IBM and many industry analysts) agree on something, Sybase is dead wrong.

3.3 Extendible Query Language

New data processing functions can be programmed by IT users in traditional programming languages (C/C++ and Java) and integrated into the DBMS. IT users must be able to add: functions that manipulate new base types and SQL-derived types; functions that manipulate pre-existing types independent of new data types.

New functions allow IT users to increase the capability, flexibility, and performance of the database without changing the database schema—applications continue to work unchanged.

INFORMIX: INFORMIX-Dynamic Server has an extendible query language.

ORACLE: Oracle 8 does not have an extendible query language because: IT users can not program routines in a 3GL and integrate them into the database server; the limited facility Oracle has in this regard requires that procedures be associated with a new SQL-derived data type. Once again, Data Cartridges are the mechanism whereby the server is extended. Unlike Informix, IT users can not write procedures to operate on pre-existing data types, thereby increasing their functionality and improving performance. Data Cartridge routines must be written in PL/SQL and are associated with the data type defined in the cartridge. Alternatively, they can be written in any language with a C call-interface, but the processes are executed external to the server. These functions can not be run in parallel and require data to be transported outside of the database server, defeating the purpose of server-side routines.

IBM: IBM DB2 Universal Server has an extendible query language.

SYBASE: Sybase Adaptive Server does not have an extendible query language because IT users can not define new base types, SQL-derived types, or functions that operate on those types. Sybase does have the ability to right stored procedures in Java, which they try to use to create FUD.

3.4 Extendible Optimizer

An exposed interface to the optimizer which allows new data types to inform the optimizer about the associated costs of reading and writing the data type and its associated indexes. This is extremely important for new base types.

An extendible optimizer is a pre-requisite for both individual query performance and scalability as the number of users and amount of data increases. Without and extendible optimizer, the DBMS query optimizer knows nothing about the I/O costs of data and index access, which is required for reorganizing queries to minimize time consuming processes. Consequently, queries accessing new data types lack any sort of optimization.

INFORMIX: INFORMIX-Dynamic Server has an extendible optimizer. Queries accessing new data types, whether derived from SQL or new base data types, fully participate in query optimization.

ORACLE: Oracle 8 does not have an extendible optimizer. There is no exposed interface to the Oracle query optimizer for IT users or 3rd party developers. Oracle tries to create FUD by referring to Verity, who is creating a Data Cartridge. An actuality, Oracle is developing the cartridge for Verity because it is impossible for anyone other than Oracle engineers to integrate new base types with the server.

IBM: IBM DB2 Universal Server does not have an extendible optimizer. The IBM Extender API does not provide a facility for supplying requisite information to the query optimizer.

SYBASE: Sybase Adaptive Server does not have an extendible optimizer.

3.5 Extendible Indexes

Enhancements to the DBMS indexing mechanism to support user-defined indexes, index new data types, support new index types, and extend the index types supported.

Traditionally, the indexes in an RDBMS are created by a simple alphabetic or nominal ordering of the data. In many business applications, these indexes do not correspond to the business needs. User-defined indexes allow the IT user to specify the rules by which the index is created, guaranteeing that the indexes created will be useful and correspond to business needs. Also, the addition of SQL-derived data types and new base types to the DBMS requires that the database implement new indexing capabilities to ensure performance and scalability with these new data types. SQL-derived data types are relatively easy to support because the same indexing mechanisms can be used with only a few adjustments. But, support for new base types requires massive changes to the indexing mechanism of a DBMS. Many new base types, such as spatial data, can not be adequately indexed using the standard mechanisms (B-Tree's).

INFORMIX: INFORMIX-Dynamic Server has extendible indexes.

ORACLE: Oracle 8 does not have extendible indexes—DBA's can not override the internal ordering or data by using a function; Oracle does not have support for any index types other than B-Tree; no one but Oracle can add index structures to the database. In short, Oracle supports none of the functionality required. Again, Oracle will try to create FUD by referencing the Verity Data Cartridge, which is being developed by Oracle, not Verity.

IBM: IBM DB2 Universal Server does not have extendible indexes—none of the requisite functionality is supported. IBM will claim to support text indexing, but the text index lives outside the database (it is not entirely clear how the index is updated, but the mechanism is apparently trigger-based). Also, IBM

marketing literature does reference the importance of supporting index structures other than B-Tree's—they specifically mention R-Tree's. Do not misconstrue this—IBM does not support extendible indexes in any way shape or form. In fact, the most likely reason for them to mention extendible indexes is because two IBM DB2 Universal White Papers were written by a consultant who also wrote a white paper for Informix.

SYBASE: Sybase Adaptive Server does not have a extendible indexes. The SQL Server engine is limited to the same indexing functionality that has been available for years. IT user's can't create SQL-derived or new base data types, so it's impossible for the server to index them. The type-specific data stores may or may not use index types other than B-Tree's.

3.6 SQL3 Enhanced

The Structured Query Language (SQL) standard is set by ANSI. The most recent approved standard is commonly referred to as ANSI-92. The ANSI committee X3H2 and others (most notably, X3H7) are in the process of creating the next standard known as SQL3. The SQL3 standard as developed by ANSI-X3H2 defines extensions for SQL-92 including: new data types list, set, and multi-set; user-defined data types including distinct types and row types; table inheritance.

SQL3 as defined by ANSI-X3H2 is not yet a standard, but it is relatively far along the process and most database vendors have implemented some SQL3 capabilities in currently available products. SQL3 greatly enhances the data modeling capabilities of the database, especially for complex data models such as parts hierarchies and bills-of-material. The object-oriented aspects of ANSI-X3H2 SQL3 provide for better integration of object-oriented development tools with database applications.

INFORMIX: INFORMIX-Dynamic Server has the most complete SQL3 implementation available.

ORACLE: Oracle 8 has limited support for the SQL3 standard. Oracle has failed to implement the new types list, set, and multi-set, Also, Oracle does not support distinct types, row types, data type inheritance, or table inheritance. Oracle deflects the issue by saying that the specification is not complete yet. But, Oracle has many vendor-specific SQL enhancements for objects that do not conform to the standard, but implement almost identical functionality. Oracle also lacks support for the list, set, and multi-set data types. Instead, Oracle offers vendor-specific VARRAYS and NESTED TABLES.

IBM: IBM DB2 Universal Server is not adequately SQL3 enhanced. IBM lacks support for the list, set, and multi-set data types. Unlike Oracle, IBM does not have a vendor-specific hack to confuse the issue. IBM also has failed to implement row types, data type inheritance, and table inheritance.

SYBASE: Sybase Adaptive Server has not implemented a single piece of SQL3.

3.7 External Data Tables

The ability to integrate external data sources with the database management system. External data sources must include, but are not limited to tertiary physical storage devices (disk hierarchies, WORM drives, near-line tape archives, …) and user-defined gateways to legacy systems. In keeping with the principle of physical abstraction, integration of these external data sources must be completely transparent to the SQL developer.

The DBMS's built-in methods for storing data on disk may not be up to the task when it comes to storing user-defined data. For example, video archiving systems use highly-optimized and extremely high capacity storage solutions. To adequately integrate video with the DBMS, integration with these systems is requisite. Similarly, while data gateways exist for common legacy systems such as IMS and IDMS, gateways do not exist for a large number of older systems and proprietary data storage mechanisms. The ability for IT users to define high-performance gateways to pre-existing data stores, real-time data feeds, and the plethora of other alternatives greatly improves the quality of applications that can be created.

INFORMIX: The Informix Server with Universal Data Option has the most complete support for user defined storage of any commercial DBMS.

ORACLE: Oracle has no support, nor do they have any plans to support, open storage managers. Instead, Oracle developers are expected to create their own storage managers and access them through a middle-ware layer.

IBM: IBM possesses an open storage manager called Data Joiner.

SYBASE: Because of its heavy middle-ware emphasis, Sybase's Adaptive Server includes features that look like the Informix and IBM open storage manager. However, access to these external data sources is not achieved through the DBMS, and important techniques like predicate and join push cannot be done.

3.8 Upgrade w/o Migration

Upgrading is the process of installing new executable and other files to make extended-relational functionality available. Migration refers to the process of converting data files from one format to another—the format changes may well be invisible to the user, but are not invisible to the executables. Users should be able to upgrade to extended-relational capabilities without migrating data.

The process of upgrading to new functionality is often nerve-wracking to IT organizations. They are putting the operational stability of mission-critical systems at risk. The requirement for data to be migrated as well as the server upgraded only compounds the problem. To ensure successful migration, much care must be taken to backup the data in a consistent state so that it can be restored if a problem occurs. Then, the IT organization must invest even more time in converting every data file to the new format. This is not necessarily a trivial or quick task—imagine an OLTP system with 100's of gigabytes of data. Data migration often has many hidden costs besides down time, such as increased temporary disk space and permanent disk space requirements. If the success of the first migration attempt can not be guaranteed, the IT organization is being put at even greater risk as they may have to repeat the process several times.

INFORMIX: INFORMIX-Dynamic Server can be upgraded to the Universal Data Option without migrating the data.

ORACLE: Oracle 8 requires a rather complex migration. The entire data set must be backed up, the amount of disk space for data segments increased (log spaces, indexes, and data spaces all require more space), the data restored, and a migration program executed. There is no deterministic way to guarantee that a DBA has completed all steps correctly and that the upgrade will work. As an added bonus, you can only upgrade from Oracle 7.1.4 or newer. Oracle users on older databases will have to go through two upgrade and migration sequences.

IBM: IBM DB2 Universal Server requires that a migration program be executed to upgrade the data files. The data files will increase in size, so sufficient space must be made available or the migration will fail. As with Oracle, there is no deterministic way to guarantee the success of the migration..

SYBASE: Sybase Adaptive Server 11.5 requires data migration. The question is why? There are effectively no changes to the SQL Server's capability.

3.9 Certified Data Extensions

Certification is to ensure that database extensions integrate with the database server in the proper manner, conform to the publish guidelines for public interfaces, and meet the rigorous quality standards required for product applications. To meet certification requirements, database extensions must have: user documentation, installation programs, technical support, and pass quality assurance tests.

INFORMIX: INFORMIX-Dynamic Server extensions are called DataBlades. Informix has a certification program for DataBlades.

ORACLE: Oracle does not support 3rd party or IT developed database extensions. Oracle Data Cartridges are simply PL/SQL libraries that define the equivalent of a SQL3 row type. Oracle does not have a Data Cartridge certification program.

IBM: IBM DB2 Universal Server extensions are called Extenders. IBM does not have a certification program for Extenders.

SYBASE: Sybase Adaptive Server 11.5 does not have database extensions. Sybase integrates separate, stovepipe data stores via their middleware engine called Adaptive Server. Sybase does not have a certification program.