







ILLUSTRATION



Illustra **Information Asset Management**



# Illustra Object-Relational Database Management System

## ILLUSTRATM: OBJECT-RELATIONAL DBMS

The Illustra Server is the industry's first DBMS that efficiently handles alpha numeric, character, text, video, images, and documents within a single repository. Built from the ground up to deliver high-performance relational and object database management, Illustra embeds Object-oriented (OO) capabilities in a relational model, providing the first breakthrough in the 25-year history of the relational DBMS. The resulting synergy delivers capabilities that neither an Object DBMS (ODMS) nor Relational DBMS (RDBMS) can provide.

Although relational DBMS has solved the problems of managing commercial data consisting of numbers and character strings, it has difficulty handling the growing predominance of more complex data. Leading application architects are seeking to expand the definition of data to include: diagrams, maps, images, sound, documents, time series, multi-dimensional data, etc.

The two predominant modern DBMS architectures, RDBMS and ODBMS, each provide a partial solution to the problem of managing rich data. Unfortunately, both suffer from serious limitations.

### RDBMS: No Rich Data Types

Repeated attempts to manage rich multi-media data using an RDBMS such as Oracle, Sybase, and Informix have failed. Because the RDBMS can store complex data only as

uninterpreted BLOBs (Binary Large Objects), a relational DBMS is just not suited to manage such information.

Many RDBMS vendors striving to capture the benefits of OO are attempting to graft an object layer on top of the old relational product. But this technique is fatally flawed as well. Because the basic engine is unable to understand how to optimize storage and access to object data, these layered products are very inefficient.

### ODBMS: No Query Language

In the 1980's, the technology that promised to manage rich multi-media data was Object Orientation (OO), which offered a new way to view and model the relationship between data and applications. OO brought the hope that applications would become easier to develop, more robust and more maintainable. The merger of the existing client/server database model with Object Orientation led to the development of Object-oriented Database Management Systems (ODBMS) that could store objects created by an OO application like C++. However, the ODBMS model suffered from a fatal flaw: it had no common query language. Much of the appeal of the RDBMS stemmed from its widely adopted query language, SQL. Lacking a query language, the ODBMSs failed to deliver the flexibility that made the RDBMS architecture so successful. In addition ODBMS systems lacked many of the high level features required by corporations, including scalability, security, server-side functions, concurrency, etc.

## OBJECT-RELATIONAL: THE BEST OF BOTH WORLDS

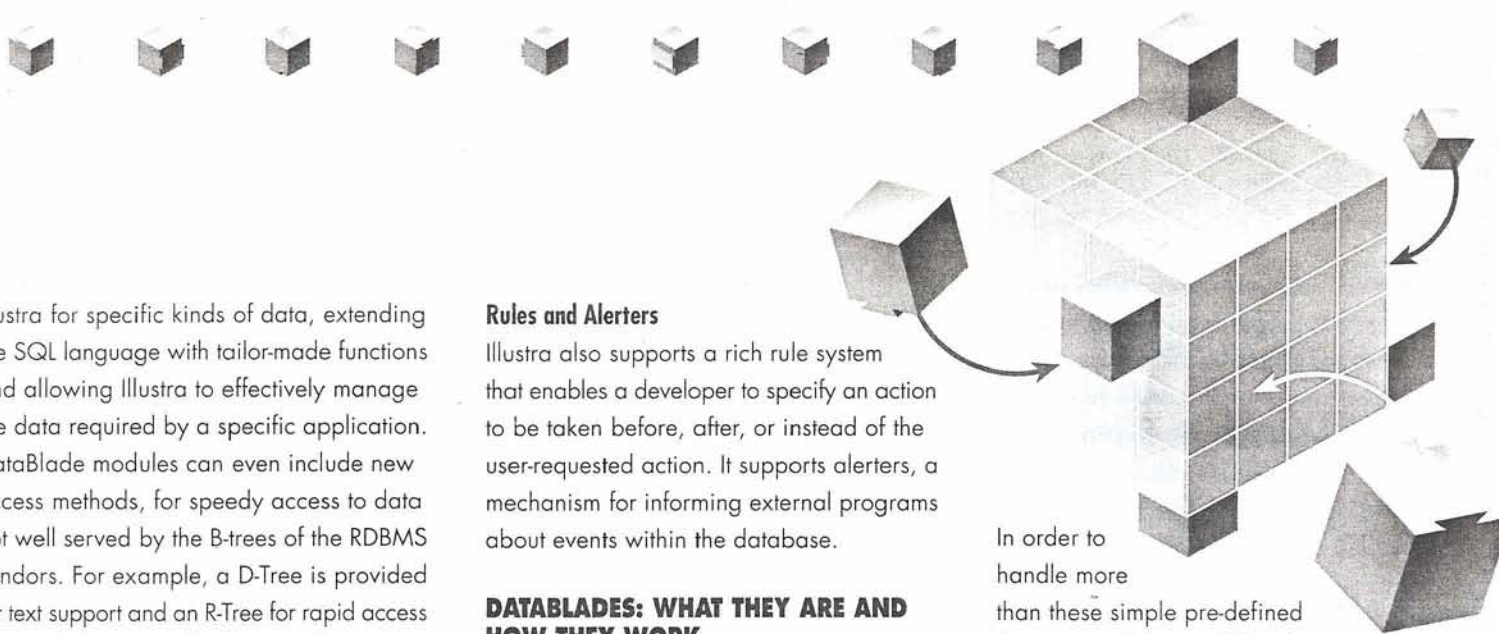
In the past, application architects were faced with an unpalatable choice: they could use either a DBMS that modeled the rich data effectively or one that provided a well-known query language, but not both. What is required is a data management system powerful enough to store and manipulate richly structured data, enforce the complex business rules that define its integrity, and is flexible enough to handle novel data structures effectively.

The Illustra Server is the world's first Object-Relational DBMS (ORDBMS). It supports Object-oriented management of rich data types, but at the same time provides an efficient query language based on extensions to industry-standard SQL. Its support for inheritance speeds application development and enhances quality. Maintenance costs are reduced because Illustra acts as a single object repository.

The essence of the problem of managing rich data types is there are so many complex and rich kinds of data to model. How can a single product hope to manage information as diverse as video images, spatial information and time-series data? Illustra does this by extending the relational database model to support new data types. In other words, the key is extensibility.

The Illustra Server™ provides a high-performance extensible platform for storage and retrieval of objects. Object extensions, called Datablade™ modules, plug intelligence into





Illustra for specific kinds of data, extending the SQL language with tailor-made functions and allowing Illustra to effectively manage the data required by a specific application. DataBlade modules can even include new access methods, for speedy access to data not well served by the B-trees of the RDBMS vendors. For example, a D-Tree is provided for text support and an R-Tree for rapid access to spatial data. DataBlades may also contain display methods to aid in the visualization of complex data.

A DataBlade embeds data intelligence into the Illustra Server, making it an intelligent repository of rich data.

#### **Illustra Server is a Relational DBMS**

It supports the key relational technologies that have proven so successful including:

- Flexible data access via standard SQL
- Standard security controls
- Full server-enforced data integrity
- Transactions and recovery
- Performance and scalability

#### **Illustra Server is an Object DBMS**

Illustra also gives relational developers access to the key OO technologies:

- Ability to create any data type
- Optimized access to rich data types
- Rapid high-quality development and low maintenance through encapsulation, inheritance, polymorphism
- Direct data access through Object IDs

#### **Rules and Alerters**

Illustra also supports a rich rule system that enables a developer to specify an action to be taken before, after, or instead of the user-requested action. It supports alerters, a mechanism for informing external programs about events within the database.

#### **DATABLADES: WHAT THEY ARE AND HOW THEY WORK**

##### **DataBlades Make BLOBs Obsolete**

Effectively managing advanced data means going beyond the limited capabilities of traditional database systems. DataBlades, a unique feature available only with Illustra, enable you to plug intelligence into your DBMS for your specific kinds of data. When fitted with a DataBlade, the Illustra system "understands" your data types, so that it can deliver fully optimized query and storage for non-traditional data.


Current RDBMS have great expertise in a limited number of data types: integers, floating point numbers, dates and character strings. Designed into the core of these traditional DBMS is a key set of capabilities for each data type:

- Content-based query capability
- Comparison operators
- Intelligent query optimization
- Efficient storage capability
- High performance
- Storage methodology (access method) appropriate to the data

In order to handle more than these simple pre-defined data types, Relational DBMS relegate rich data types to second class status, storing them as Binary Large Objects (BLOBs). Because BLOBs are uninterpreted bit patterns, the RDBMS does not know how to perform content-based queries on them. It has no sensible comparison operators, so it cannot intelligently build an optimal query plan and it cannot provide high performance storage and retrieval. In fact, the RDBMS cannot provide for BLOBs any of the features that make the RDBMS so useful for traditional data, making the database management system incapable of managing the more complex data. The hard work is left up to the individual application developer – the work of understanding the contents, format, and methodologies required of each data type – and this work needs to be re-invented and re-embedded in each application. Worse, because the RDBMS cannot understand BLOB contents, BLOBs have to be shipped across the network to the client for processing, placing a heavy burden on scarce network bandwidth.



**ILLUSTR**



DataBlades represent a completely new approach to supporting rich data, passing well beyond simple BLOB support. A DataBlade is a snap-in module that adds expertise and intelligence to the DBMS system, extending the system with the key functionality required for native support of specific data types. Just as a general purpose utility knife can be extended to perform different cutting jobs by inserting special-purpose blades, so the Illustra Server can be extended to manage new data types by snapping in the required DataBlade.

An Illustra DataBlade™ moves object orientation beyond a programming methodology to a data management strategy. DataBlades include new data types and functions, and may also include display methods, and new access methods to allow intelligent query support for the new data types.

#### **Relational Supercharged with Objects**

The Illustra Server combined with DataBlade modules supercharges the relational model, extending the SQL query language to efficiently process rich data. It provides content-based queries, optimized query plans and high performance storage and retrieval tailored to the individual data type. Illustra:

- Is the first DBMS that enables integration, movement, and management of multiple data types
- Provides the power of objects, the flexibility of relational
- Features server-enforced integrity for complex data
- Uses DataBlades that make BLOBs obsolete

**For more information about Illustra please call:**

#### **NEW YORK**

One World Trade Center, Suite 7967  
New York, NY 10048  
Tel. 212.432.3188 Fax 212.432.3155

#### **WASHINGTON, DC**

3914 Stoconga Drive, Beltsville, MD 20705  
Tel. 301.937.6245 Fax 301.937.4908

#### **UNITED KINGDOM**

150 Minorities, London, EC3N 1LS  
Tel +44.171.264.2185 Fax +44.171.264.2193

#### **CORPORATE HEADQUARTERS**

1111 Broadway, Suite 2000, Oakland, CA 94607  
Tel 510.652.8000 Fax 510.869.6388  
Email [info@illustra.com](mailto:info@illustra.com)  
URL <http://www.illustra.com>



**ILLUSTR**



































