

The Montage Extensible DataBlade™ Architecture

Michael Ubell
Montage Software Inc.
1111 Broadway, #2000
Oakland, Ca. 94607
michael@montage.com

Relational database systems have supported business applications for over a decade, providing capabilities such as optimization, concurrency, recovery and data independence. Customers now seek to integrate more areas of their business into a database framework. New applications need data management systems that handle a variety of data types, such as images, drawings, spatial information, time series and documents. These data types must be integrated with traditional data and with each other in a system that provides the same capabilities as relational systems. Relational systems were initially limited to simple data types: numbers, character strings, dates and money. Although relational database vendors have added some new data types, they lack the extensible architecture needed for an integrated solution. Often these extensions appear as wrappers with restricted capabilities and limited integration with other data types.

The Montage database system supports general extensibility of integrated data types. User definable capabilities include base and composite types, C and SQL functions, and functional indices. These features are dynamic: they are added to the system using SQL-3 data definition language. New types of access methods and storage managers are also easily added to the system. The optimizer can use new statistics and formulas to determine when to choose a new access method and the cost of executing a user defined function.

In this architecture the database management system can be thought of as a “razor” into which technology specific “DataBlades” are inserted. DataBlades are packages of data types, functions (or methods) defined on those types, and access methods customized for those types and functions. DataBlades range in specificity and can address technical areas (images), industry sectors (medical) or specific companies (Medical Widgets Corporation). DataBlades may be written by Montage, third parties, or end users.

Montage offers DataBlades for textual document, spatial, image and statistical applications. DataBlades for OCR, time series and isochronous (real-time) data are under development. The latter DataBlade will include a special storage manager to optimize the data delivery. Other ven-

dors can easily add DataBlades that exploit their own expertise.

For example, Montage offers a textual document blade that integrates a third party format conversion suite with a document data type and a full text indexing access method. Queries involving full text search may be combined in relational queries with other data types:

```
SELECT * FROM employees WHERE
    Contains(resume,
        'sql or quel and programming')
    AND dept = 'Marketing';
```

The Montage data manager implements SQL-92 with data type and function extensions from SQL3, including type and table hierarchies, multiple inheritance, object identifiers and function overloading. It also implements a general rules and alerters system. Rules may be placed on data retrieval as well as on data modification statements. Alters can be used to signal other processes when a rule has fired. Dump, restore, rollforward and copy utilities are also provided.

The Montage system runs as a client-server application. Servers are supported in various UNIX environments. Clients may be on UNIX, Windows or Macintosh systems. C-functions registered in the database may either be dynamically linked to the server or be run in the client application during development (debugging) or production (performance, security).

Applications may employ multiple DataBlades. For example an integrated insurance application combines image, document, OCR, spatial, statistical and financial data. By storing all insurance data, Montage encapsulates the knowledge of business practices to a centralized repository where it is automatically applied to all application programs. The true power of the system is in managing all the data, rules and functions for a business in a uniformly accessible, integrated extensible database.

References

- 1 - “The Miro DBMS”, M. Stonebraker, Proceedings 1993 ACM SIGMOD, Washington DC. Page 439.
- 2 - Montage User’s Guide Release 2.0, Montage Software, March 1994.

DataBlade is a trademark of Montage Software, Inc.