

IBM's DB2 Universal Database demonstrations at VLDB '98

Berni Schiefer
IBM Toronto

Jim Kleewein
IBM Santa Teresa

Karen Brannon
IBM Almaden

Guy Lohman
IBM Almaden

Gene Fuh
IBM Santa Teresa

Abstract

Today's competitive business climate dictates that companies derive more information out of their databases. Analysts looking for business trends in their company's database pose increasingly complex queries, often through query generator front-end tools. Businesses must extract as much useful information as possible from the large volumes of data that they keep, making parallel database technology a key component of such business intelligence applications. Enterprises and independent software vendors continue to require support for more application productivity and capability. And many growing enterprises have data stored in many systems, often both file systems and database systems from a variety of vendors. All of these areas contribute to high performance at low cost. Being able to access and manage this data with high performance, fast response time and low total cost of ownership is a compelling advantage in business today.

Scalability

DB2 Universal Database Enterprise - Extended Edition (DB2 UDB EEE) was designed to support the very large databases that business intelligence applications often require. Based on a highly scalable shared-nothing software architecture which exploits clusters of SMP systems, DB2 UDB EEE supports a variety of hardware including the IBM Netfinity servers running Windows NT connected with Gigaset cLAN technology that supports the VI Architecture. VI enables a high bandwidth, low latency, and fault tolerance interconnect. Using this hardware architecture and interconnect, DB2 UDB EEE has linear scalability on standard queries accessing a large decision support database.

Performance

DB2 UDB draws on 20 years of experience in cost-based SQL query optimization to provide sophisticated industry-leading SQL compiler technology to ensure that high performance is achieved for a broad class of queries and functions. DB2 provides an easy-to-use tool called Visual Explain for graphically displaying the optimizer's plan to the user for analysis and tuning. We will use Visual Explain to show DB2 UDB's ability to optimize and apply query rewrite technology to complex OLAP and data warehousing queries, including star joins, databases partitioned for SMP or MPP parallelism, CUBE queries, and even recursive SQL.

File management

IBM offers DataLinks technology to extend DB2 UDB to manage arbitrary data stored in external files. Using DataLinks, DB2 UDB provides referential integrity, access control, and coordinated backup and recovery for data stored in file systems as if the files were stored in DB2. We will demonstrate a web-based client using Java applets or servlets to access DB2 data with linked data in files stored in an external filesystem, illustrating the use of DataLinks for access control, data integrity and automated backup and recovery.

Heterogeneous Data Access

IBM has been incorporating new features and functions into DB2 UDB using object-relational features. These provide a higher level of abstraction, making applications easier to build. One type of extension is the ability to define, use, and manipulate geospatial data. DB2's DataJoiner product makes this new capability available everywhere in an enterprise by providing the ability to access data stored in a variety of databases. We will exploit geospatial capability using object-relational extensions, not only against data stored in DB2, but also data stored in Oracle and IMS (IBM's hierarchical database).

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

**Proceedings of the 24th VLDB Conference
New York, USA, 1998**