

Real World Requirements for Decision Support - Implications for RDBMS

Sanju K. Bansal
Executive Vice-President
MicroStrategy

There are a number of data manipulation requirements common to most decision support (DSS) systems; these requirements are currently being addressed through the use of sophisticated decision support engines working in conjunction with the RDBMS. Since RDBMSs are the preferred data storage layer for DSS, there is an opportunity to provide many of the most frequent data manipulation requirements within the database. Recommendations on how decision support functionality should be partitioned between the RDBMS engine and the DSS application will be presented.

The session will begin with an overview of decision support architectures, covering the role of the RDBMS, the DSS Engine, and the DSS Application as well as the differences between two tier vs. three tier decision support architectures. End-user requirements such as Multidimensional Conceptual Data Views and the need for On Line Analytical Processing (OLAP) will then be addressed, with real world examples of OLAP requirements.

Finally, requirements for interactive decision support such as Query Governing and proactive Cost Feedback Mechanisms will be covered. The speaker will conclude with thoughts on which DSS features are best implemented in the RDBMS and which DSS features are best provided by a DSS application engine in order to maximize efficiency and the power of the application.

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 ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association of Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

SIGMOD '95, San Jose, CA USA
© 1995 ACM 0-89791-731-6/95/0005..\$3.50