

1. Record Nr.	UNINA9910823892403321
Titolo	DB2 UDB's high function business intelligence in e-business // [Nagraj Alur ... et al.]
Pubbl/distr/stampa	San Jose, Calif., : IBM Corp., 2002
Edizione	[1st ed.]
Descrizione fisica	xxii, 252 p. : ill
Collana	IBM redbooks
Altri autori (Persone)	AlurNagraj
Soggetti	Electronic commerce Business intelligence Database management
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Bibliographic Level Mode of Issuance: Monograph
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Front cover -- Contents -- Figures -- Tables -- Examples -- Notices -- Trademarks -- Preface -- The team that wrote this redbook -- Notice -- Comments welcome -- Chapter 1. Business Intelligence overview -- 1.1 e-business drivers -- 1.1.1 Impact of e-business -- 1.1.2 Importance of BI -- 1.2 IBM's BI strategy and offerings -- 1.2.1 BI and analytic enhancements in DB2 UDB -- 1.2.2 Advantages of BI functionality in the database engine -- 1.3 Redbook focus -- 1.3.1 Materialized views -- 1.3.2 Statistics, analytic and OLAP functions -- Chapter 2. DB2 UDB's materialized views -- 2.1 Materialized view overview -- 2.1.1 Materialized view motivation -- 2.1.2 Materialized view concept overview -- 2.1.3 Materialized view usage considerations -- 2.1.4 Materialized view terminology -- 2.2 Materialized view CREATE considerations -- 2.2.1 Step 1: Create the materialized view -- 2.2.2 Step 2: Populate the materialized view -- 2.2.3 Step 3: Tune the materialized view -- 2.3 Materialized view maintenance considerations -- 2.3.1 Deferred refresh -- 2.3.2 Immediate refresh -- 2.4 Loading base tables (LOAD utility) -- 2.5 Materialized view ALTER considerations -- 2.6 Materialized view DROP considerations -- 2.7 Materialized view matching considerations -- 2.7.1 State considerations -- 2.7.2 Matching criteria considerations -- 2.7.3 Matching permitted -- 2.7.4 Matching inhibited -- 2.8 Materialized view design considerations -- 2.8.1 Step 1: Collect queries &

prioritize -- 2.8.2 Step 2: Generalize local predicates to GROUP BY -- 2.8.3 Step 3: Create the materialized view -- 2.8.4 Step 4: Estimate materialized view size -- 2.8.5 Step 5: Verify query routes to "empty" the materialized view -- 2.8.6 Step 6: Consolidate materialized views -- 2.8.7 Step 7: Introduce cost issues into materialized view routing -- 2.8.8 Step 8: Estimate performance gains. 2.8.9 Step 9: Load the materialized views with production data -- 2.8.10 Generalizing local predicates application example -- 2.9 Materialized view tuning considerations -- 2.10 Refresh optimization -- 2.11 Materialized view limitations -- 2.11.1 REFRESH DEFERRED and REFRESH IMMEDIATE -- 2.11.2 REFRESH IMMEDIATE and queries with staging table -- 2.12 Replicated tables in nodegroups -- Chapter 3. DB2 UDB's statistics, analytic, and OLAP functions -- 3.1 DB2 UDB's statistics, analytic, and OLAP functions -- 3.2 Statistics and analytic functions -- 3.2.1 AVG -- 3.2.2 CORRELATION -- 3.2.3 COUNT -- 3.2.4 COUNT_BIG -- 3.2.5 COVARIANCE -- 3.2.6 MAX -- 3.2.7 MIN -- 3.2.8 RAND -- 3.2.9 STDDEV -- 3.2.10 SUM -- 3.2.11 VARIANCE -- 3.2.12 Regression functions -- 3.2.13 COVAR, CORR, VAR, STDDEV, and regression examples -- 3.3 OLAP functions -- 3.3.1 Ranking, numbering and aggregation functions -- 3.3.2 GROUPING capabilities ROLLUP & CUBE -- 3.3.3 Ranking, numbering, aggregation examples -- 3.3.4 GROUPING, GROUP BY, ROLLUP and CUBE examples -- Chapter 4. Statistics, analytic, OLAP functions in business scenarios -- 4.1 Introduction -- 4.1.1 Using sample data -- 4.1.2 Sampling and aggregation example -- 4.2 Retail -- 4.2.1 Present annual sales by region and city -- 4.2.2 Provide total quarterly and cumulative sales revenues by year -- 4.2.3 List the top 5 sales persons by region this year -- 4.2.4 Compare and rank the sales results by state and country -- 4.2.5 Determine relationships between product purchases -- 4.2.6 Determine the most profitable items and where they are sold -- 4.2.7 Identify store sales revenues noticeably different from average -- 4.3 Finance -- 4.3.1 Identify the most profitable customers -- 4.3.2 Identify the profile of transactions concluded recently -- 4.3.3 Identify target groups for a campaign. 4.3.4 Evaluate effectiveness of a marketing campaign -- 4.3.5 Identify potential fraud situations for investigation -- 4.3.6 Plot monthly stock prices movement with percentage change -- 4.3.7 Plot the average weekly stock price in September -- 4.3.8 Project growth rates of Web hits for capacity planning purposes -- 4.3.9 Relate sales revenues to advertising budget expenditures -- 4.4 Sports -- 4.4.1 For a given sporting event -- 4.4.2 Seed the players at Wimbledon -- Appendix A. Introduction to statistics and analytic concepts -- A.1 Statistics and analytic concepts -- A.1.1 Variance -- A.1.2 Standard deviation -- A.1.3 Covariance -- A.1.4 Correlation -- A.1.5 Regression -- A.1.6 Hypothesis testing -- A.1.7 HAT diagonal -- A.1.8 Wilcoxon rank sum test -- A.1.9 Chi-Squared test -- A.1.10 Interpolation -- A.1.11 Extrapolation -- A.1.12 Probability -- A.1.13 Sampling -- A.1.14 Transposition -- A.1.15 Histograms -- Appendix B. Tables used in the examples -- DDL of tables -- Appendix C. Materialized view syntax elements -- Materialized view main syntax elements -- Related publications -- IBM Redbooks -- Other resources -- Referenced Web sites -- How to get IBM Redbooks -- IBM Redbooks collections -- Index -- Back cover.

Sommario/riassunto

This IBM Redbooks publication deals with exploiting DB2 UDB's materialized views (also known as ASTs/MQTs), statistics, analytic, and OLAP functions in e-business applications to achieve superior performance and scalability. This book is aimed at a target audience of DB2 UDB application developers, database administrators (DBAs), and

independent software vendors (ISVs). We provide an overview of DB2 UDB's materialized views implementation, as well as guidelines for creating and tuning them for optimal performance. We introduce key statistics, analytic, and OLAP functions, and describe their corresponding implementation in DB2 UDB with usage examples. Finally, we describe typical business level queries that can be answered using DB2 UDB's statistics, analytic, and OLAP functions. These business queries are categorized by industry, and describe the steps involved in resolving the query, with sample SQL and visualization of results.
