

1. Record Nr.	UNINA9910825287103321
Titolo	Database strategies : using Informix XPS and DB2 Universal Database / / Chuck Ballard ... [et al.]
Pubbl/distr/stampa	[San Jose, Calif., : IBM Corp., International Technical Support Organization], c2005
Edizione	[1st ed.]
Descrizione fisica	xx, 457 p. : ill
Collana	Redbooks
Soggetti	Database management Informix software
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	"This edition applies to Version 8.2 of DB2 Universal Database (UDB) and Version 8.5 of Informix Extended Parallel Server (XPS)." "August 2005."
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Front cover -- Contents -- Notices -- Trademarks -- Preface -- The team that wrote this redbook -- Become a published author -- Comments welcome -- Chapter 1. Introduction to this redbook -- 1.1 Understanding strategic directions for XPS -- 1.2 Objective of this redbook -- 1.3 Chapter abstracts -- Chapter 2. XPS and DB2 UDB architectures -- 2.1 High-level product overviews -- 2.1.1 IBM Informix Extended Parallel Server -- 2.1.2 DB2 Data Warehouse Edition -- 2.2 Understanding the architectures -- 2.3 Defining an instance -- 2.3.1 Informix XPS instance architecture -- 2.3.2 DB2 Universal Database instance architecture -- 2.4 Storage architecture -- 2.4.1 Pages -- 2.4.2 Containers and chunks -- 2.4.3 Logical disks -- 2.4.4 Logging -- 2.4.5 Storage architecture summary -- 2.5 Parallelism -- 2.5.1 The process model of XPS -- 2.5.2 The process model of DB2 -- 2.5.3 Intra-node parallelism -- 2.5.4 Inter-node parallelism -- 2.6 Memory management -- 2.6.1 XPS memory model -- 2.6.2 DB2 memory model -- 2.7 Partitioning -- 2.7.1 Fragmentation in XPS -- 2.7.2 Partitioning in DB2 -- 2.8 Terminology -- Chapter 3. Configuration -- 3.1 XPS and DB2 configuration -- 3.1.1 Knobs (configuration files and tuning parameters) -- 3.1.2 Commands -- 3.1.3 Granularity -- 3.1.4 Database manager -- 3.1.5 Dynamic

parameters -- 3.1.6 Cataloging -- 3.1.7 Client access to DB2 instances -- 3.2 Configuration methods -- 3.2.1 DB2 configuration methods -- 3.2.2 Configuration Advisor and the AUTOCONFIGURE command -- 3.3 Configuration files and objects overview -- 3.3.1 Environment variables and the profile registry -- 3.3.2 Setting registry and environment variables -- 3.3.3 DB2 configuration files and objects -- 3.4 Configuring the instance -- 3.4.1 Page size(s) -- 3.4.2 Table spaces -- 3.4.3 Bufferpools -- 3.4.4 Physical and Logical Logs.

Chapter 4. Instance and database operations -- 4.1 Instance operation modes -- 4.1.1 Online mode -- 4.1.2 Offline mode -- 4.1.3 Quiescent mode -- 4.1.4 Creating and dropping the instance -- 4.2 Modifying the configuration -- 4.2.1 Working with the DAS -- 4.2.2 Viewing or updating the configuration using Control Center -- 4.2.3 Managing database partition groups -- 4.2.4 Managing buffer pools -- 4.3 Managing database storage -- 4.3.1 Table spaces and containers -- 4.3.2 Monitoring table space and container storage -- 4.3.3 Transactions and logs -- 4.4 Backup and recovery -- 4.4.1 Recovery types -- 4.4.2 Backup and restore methods -- 4.4.3 Table level restore -- 4.5 High availability -- 4.5.1 Log mirroring -- 4.5.2 Replication -- 4.5.3 Online split mirror and suspended I/O support -- 4.6 Security -- 4.6.1 Authorization and privileges -- 4.6.2 Roles and groups -- 4.6.3 Security levels -- 4.6.4 Client/server security -- 4.6.5 Authentication methods -- Chapter 5. Data types -- 5.1 Object names -- 5.2 Data type mapping -- 5.3 NULL values -- 5.4 Disk considerations -- 5.5 Character types -- 5.5.1 Truncation -- 5.5.2 NCHAR data type -- 5.5.3 VARCHAR data type -- 5.5.4 TEXT data type -- 5.6 Numerical data types -- 5.6.1 Numerical limits -- 5.7 DECIMAL -- 5.7.1 MONEY data type -- 5.7.2 SERIAL and SERIAL8 -- 5.8 Date and time types -- 5.8.1 DATE data type -- 5.8.2 DATETIME, TIME, and TIMESTAMP data types -- 5.8.3 INTERVAL data type -- 5.9 FLOAT -- 5.10 REAL or SMALLFLOAT -- 5.11 LOB data types -- 5.12 Sequence objects -- 5.13 Other object limits in DB2 -- 5.14 DB2 manuals -- Chapter 6. Data partitioning and access methods -- 6.1 Benefits of data partitioning -- 6.2 Hash fragmentation -- 6.3 Round robin fragmentation -- 6.4 Expression and range fragmentation -- 6.5 Hybrid fragmentation -- 6.6 Range partitioning using MDC -- 6.6.1 Benefits of MDC. 6.6.2 Design considerations for MDC tables -- 6.6.3 Operations on MDC tables -- 6.6.4 Space requirement for MDC -- 6.7 Range-clustered tables in DB2 -- 6.8 Roll-in and roll-out of data using UNION ALL views -- 6.8.1 Query optimization of UNION ALL views -- 6.8.2 Benefits of UNION ALL views -- 6.8.3 Limitations of UNION ALL views -- 6.9 MDC and UNION ALL views for roll-in and roll-out -- 6.10 Indexing strategies -- 6.10.1 Syntax for index creation -- 6.10.2 DB2 index expansions -- 6.10.3 Index types and access methods -- 6.10.4 Space requirements for indexes -- 6.10.5 Table and Index reorganization on DB2 -- 6.11 Joins -- 6.11.1 Join syntax -- 6.11.2 Join methods (generic) -- 6.11.3 Join strategies in a partitioned database -- 6.11.4 MERGE, UPDATE, and DELETE joins -- 6.12 Optimizer -- 6.12.1 The role of query optimizer -- 6.12.2 LEO: Learning Optimizer -- 6.12.3 Push-down hash join -- 6.12.4 Optimization strategies for intra-partition parallelism -- 6.12.5 Directives -- 6.12.6 Optimization classes -- 6.13 Performance enhancements in DB2 UDB V8.1 -- 6.13.1 Distributed catalog cache -- 6.13.2 Prefetch -- 6.13.3 Page cleaner I/O improvements -- 6.13.4 Multi-threading of Java-based routines -- 6.13.5 Join variations -- 6.13.6 Increased opportunity for selection of bit-filters -- 6.13.7 Informational constraints -- 6.13.8 Uniform page size -- Chapter 7. SQL considerations -- 7.1 SELECT issues -- 7.1.1 Selectivity -- 7.1.2

Statistical sampling -- 7.1.3 SELECT cursors -- 7.1.4 Joins -- 7.2 MATCHES predicate -- 7.3 Comments -- 7.4 SQLCODE and SQLSTATE -- 7.5 Built-in functions -- 7.6 SQL access to system catalogs -- 7.7 Quotations and character strings -- 7.8 Concatenation behavior -- 7.9 Implicit casting -- 7.10 Deferred constraint checking -- 7.11 Set Operators: UNION, INTERSECT, and MINUS -- 7.12 Multi-database access -- 7.13 Temporary tables.

7.13.1 Implicit -- 7.13.2 Explicit -- 7.14 Compound SQL -- 7.15 INSERT cursors -- 7.16 MERGE INTO -- 7.17 Online analytical processing SQL -- 7.18 Isolation levels -- 7.19 Optimizer directives -- 7.20 DDL issues -- 7.20.1 Creating and altering tables -- 7.20.2 Synonyms -- 7.20.3 Primary key definitions -- 7.20.4 Constraint naming -- 7.21 Triggers -- 7.21.1 SELECT triggers -- 7.21.2 BEFORE-statement triggers -- 7.21.3 Disabling triggers -- 7.22 Multidimensional Clustering in DB2 -- 7.23 DB2 Materialized Query Tables -- 7.23.1 Using and Configuring MQTs -- 7.24 System commands -- 7.24.1 CREATE DATABASE -- 7.24.2 Administrative commands -- 7.25 Statistics for table and indexes -- 7.26 Query Monitoring -- Chapter 8. Loading and unloading data -- 8.1 Loading and inserting data in a single stream -- 8.2 Parallel bulk loading -- 8.2.1 Handling bad rows -- 8.2.2 Performance and tuning considerations for loading with DB2 -- 8.3 Parallel unloading -- 8.3.1 XPS unloading -- 8.3.2 DB2 unloading -- 8.3.3 Parallel exports -- 8.4 Specific issues -- Chapter 9. Administration tools and utilities -- 9.1 Resource management -- 9.2 Performance tuning -- 9.3 Tools and wizards that are included with DB2 -- 9.3.1 Control Center -- 9.3.2 Command Editor -- 9.3.3 Task Center -- 9.3.4 SQL Assist -- 9.3.5 Visual Explain -- 9.3.6 Configuration Assistant -- 9.3.7 Journal -- 9.3.8 Health Center -- 9.3.9 Replication Center -- 9.3.10 License Center -- 9.3.11 Information Catalog Center -- 9.3.12 Data Warehouse Center -- 9.3.13 Web administration -- 9.3.14 Wizards, advisors, and launchpads -- 9.4 Optional tools -- 9.4.1 DB2 Performance Expert -- 9.4.2 DB2 Recovery Expert -- 9.4.3 DB2 High Performance Unload -- 9.4.4 DB2 Test Database Generator -- 9.4.5 DB2 Table Editor -- 9.4.6 DB2 Web Query Tool -- 9.4.7 Query Patroller -- 9.5 Utilities -- 9.5.1 Database reorganization. 9.5.2 Database statistics -- 9.5.3 Schema extraction -- 9.5.4 Maintaining database integrity -- 9.5.5 Throttling utilities -- 9.5.6 Validating a backup -- 9.6 Other administrative operations -- 9.6.1 Configuring automatic maintenance -- 9.7 Monitoring tools and advisors -- 9.7.1 Health check tools -- 9.7.2 Memory Visualizer -- 9.7.3 Storage Manager -- 9.7.4 Event monitor -- 9.7.5 Snapshots -- 9.7.6 Activity Monitor -- 9.7.7 DB2 Performance Expert -- 9.7.8 The db2pd utility, an onstat equivalent -- 9.7.9 Diagnostic files -- 9.7.10 Error message and command help -- Chapter 10. Planning the transition -- 10.1 Tasks and activities -- 10.1.1 Readiness assessment and scope -- 10.1.2 Tool evaluation -- 10.1.3 Estimating project duration -- 10.2 Data conversion -- 10.2.1 Preparation overview -- 10.2.2 Data conversion process -- 10.2.3 Time planning -- 10.2.4 The database structure -- 10.2.5 Data movement approaches -- 10.2.6 WebSphere Information Integrator -- 10.2.7 Modifying the application -- 10.2.8 Database objects and interfaces -- 10.3 After the transition -- Chapter 11. Application conversion considerations -- 11.1 Key considerations -- 11.2 Application transitioning from XPS to DB2 -- 11.3 Transactions -- 11.4 Savepoints -- 11.5 Locks and isolation levels -- 11.5.1 Lock escalation -- 11.5.2 Deadlocks -- 11.5.3 Isolation levels -- 11.6 Packages -- 11.6.1 Static versus Dynamic SQL -- 11.6.2 Binding -- 11.7 Cursors -- 11.8 Stored procedures -- 11.9

Programming languages -- 11.9.1 ESQL/C -- 11.9.2 JDBC -- 11.9.3 ODBC/CLI -- 11.9.4 C++ -- 11.9.5 Large objects -- 11.9.6 SQL Communications Area -- 11.9.7 SQLDA -- Chapter 12. DB2 Migration ToolKit for Informix -- 12.1 Features and functionality -- 12.2 Recommendations for Use -- 12.2.1 MTK installation and configuration -- 12.2.2 MTK Configurations -- 12.3 Technical overview of MTK -- 12.3.1 The MTK GUI. 12.3.2 The migration process.

Sommario/riassunto

The acquisition of Informix by IBM has provided the opportunity for Informix customers to consider new alternatives to further enrich their data management systems infrastructure. They can now more easily take advantage of available products, services, and capabilities as they grow and change. This IBM Redbooks publication focuses on strategies, techniques, capabilities, and considerations for using Informix Extended Parallel Server (XPS) and DB2 Universal Database (UDB). It provides detailed discussions and data to give a good understanding of the two products, their capabilities, and their similarities. XPS customers can choose to adopt a database strategy of coexistence or consider transitioning to DB2 UDB. The features and functionality of each DBMS are briefly described for a better understanding, in areas such as architecture, partitioning techniques, SQL considerations, configuration, indexing, data types, DML, and DDL. It also discusses products and tools to complement these database management systems. With this information, you can better decide which products satisfy your particular requirements, and better plan on how to achieve your objectives as you develop your database management system strategy. You will be better positioned to make informed decisions that can give you the best return on your DBMS investment.
