Oracle Test-to-Scale Program: Data Warehousing

Dramatic Performance Improvements in Oracle7 releases

An Oracle White Paper
October 1995

ORACLE®
Enabling the Information Age

Part #: C10240

Oracle7 Server: Scalable Parallel Architecture for Open Data Warehousing

Part#: C10240

October 1995

Contributing Authors: Andy Mendelsohn, Ray Roccaforte, Steve Faris, Mike Hartstein

Oracle Test-to-Scale Program: Data Warehousing

Dramatic Performance Improvements in Oracle7 releases

INTRODUCTION

This report outlines the results of a set of 'tests-to-scale' performed as part of Oracle's on-going efforts at characterizing the performance and scalability of the Oracle7 server in complex, data-intensive operations typical of decision support applications. The tests described here model a real-world business analysis situation involving drill-down analysis and complex queries against a large 100 gigabyte dataset and demonstrate the performance and scalability improvements engineered into successive releases of the Oracle7 server. These tests illustrate improvements in three important areas: *parallel index creation, complex queries* and *drill-down analysis*. Oracle7 Release 7.2 (available now) delivers a two-fold improvement over Release 7.1 in creating large indexes and running complex queries. Drill-down analysis is nearly 10 times faster with Release 7.2. Oracle7 Release 7.3 delivers three-fold indexing, five-fold complex query and 20-fold drill-down performance gains over Release 7.1 in these tests.

REAL-WORLD BUSINESS SCENARIO

The tests outlined in this report model a real-world business scenario typical in sales analysis problems. Sales Executives and Brand Managers often perform an in-depth examination of regional variations in sales, attempting to identify laggard brands that performed very poorly during a given period. A typical question that characterizes such analysis is : "What was the worst performing product during the year 1994?". Identification of such poor performers is followed by a detailed analysis of the given product-market, looking into details such as geography, pricing, distribution, sales promotion and packaging, in an attempt at identifying factors that resulted in the poor performance. In other words, this phase deals with the question "Why did this product do poorly?". This analysis eventually feeds into the formulation of product strategy aimed at correcting the problem. Typical characteristics of such analysis that are of interest to us are:

- Data-intensive: The data being examined is typically very large—in the tens or hundreds of gigabytes or sometimes in the terabyte range. The tests described here were run against a 100GB dataset. A 10GB dataset was also used as a reference, primarily to demonstrate the scaling.

- Successive refinement: The analysis is often performed at different levels of summarization, beginning with highly aggregated data, and successively reducing the level of aggregation as more patterns become evident from the data. The factors identified from summary data are typically used as constraints to restrict the amount of data that needs to be examined, as the analyst proceeds to look at more details.

DATA DESCRIPTION

The tests use a simplified order entry schema involving the following tables: LINEITEM, ORDER, PART, CUSTOMER, SUPPLIER, NATION. The tables describe customers who place orders and suppliers who provide the necessary parts from different nations in the world. The tables were populated with randomly generated data. We chose to introduce variations and skew in the generated data in order to facilitate drill-down analysis to identify anomalies.

The total size of the tables was approximately 100GB. The entire database, including indexes and temporary tablespaces, was over 200GB in size. Here is an indication of the table sizes used in our tests:

Table	Rows
LINEITEM	600 M
ORDER	150 M
CUSTOMER	15 M
SUPPLIER	1 M
PART	20 M
NATION	25

The LINEITEM and ORDER tables contain most of the table data and account for over 90% of the 100GB data. NATION is a very small look-up table.

TEST ENVIRONMENT

The tests were performed using an 18-processor Sequent Symmetry machine.

Processor: 66 MHz Pentium
No. of CPUs: 18
Total Disk capacity: 300 GB (150 x 2GB)

DATA ANALYSIS METHODOLOGY

Data analysis was performed using a drill-down approach commonly employed in real-world decision support applications.

- At the first level, we start with a highly summarized 'data cube' that contains sales figures for each month during the given year (say 1994).

An examination of this data reveals the month with the lowest sales. In our example, we may have identified May as that month.

- At the next level, we add a further dimension to the data cube, to include regional sales data. A query of this summary tells us the nation that had the worst sales during May 1994. We may now have (May, France) as the month-market combination in our example.
- At the next level, we expand the cube further, adding data on products. Analysis of this data reveals the product-market-month combination that performed the worst during 1994. Let's say we end up with (brand 43, France, May) as that combination.
- We continue along the same lines, adding further details in terms of package type, and supplier nation to successively reduce the level of data aggregation and examine larger subsets of the data.
- Finally, we query the entire 100GB dataset, using factors identified in earlier steps as constraints, and extract details for the (product, market, month, container, supplier_nation) combination with the worst sales volume during 1994. This is typically followed in real-life by further analysis of the fact data or a qualitative examination of the business issues, to provide insights for strategy formulation.

TEST RESULTS

The tests were performed using maintenance releases 7.1, 7.2 and 7.3 of the Oracle7 server as a demonstration of the performance and scalability improvements engineered into these releases. Due to time constraints, all of the tests could not be run on releases 7.1 and 7.2. Some estimations and extrapolations based on available data were used to fill out the performance matrix.

The test results illustrate dramatic performance gains for all the major operations involved—summary creation, index builds and complex queries. Release 7.2 provides a ten-fold improvement in summary creation performance over release 7.1. Index builds and complex query execution are about twice as fast. These performance gains are further extended with release 7.3: summary creation exhibits an amazing 20-fold improvement over release 7.1. Complex query execution is 5 times faster and index creation delivers over a 3-fold improvement.

We also provide a 'size-up' comparison for release 7.3, contrasting the 100GB performance numbers against a 10GB dataset as the reference point. These results illustrate Oracle7's ability to deliver scalable performance with substantial growth in data volumes.

Performance Comparison: Release 7.1 vs. 7.2 vs. 7.3

Time (minutes)

Operation	Release 7.1	Release 7.2 over 7.1	Improvement	Release 7.3	Improvement over 7.1
Summary Creation (6-D Cube)	13,400.0**	1320.8	1015%	647.4	2070%
Index Creation Single key Concatenated key	 967.7 1719.6	 398.6 620.0	 243% 277%	 304.3 389.9	 318% 441%
Complex Query	149.1	75.4	198%	29.9	498%

** *This number is an estimate. The operation was not run to completion.*

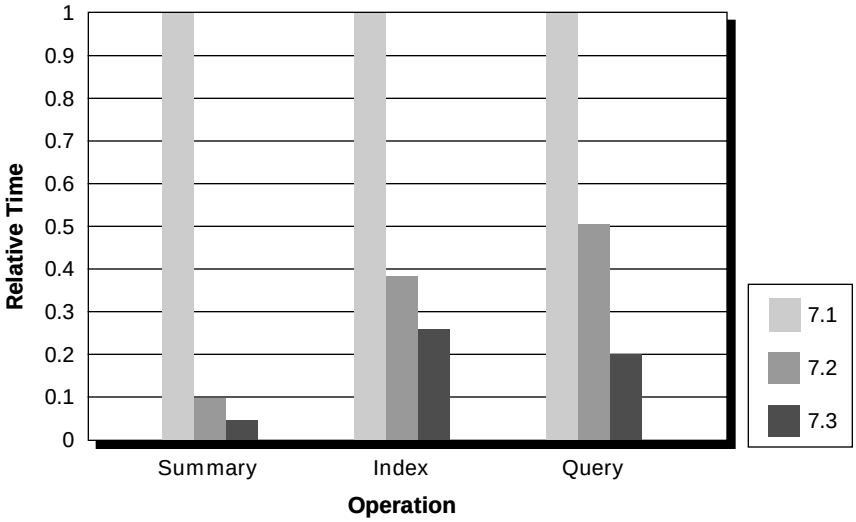

Release 7.3 Tests: 10GB vs. 100GB table data

Time (seconds)

	10 GB (T_{10})	100 GB (T_{100})	Size-up Factor
Summary Creation			
T_123456 (6-D cube)	3722.82	38846.21	95.65%
Complex Query	192.11	1796.25	106.50%
Index Creation			
Single key	1606.42	18259.76	86.33%
Concatenated key	1951.99	23396.38	80.14%
ANALYZE TABLE			
lineitem	1932.48	3501.8	181.88%
order	585.31	1616.32	172.39%
ANALYZE INDEX			
Single key	712.93	6089.76	114.58%
Concatenated key	1252.75	12683.08	98.76%

Note: Size-up Factor is defined as $[1 - (T_{100} - 10 * T_{10}) / (10 * T_{10})] * 100\%$

PERFORMANCE IMPROVEMENTS IN ORACLE7 RELEASES

In this section, we provide an outline of the performance improvements engineered into successive releases of the Oracle7 server, with specific emphasis on the enhancements that play a role in our tests.

Release 7.1: Performance Features

Parallel query execution technology was introduced in Oracle7 Release 7.1 to serve as the basis for enabling data-intensive decision support applications on cost-effective open systems. Parallel execution capabilities were architected as core, internal facilities designed to achieve highly scalable performance on all parallel hardware architectures—Symmetric Multi-processor systems (SMP), clusters and Massively Parallel systems (MPP). The initial release provided support for parallel execution of most operations involved in query execution including table scans, sorts, joins, aggregations and ordering. In addition, this release included parallel execution of data loads, index creation and recovery operations. A direct read mechanism was included in patch release 7.1.5 to enable data-intensive parallel table scan operations to bypass the buffer cache, resulting in improved performance for all parallel query operations. This facility effectively eliminates contention for the buffer cache between concurrent DSS queries and isolates DSS queries from concurrent transaction processing (OLTP) activities on the system. The matrix below indicates the operations in our tests that benefit from the performance features in release 7.1:

	Operation		
Performance Feature	Index Creation	Complex Query	Drill-down
Parallel Query operations: scan, join, sort, Group By, Distinct, Order By		✓	✓
Parallel index creation	✓		
Direct Reads (7.1.5)	✓	✓	✓

Release 7.2: Performance Features

Release 7.2 includes further enhancements to achieve improved scalability and performance for common data warehouse operations. Parallel execution of the CREATE TABLE ... AS SELECT operation provides a fast, scalable means to efficiently create intermediate 'rollup' tables. In addition to executing the underlying SELECT query in parallel, release 7.2 populates the new table in parallel, achieving substantial performance gains. Further, table and index creation operations use efficient direct reads and asynchronous, direct writes in parallel to provide significant gains in throughput. A new UNRECOVERABLE option further improves the performance and scalability of these operations by disabling the generation and logging of redo information.

The performance of sorts is dramatically improved in release 7.2 using a direct write mechanism for transferring intermediate sort runs to temporary segments on disk. All operations that use sorts–sort/merge joins, index creation, ordering, aggregation, ANALYZE–benefit from this improvement.

The cost-based optimizer in Release 7.2 supports faster execution of star queries commonly used in decision support applications, delivering performance competitive with specialized query processing systems from niche vendors. Release 7.2 also incorporates several other improvements in query optimization.

	Operation		
Performance Feature	Index Creation	Complex Query	Drill-down
Direct writes for sorts	✓	✓	✓
Parallel table creation			✓
Direct reads	✓	✓	✓
Index direct write	✓		
No-logging option for Create Table, Create Index	✓		✓
Faster execution of Star queries*		✓	

**Does not a play a role in our tests*

Release 7.3: Performance Features

Release 7.3 further extends the performance and scalability of the Oracle7 server in common data warehouse operations with significant improvements in query execution and optimization. Release 7.3 supports parallel execution of UNION and UNION ALL operations frequently used in decision support applications. There is a new join algorithm, hash join, that significantly outperforms current join methods in many complex queries and is ideally suited for scalable parallel execution. The cost-based optimizer incorporates the intelligence to choose the appropriate join method for a given query. In addition, application developers can also explicitly specify the use of hash join through a new hint. Parallel table scans employ an asynchronous read-ahead mechanism for queries involving large tables. Using this mechanism, additional data is asynchronously retrieved from disk while the current set is being processed, achieving an overlap between I/O and processing operations and delivering much higher throughput.

Release 7.3 improves the performance of parallel query execution on multi-computer systems based on 'shared-nothing' hardware architectures by exploiting the affinity of disks to individual processing nodes. Using this knowledge, the system optimally allocates work to query slaves located on nodes where the data resides and minimizes data movement, achieving substantial performance gains. Space management for sort operations is performed more efficiently using TEMPORARY tablespaces designated exclusively for sorts.

The cost-based optimizer utilizes data value histograms to deliver improved query optimization in the presence of data skew. The optimizer incorporates parallel execution considerations as a fundamental component, producing query execution plans optimized for parallel execution. Release 7.3 includes further performance improvements in the processing of Star queries. The optimizer also incorporates numerous other enhancements to achieve significant gains in the processing of operations such as 'anti-joins' involving the NOT IN condition, and the evaluation of OR conditions.

Release 7.3 incorporates bit-mapped indexing as an integrated server capability, delivering the dramatic performance benefits of this technology to decision support applications, without any compromises in server functionality. Bit-mapped indexes will coexist with and complement other available indexing schemes.

	Operation		
Performance Feature	Index Creation	Complex Query	Drill-down
Hash join		✓	
Asynch. read-ahead	✓	✓	✓
Optimizer histograms		✓	
Parallel Union, Not In*		✓	
MPP device affinity*	✓	✓	✓
STAR query improvements*		✓	✓
Bit-mapped indexing*		✓	
Improved space mgmt. for sorts	✓		✓

**Does not a play a role in our tests*

SUMMARY

Results from the tests-to-scale presented in this report reaffirm the superior performance and scalability of the industry-leading Oracle7 server in enterprise data warehouse environments. The tests offer substantial proof of the dramatic performance improvements Oracle customers can realize as they move from Oracle7 Release 7.1 to Releases 7.2 and 7.3.

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
USA

Worldwide Inquiries:
415.506.7000
Fax 415.506.7200

Part #: C10240

Printed on recycled paper.