

Oracle Partitioning

18^c ORACLE[®]
Database

KEY BENEFITS

- Faster query performance without changing application code
- Unlimited pruning with zone maps on Engineered Systems
- Efficient data management through partition-level administrative operations
- Reduced storage costs using tiered storage and compression
- Higher availability of partitioned tables and indexes

Partitioning enables large database tables and indexes to be subdivided into smaller objects that can be managed and accessed at a finer level of granularity. Oracle Partitioning enhances the manageability, performance, and availability of large databases. It offers a comprehensive choice of partitioning methods to address different business requirements and can be applied to OLTP and Data Warehousing applications without having to change any code. Zone maps, available on Oracle Engineered Systems, furthermore allow breaking down tables and partitions in even smaller zones for unlimited data pruning. Oracle's comprehensive data management capabilities complement Oracle Database 18c storage management and compression features to help customers reduce storage costs.

Improving Database Manageability, Performance and Availability

Now in its 12th generation with Oracle Database 18c, Oracle Partitioning is widely proven to enhance the manageability, performance and availability of multi-terabyte OLTP and Data Warehousing systems. It offers a comprehensive choice of partitioning methods (e.g. range, interval, hash, list and composite) to determine how the data is placed into partitions.

Partition for Manageability

Oracle Partitioning enables database administrators to take a "divide and conquer" approach to data management, and perform maintenance operations such as indexing, loading data, compressing data, and purging data on a per-partition basis. For example, a database administrator could merge and compress multiple older partitions in a single operation, moving data onto a low cost storage tier without disrupting access to data. By partitioning large table and indexes into smaller more easily managed chunks, such maintenance operations are applied to the relevant database objects (i.e. partitions) only.

With ever increasing data volumes and decreasing maintenance windows towards 24x7 uninterrupted business operations, transparent and efficient data maintenance operations in complete online fashion are becoming equally crucial and important than achieving extreme performance.

Partition for Performance

Oracle Partitioning address the challenge of performance degradation when faced with growing data volumes by limiting the amount of data to be examined or operated on, thus

RELATED PRODUCTS

Oracle Partitioning is included in Oracle Database Cloud Service High Performance Edition and Extreme Performance Edition and also available as option for Enterprise Edition on-premise

Oracle Partitioning complements many related products including:

- Database In-Memory
- Advanced Compression
- Real Application Cluster
- Oracle Engineered Systems

significantly improving performance beyond what is possible with a non-partitioned table. It fully complements Oracle Database 18c performance features, and is used in conjunction with any indexing technique, join technique, or parallel access method. Plus, partitioning is implemented at the database level and doesn't require any changes to application code or query statements in order to easily take advantage of performance benefits, including:

Partitioning Pruning to automatically use partitioning metadata and only touch data of relevance for a SQL operation. This ensures that the database only accesses relevant partitions instead of accessing complete tables or indexes to satisfy certain queries.

Zone Maps to enable pruning of "zones of blocks" and full partitions for any column defined in a zone map, regardless of whether the column is included in the table's partitioning key. Partitioning, together with zone maps, provides the most comprehensive data pruning on the market. Zone maps are only available on Oracle Engineered Systems.

Partition-wise Joins that can be applied with two tables being joined together and both tables are partitioned on the join key. This automatically breaks large joins into smaller joins that occur between each of the partitions, completing the overall join in less time and using less resources. This offers significant performance benefits for both serial and parallel query execution.

Together with Database In-Memory, the autonomy of individual partitions allows to populate only the most business- and performance-critical data as In-Memory Column store while less critical data still resides on lower cost flash or disk. Thus, even the largest data warehouse can see considerable performance benefits from In-Memory without being constrained by the available memory.

Partition for Availability

Partitioned database objects provide partition independence, which is an important part of a high-availability strategy. For example, if one partition in a table is unavailable, all of the other partitions of the table remain online and accessible. Applications can continue to execute queries and transactions against this partitioned table, and these database operations that do not need to access the unavailable partition will run successfully. Moreover, partitioning can reduce scheduled downtime by enabling database administrators to perform maintenance operations on large database objects in relatively short time windows.

Information Lifecycle Management with Oracle Partitioning

IT organizations challenged with storing vast quantities of data for the lowest possible cost can be optimally addressed using Oracle Partitioning. By understanding of how data is accessed, the independence of individual partitions is a key enabler for implementing a tiered storage strategy. Specifically in tables containing historical data, partitioning enables individual partitions (or groups of partitions) to be stored on different storage tiers, providing different physical attributes, such as compression, and price points. For example an Orders table containing 2 years worth of data could have only the most recent quarter being stored on an expensive high-end storage tier and keep the rest of the table (almost 90% of the data) on an inexpensive low cost storage tier stored as read only partitions. Through Oracle Partitioning, the storage costs are reduced by factors (cost savings of 50% or more are common), without impacting the end user access, thus optimizing the cost of ownership for the stored information. Together with Automatic Data Optimization and Heat Map – functionality in the Advanced Compression option – the process of data usage analysis and maintenance of storage tiering and compression can be fully automated without the need of manual

intervention or coding.

Oracle Partitioning for all Database Applications

Partitioning large tables and indexes in Oracle Database 18c can help ensure that an optimal method is available for every business requirement, and can enhance the manageability, performance, and availability of almost any database application. It is a very straightforward process to implement partitioning and can be applied to any transaction processing and data warehousing application to ensure success - both in terms of business user's performance and availability expectations, and in terms of reducing IT capital and operational costs.



CONTACT US

For more information about Oracle Partitioning, visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.

CONNECT WITH US



Hardware and Software, Engineered to Work Together

Copyright © 2016, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0317

