

# RDF Semantic Graph Prerequisites, and Advanced Performance and Scalability for Semantic Web Applications

ORACLE WHITE PAPER | SEPTEMBER 2014



ORACLE®

## Introduction

RDF Semantic Graph, a feature of the Oracle Spatial and Graph option for Oracle Database Enterprise Edition, delivers an advanced semantic data management capability not found in any other commercial or open source triple store. RDF Semantic Graph exploits key performance and scalability features of Oracle Database and Oracle Exadata Database Machine to address the most demanding enterprise-class semantic web solutions. For this reason, users of the RDF/OWL features of Oracle Spatial and Graph are required to license Oracle Database Enterprise Edition and the Partitioning option.

## Partitioning

RDF Semantic Graph uses Oracle Partitioning, an option for Oracle Database Enterprise Edition to enhance the manageability, performance, and availability of RDF graph applications. Partitioning allows graphs and their indexes to be subdivided into smaller pieces, enabling these database objects to be managed and accessed at a finer level of granularity. RDF graphs and indexes are range partitioned in Oracle Database.

Partitioning offers significant performance, scalability, and manageability benefits, including the following:

The white paper templates contain detailed instructions to help content owners accurately use the template. The instructions must be read to ensure the format is maintained.

- » Reduced response times for long-running queries; partitioning reduces disk I/O operations.
- » Reduced response times for concurrent queries; I/O operations run concurrently on each partition.
- » Easier index maintenance, because of partition-level create and rebuild operations.
- » Ability to rebuild indexes on partitions without affecting the queries on other partitions.
- » Ability to change storage parameters for each local index independent of other partitions. Partitions can also be split, merged, and exchanged.

## Compression

The RDF Semantic Graph transparently and automatically uses the table compression feature of Oracle Database. This is known as `COMPRESS BASIC`. It compresses data by eliminating duplicate values in a database block. All database features and functions that work on regular database blocks also work on compressed database blocks.

Some key benefits that compression brings to Semantic Web solutions include:

- » Up to a sixty percent reduction in storage costs. These savings also extend to test, development, backup and disaster recovery environments, further magnifying the cost savings.
- » Improved query performance, as queries read the compressed version of data directly without having to decompress it. Compression improves disk scan rate and reduction in the number of I/Os.
- » Enhanced memory efficiency, as data in memory is in a compressed format. This allows more data to be stored in memory and reduces the amount of I/O, which can improve performance.

## Real Application Clusters (RAC)

RDF Semantic Graph is compatible with Oracle Real Application Clusters (Oracle RAC), an option for Oracle Database Enterprise Edition. Oracle RAC allows the transparent deployment of a single graph database across a cluster of servers, providing fault tolerance from hardware failures or planned outages, performance, and scalability with no application changes necessary.

Oracle RAC provides high availability for applications by removing the single server point of failure. If a node in the cluster fails, the Oracle Database continues running on the remaining nodes. Individual nodes can be shutdown for maintenance while application users continue to work. Oracle RAC provides flexibility for scaling applications. To keep costs low, clusters can be built from standardized, commodity-priced processing, storage, and network components. When more processing power is needed, another server can be added without taking users offline to gain horizontal scalability.

## Oracle Exadata Database Machine

RDF Semantic Graph supports Oracle Exadata Database Machine. RDF graph querying, inferencing and loading can benefit from running on an Oracle Exadata Database Machine.


- » Oracle Exadata provides a balanced system that prevents graph applications from becoming I/O bound. Graph applications can take advantage of extra-large SGA, Smart Flash Cache, Smart scan, high performance I/O, and InfiniBand fabrics to increase the buffer cache and improve performance.
- » Graph applications can specify a degree of parallelism on Oracle Exadata to automatically break a complex graph processing task into smaller tasks and distribute them to multiple compute nodes.
- » Graph queries can run faster using Oracle Exadata OLTP index compression, which improves the compression ratio and memory residence for B-tree indexes used to query a graph.
- » Inferencing can be accelerated using Oracle Exadata Hybrid Columnar Compression (HCC), especially when large data sets or large rules sets or both are involved. HCC high compression ratios increase memory residence and therefore the accessibility of data and rules being processed.



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

Worldwide Inquiries  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

CONNECT WITH US



**Hardware and Software, Engineered to Work Together**

Copyright © 2014, 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. 0914