

# ORACLE DATABASE SEMANTIC TECHNOLOGIES

## SEMANTIC DATABASE FEATURES

### NEW IN ORACLE SPATIAL 11.2

- Triple-level security with Virtual Private Database and Label Security support
- Semantic indexing for documents based on popular natural language tools
- Change management to version triples within an RDF graph
- Parallel and incremental inference, and owl:sameAs optimization
- Support for the NIH SNOMED comprehensive clinical ontology
- W3C SKOS inference support
- More OWL constructs: union, intersection, oneOf, OWL 2 property chains, disjoint properties, keys and OWL 2 RL/RDF rules support
- SPARQL Filter and Union clause
- Cytoscape plugin visualization

### EXISTING FEATURES

- RDF/OWL data and ontology support
- High performance bulk load
- Native inference engine for OWL, RDFS, and user-defined rules
- Querying RDF/OWL via SQL, using SPARQL-like graph patterns
- SPARQL Optional clause
- Support for Virtual Models
- Java API and SPARQL query support via Jena Adaptor
- Ontology-assisted querying of relational data using SQL operators
- Compression reduces storage by 75%, enhances query, load, and inference performance
- Automatic partitioning enhances load and inference performance and concurrent operations for large RDF data sets
- Proven scalability to over 10 billion triples

*As part of Oracle Spatial 11g, an option for Oracle Database 11g Enterprise Edition, Oracle delivers an advanced semantic data management capability not found in any other commercial or open source triple store. With native support for RDF/RDFS/OWL/SKOS standards, this semantic data store enables application developers to benefit from an open, scalable, secure, integrated, efficient platform for RDF and OWL-based applications. These semantic database features enable storing, loading, and DML access to RDF/OWL data and ontologies, inference using RDFS, OWL and SKOS semantics and user-defined rules, querying of RDF/OWL data and ontologies using SPARQL-like graph patterns embedded in SQL, and ontology-assisted querying of enterprise (relational) data.*

### Store, Load, and DML operations on the Semantic Data Store

Semantic Technologies features support storing, loading and DML operations on RDF/OWL models. Each model contains a set of subject – object – relationship triples organized as an RDF/OWL graph of directed labeled edges. The edge is the link (or relationship) that connects a subject node to an object node and is labeled by a predicate. A normalized, compressed and partitioned storage architecture manages the complexity arising from repeated usage of typically long URIs and literal values associated with the subjects, objects and predicates across triples. This provides space-efficient storage that requires 75% less disk hardware than uncompressed semantic data, and scalable and performant loading, querying, and inferencing.

### Native Inference engine for OWL, RDFS, and user-defined rules


Semantic Technologies features include a native inference engine for efficient and scalable inference using the most used subset of OWL semantics. This OWL inference engine makes the existing native inference for RDF, RDFS, and user-defined rules (used for additional specialized inference capabilities) more efficient and scalable. Inference can be done using any combination of these supported entailment regimes.

New in release 11.2, are features to optimize inference performance for large owl:sameAs cliques with a compact data structure for inference, perform incremental inference to update entailments efficiently after triple inserts, and parallel inference on multi-core or multi-CPU architectures. Additional OWL constructs are supported for: union, intersection, oneOf, OWL 2 property chains, disjoint properties, and keys, and more OWL 2 RL/RDF rules are supported. A new

**SQL-BASED SEMANTIC QUERIES**

Find Tom's grandfather:

```
select y from
TABLE(SEM_MATCH
('{:Tom :hasParent ?x .
?x :hasFather ?y}',
SEM_Models('family'),
...));
```

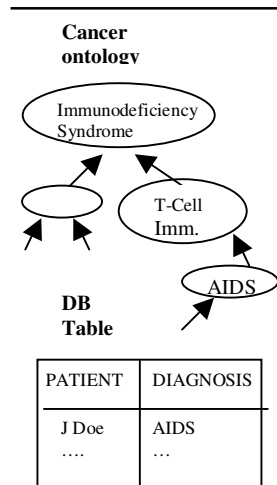


Add rule:

```
A :hasParent B,
B :hasFather C,
=> A :hasGrandFather C
```

Now find Tom's grandfather:

```
select x from
TABLE(SEM_MATCH
('{:Tom :hasGrandFather
?x}',
SEM_Models('family'), ...
```



Use the ontology to retrieve the row with 'AIDS' when 'Immunodeficiency Syndrome' is specified in the query:

```
select patient,
diagnosis from patients
where sem_related(
diagnosis,
'rdfs:subClassOf'
'Immunodeficiency
Syndrome',
'Cancer_ontology')= 1
```

CREATE\_ENTAILMENT component supports the US NIH comprehensive clinical Systematized Nomenclature of Medicine (SNOMED) ontology. A new rulebase supporting the emerging W3C Simple Knowledge Organization System (SKOS) standard on RDF enables easy sharing of controlled and structured vocabularies such as thesauri, taxonomies, and classification schemes.

**Query Semantic Data in Oracle Database**

RDF/OWL data can be queried using SQL. The SEM\_MATCH table function, which can be embedded in a SQL query, has the ability to search for an arbitrary pattern against the RDF/OWL models, and optionally, data inferred using RDFS, OWL, and user-defined rules. The SEM\_MATCH function meets most of the requirements identified by W3C SPARQL standard for graph queries. Support for virtual models, a view-like feature for combining models and optionally corresponding entailments from a UNION or UNION ALL operation, can be used in a SEM\_MATCH query. New in release 11.2, the SPARQL FILTER, UNION, and OPTIONAL keywords are supported in the SEM\_MATCH table function.

**Ontology-assisted Query for Relational Data**

Queries can extract more semantically complete results from relational data by associating relational data with ontologies that organize the domain knowledge of the relational data. Semantic Technologies include support for the semantic operator SEM\_RELATED (and optionally its SEM\_DISTANCE ancillary operator) for efficient ontology-assisted querying of relational data.

**Fine-grained Security**

The default control of access to semantic data in Oracle Database is at the model level. New in release 11.2, security administrators can define policies that conditionally restrict a user's access to triples that involve instances of a specific RDF class or property using the Virtual Private Database feature of Oracle Database, and associate sensitivity labels with individual triples stored in an RDF model using the Label Security Option.

**Semantic Indexing for Documents**

New in release 11.2, semantic indexing for documents introduces an index type that makes use of third party information extractors and annotators that locate and extract meaningful information from unstructured documents to semantically index documents stored in relational tables. Documents indexed semantically can be searched using SEM\_CONTAINS operator within a standard SQL query. The search criteria for these documents are expressed using SPARQL query patterns that operate on the information extracted from the documents.

**Change Management for RDF Models and OWL Ontologies**

New in release 11.2, RDF data stored in the Oracle semantic data store can be version-enabled using the Workspace Manager feature Oracle Database 11g. This allows users to create new versions of RDF data through data modification operations while maintaining a copy of the old data. The unit of versioning for the

**RESOURCES**

go to:

SEARCH.ORACLE.COM

"SEMANTIC TECHNOLOGIES"

semantic data store is a model, which is in turn associated with an application table that resides in a user schema.

### **Advanced Performance and Scalability for Semantic Web Applications**

Semantic Technologies supports compression, partitioning and Real Applications Clusters (RAC) for performance and scalability to address the most demanding enterprise-class semantic web solutions. **Users of the RDF/OWL/SKOS features of the Oracle Spatial option are required to license Oracle Database Enterprise Edition and the Partitioning option.**

### **Incorporate Leading Partner Tools Into An Open Data Management Solution**

The Semantic Database features in Oracle Spatial are directly integrated with the leading third party semantic technology tools and applications. New in release 11.2 is a plug-in that allows the open source Cytoscape tool to visualize semantic data in Oracle Database. Since Oracle's RDF and OWL data type is compliant with open W3C standards, Oracle Database can serve as an interoperable knowledge base. Semantic data can be shared more easily within organizations, and across the enterprise, so you can realize increased return on knowledge bases while reducing costs.

Oracle consistently works to help shape, drive, implement and support the latest open standards for the Semantic Web. Oracle is W3C member and actively participates in various technical working groups, like the OWL WG and the Data Access WG. As a result, Oracle is also committed to supporting standards specifications for RDF, RDFS, OWL, SKOS and SPARQL.

With Oracle Database Semantic Technologies features, Oracle brings the power and value of semantic analysis to your business applications. These advanced knowledge management features support semantic applications in domains ranging from national intelligence and financial fraud detection to the data integration in life sciences. Only Oracle provides world-class performance, scalability, security, and manageability to your semantic data assets, while reducing costs, with support from the leading tools vendors.

Copyright 2009, Oracle. 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 is it 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, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.