19c ORACLE Database

Graph Technology at Oracle

Oracle Database delivers high performance, scalable graph data management, query, and analytics for enterprise applications. State-of-the-art graph features are available along with functionality required for enterprise grade applications: fine-grained security, high availability, easy manageability, and integration with other data in an application. Oracle's converged, multi-model database natively supports graphs, spatial, XML, JSON and other types of data.

The Property Graph feature enables high performance graph analytics in memory using over fifty pre-built algorithms, and graph queries with a SQL-like language PGQL. Customers use Property Graphs in fraud analytics, vulnerability analysis, recommendation systems, and more.

The RDF Graph feature includes native support for W3C standards RDF, OWL, and SPARQL, enabling high performance, scalable SPARQL and SPARQL embedded in SQL queries. Customers use RDF Graphs in linked data and data sharing applications in pharma, publishing, public sector and more.

GRAPH ANALYSIS

Oracle's mission is to help people see data in new ways, discover insights, and unlock endless possibilities. Graph analysis is about understanding relationships and connections in data, and detecting patterns that identify new insights. With Oracle's Graph offerings developers can use a comprehensive suite of graph query and analytics tools to integrate graphs into applications on enterprise grade data management infrastructure. Innovative technologies of Oracle Cloud Gen 2 and Oracle Autonomous Database, the industry's only self-driving, self-securing, and self-repairing database, are available to graph applications. Enterprise infrastructure enables use of graph analysis functionality such as community detection, clustering, influencer analysis, anomaly detection, path analysis, and pattern matching in a highly secure, highly available environment. Graph analysis is used by a wide variety of enterprise customers in financial, manufacturing, public sector, pharma and other industries.

Advanced Graph Data Management for the Enterprise

"Identifying suspicious activity is mission critical for our business. Oracle Graph feature in Oracle Database provides outstanding performance benefits for this. Queries that used to take minutes, hours, or even days now run in sub-seconds with Oracle's graph features. The improvement in delivering anti-fraud alerts on time is in orders of magnitude. This brings a lot of real business value."

Yavor I Ivanov
Head of Database Administration
Paysafe Group

ORACLE

Oracle's support includes two graph data models: Property Graph and RDF Graph.

PROPERTY GRAPH FEATURES

Much of the data generated and collected today contains inherent relationships between data entities. These relationships can be easily structured as a Property Graph – a set of connected entities. The property graph vertices denote entities, the edges denote relationships, and the associated properties (attributes) are stored as key-value pairs for both.

Property Graphs are analyzed in an in-memory analytics server (PGX) using Java APIs that implement over fifty pre-built algorithms. The highly scalable, parallel, in-memory analytics server is an innovation from researchers at Oracle. Graphs can be queried with the built-in declarative graph query language PGQL, a powerful, easy-to-use SQL-like graph pattern matching query language that matches patterns and returns vertices, edges, vertex and edge properties, and subgraphs in the graph. PGQL queries can run in the in-memory analytics server or directly on graph data in Oracle Database.

Property Graph Analytics: Pre-built analytics graph algorithms include including ranking, centrality, recommendation, community detection, and path finding routines. Designed for high performance and extreme scale, PGX can work with graphs with billions of vertices and edges. Users can also build their own algorithms for specialized graph analytics. The results of analytics algorithms are stored as transient properties of vertices and edges in the graph and inspected with pattern matching queries in PGQL. The combination of pattern-matching and analytics leads to a highly expressive and flexible interface for graph analytics. Graph analytics can be executed within a Java application or in PGX deployed in Oracle WebLogic Server and Apache Tomcat.

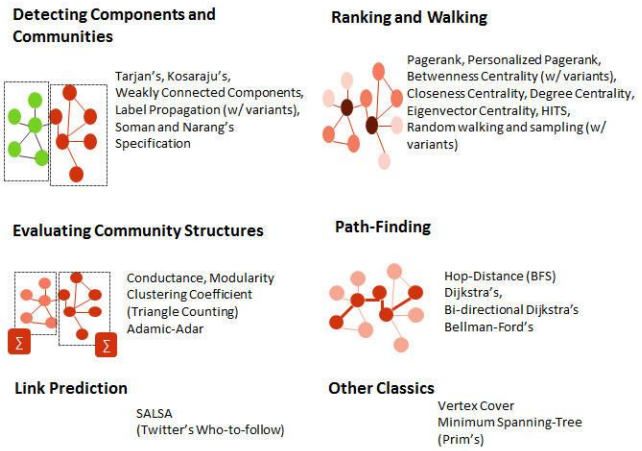

Property Graph Query Language PGQL: Alongside SQL constructs like SELECT, FROM, WHERE, GROUP BY and ORDER BY, and functions like ABS, CEIL, FLOOR and ROUND, PGQL allows for matching graph patterns. Fixed-length graph patterns match a fixed number of vertices and edges per query. Variable-length graph patterns contain one or more quantifiers like *, +, ? or {2,4} for matching vertices and edges in a recursive fashion. PGQL integrates graph pattern matching with subquery capabilities so that vertices and edges matched in one query can be passed to another query for continued joining or pattern matching. Since PGQL is built on SQL's foundation, it benefits from all existing SQL features and any SQL features that will be added to the standard over time.

Key Benefits

- Oracle Database scalability, security and manageability for enterprise graph applications
- Extreme performance for critical enterprise graph datasets with Oracle Engineered Systems
- Commercial strength scalable, and comprehensive support for property graph query and analytics
- W3C RDF semantic graphs with advanced support for Linked Data applications

Key Property Graph Features

- 50+ powerful parallel, in-memory graph analytics
- Java API to execute analytics in parallel in the in-memory analytics server
- Optimized Property Graph storage model
- PGQL, powerful SQL-like language to run graph queries
- Ease of development with rich UI – interpreters for notebook, shell UI, and graph visualization
- Analysis results can flow into a bipartite, filtered, undirected, sorted or simplified edge output graph
- Property graph analysis of RDF graphs using property graph views

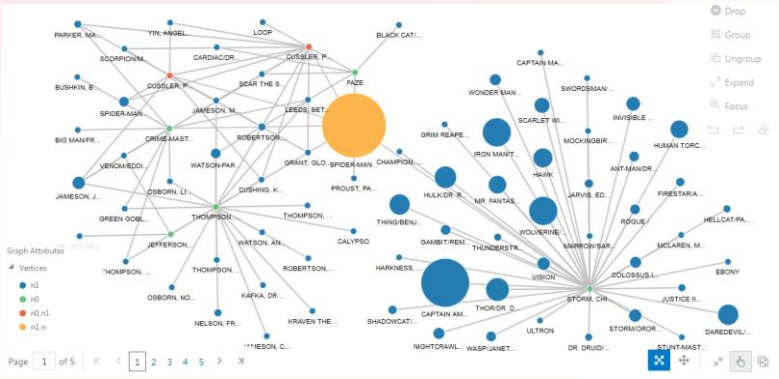

Enterprise Features: Graph data management is part of the converged, multi-model Oracle Database. Enterprise features of the database all extend to graph data management: fine-grained security, high availability, easy manageability, integration with all other data in business applications, and more.

Ease of Development and other Features: Ease of development and management is provided through a set of Java APIs to create and drop property graphs, add and remove vertices and edges, search for vertices and edges using key-value pairs, create text indexes, and perform other manipulations. The rich user interface includes PGX and PGQL interpreters for notebook, a shell UI, and graph visualization.

RDF GRAPH FEATURES

RDF graph model in Oracle Database provides native support for RDF and OWL, Worldwide Web Consortium (W3C) standards for representing and defining knowledge graphs, semantic data, and SPARQL, a graph query language that enables comprehensive RDF query, reasoning, and analytics. RDF graphs create a unified metadata layer for disparate applications for identification, integration, and discovery of data. RDF graphs are central to knowledge graphs and linked data applications common in the healthcare and life sciences, finance, media and intelligence communities.

RDF Graph Queries: Oracle's RDF graph scales to billions of triples, with parallel loading, querying and inferencing. RDF graphs are stored, loaded, and queried directly in Oracle Database. Pattern-matching queries can be executed with SPARQL 1.1 in SQL, with Apache Jena Java APIs, and with the Fuseki SPARQL end-point web services. Built-in reasoning uses forward-chaining rules and inference using OWL 2, SKOS, and user-defined rules. Open Geospatial Consortium standards-based GeoSPARQL evaluates spatial data in an RDF graph. RDF views defined on tables enable SPARQL queries on relational data. Property graph views enable graph analytics. A framework for third party natural language processing enables semantic text indexing.

Enterprise Features: Enterprise applications benefit from RDF Graph's extensive integration with Oracle Database and Oracle tools. SQL Developer provides an easy-to-use interface to work with and query RDF graph data. Oracle Database management utilities and tuning can be applied to RDF graphs, including Enterprise Manager, Oracle optimizer hints, SQL*Loader direct path load, Data Guard physical standby, Data Pump import/export, Recovery Manager and external tables. Access controls can be applied at the model level and optionally at the triple level using the Oracle Label Security option. RDF graphs can be analyzed using SPARQL 1.1 path expressions, Apache Jena compliant graph visualization tools, Oracle SQL Developer, and Advanced Analytics.

- Fast load of graphs into in-memory analytics server
- Fine-grained security for graph elements
- Collaborative filtering with SQL-based property graph analytics

Key RDF Graph Features

- Support for W3C and OGC standards: RDF, SPARQL, R2RML, OWL, SKOS
- RDF Triple store in compressed, partitioned tables
- SPARQL 1.1: including Update and ORDER BY, and query functions
- Fine-grained security: schema private RDF graphs, label security, data vault, and more
- Parallel load, query and inference
- RDF views on property graph data
- Property graph views on RDF data
- Oracle Flashback Query support
- Semantic indexing of text
- Ontology-assisted relational querying
- Oracle Database In-Memory support
- Full integration with SQL Developer
- Load from a variety of RDF file formats into RDF graph model in database

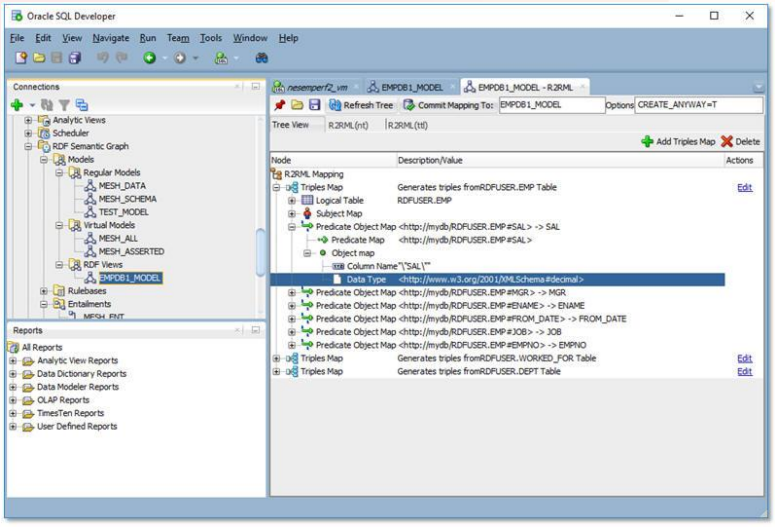

GRAPH FEATURES IN THE WORLD'S LEADING IT PLATFORM

Oracle's graph features are a native component of Oracle Database – and of the world's leading information technology platform for Oracle Cloud, on premise, and big data deployments. Applications benefit from the enterprise grade performance, scalability, security, and data integration capabilities of Oracle Database. They can exploit the extreme processing power and bandwidth of Oracle Exadata Database Machine. Developers can easily incorporate these capabilities in their solutions using modern development frameworks. Large enterprises worldwide in financial services, retail, life sciences, national and local government rely on Oracle Graph technologies.

Related Products

- Oracle Cloud
- Oracle Database 19c
- Oracle Advanced Analytics
- Oracle Big Data Spatial and Graph
- Oracle Exadata Database Machine

Resources

For more information, visit

Spatial and Graph on Oracle.com

Product overviews, blogs, videos, software downloads, documentation, tutorials, white papers, use case studies:

- www.oracle.com/database/technologies/spatialandgraph.html

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com.
Outside North America, find your local office at oracle.com/contact.

blogs.oracle.com/oraclespatial Oracle Spatial and Graph twitter.com/SpatialHannes

Integrated Cloud Applications & Platform Services

ORACLE