
ORACLE®

GeoSPARQL Support and Other Cool Features in Oracle 12c Spatial and Graph

Linked Data Seminar – Culture, Base Registries & Visualisations

Hans Viehmann
Product Manager EMEA
Oracle Corporation
December 2, 2016


@SpatialHannes

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Strategic Vision: The platform

Database and Big Data Platform support, both On-premise and in the Cloud

Oracle Big Data Spatial and Graph

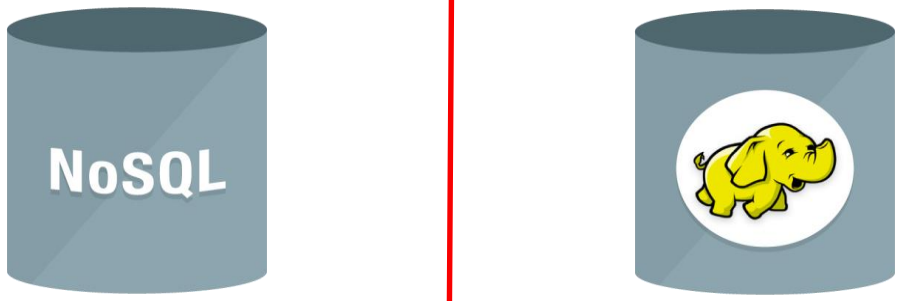

Oracle Database Spatial and Graph

Spatial and Graph in Cloud Offerings

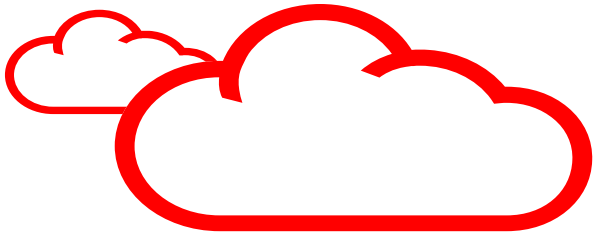

Big Data:
Single Model Data Store

Database 12c:
Polyglot (Multi-model) Data Store

Using RDF to create virtual data warehouses of linked data

RDF-based Metadata Layer

- W3C standard, flexible model for sparse and evolving data
- Common vocabulary enables data integration & app development
- Relational data stays in place, apps don't need to change

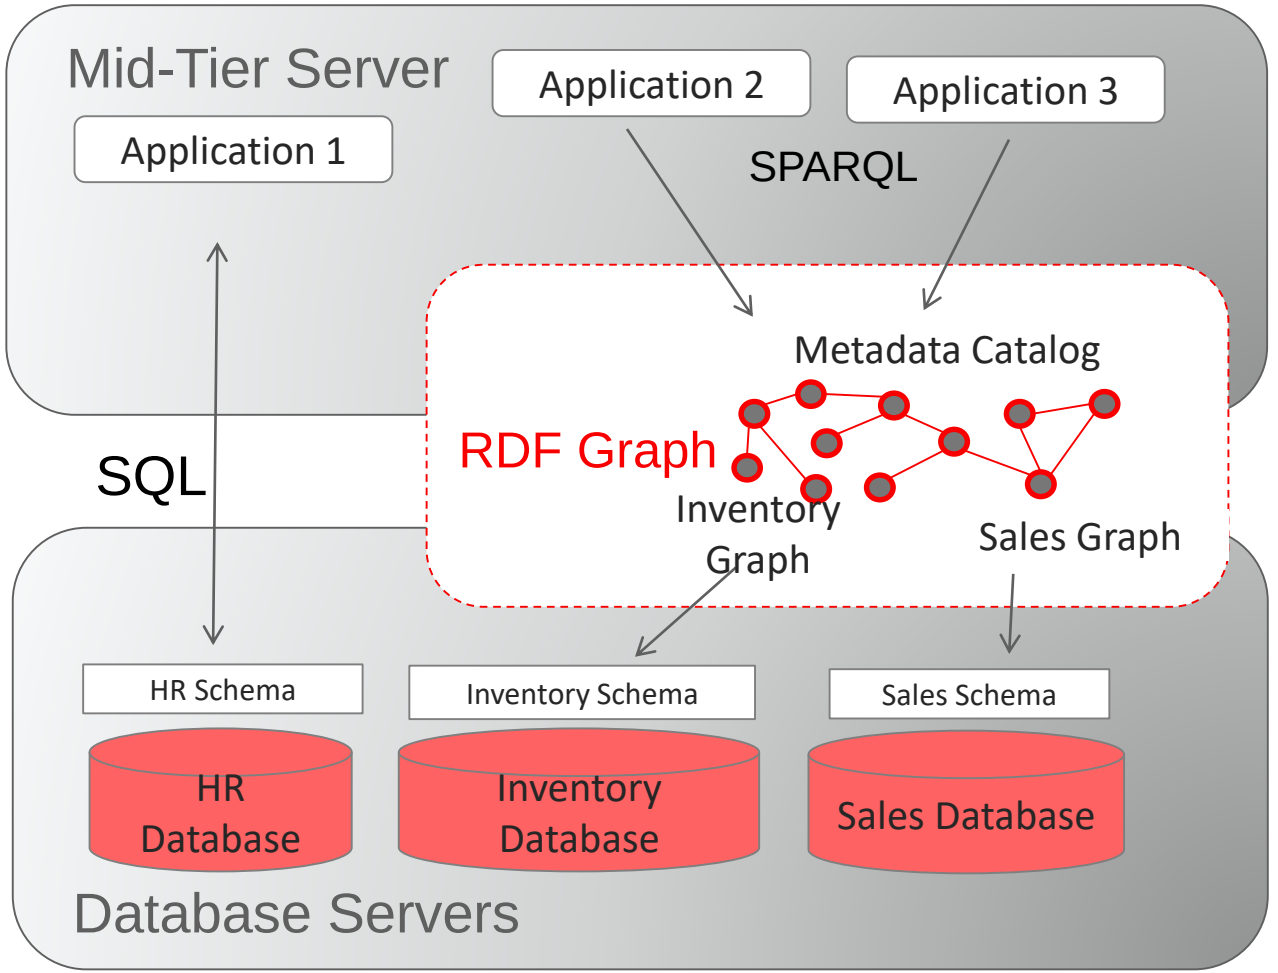

Industries Actively Deploying Linked Data Systems

Industries

- Life Sciences
- Health Care
- Finance
- Media
- Networks & Communications
- Defense & Intelligence

EU Publications Office

Linked Metadata Platform for European Union

Objectives

- Common metadata model supports:
  - Search and discovery of EU Publications
  - Multiple domains and languages

Solution

- Validate and tag EU law, tenders, and publicity to standardized vocabularies
- Unified RDF graph metadata model
- Supports discovery of content through user's terminology and language
- Provides variety of dissemination modes

Benefits

- Evolving data model that flexibly supports a variety of business use cases
- Scalability:
  - Over billion RDF triples in Oracle Graph DB
  - 2.5 TB of compressed data in Oracle DB
  - Links to 3.9 TB (60M) files of EU pubs
- Reliability and maintainability
  - Oracle ASM (Automatic Storage Management)
  - Two failover systems

ORACLE

Geospatial Linked Data Services

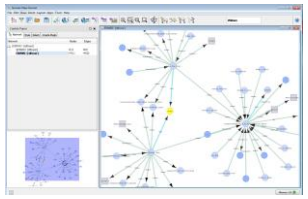

Web Analyst 1

Web Analyst 2

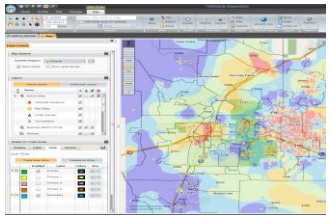

GeoSPARQL, REST, JSON

Linked Data Graphs

Statistics Terminologies

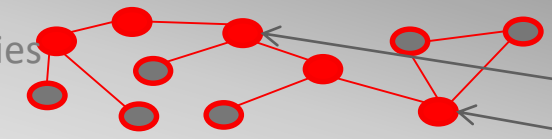

Pop_Stat_Graph

Geometries (or views)

Relational Schema

2D Feature Schema

Population Statistics Database

Census Block Geometry

Enterprise Data Servers

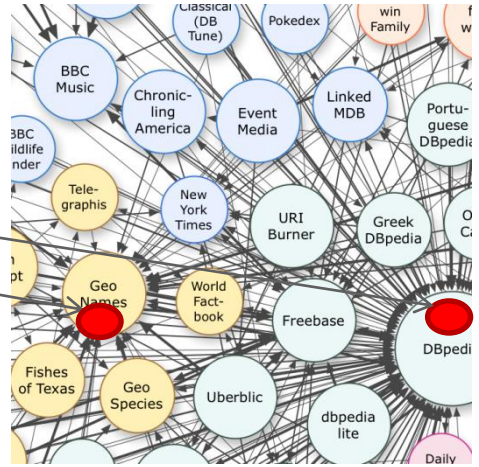

Public LOD endpoints (DBPedia & GeoNames)

OGC GeoSPARQL Support

Oracle Spatial and Graph 12c

ORACLE

OGC GeoSPARQL

- GeoSPARQL – A Geographic Query Language for RDF Data
  - OGC Standard (document 11-052r4)
  - Published in June 2012
  - Submitting Organizations

Australian Government Bureau of Meteorology

Bentley

CSIRO

GeoConnections

Defence Geospatial Information Working Group

USGS science for a changing world

interactive instruments

ORACLE

Ordnance Survey

Raytheon

Traverse Technologies, Inc.

ORACLE

Why GeoSPARQL? – Linked Geo Data

- Many Linked Open Data (LOD) datasets have geospatial components
- Barriers to integration
  - Vendor-specific geometry support
  - Different vocabularies
    - W3C Basic Geo, GML XMLLiteral, Vendor-specific
  - Different spatial reference systems
    - WGS84 Lat-Long, British National Grid

OpenStreetMap
The Free Wiki World Map

Why GeoSPARQL? – Semantic GIS

- GIS applications with semantically complex thematic aspects
  - Logical reasoning to classify features
    - Land cover type, suitable farm land, etc.
  - Complex Geometries
    - Polygons and Multi-Polygons with 1000's of points
  - Complex Spatial Operations
    - Union, Intersection, Buffers, etc.

Find parcels with an area of at least 3 sq. miles that touch a local feeder road and are inside an area of suitable farm land.

From SPARQL to GeoSPARQL

RDF Data

```
:res1 rdf:type   :House .
:res1 :baths       "2.5"^^xsd:decimal .
:res1 :bedrooms "3"^^xsd:decimal .

:res2 rdf:type     :Condo .
:res2 :baths       "2"^^xsd:decimal .
:res2 :bedrooms "2"^^xsd:decimal .

:res3 rdf:type   :House
:res3 :baths       "1.5"^^xsd:decimal .
:res3 :bedrooms "3"^^xsd:decimal .
```

SPARQL Query

```
SELECT ?r ?ba ?br
WHERE { ?r rdf:type :House .
        ?r :baths ?ba .
        ?r :bedrooms ?br }
```

Result Bindings

```
?r     | ?ba    | ?br
====================
:res1 | "2.5" | "3"
:res3 | "1.5" | "3"
```

From SPARQL to GeoSPARQL

RDF Data

```
:res1 rdf:type   :House .
:res1 :baths       "2.5"^^xsd:decimal .
:res1 :bedrooms "3"^^xsd:decimal .

:res2 rdf:type     :Condo .
:res2 :baths       "2"^^xsd:decimal .
:res2 :bedrooms "2"^^xsd:decimal .

:res3 rdf:type   :House
:res3 :baths       "1.5"^^xsd:decimal .
:res3 :bedrooms "3"^^xsd:decimal .
```

SPARQL Query

```
SELECT ?r ?ba ?br
WHERE { ?r rdf:type :House .
        ?r :baths ?ba .
        ?r :bedrooms ?br
        FILTER (?ba > 2) }
```

Result Bindings

?r	?ba	?br
:res1	"2.5"	"3"

From SPARQL to GeoSPARQL

Spatial RDF Data

```
:res1   rdf:type          :House .
:res1   :baths              "2.5"^^xsd:decimal .
:res1   :bedrooms         "3"^^xsd:decimal .
:res1   ogc:hasGeometry :geom1 .
:geom1 ogc:asWKT          "POINT(-122.25 37.46)"^^ogc:wktLiteral .

:res3   rdf:type          :House
:res3   :baths              "1.5"^^xsd:decimal .
:res3   :bedrooms         "3"^^xsd:decimal .
:res3   ogc:hasGeometry :geom3 .
:geom3 ogc:asWKT          "POINT(-122.24 37.47)"^^ogc:wktLiteral .
```

This is what GeoSPARQL standardizes

Vocabulary & Datatypes

Find houses within a search polygon

GeoSPARQL Query

```
SELECT ?r ?ba ?br
WHERE { ?r rdf:type :House . ?r :baths ?ba . ?r :bedrooms ?br .
        ?r ogc:hasGeometry ?g . ?g ogc:asWKT ?wkt
        FILTER(ogcf:sfWithin(?wkt, "POLYGON(…)"^^ogc:wktLiteral))
}
```

Extension Functions

ORACLE

RDB2RDF for viewing Spatial Data as RDF

Relational Data

HOUSE table

id int	baths number	bedrooms number	geom SDO_GEOMETRY
1	2.5	3	POINT(-122.25 37.46)
3	1.5	3	POINT(-122.24 37.47)

RDB2RDF: Direct Mapping

```
sem_apis.CREATE_RDFVIEW_MODEL (
   'House_Model',
   sys.odcivarchar2list('HOUSE'),
   'http://dm/');
```

RDF View (of Relational Data)

```
<http://dm/RDFUSER.HOUSE/ID=1>
  rdf:type
    <http://dm/RDFUSER.HOUSE>;
  :baths
"2.5"^^xsd:decimal;
  :bedrooms "3"^^xsd:decimal;
  :geom
    "POINT(…)"^^ogc:wktLiteral.
<http://dm/RDFUSER.HOUSE/ID=3>
  rdf:type
    <http://dm/RDFUSER.HOUSE>;
  :baths
"1.5"^^xsd:decimal;
  :bedrooms "3"^^xsd:decimal;
  :geom
    "POINT(…)"^^ogc:wktLiteral.
```

ORACLE

RDB2RDF for viewing Spatial Data as RDF

Relational Data

House table

id int	baths number	bedrooms number	geom SDO_GEOMETRY
1	2.5	3	POINT(-122.25 37.46)
3	1.5	3	POINT(-122.24 37.47)

RDB2RDF: Direct Mapping

```
sem_apis.CREATE_RDFVIEW_MODEL (
   'House_Model',
   sys.odcivarchar2list('HOUSE'),
   'http://dm/');
```

Querying RDF View

```
PREFIX :
 <http://dm/RDFUSER.HOUSE#>.
SELECT ?r ?ba ?br
WHERE {
 ?r rdf:type
  <http://dm/RDFUSER.HOUSE>;
 :baths ?ba;
 :bedrooms ?br;
 :geom ?wkt.
FILTER
(ogcf:sfWithin(?wkt,
"POLYGON(…)"^^ogc:wktLiteral)
)
}
```

ORACLE

GeoSPARQL Support in Oracle

- Oracle Spatial and Graph supports the following conformance classes for GeoSPARQL
  - Core
  - Topology Vocabulary Extension (Simple Features)
  - Geometry Extension (WKT, 1.2.0)
  - Geometry Topology Extension (Simple Features, WKT, 1.2.0)
  - RDFS Entailment Extension (Simple Features, WKT, 1.2.0)

Builds on the power of Oracle Spatial

- Efficient Spatial Indexing
- Spatial Reference Systems
  - Built-in support for 1000's of SRS
  - Coordinate system transformations applied transparently during indexing and query
- Geometry Types
  - Support OGC Simple Features geometry types
    - Point, Line, Polygon
    - Multi-Point, Multi-Line, Multi-Polygon
    - Geometry Collection
  - Up to 500,000 vertices per Geometry

GeoSPARQL – New 12.2 Features

- New utility functions
- Support for EPSG SRID URIs
- Revised Geometry Storage Scheme – big performance gain
- SDO_JOIN
- Spatial Aggregates
- 3D Support

Semantic technologies platform

Oracle Spatial and Graph 12c

ORACLE

Oracle Database 12c Spatial and Graph Tooling

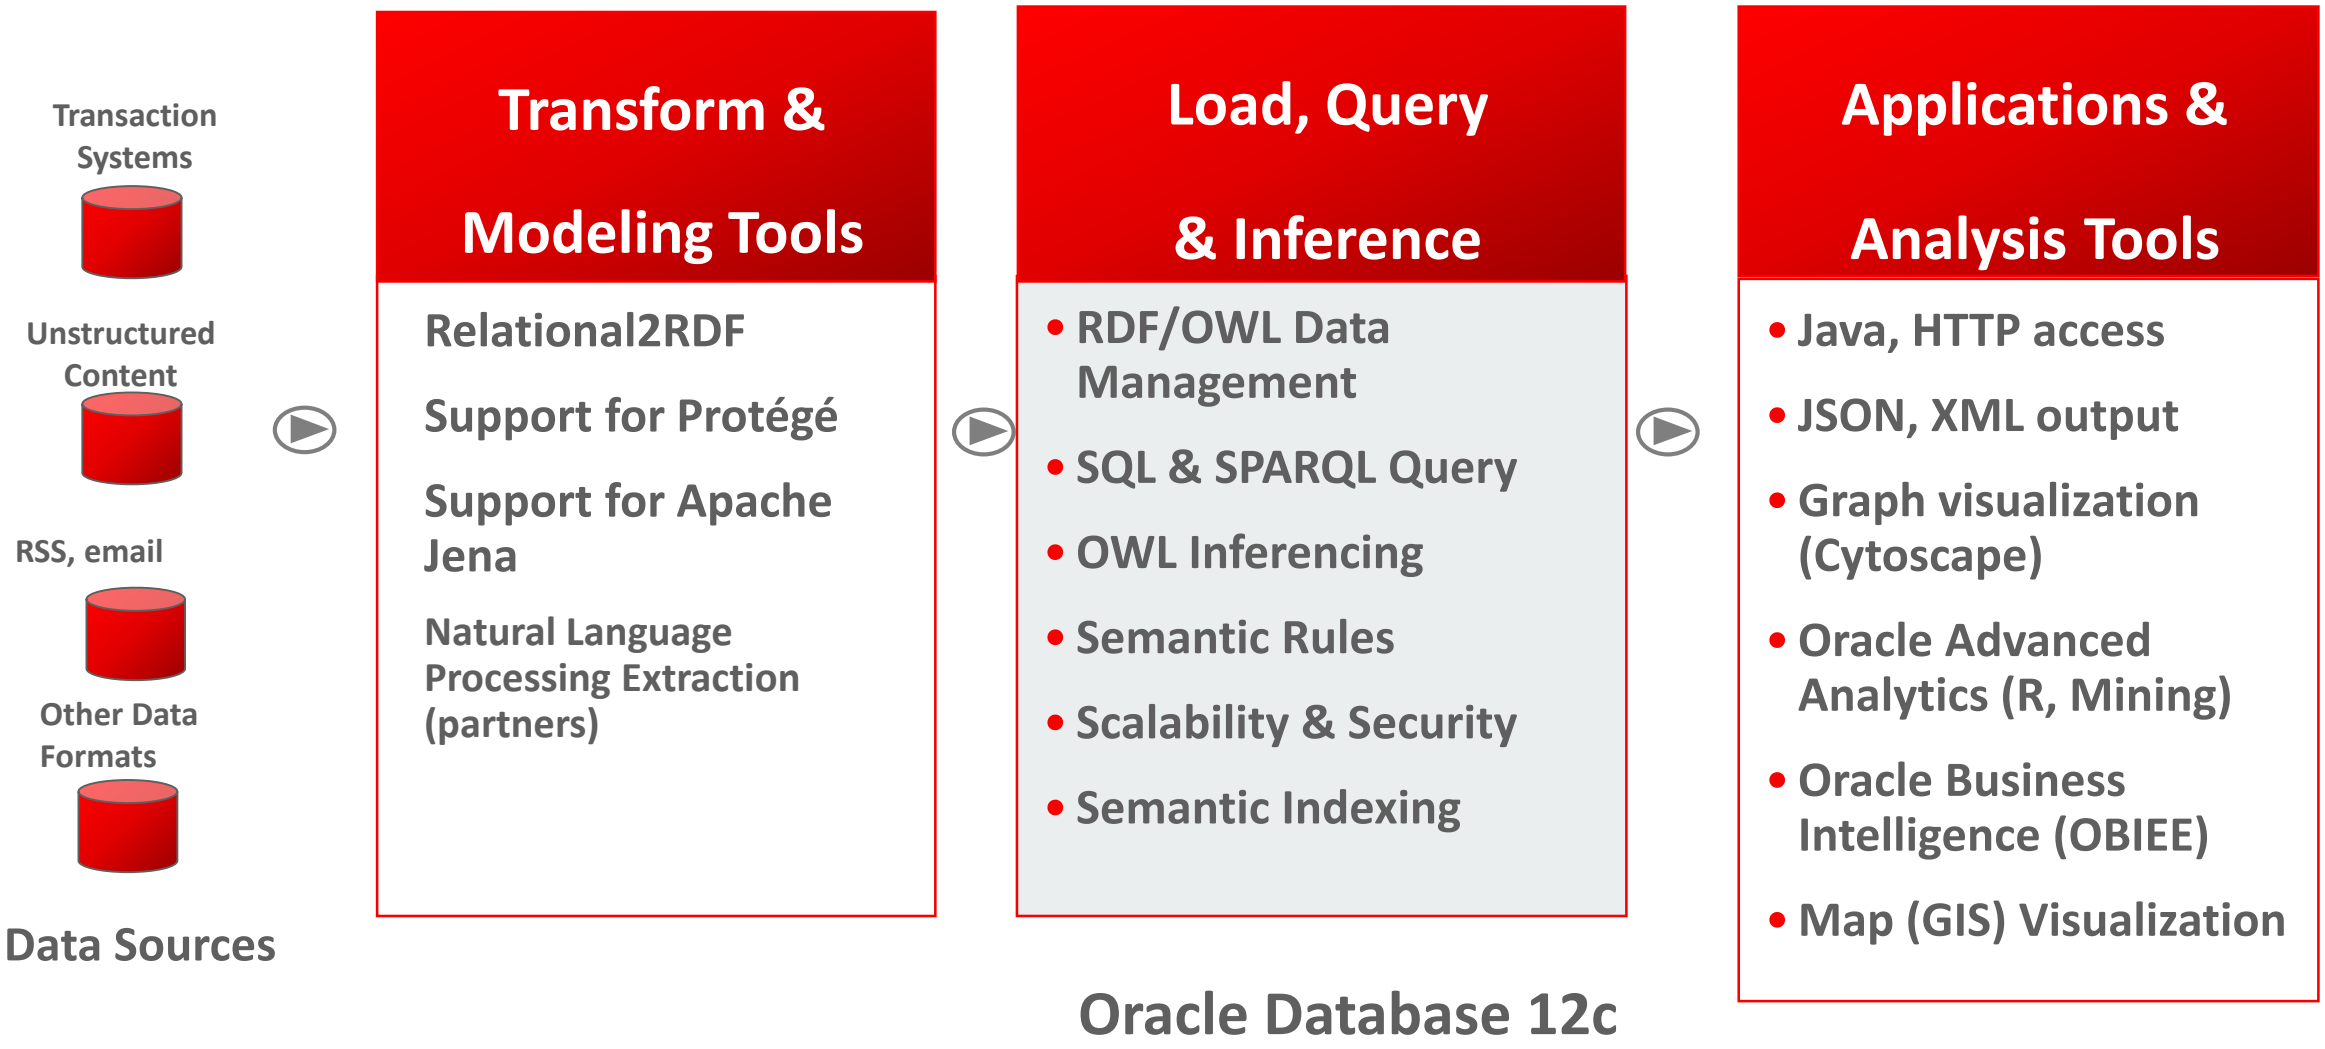

R2RML: From Relational to RDF
RDF Views on Relational Tables

- Pattern matching on relational tables
- Supports W3C RDF & SPARQL standard
- Automatic and custom mapping
- RDF views: on tables, views, SQL query results
- No duplication of data and storage
- Direct Mapping – Automatic
- R2RML - express customized mappings

EmpNo	Ename	Job	Mgr	DeptNo
7521	Ward	Salesman	7698	10
7698	Blake	Manager	7839	10
7839	King	President		30

DeptNo	LOC
10	NYC
30	CHI

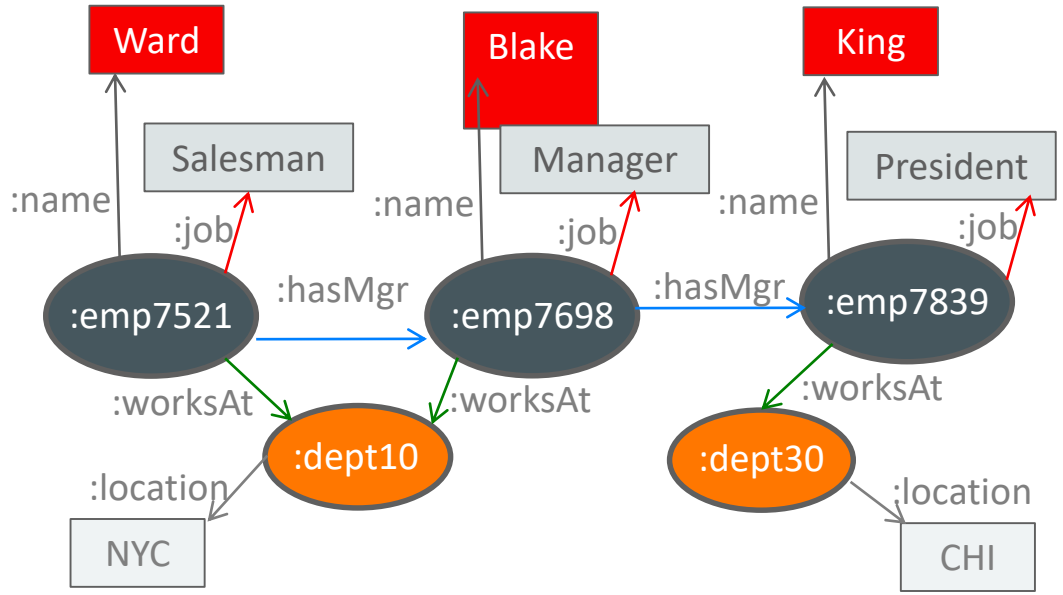

Core Inferencing Features

- Forward-chaining based inference engine in the database
  - Removes on-the-fly reasoning and results in fast query times
- Native rulebases: RDFS, OWL 2 RL, OWL 2 EL, SKOS
  - SNOMED (subset of OWL 2 EL)
- Validation of inferred data
- Proof generation
- User defined inferencing
  - Temporal reasoning, Spatial reasoning
- Ladder Based Inference
  - Fine grained security for inference graph
- Integration with external OWL 2 reasoners (TrOWL)

Integration with Tom Sawyer Graph Visualization

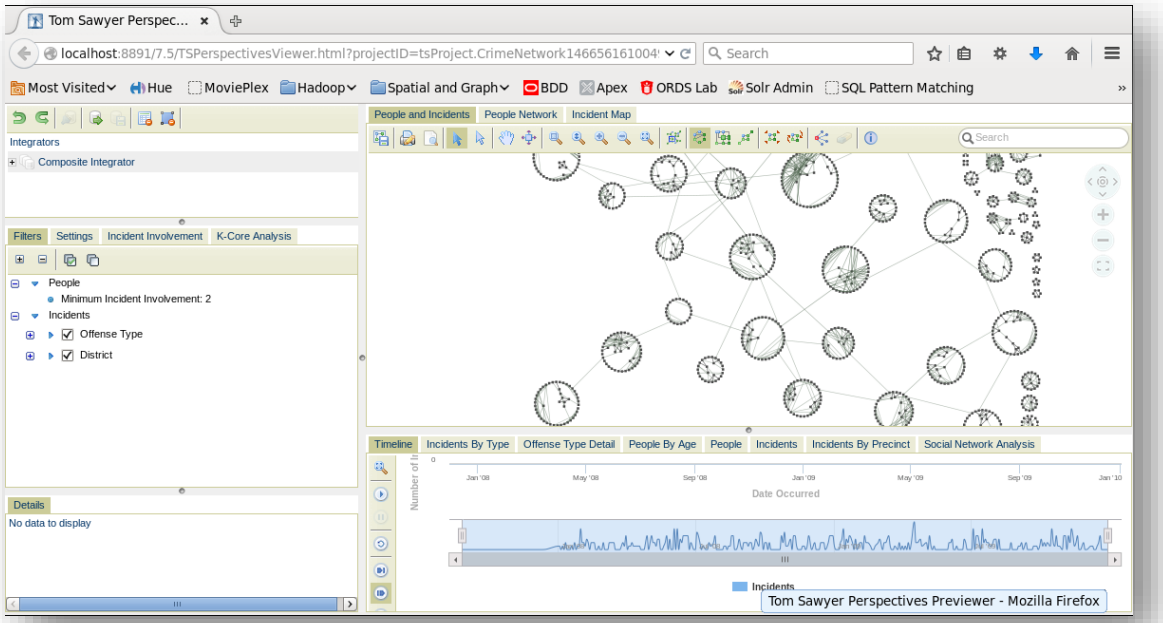

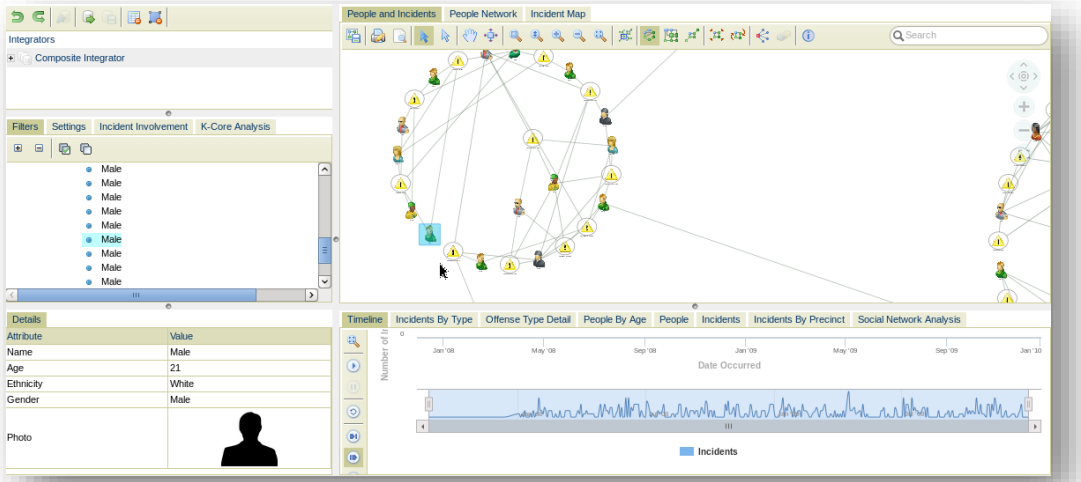

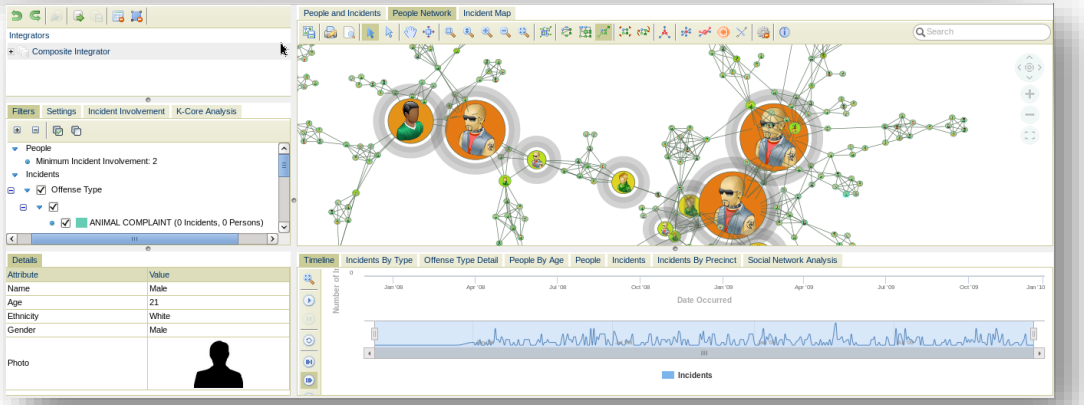

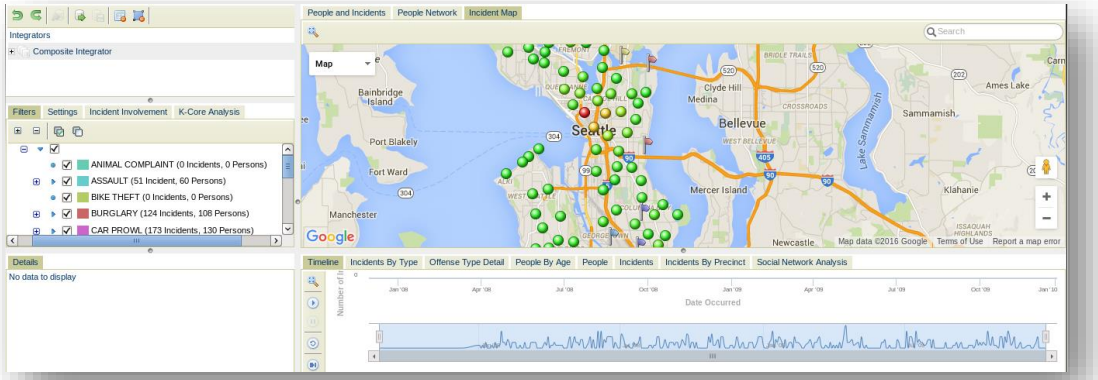

Integration with IO Informatics

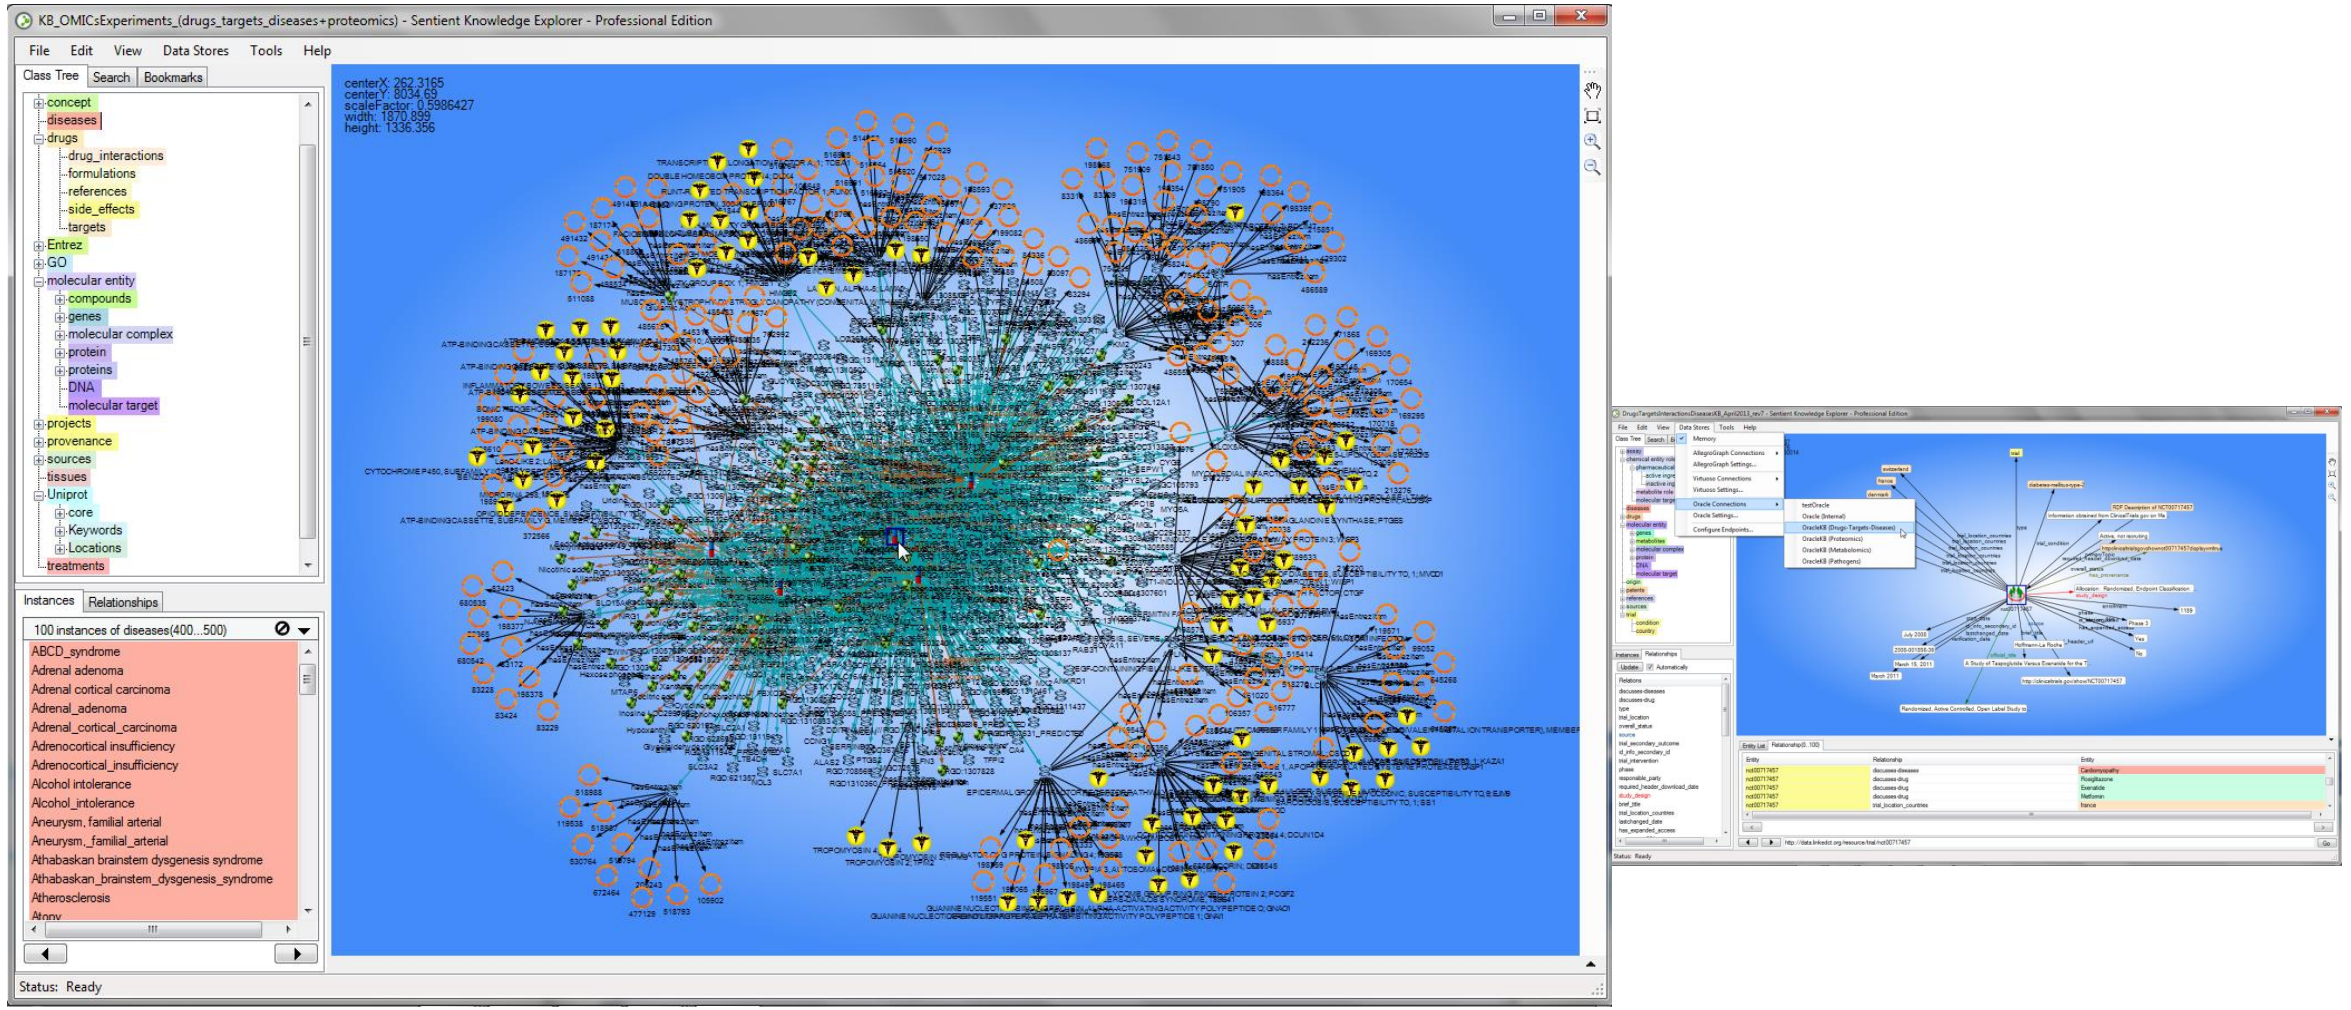

Integration with Cytoscape Open Source Visualization

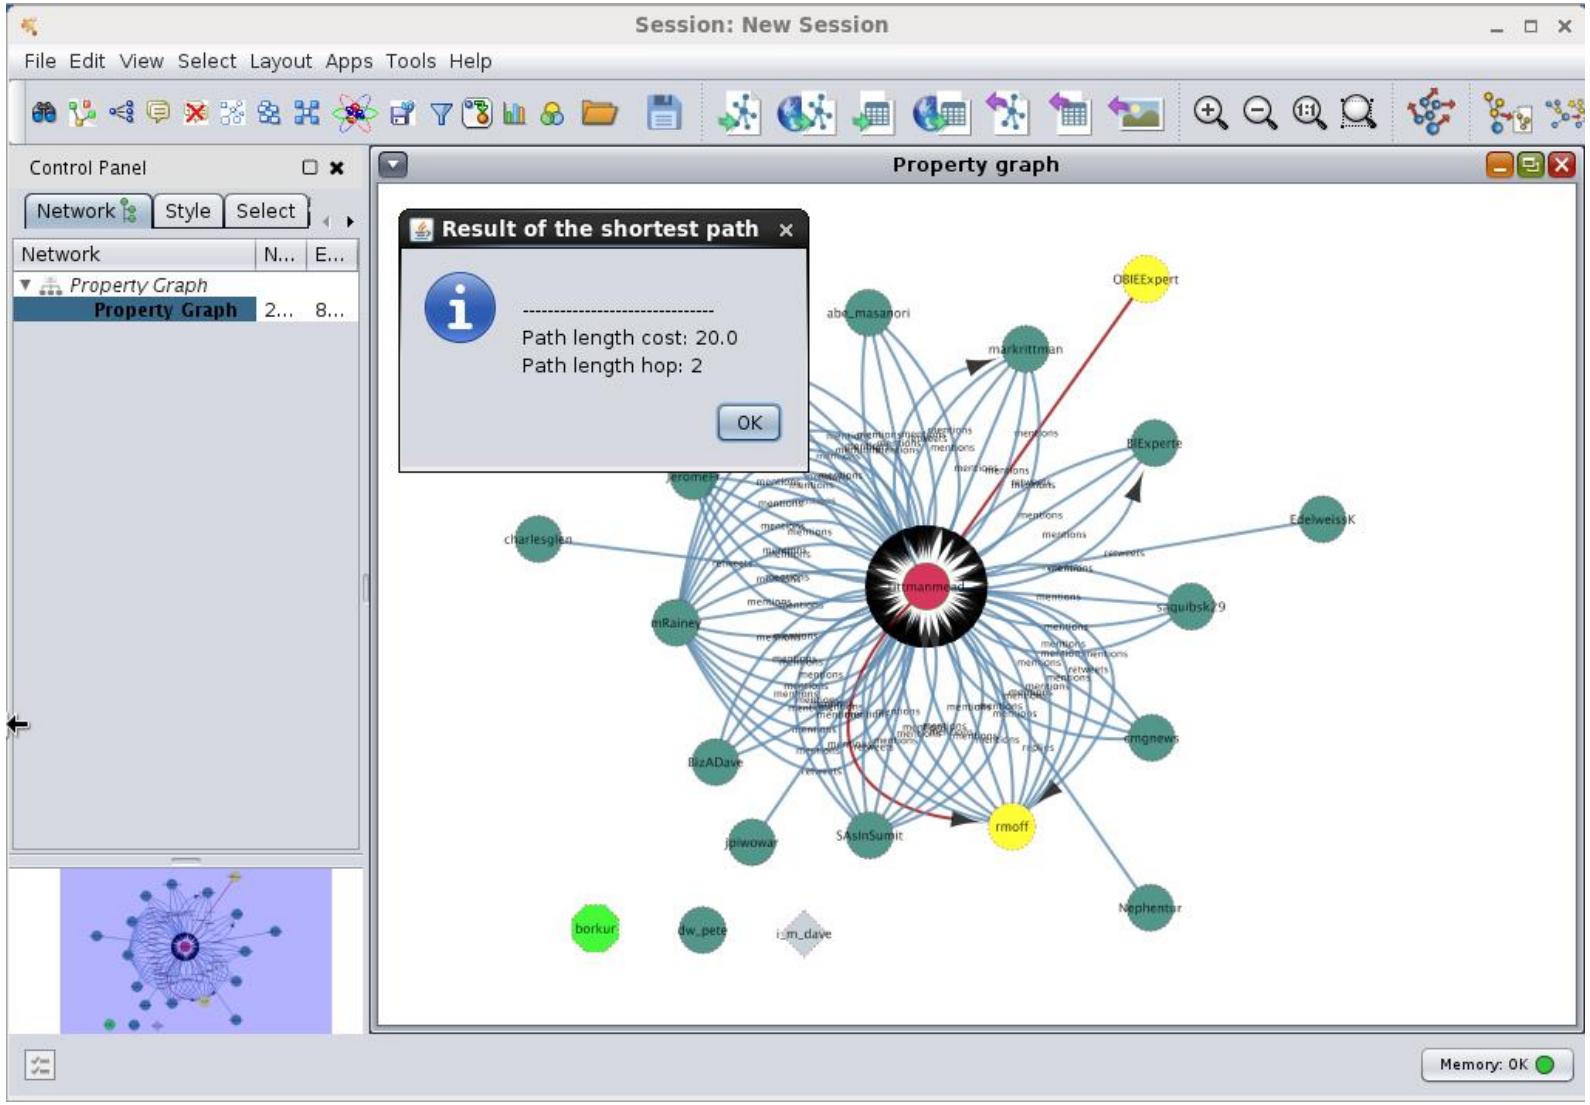

Integration with Protégé

Semantic Modeling

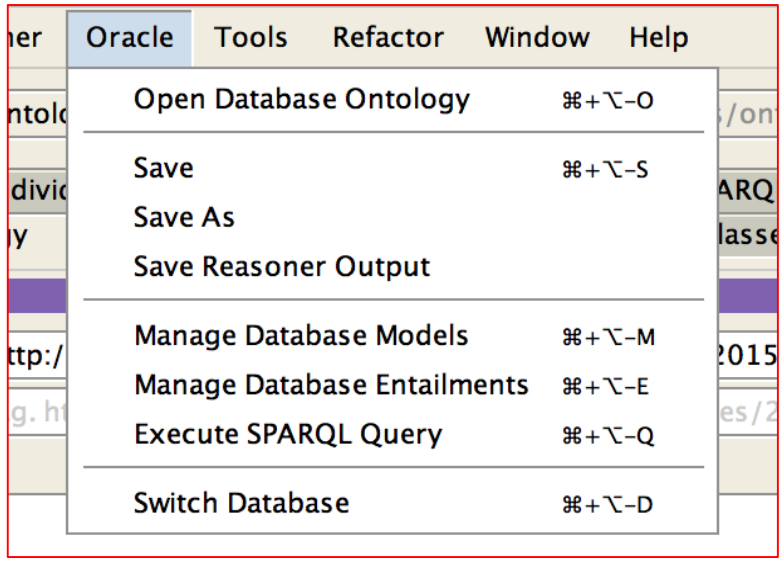

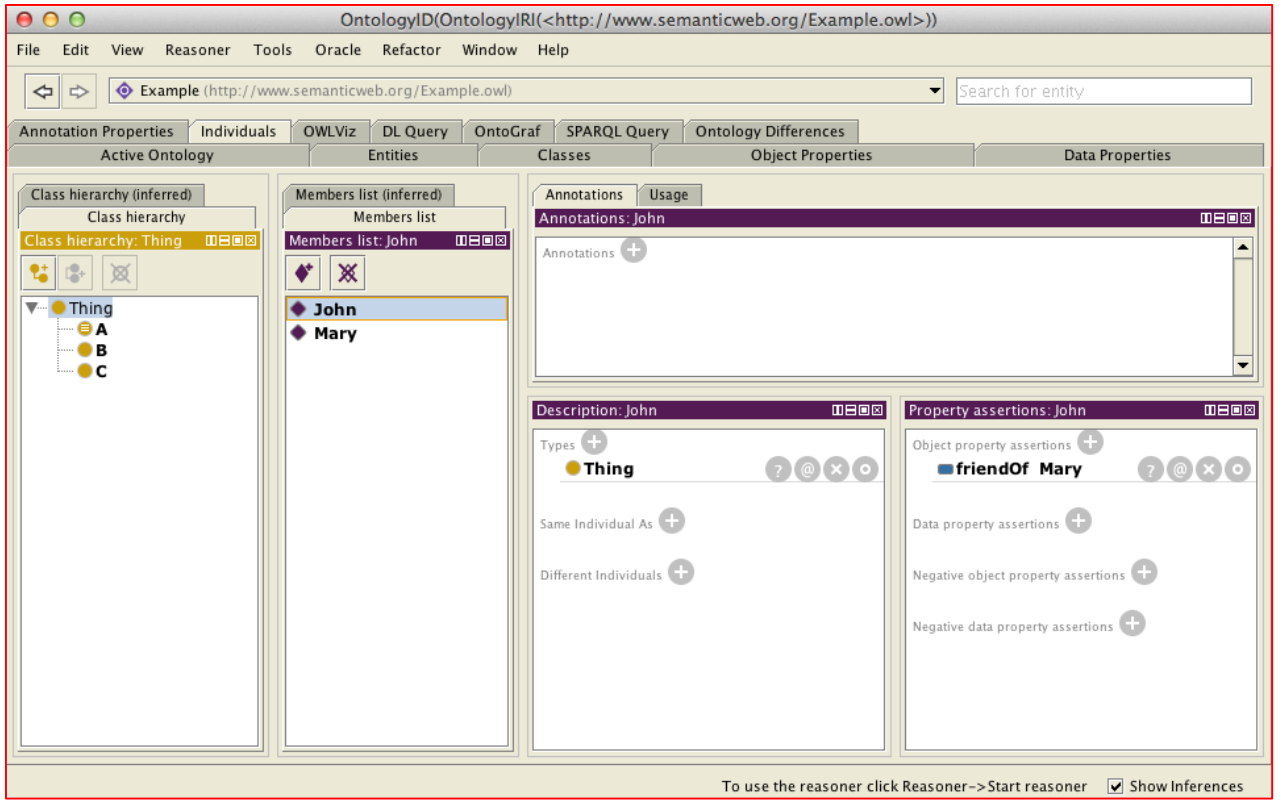

Integration with TopQuadrant TopBraid Composer

Semantic Modeling

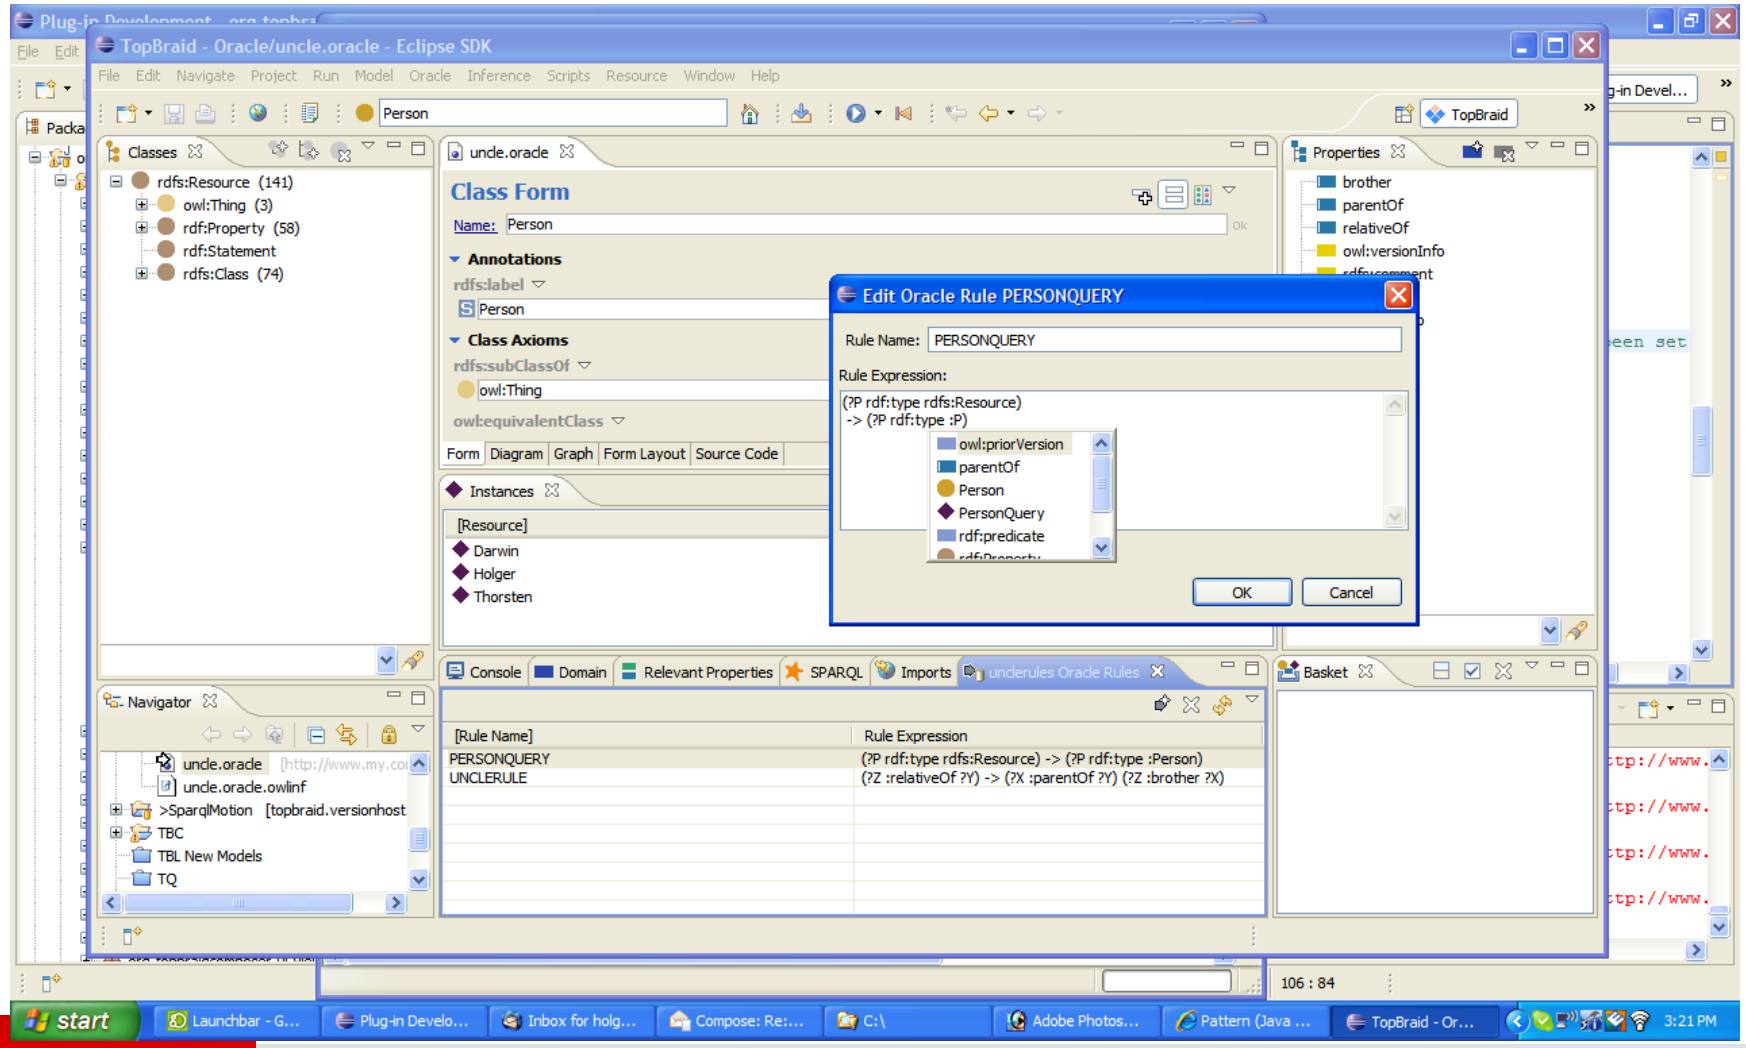

Reporting RDF Data with Oracle BI EE

- Powerful BI dashboards
  - Visually appealing
  - 100% thin client
- Across all styles of analysis
  - R-OLAP, M-OLAP, Scorecards, Reporting, Collaboration, Actions
- Across all data sources
  - Federated data access
  - Share, collaborate, & publish

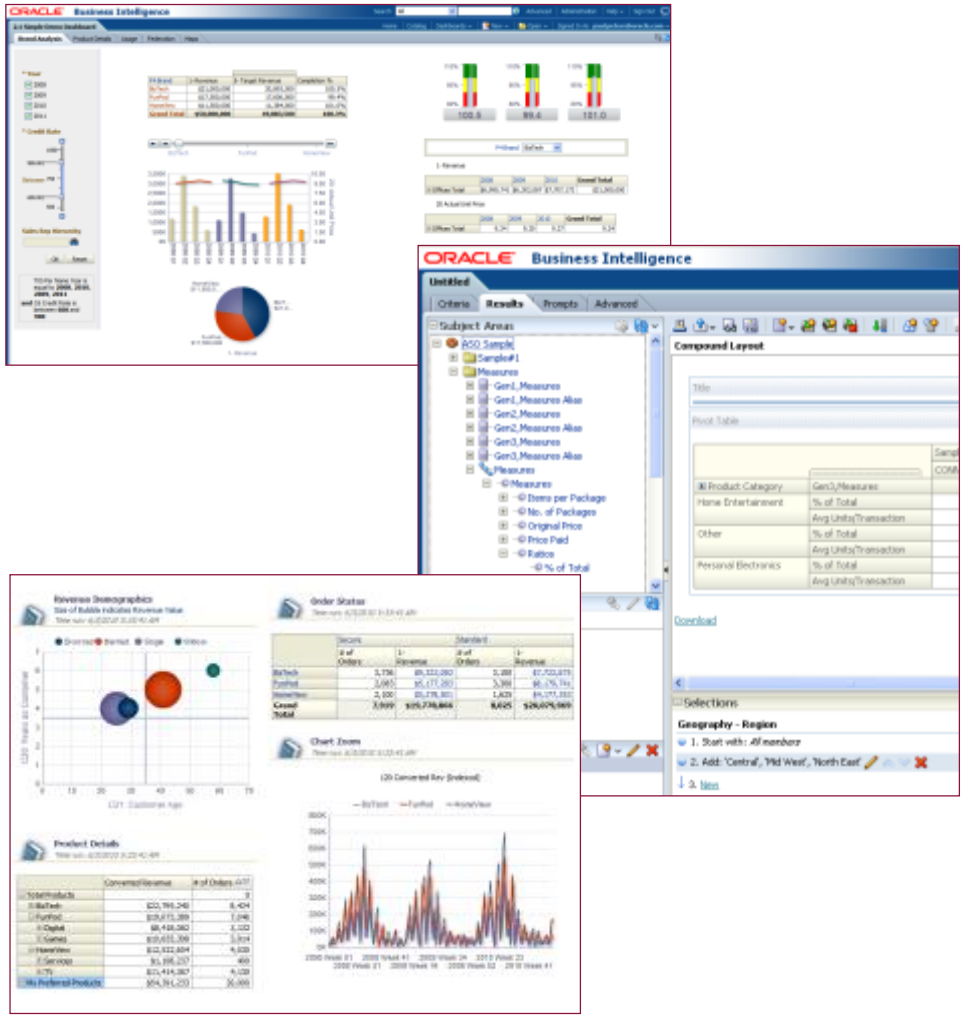

In-Database Graph and Predictive Analytics

Oracle Advanced Analytics – Data Mining

Problem Classification	Sample Problem
Anomaly Detection 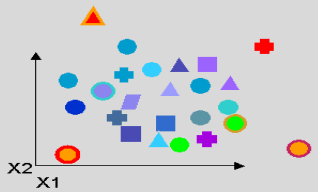	Given demographic data about a set of customers, identify customer purchasing behavior that is significantly different from the norm
Association Rules 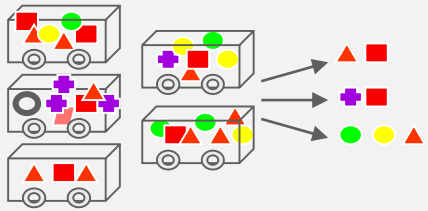	Find the items that tend to be purchased together and specify their relationship – market basket analysis
Clustering 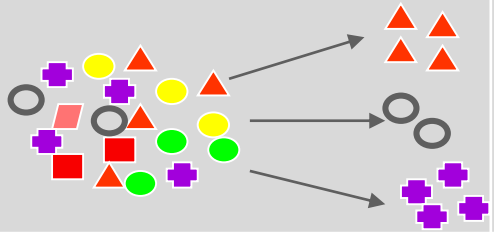	Segment demographic data into clusters and rank the probability that an individual will belong to a given cluster
Feature Extraction 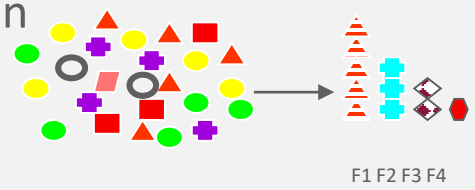	Given demographic data about a set of customers, group the attributes into general characteristics of the customers

Performing Statistical Graph Analytics

Oracle Advanced Analytics – Oracle R Enterprise

- Open source language
- Statistical computing and chart for graph data
- Produces publication quality plots
- Highly extensible with open source R packages

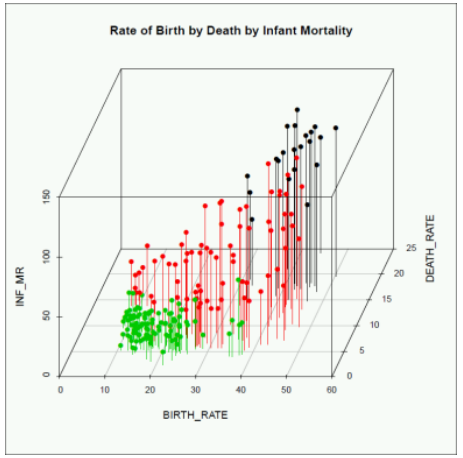

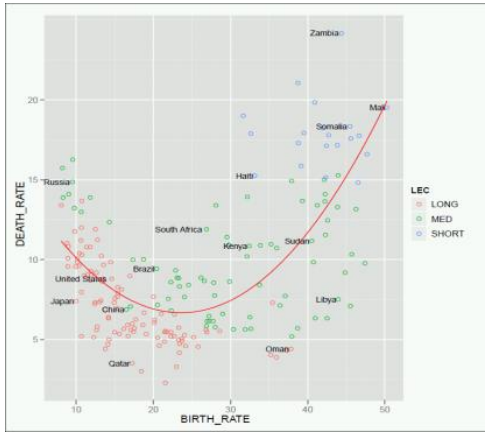

World's Fastest Big Data Graph Benchmark

1 Trillion Triple RDF Benchmark with Oracle Spatial and Graph

- **World's fastest data loading performance**
- **World's fastest query performance**
- **Worlds fastest inference performance**
- **Massive scalability: 1.08 trillion edges**

- **Platform**: Oracle Exadata X4-2 Database Machine
- **Source**: w3.org/wiki/LargeTripleStores, 9/26/2014

Oracle Database 12*c* can load, query and inference millions of RDF graph edges per second

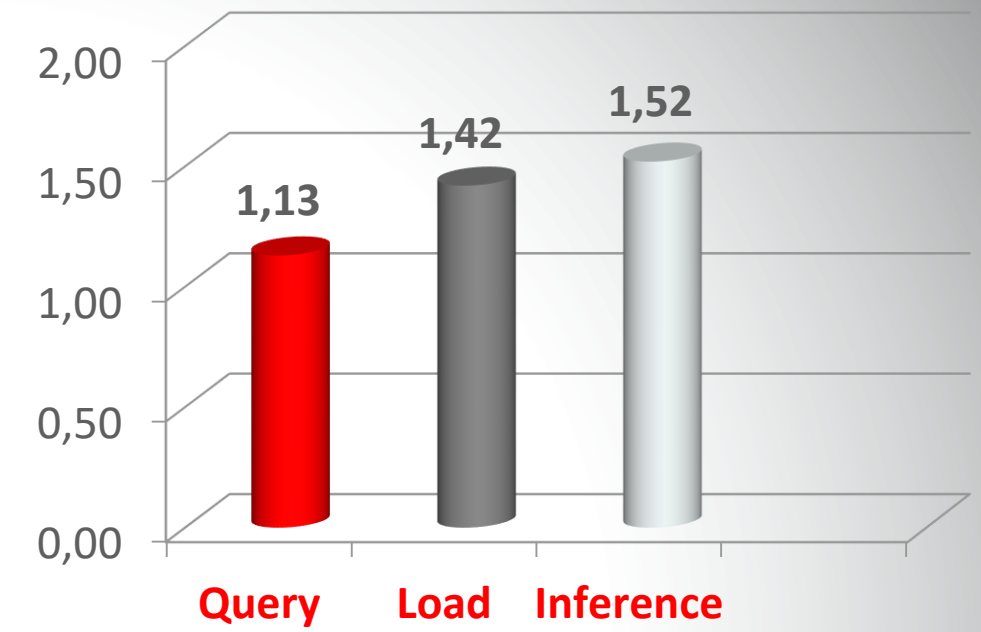

Oracle Database 12*c* RDF Semantic Graph Database

- Exadata ready
- Compression & partitioning
- Parallel load, inference, query
- High availability
- Label security: triple-level
- W3C standards compliance
- Semantic Indexing of text
- Enterprise Manager
- Support for Open Source
  - Development framework, ontology editing, visualization

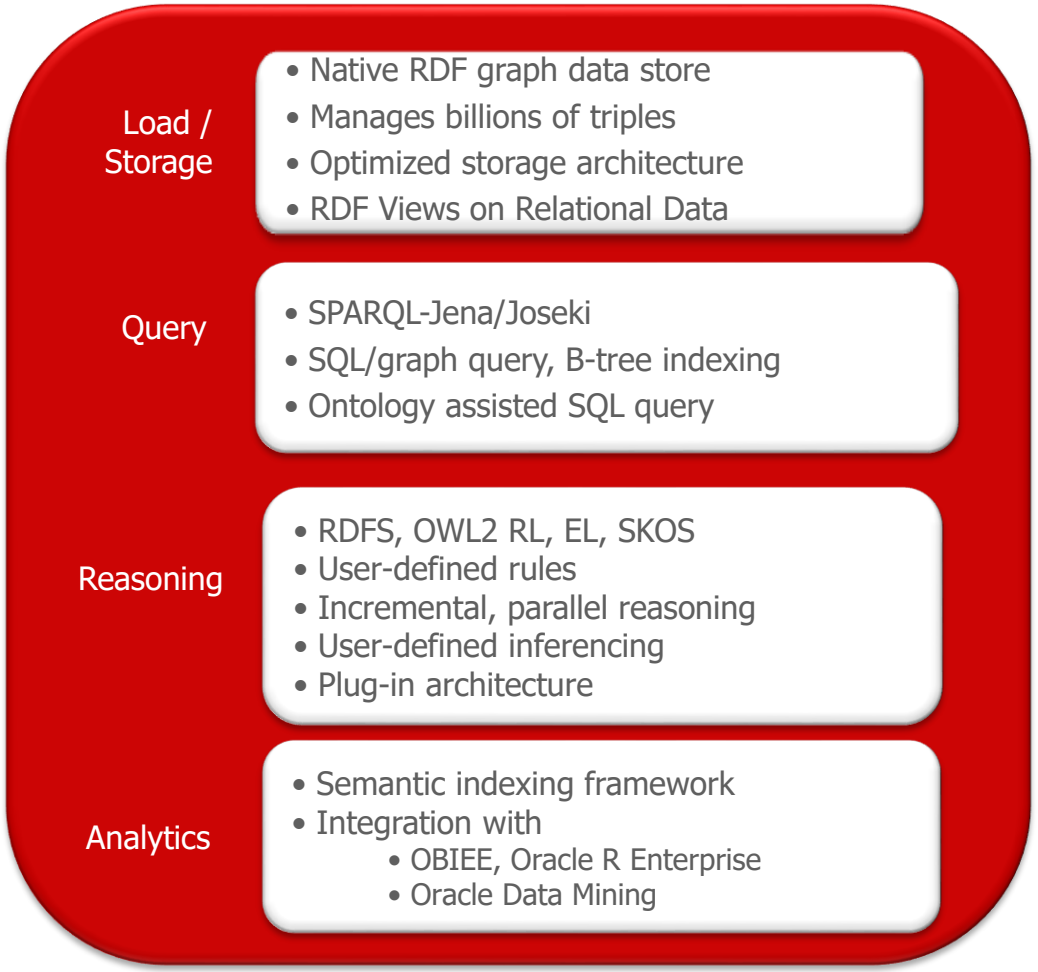

Oracle Participation in W3C & OGC Standards

- W3C Web Semantic Activities:
  - W3C RDF Working Group (Data Model)
  - W3C SPARQL Working Group (Query)
  - W3C RDB2RDF Working Group (Data Mapping)
  - W3C OWL Working group (Rules / Inferencing)
  - W3C Semantic Web Rules Language (SWRL)
  - Spatial Data on the Web (Community Group)
- OGC GeoSPARQL Standard Working Group

Integrated Cloud

Applications & Platform Services


ORACLE®