

ORACLE®

How to Analyze Data Warehouse Data as a Graph

How graph databases can complement relational technology

Hans Viehmann
Product Manager EMEA
ORACLE Corporation
HOUG Conference 2019

 @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.

Graph – an important growth area for data & analytics

Gartner Identifies Top 10 Data and Analytics Technology Trends for 2019



Trend No. 5: Graph

Graph analytics is a set of analytic techniques that allows for the exploration of relationships between entities of interest such as organizations, people and transactions.

The application of graph processing and graph DBMSs will grow at 100 percent annually through 2022 to continuously accelerate data preparation and enable more complex and adaptive data science.

Graph data stores can efficiently model, explore and query data with complex interrelationships across data silos, but the need for specialized skills has limited their adoption to date, according to Gartner.

Graph analytics will grow in the next few years due to the need to ask complex questions across complex data, which is not always practical or even possible at scale using SQL queries.

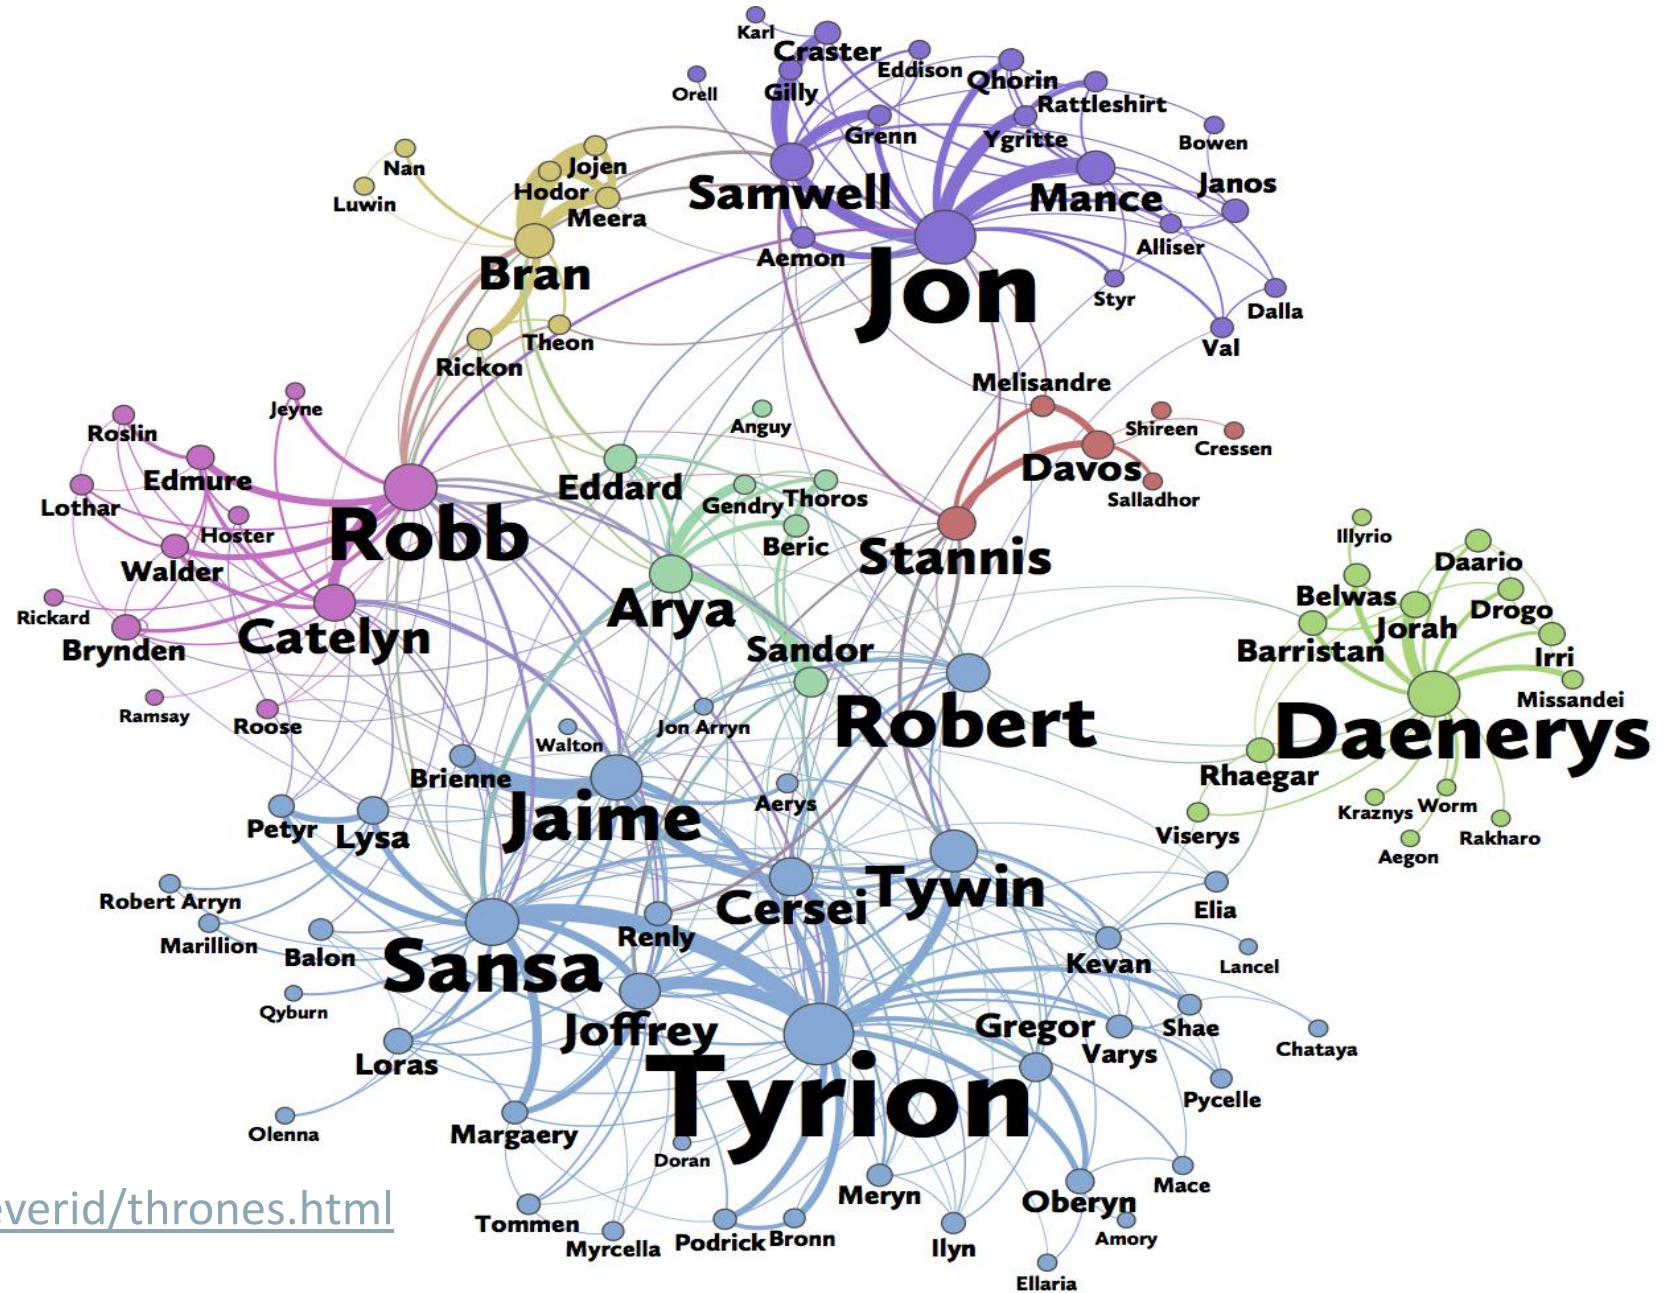
Source: Gartner press release, 2/18/2019, www.gartner.com/en/newsroom/press-releases/2019-02-18-gartner-identifies-top-10-data-and-analytics-technolo

A social network generated from Game of Thrones.

The color of a vertex indicates its community.

The size of a vertex corresponds to its PageRank value, and the size of its label corresponds to its betweenness centrality.

An edge's thickness represents its weight.




<https://www.macalester.edu/~abeverid/thrones.html>




Graph Data Model

- What is a graph?
 - Data model representing entities as vertices and relationships as edges
 - Optionally including attributes
 - Also known as „linked data“
- What are typical graphs?
 - Social Networks
 - LinkedIn, facebook, Google+, ...
 - IP Networks, physical networks, ...
 - Knowledge Graphs
 - Apple SIRI, Google Knowledge Graph, ...



Graph Data Model

- Why are graphs popular?
 - Easy data modeling
 - „whiteboard friendly“
 - Flexible data model
 - No predefined schema, easily extensible
 - Particularly useful for sparse data
 - Insight from graphical representation
 - Intuitive visualization
 - **Enabling new kinds of analysis**
 - Overcoming some limitations in relational technology
 - Basis for Machine Learning (Neural Networks)



Oracle's Spatial and Graph Strategy


Enabling Spatial and Graph use cases on every platform

Oracle Database
Spatial and Graph Option



Exadata
Non-Engineered Systems

Oracle Big Data
Spatial and Graph



Big Data Appliance
Commodity Hadoop
Spark

Cloud
Services



Database Cloud Service
Exadata Cloud Service

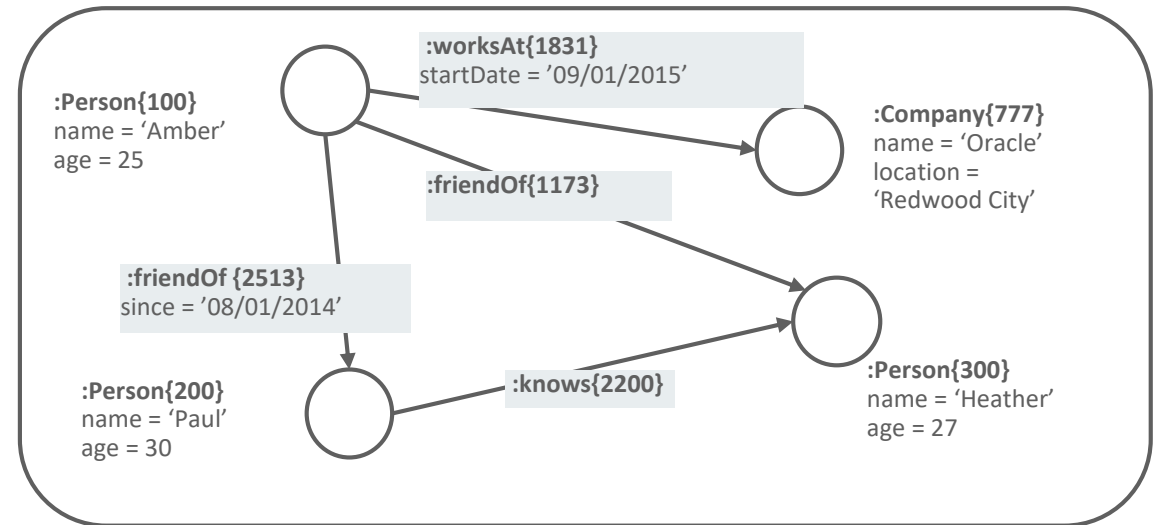
Categories of Graph Analysis

Computational Graph Analytics

- Compute values on vertices and edges
- Traversing graph or iterating over graph (usually repeatedly)
- Procedural logic
- Examples:
 - Shortest Path, PageRank, Weakly Connected Components, Centrality, ...

Graph Pattern Matching

- Based on description of pattern
- Find all matching sub-graphs

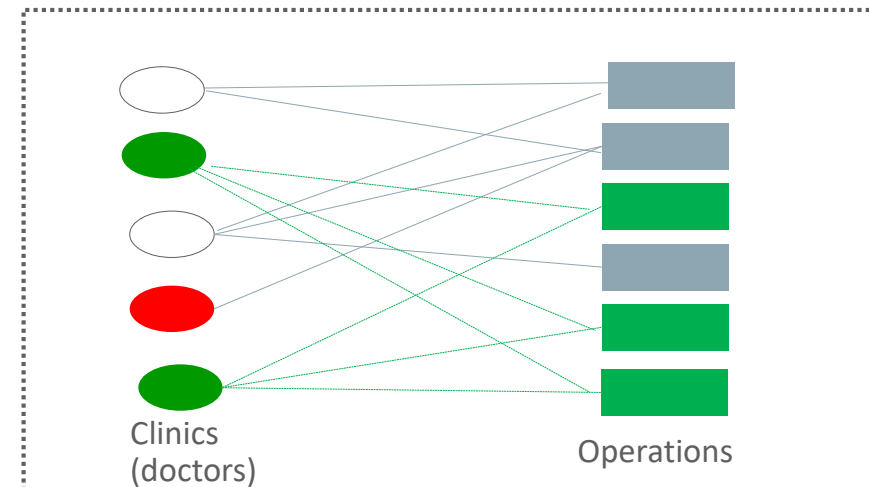


Example: Fraud Detection in Healthcare

- Finding anomalies in healthcare billing data
 - Public domain dataset
 - Medical providers and their operations
- Question
 - Are there any medical providers that are suspicious
 - medical providers that perform different operations than their fellows (e.g. eye doctors doing plastic surgery)

- Approach

- Create graph between doctors and operations
- Apply personalized pagerank (a.k.a equivalent to random walking)
- Identify doctors that are *far* from their fellows



Pattern matching in Property Graphs using PGQL

- Finding a given pattern in graph
 - Fraud detection
 - Anomaly detection
 - Subgraph extraction
 - ...
- SQL-like syntax but with graph pattern description and property access
 - Interactive (real-time) analysis
 - Supporting aggregates, comparison, such as max, min, order by, group by

- Proposed for standardization by Oracle
 - Specification available on-line
 - Open-sourced front-end (i.e. parser)



<https://github.com/oracle/pgql-lang>

PGQL · Property Graph Query Language

An SQL-like query language for graphs

Try It

Graphs + SQL

PGQL is a graph pattern matching query language for the [property graph data model](#), inspired by [Cypher](#), [SQL](#), and [G-CORE](#). PGQL combines Cypher-like [ASCII art syntax](#) with familiar constructs from SQL, such as `SELECT`, `FROM` and `WHERE`. PGQL also provides powerful constructs for matching regular path expressions (e.g. `PATH`).

An example PGQL query is as follows:

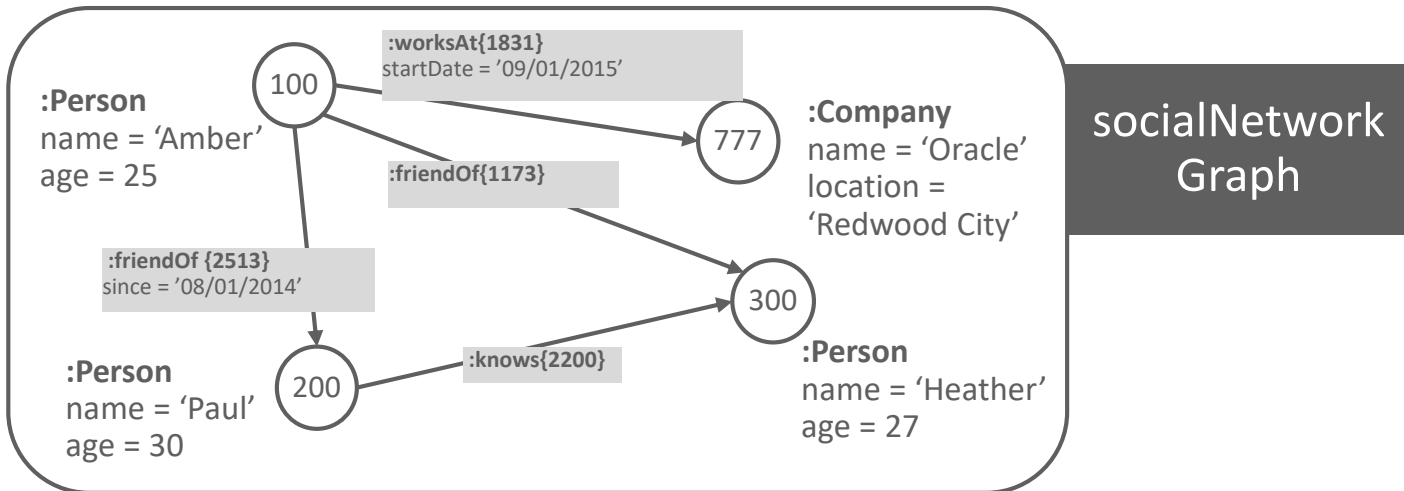
```
SELECT p2.name AS friend_of_friend
FROM facebook_graph
MATCH (p1:Person)-/:friend_of{2}/->(p2:Person) /* ..match two-hop friends.. */
WHERE p1.name = 'Mark' /* ..of Mark. */
```

See [PGQL 1.1 Specification](#) for a detailed specification of the language.

Basic graph pattern matching

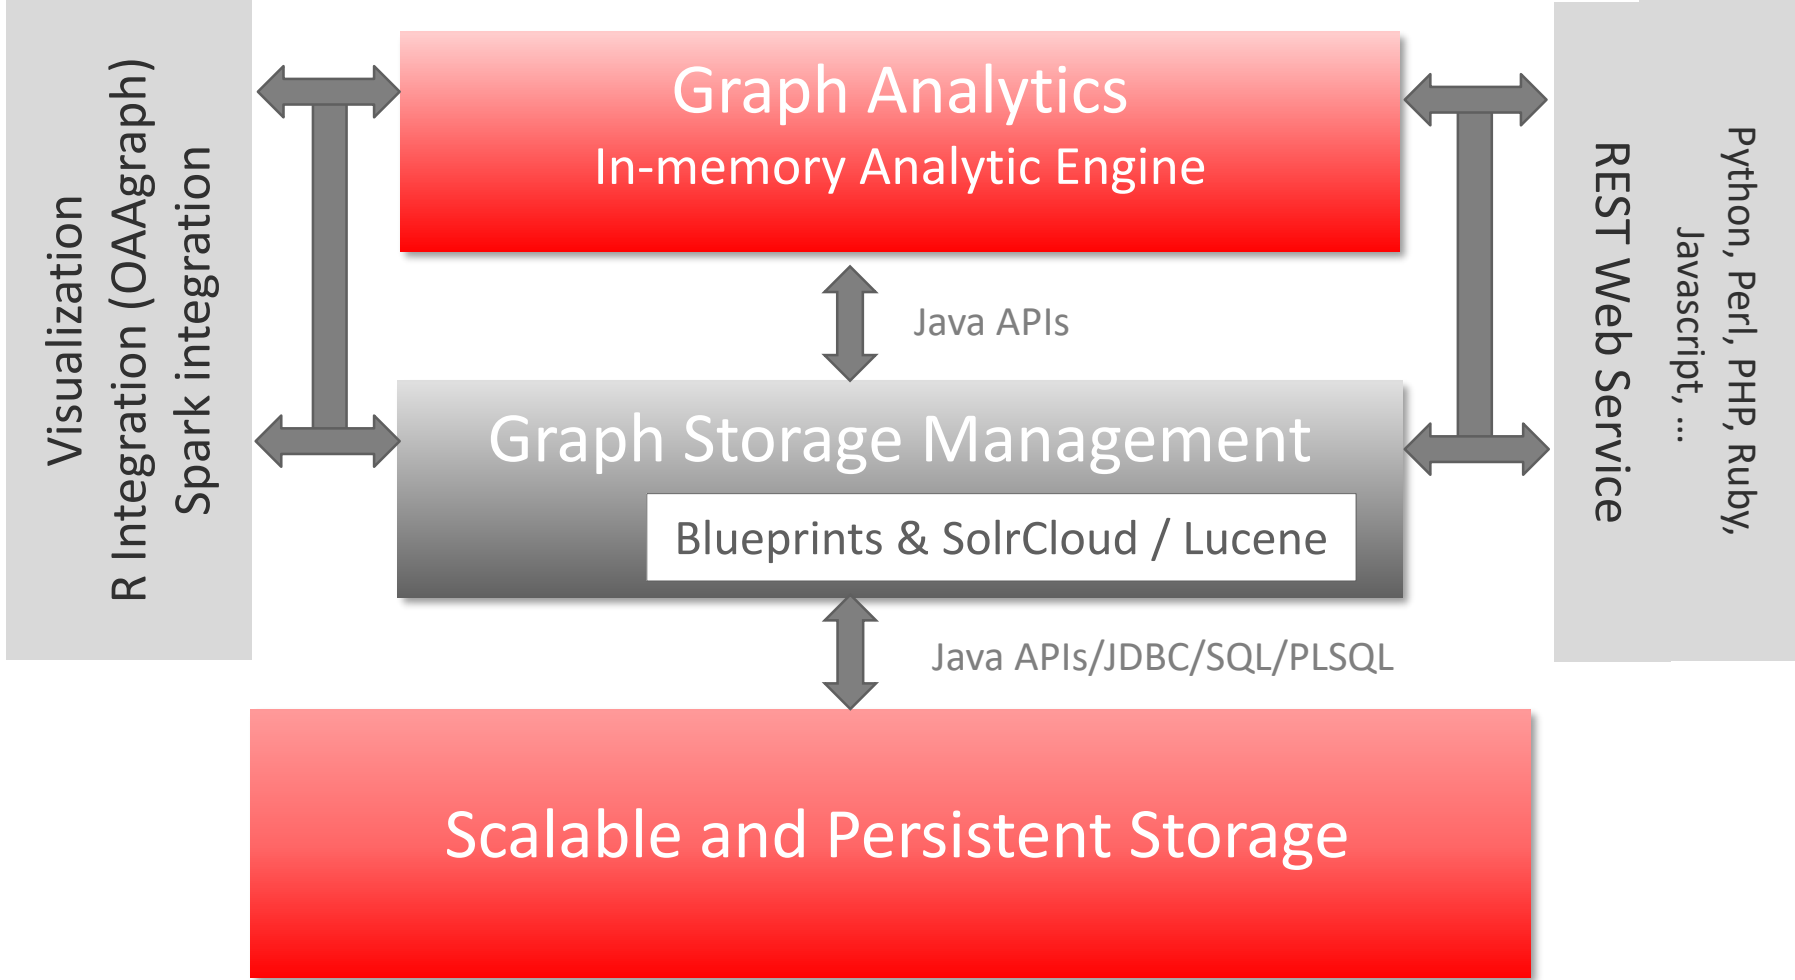
- Find all instances of a given pattern/template in the data graph

```
SELECT v3.name, v3.age
FROM socialNetworkGraph
MATCH (v1:Person) -[:friendOf]-> (v2:Person) -[:knows]-> (v3:Person)
WHERE v1.name = 'Amber'
```

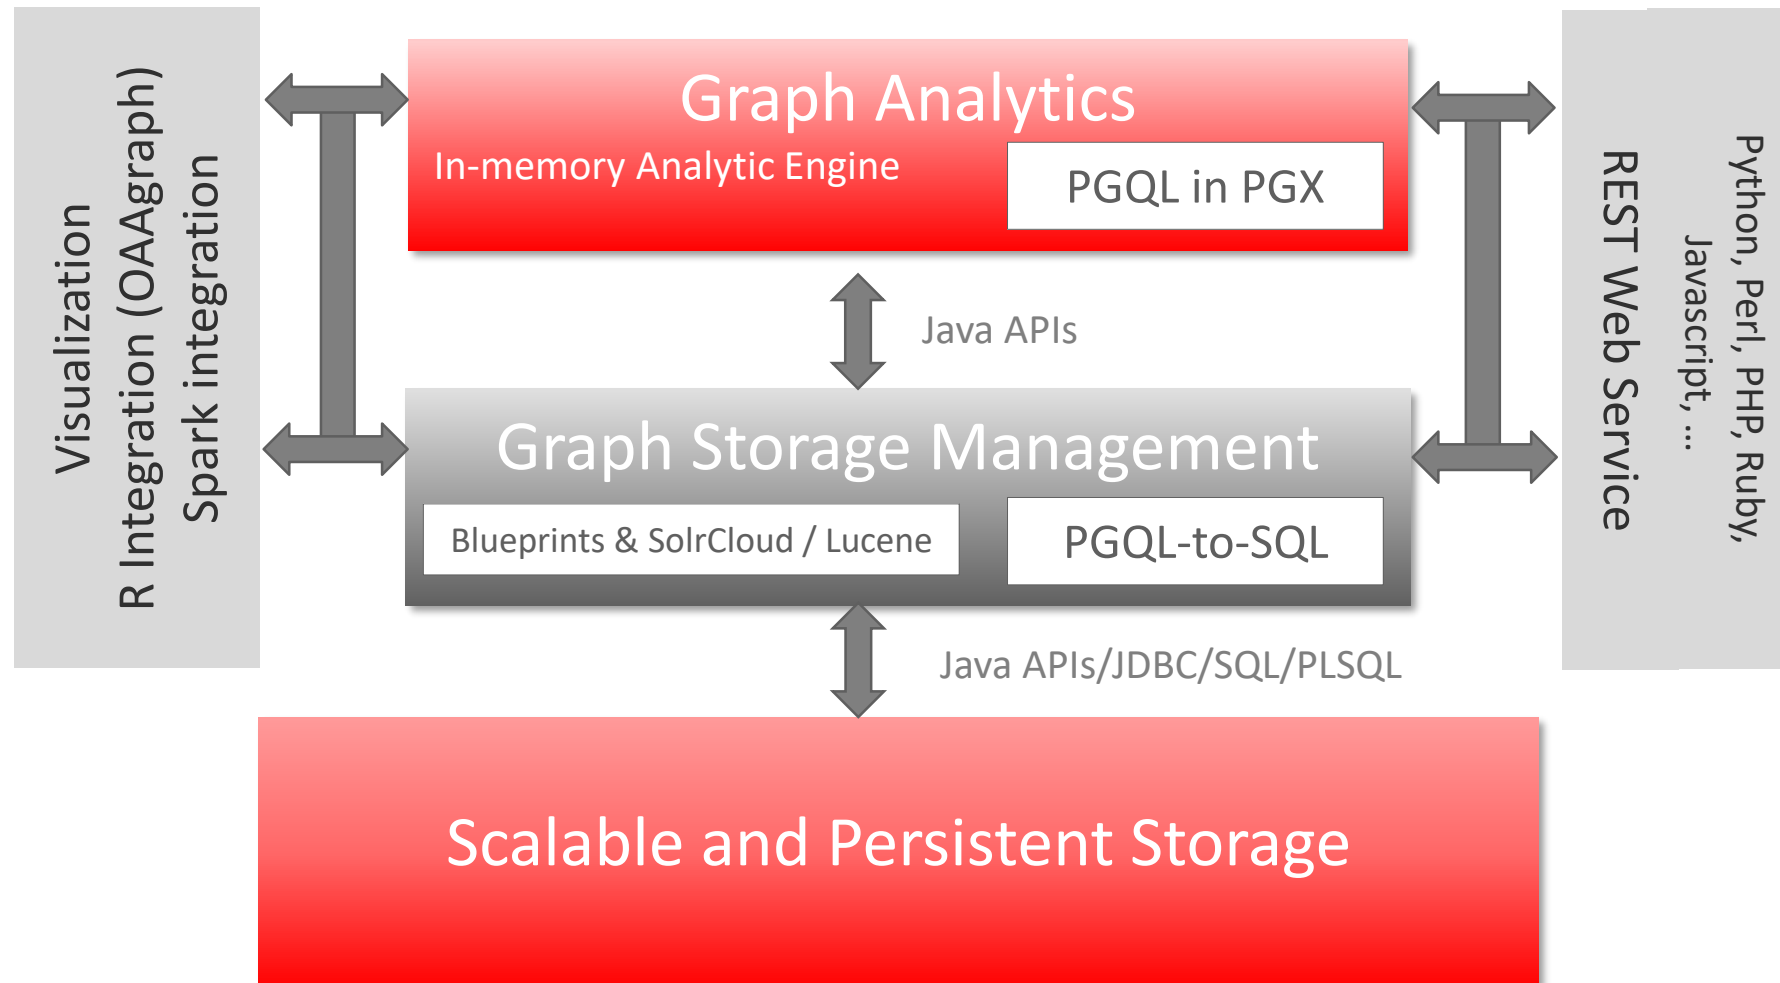


Query: Find all people who are known by friends of 'Amber'.

Oracle Graph Analytics Architecture



Support for Graph Pattern Matching



Path Query (Parallel Recursive With)

PGQL:

```
PATH knows_path := () -[:knows]-> ()
SELECT s1.fname, s2.fname
WHERE (s1) -/[:knows_path*]/-> (o) <-/[:knows_path*]/-(s2)
ORDER BY s1,s2
```

Find the **pairs of people** who are connected to a common person through the **“knows”** relation

SQL:





```
SELECT T2.T AS "s1.fname$T",T2.V AS "s1.fname$V",T2.VN AS "s1.fname$VN",T2.VT AS "s1.fname$VT",
       T3.T AS "s2.fname$T",T3.V AS "s2.fname$V",T3.VN AS "s2.fname$VN",T3.VT AS "s2.fname$VT"
FROM (/*Path[*/SELECT DISTINCT SVID, DVID FROM ( SELECT VID AS SVID, VID AS DVID FROM "GRAPH1VT$" UNION ALL SELECT SVID,DVID
FROM (WITH RW (ROOT, SVID, DVID, LVL) AS ( SELECT ROOT, SVID, DVID, LVL FROM (SELECT SVID ROOT, SVID, DVID, 1 LVL
FROM (SELECT T0.SVID AS SVID, T0.DVID AS DVID FROM "GRAPH1GT$" T0 WHERE (T0.EL = n'knows'))
) UNION ALL SELECT DISTINCT RW.ROOT, R.SVID, R.DVID, RW.LVL+1 FROM (SELECT T1.SVID AS SVID,
T1.DVID AS DVID FROM "GRAPH1GT$" T1 WHERE (T1.EL = n'knows')) R, RW WHERE RW.DVID = R.SVID )
CYCLE SVID SET cycle_col TO 1 DEFAULT 0 SELECT ROOT SVID, DVID FROM RW )/*]Path*/) T6,
/*Path[*/SELECT DISTINCT SVID, DVID FROM ( SELECT VID AS SVID, VID AS DVID FROM "GRAPH1VT$" UNION ALL SELECT SVID,DVID
FROM (WITH RW (ROOT, SVID, DVID, LVL) AS ( SELECT ROOT, SVID, DVID, LVL FROM (SELECT SVID ROOT, SVID, DVID, 1 LVL
FROM (SELECT T4.SVID AS SVID, T4.DVID AS DVID FROM "GRAPH1GT$" T4 WHERE (T4.EL = n'knows'))
) UNION ALL SELECT DISTINCT RW.ROOT, R.SVID, R.DVID, RW.LVL+1 FROM (SELECT T5.SVID AS SVID,
T5.DVID AS DVID FROM "GRAPH1GT$" T5 WHERE (T5.EL = n'knows')) R, RW WHERE RW.DVID = R.SVID )
CYCLE SVID SET cycle_col TO 1 DEFAULT 0 SELECT ROOT SVID, DVID FROM RW )/*]Path*/) T7,
"GRAPH1VT$" T2, "GRAPH1VT$" T3
WHERE T2.K=n'fname' AND T3.K=n'fname' AND T6.SVID=T2.VID AND T6.DVID=T7.DVID AND T7.SVID=T3.VID
ORDER BY T6.SVID ASC NULLS LAST, T7.SVID ASC NULLS LAST
```

Browser: Datasources | Not secure | slc14rui:32000/graphstudio/?root=ds-datasources

ORACLE Graph Cloud | Search | ogcs-dev-user

Data Sources Create

Search Sort by:

Type	Name	Description	Owner	Status	Last Applied
	Financial Demo Dataset	Financial Demo Dataset	ogcs-dev-user	Tue, 16 Oct 2018 20:18:04 GMT	
	slcai607	Sales Demo Dataset	ogcs-dev-user	Fri, 12 Oct 2018 21:30:25 GMT	



A woman with long brown hair and glasses is sitting at a wooden table in a bright, modern office or cafe. She is wearing a brown leather jacket over a blue patterned scarf. She is holding a black mobile phone to her ear with her left hand and looking down at a newspaper or magazine on the table with her right hand. The background is slightly blurred, showing other people and large windows.

Summary

Summary

Graph capabilities in Oracle Big Data Spatial and Graph





- Graph databases are powerful tools, complementing relational databases
 - Especially strong for analysis of graph topology and connectedness
- Graph analytics offer new insight
 - Especially relationships, dependencies and behavioural patterns
- Oracle Property Graph technology offers
 - Comprehensive analytics through various APIs, integration with relational database
 - Scalable, parallel in-memory processing
 - Secure and scalable graph storage using Hadoop platform or Oracle Database
- Available both on-premise or in the Cloud already today

Q&A



More information

- Oracle Big Data Spatial and Graph OTN product page:
www.oracle.com/technetwork/database/database-technologies/bigdata-spatialandgraph
– White papers, software downloads, documentation and videos
- Oracle Big Data Lite Virtual Machine - a free sandbox to get started:
www.oracle.com/technetwork/database/bigdata-appliance/oracle-bigdatalite-2104726.html
- Hands On Lab included in `/opt/oracle/oracle-spatial-graph/`
– Content also available on GitHub under <http://github.com/oracle/BigDataLite/>
- Blog – examples, tips & tricks: blogs.oracle.com/bigdataspatialgraph
-  @OracleBigData, @SpatialHannes, @Jeanlhm  Oracle Spatial and Graph Group

Introduction to Graph analytics

Youtube videos

- What is Oracle Big Data Spatial and Graph?
<https://youtu.be/t9pJJhzZKOE>

How can graph analytics help my business?
<https://youtu.be/0dJNzBi7B-k>

Detecting anomalies with Oracle Big Data Spatial and Graph
<https://youtu.be/nfP6HD0ImjY>

Generating recommendations with Oracle Big Data Spatial and Graph
<https://youtu.be/9LRIF3of-Hs>



Integrated Cloud

Applications & Platform Services

ORACLE®