











Graph Databases, including Knowledge Graphs

Dr Mercedes Arguello Casteleiro
Lecturer, University of Southampton, UK
Visiting academic, University of Manchester, UK

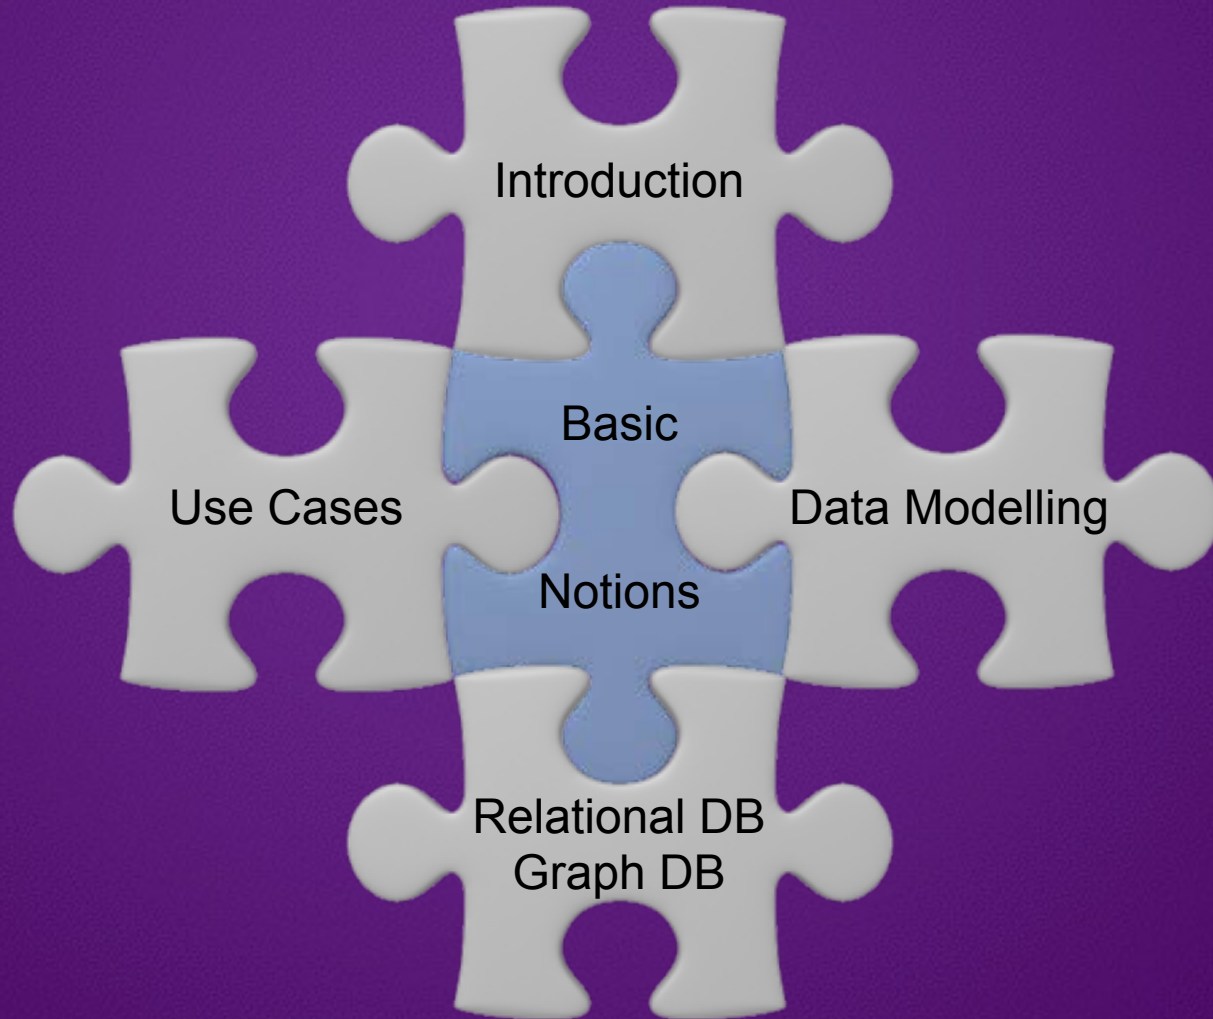
Top-10 Graph databases

<https://db-engines.com/en/ranking/graph+dbms>

39 systems in ranking, June 2023

Rank			DBMS	Database Model	Score		
Jun 2023	May 2023	Jun 2022			Jun 2023	May 2023	Jun 2022
1.	1.	1.	Neo4j +	Graph	52.77	+1.66	-6.76
2.	2.	2.	Microsoft Azure Cosmos DB +	Multi-model 	36.57	+0.58	-4.41
3.	3.	3.	Virtuoso +	Multi-model 	5.24	-0.33	-0.93
4.	4.	4.	ArangoDB +	Multi-model 	4.89	+0.01	-0.61
5.	5.	5.	OrientDB	Multi-model 	4.53	+0.03	-0.33
6.	6.	6.	Amazon Neptune	Multi-model 	3.03	+0.13	+0.21
7.	7.	 8.	JanusGraph	Graph	2.83	+0.15	+0.43
8.	8.	 19.	Memgraph +	Graph	2.82	+0.18	+2.38
9.	9.	 15.	NebulaGraph +	Graph	2.72	+0.11	+1.61
10.	10.	 7.	GraphDB +	Multi-model 	2.55	+0.07	-0.07

Part 1: Graph databases

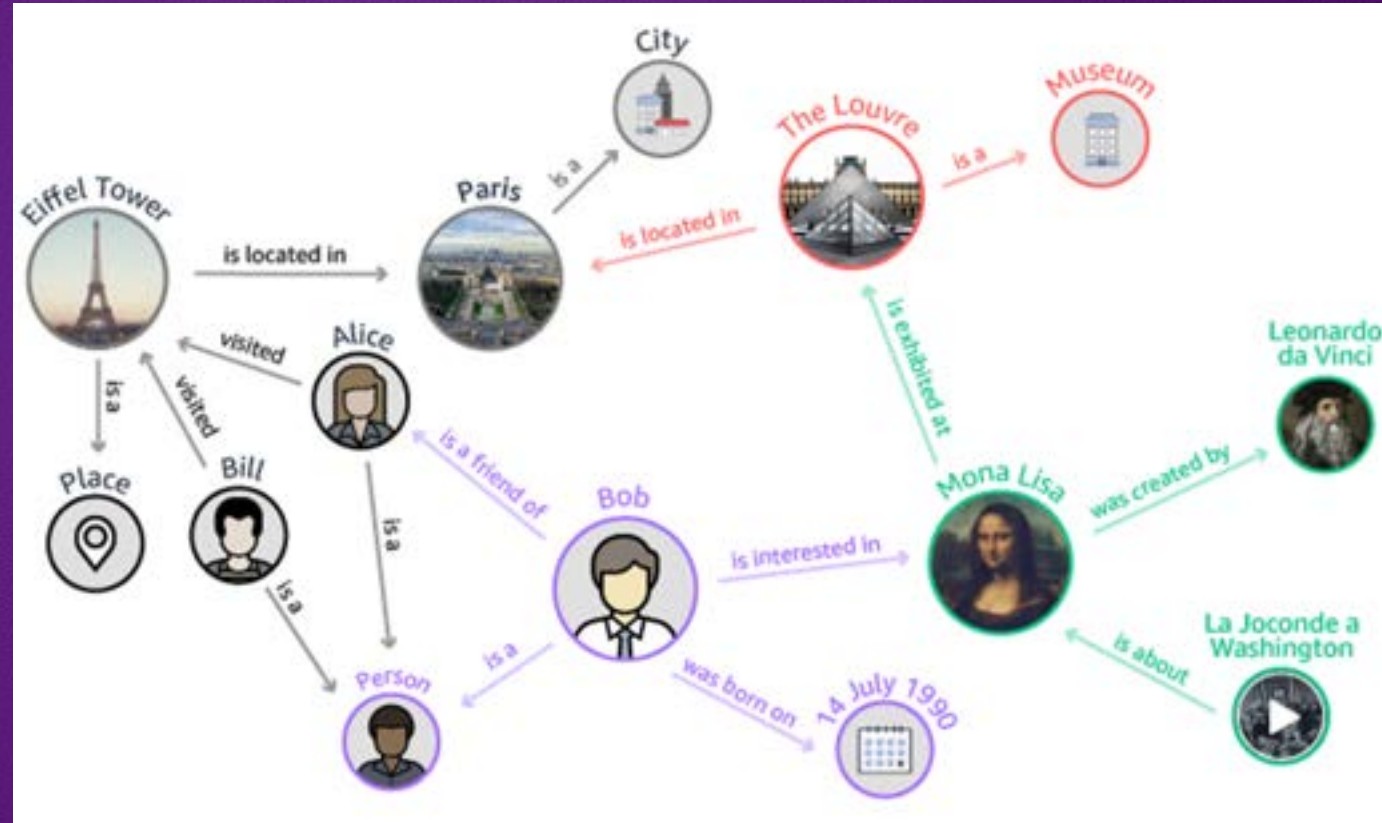


Using Amazon Neptune to build an Enterprise Knowledge Graph

<https://aws.amazon.com/neptune/knowledge-graphs-on-aws/>



Knowledge
Graph




Graphs as networks

A knowledge graph captures the semantics of a particular domain

Neo4j graph platform

Neo4j is the company behind Neo4j graph platform



Neo4j database is designed to store, reveal, and query relationships

Customers

LinkedIn

UBS

COMCAST

CISCO

Lufthansa

Walmart

ebay

Partners



Google Cloud Platform

amazon
web services

IBM

Microsoft
Azure



TATA

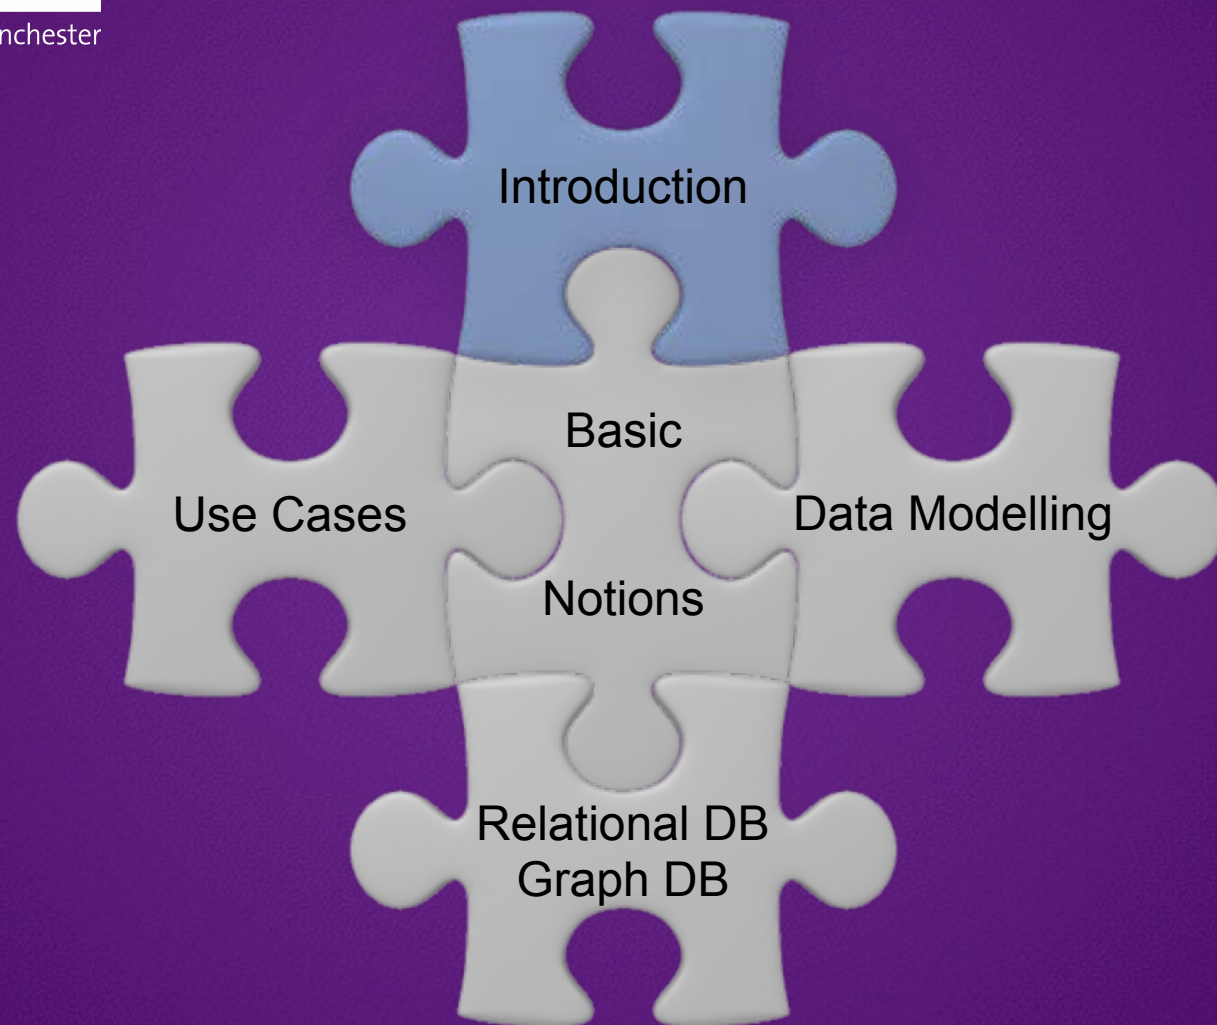
LINKURIUS
VISUALIZE GRAPH DATA EASILY



pitney bowes




GraphAware
structr




Evolution of DBs

SQL database

Relational




Analytical (OLAP)




online
analytical
processing
(OLAP)

NoSQL database


Column-Family




Graph



Document



Key-Value




...so... why using graph DBs?


For highly interconnected data, graph model seems the most natural.

One to Many

HIERARCHICAL




RELATIONAL



Any to Any

NETWORK



What problems the graph DBs were created to solve?

Data Measurement Chart

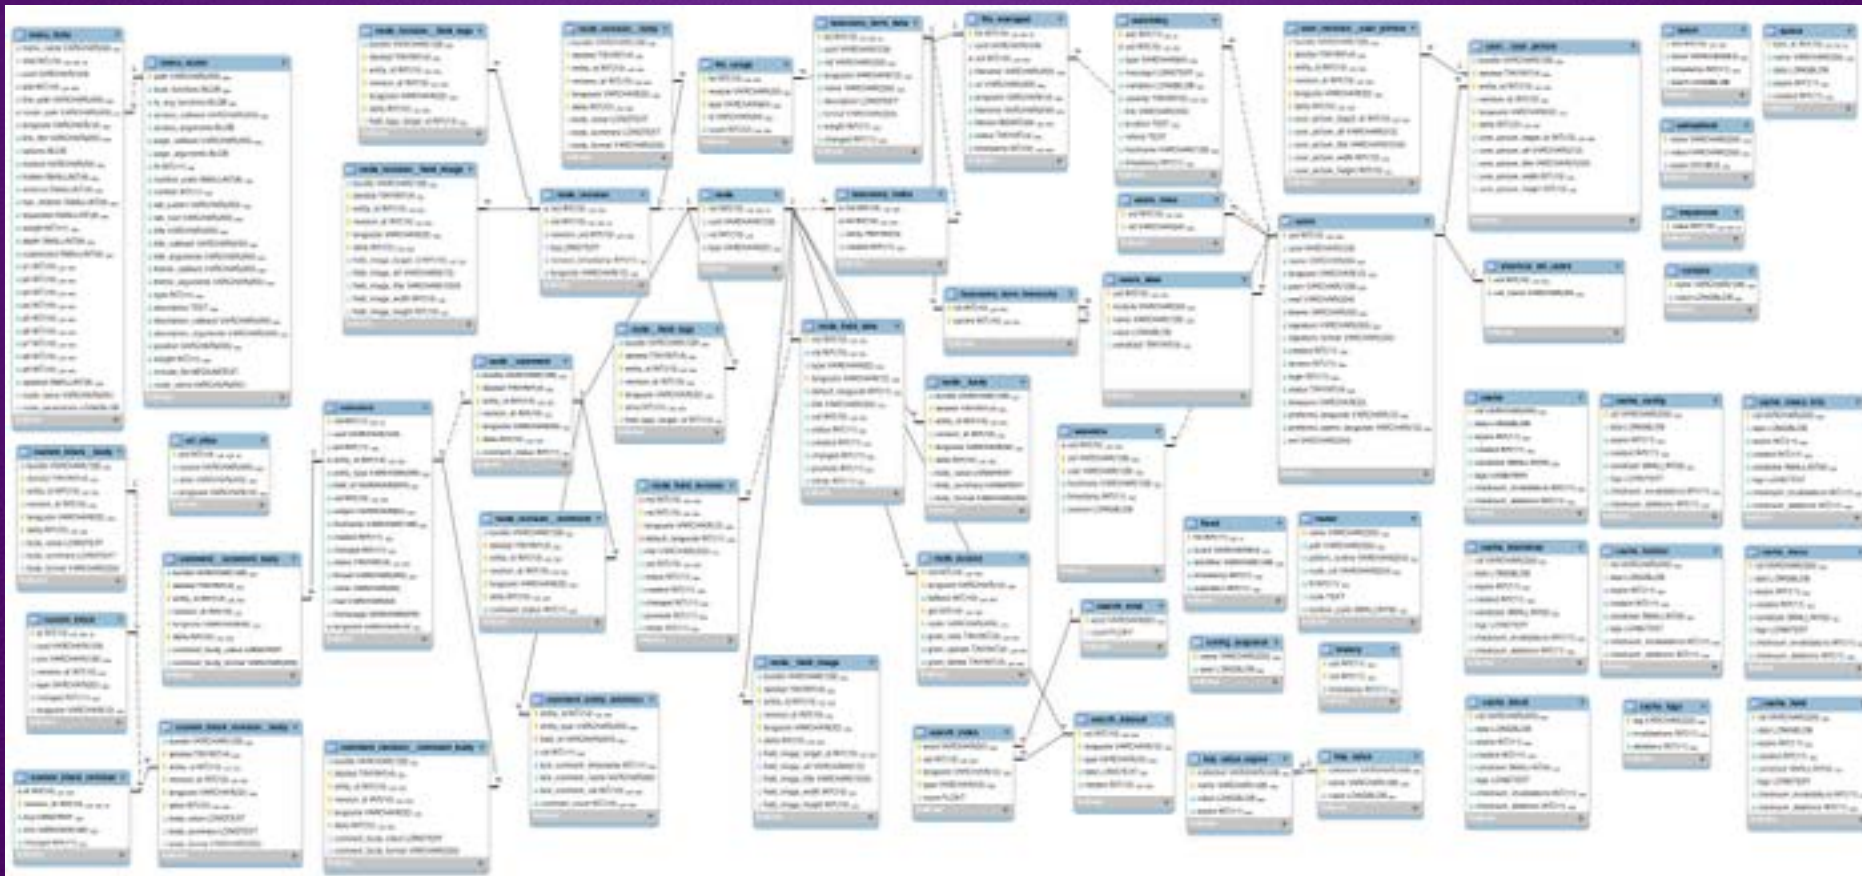
Data Measurement	Size
Bit	Single Binary Digit (1 or 0)
Byte	8 bits
Kilobyte (KB)	1,024 Bytes
Megabyte (MB)	1,024 Kilobytes
Gigabyte (GB)	1,024 Megabytes
Terabyte (TB)	1,024 Gigabytes
Petabyte (PB)	1,024 Terabytes
Exabyte (EB)	1,024 Petabytes

With the advent of the Cloud, there has been a data explosion with Exabytes of data available



What problems the graph DBs were created to solve?

The normalisation of data gets more and more expensive as the data size grows



What problems the graph DBs were created to solve?

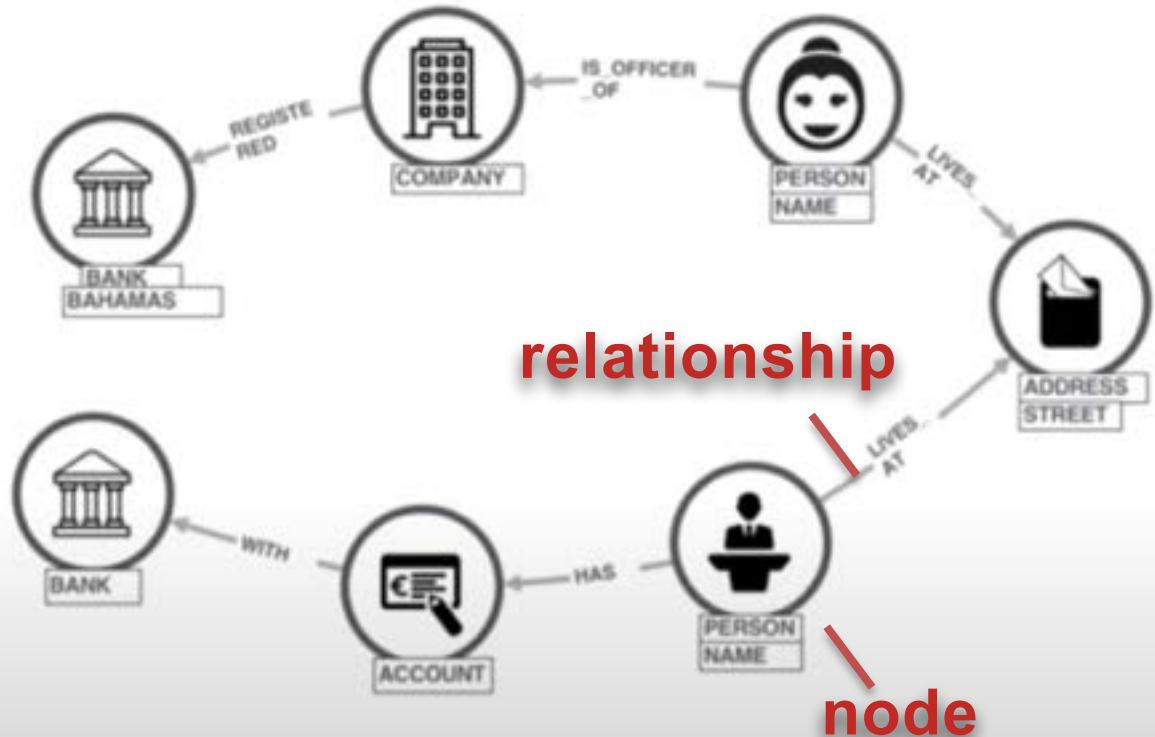
...so... how the graph DBs deal with data explosion?

11,5 million documents

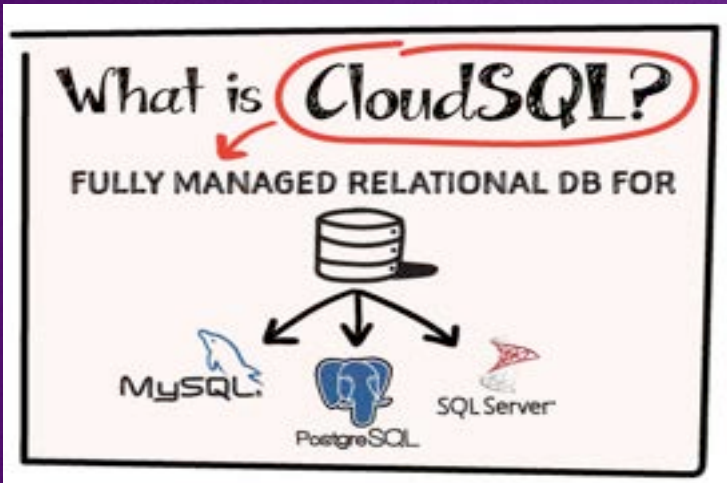
Emails, Scanned Documents,
Bank Statements etc...



2.6 TB



What problems the graph DBs were created to solve?



With the advent of the Cloud, there has been a data explosion with Exabytes of data available



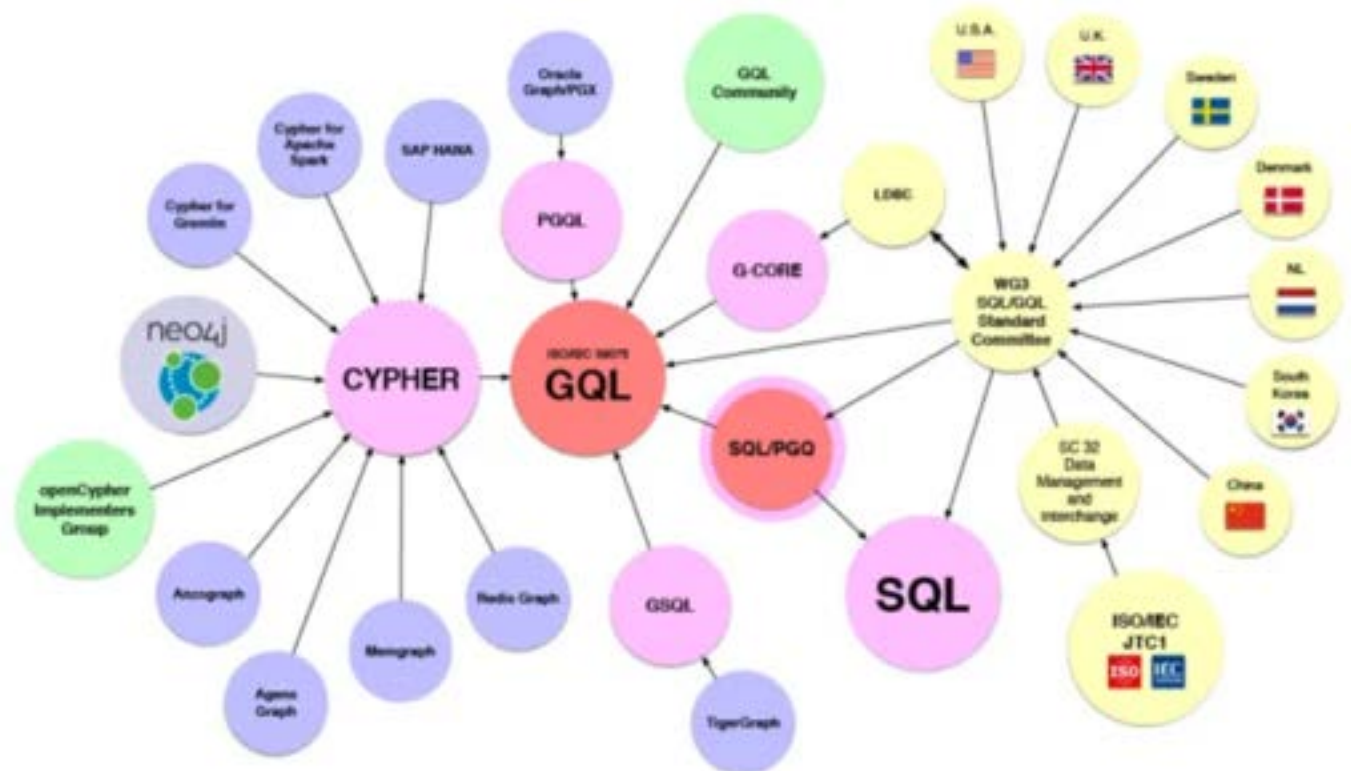
Another difficulty that is equally important is that, the data becomes also difficult to maintain needing complicated joined queries

GQL Manifesto

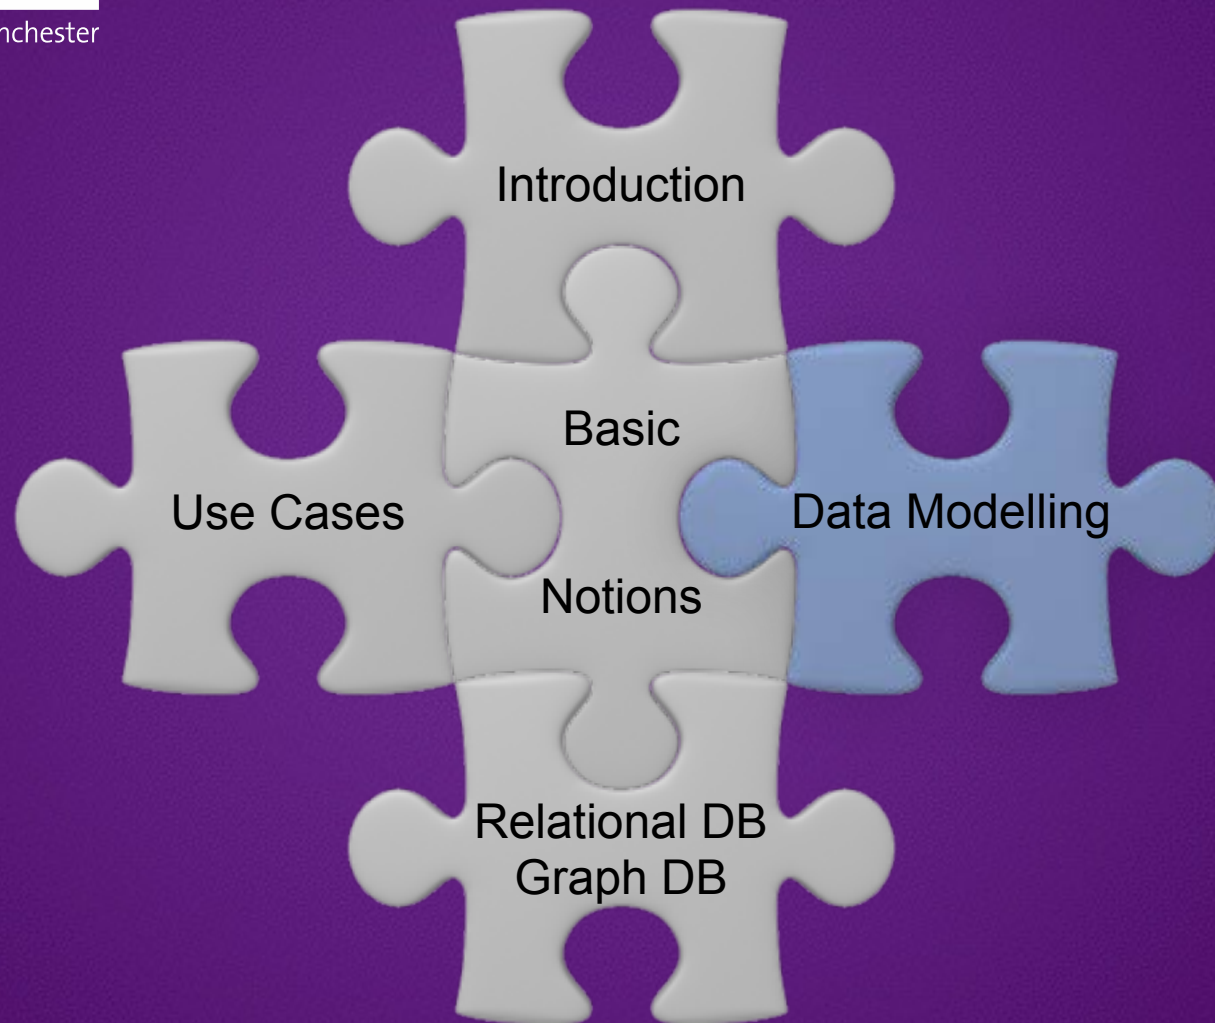
GQL stands for Graph Query Language

GQL Manifesto similar to SQL or Third manifesto

GQL Is Now a Global Standards Project alongside SQL



Like SQL, the new GQL (Graph Query Language) needs to be an industry standard.

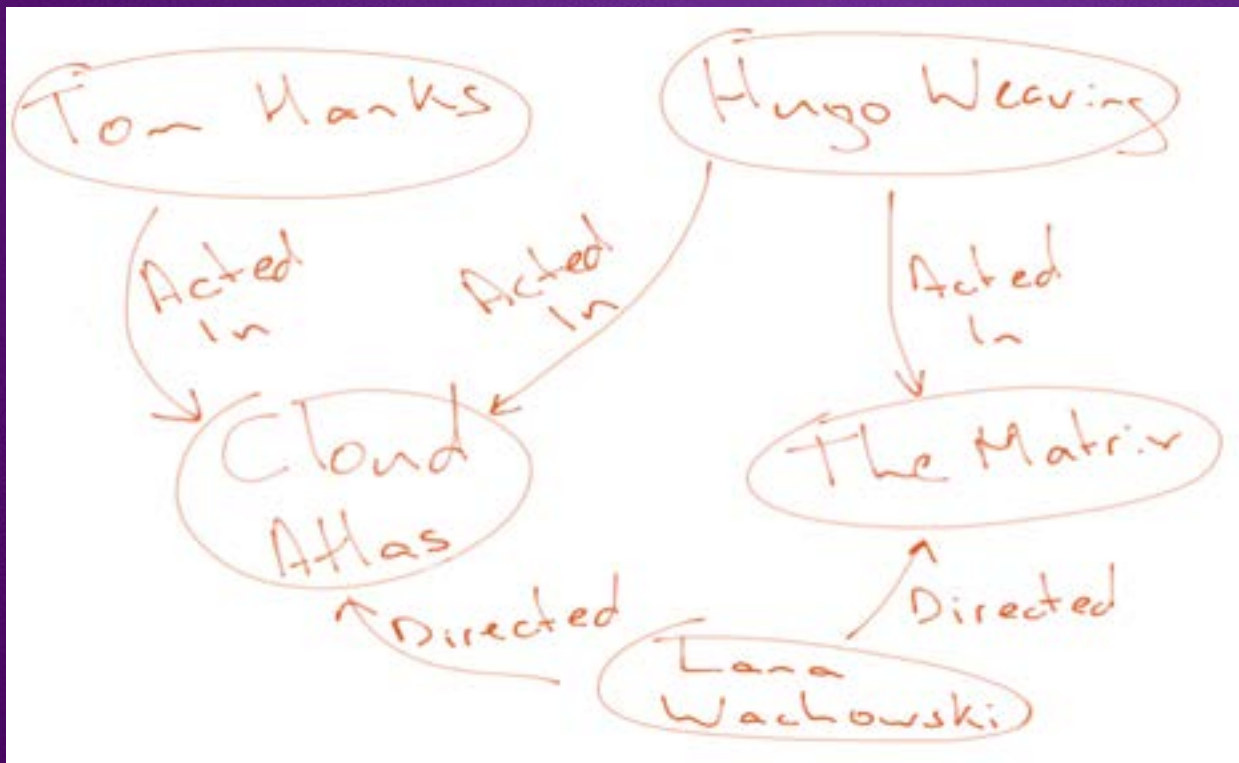


Properties of graph DBs

For a developer, attractive properties of a graph DB are

Intuitiveness = whiteboard-friendly

The graph model is intuitive and easy to understand.



Intuitiveness

Speed

Agility

Properties of graph DBs

For a developer, attractive properties of a graph DB are

Intuitiveness

Speed

Agility

Speed in development and execution



...create data...

```

CREATE (a:Person { name:"Tom Hanks",
  born:1956 })-[r:ACTED_IN { roles: ["Forrest"]}]>(m:Movie { title:"Forrest Gump",released:1994 })
CREATE (d:Person { name:"Robert Zemeckis", born:1951 })-[:DIRECTED]->(m)
RETURN a,d,r,m
  
```

...find interesting connections...

```

MATCH (p:Person { name:"Tom Hanks" })-[r:ACTED_IN]->(m:Movie)
RETURN m.title, r.roles
  
```


Properties of graph DBs

For a developer, attractive properties of a graph DB are

Intuitiveness

Speed

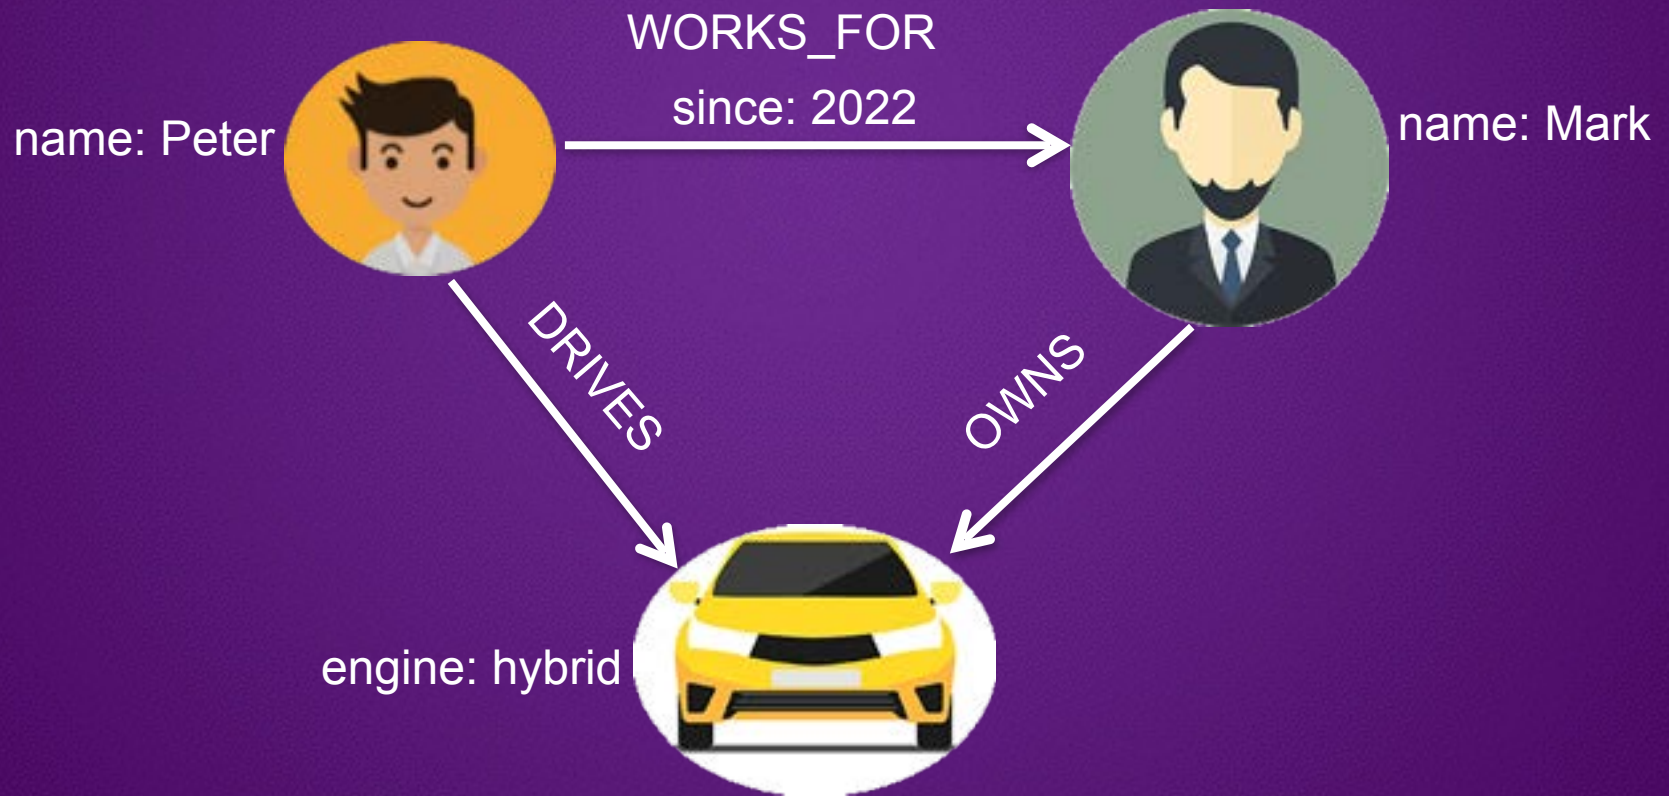
Agility

Agility = A Naturally Adaptive Model +
A Query Language Design for Connectedness

How easily and quickly your code adapt of the changing business.

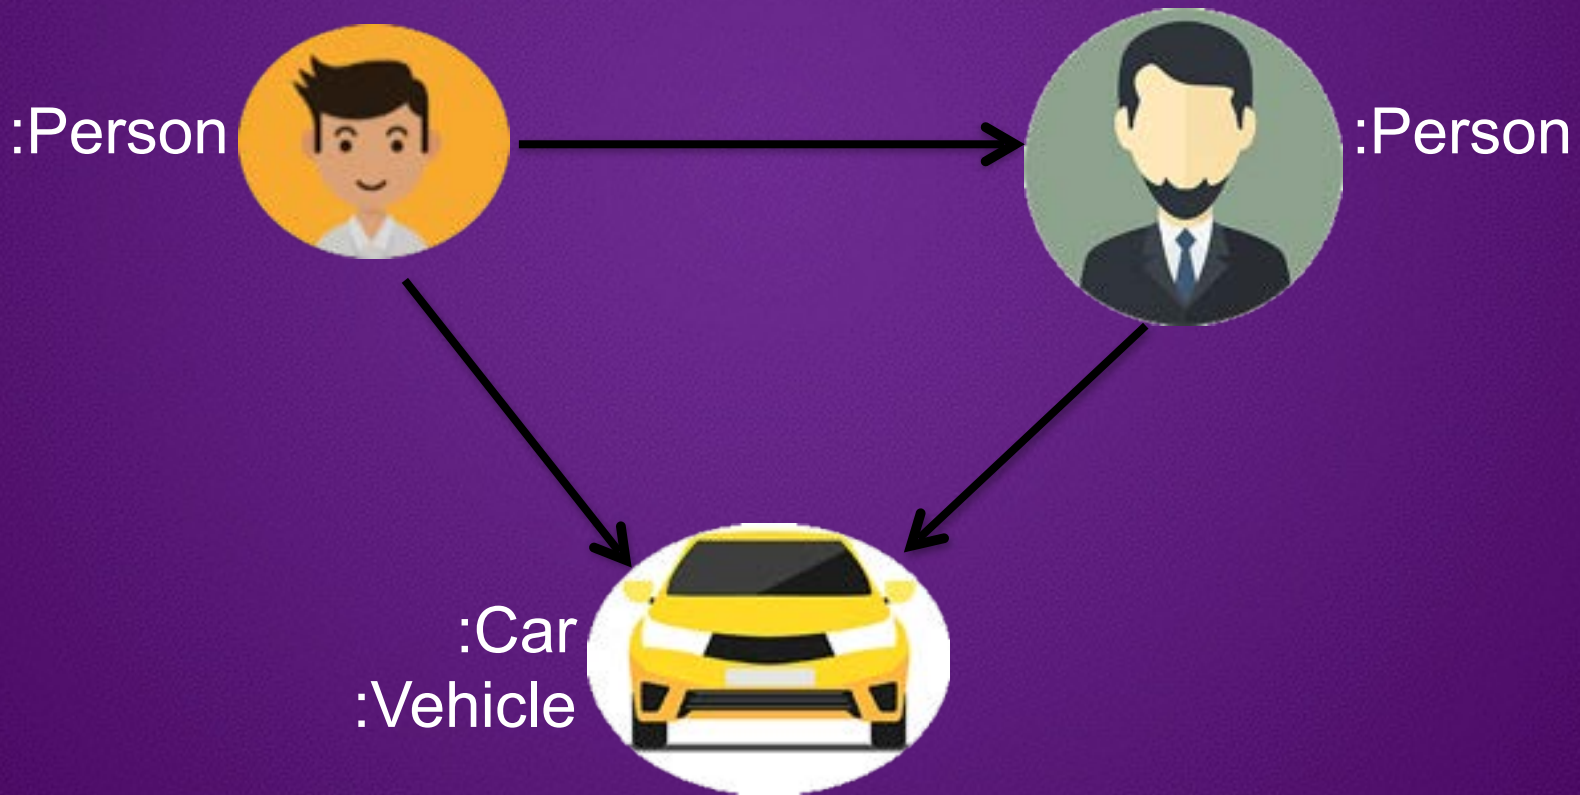
Detailed Property Graph

Let's start with a Detailed Property Graph



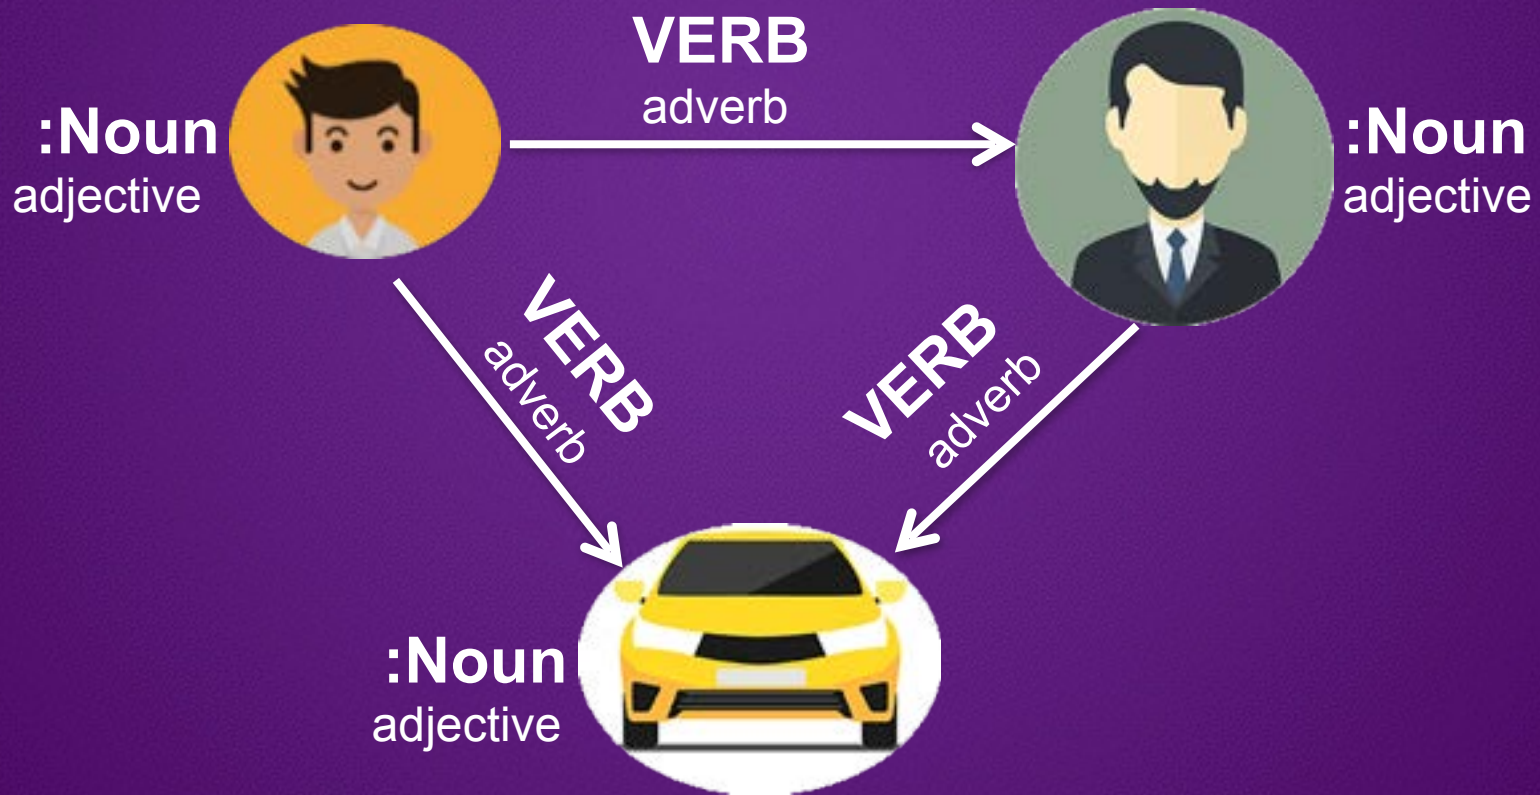
Detailed Property Graph

A Detailed Property Graph
is also called a Labeled Property Graph

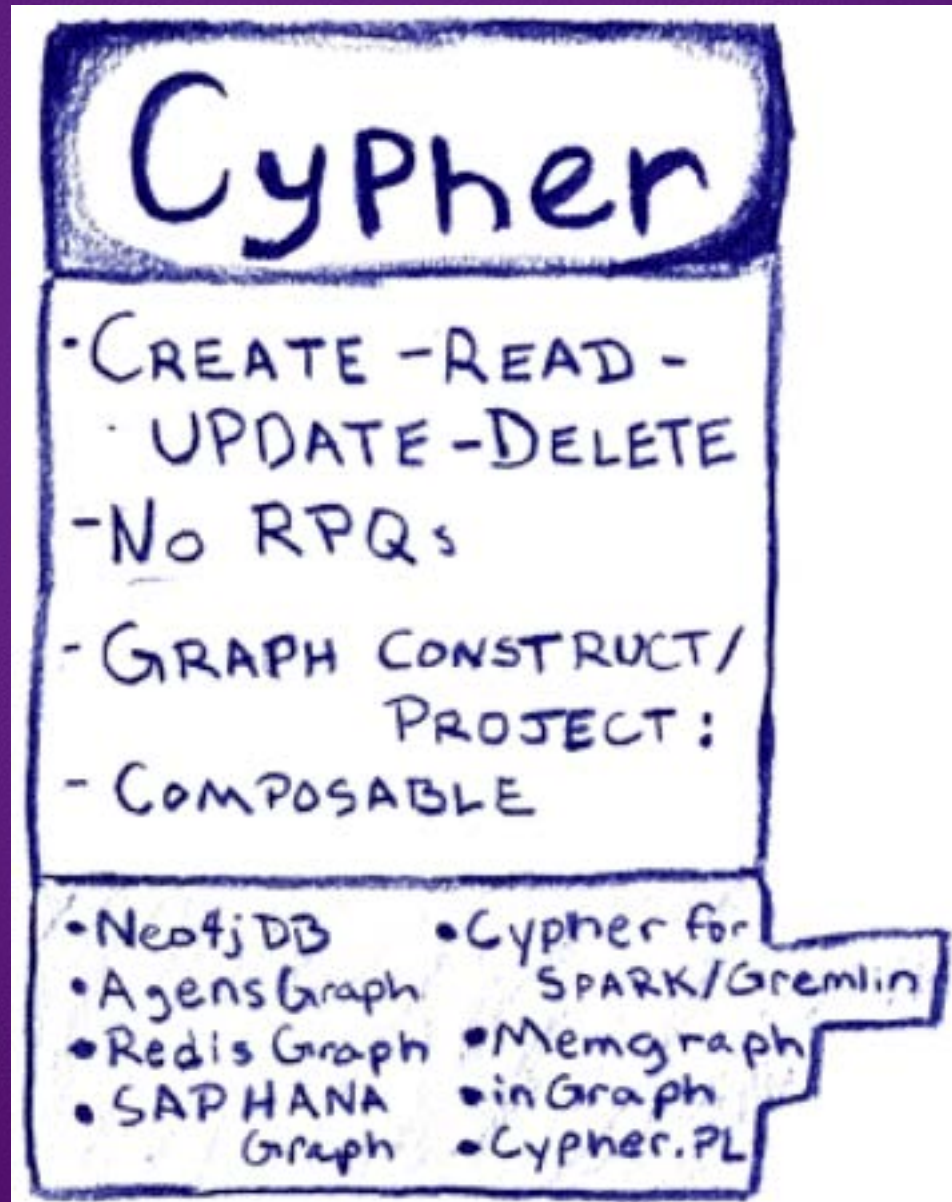


Mapping to Languages

How to map a Detailed Property Graph to the English language?



Declarative query language



Cypher versus SQL

Typical Complex SQL Join

```

(SELECT T.directReportees AS directReportees, sum(T.count) AS count
FROM T
SELECT manager_pid AS directReportees, 0 AS count
FROM person_reportee manager
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
UNION
SELECT manager_pid AS directReportees, count(manager.directly_manages) AS count
FROM person_reportee manager
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
GROUP BY directReportees
UNION
SELECT manager_pid AS directReportees, count(reportee.reportee) AS count
FROM person_reportee manager
JOIN person_reportee reportee
ON manager.directly_manages = reportee_pid
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
GROUP BY directReportees
UNION
SELECT manager_pid AS directReportees, count(L2Reportees.directly_manages) AS count
FROM person_reportee manager
JOIN person_reportee L1Reportees
ON manager.directly_manages = L1Reportees_pid
JOIN person_reportee L2Reportees
ON L1Reportees.directly_manages = L2Reportees_pid
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
GROUP BY directReportees
) AS T
GROUP BY directReportees)
UNION
(SELECT T.directReportees AS directReportees, sum(T.count) AS count
FROM T
SELECT reportee.directly_manages AS directReportees, 0 AS count
FROM person_reportee manager
JOIN person_reportee reportee
ON manager.directly_manages = reportee_pid
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
GROUP BY directReportees
UNION
SELECT L2Reportees_pid AS directReportees, count(L2Reportees.directly_manages)
AS count
FROM person_reportee manager
JOIN person_reportee L1Reportees
ON manager.directly_manages = L1Reportees_pid
JOIN person_reportee L2Reportees
ON L1Reportees.directly_manages = L2Reportees_pid
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
GROUP BY directReportees
) AS T
GROUP BY directReportees)
UNION
(SELECT L2Reportees.directly_manages AS directReportees, 0 AS count
FROM person_reportee manager
JOIN person_reportee L1Reportees
ON manager.directly_manages = L1Reportees_pid
JOIN person_reportee L2Reportees
ON L1Reportees.directly_manages = L2Reportees_pid
WHERE manager_pid = (SELECT id FROM person WHERE name = "Name Name")
GROUP BY directReportees
UNION
)

```

The same query
using Cypher

```

MATCH (boss)-[:MANAGES*0..3]->(sub),
      (sub)-[:MANAGES*1..3]->(report)
WHERE boss.name = "John Doe"
RETURN sub.name AS Subordinate,
       count(report) AS Total

```

Cypher: declarative query language

Cypher is designed to be easily read and understood by developers.

Cypher in action: creating data

```
CREATE (a:Person { name:"Tom Hanks",  
  born:1956 })-[r:ACTED_IN { roles: ["Forrest"]}]>(m:Movie { title:"Forrest Gump", released:1994 })  
CREATE (d:Person { name:"Robert Zemeckis", born:1951 })-[:DIRECTED]->(m)  
RETURN a,d,r,m
```

Cypher in action: finding interesting connections

```
MATCH (p:Person { name:"Tom Hanks" })-[r:ACTED_IN]->(m:Movie)  
RETURN m.title, r.roles
```

Cypher for creating data

ASCII art representation in Cypher



```
CREATE (:Person {name:"Peter"}) - [[:WORKS_FOR]] -> (:Person {name:"Mark"})
```

label

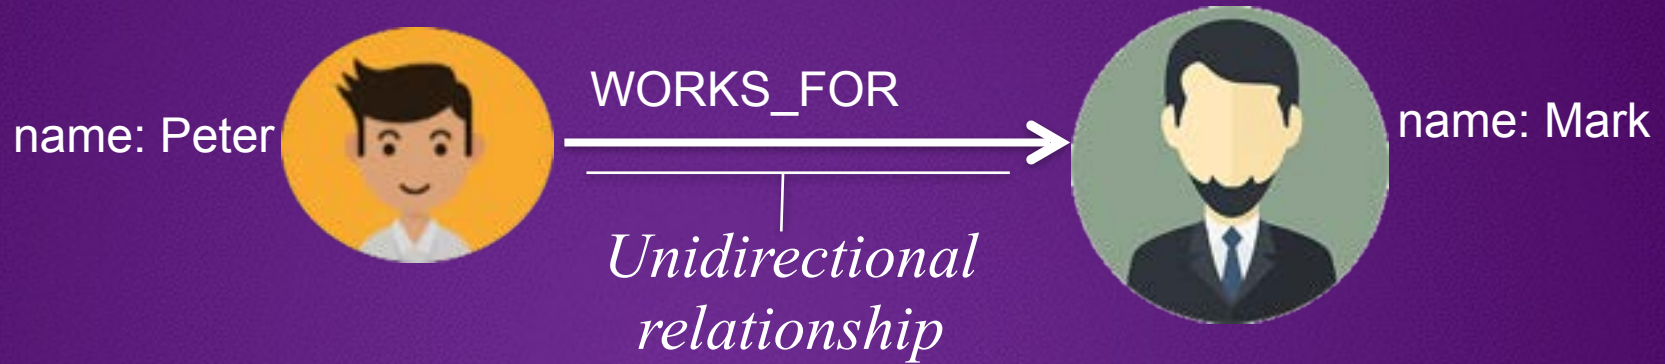
property

Drawing the pattern of the graph
in your query through
ASCII characters ->

label

property

Cypher for representing Bi-Directionality



```
MATCH (:Person {name:"Peter"} - [:WORKS_FOR] -> (p:Person))
```

label

property

*Unidirectional
relationship*

label

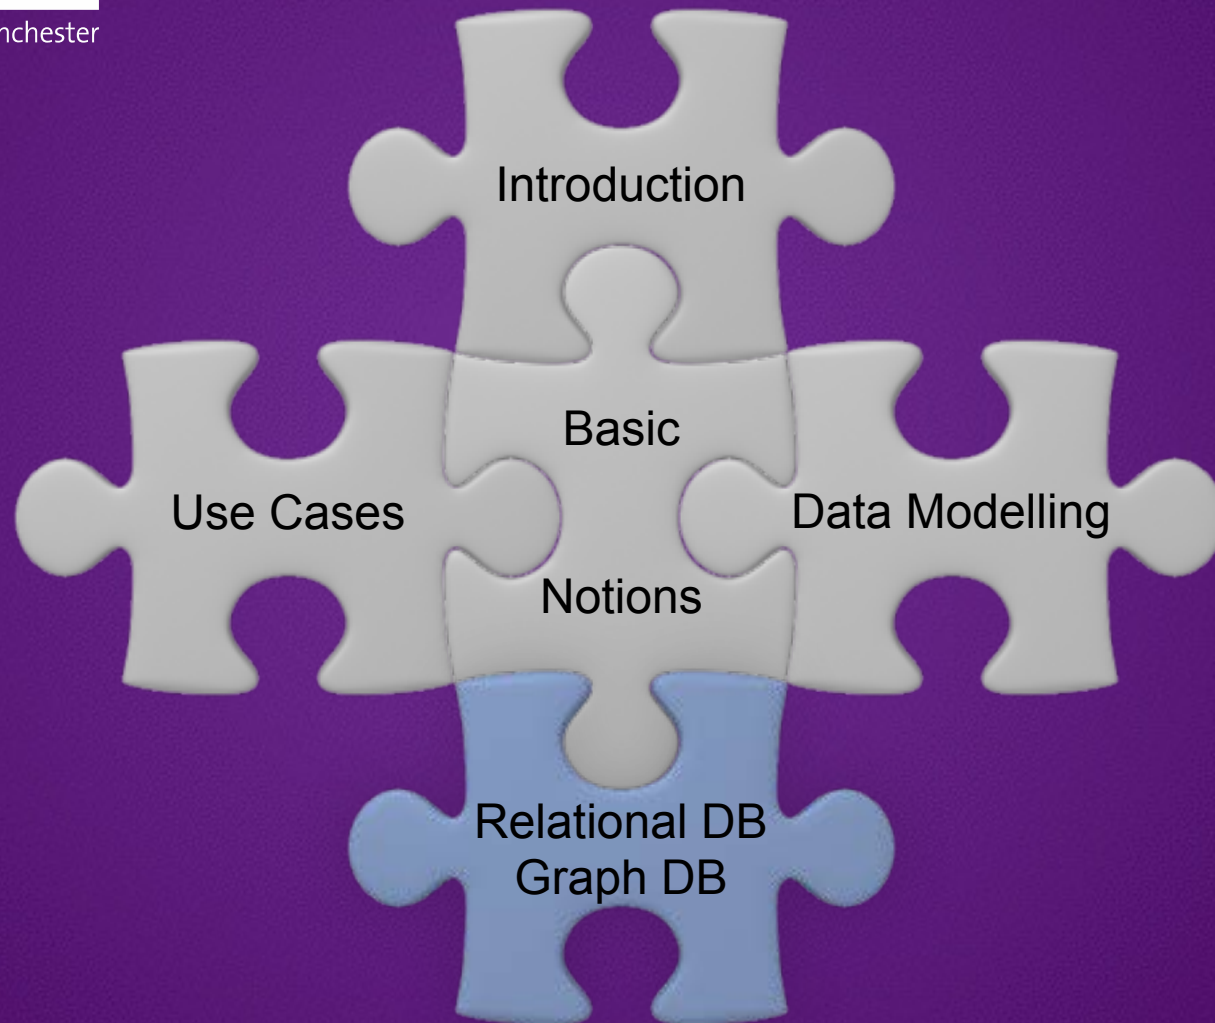
```
MATCH (:Person {name:"Peter"} - [:WORKS_WITH] - (p:Person))
```

label

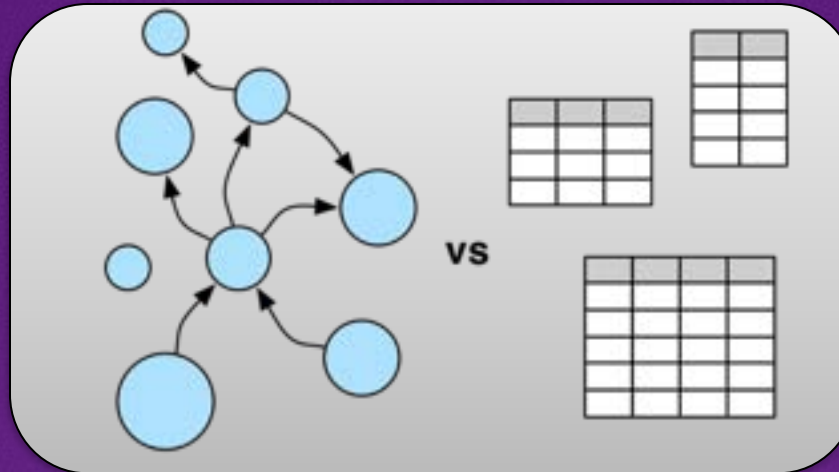
property

*Bidirectional
relationship*

label





Relational DB vs Graph DB



Tabular data may be best stored in a relational DB


Interconnected data



A is not directly related to C
A is transitively related to C

How to use Neo4j

There are 3 basic steps to using Neo4j.




Step 1: creating a model

Creating in advance
labels; nodes; relationships; and properties

How to use Neo4j

There are 3 basic steps to using Neo4j.



Step 1: creating a model

Creating in advance
labels; nodes; relationships; and properties




Step 2: loading data

The easiest way to load data from a relational DB in Neo4j is by exporting the data as a CSV file

How to use Neo4j

There are 3 basic steps to using Neo4j.



Step 1: creating a model

Creating in advance
labels; nodes; relationships; and properties



Step 2: loading data

The easiest way to load data from a relational DB in Neo4j is by exporting the data as a CSV file



Step 3: querying data

Neo4j has a built-in web application for querying data, so you can explore your data

Moving from Relational DB to Graph DB

Not all applications are the same!

Depending on data interconnections,
you may consider

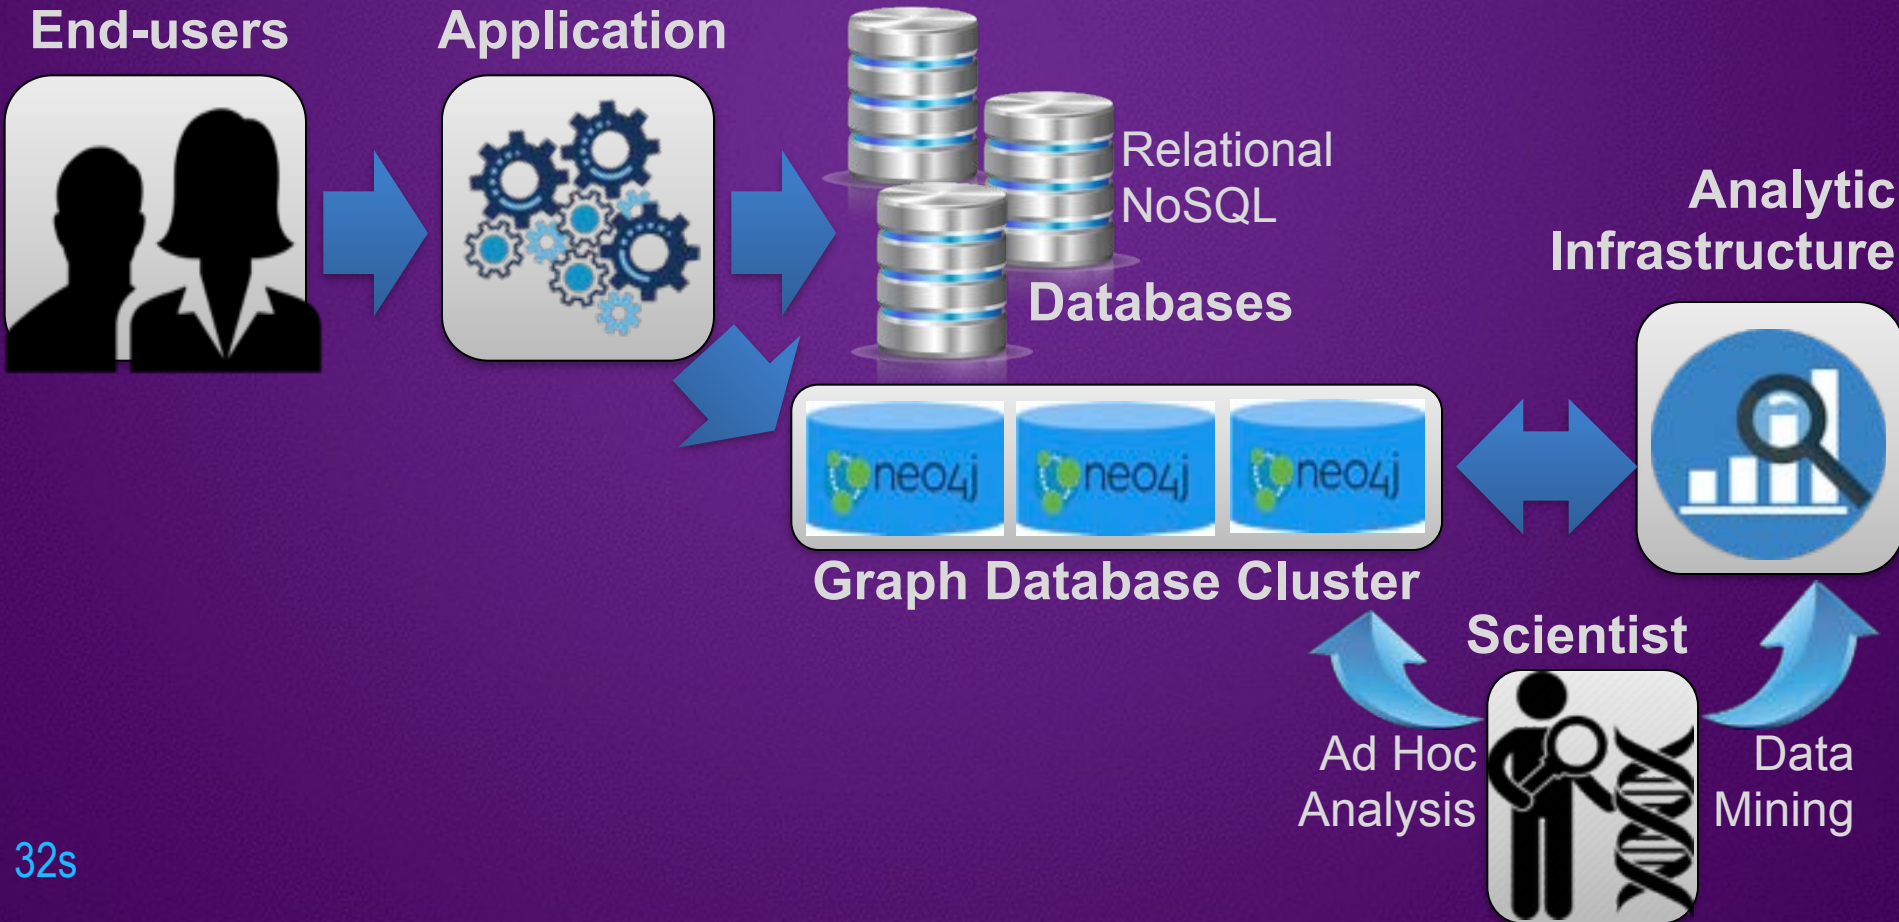
Option 1: migrate all data

Option 2: migrate a subset of data

Option 3: duplicate a subset of data

Moving from Relational DB to Graph DB

Option 2: migrate a subset of data
Architectural overview



Relational DB vs Graph DB

PersonID	Person [Name]	Nationality [Country]	UK University	University location [UK country]
1024	<i>Luca</i>	<i>Italy</i>	<i>Queen's University Belfast</i>	<i>Northern Ireland</i>
12	<i>Lisa</i>	<i>UK</i>	<i>University College London</i>	<i>England</i>
2048	<i>Basil</i>	<i>UK</i>	<i>The University of Manchester</i>	<i>England</i>
24	<i>Alice</i>	<i>USA</i>	<i>Cardiff University</i>	<i>Wales</i>
212	<i>Maria</i>	<i>Spain</i>	<i>The University of Edinburgh</i>	<i>Scotland</i>



...difficult
to read

Relational DB vs Graph DB

Name	Country	University	UK country
<i>Luca</i>	<i>84</i>	<i>32</i>	<i>1</i>
<i>Lisa</i>	<i>186</i>	<i>3</i>	<i>2</i>
<i>Basil</i>	<i>186</i>	<i>6</i>	<i>2</i>
<i>Alice</i>	<i>187</i>	<i>25</i>	<i>3</i>
<i>Maria</i>	<i>165</i>	<i>4</i>	<i>4</i>

ID	Country name
<i>84</i>	<i>Italy</i>
<i>186</i>	<i>UK</i>
<i>187</i>	<i>USA</i>
<i>165</i>	<i>Spain</i>

ID	University name
<i>32</i>	<i>Queen's University Belfast</i>
<i>3</i>	<i>University College London</i>
<i>6</i>	<i>The University of</i>
<i>25</i>	<i>Cardiff University</i>
<i>4</i>	<i>The University of Edinburgh</i>

ID	UK Country name
<i>1</i>	<i>Northern Ireland</i>
<i>2</i>	<i>England</i>
<i>3</i>	<i>Wales</i>
<i>4</i>	<i>Scotland</i>

Relational DB vs Graph DB



Name	Country	University	UK country
Luca	84	32	1
Lisa	186	3	2
Basil	186	6	2
Alice	187	25	3
Maria	165	4	4

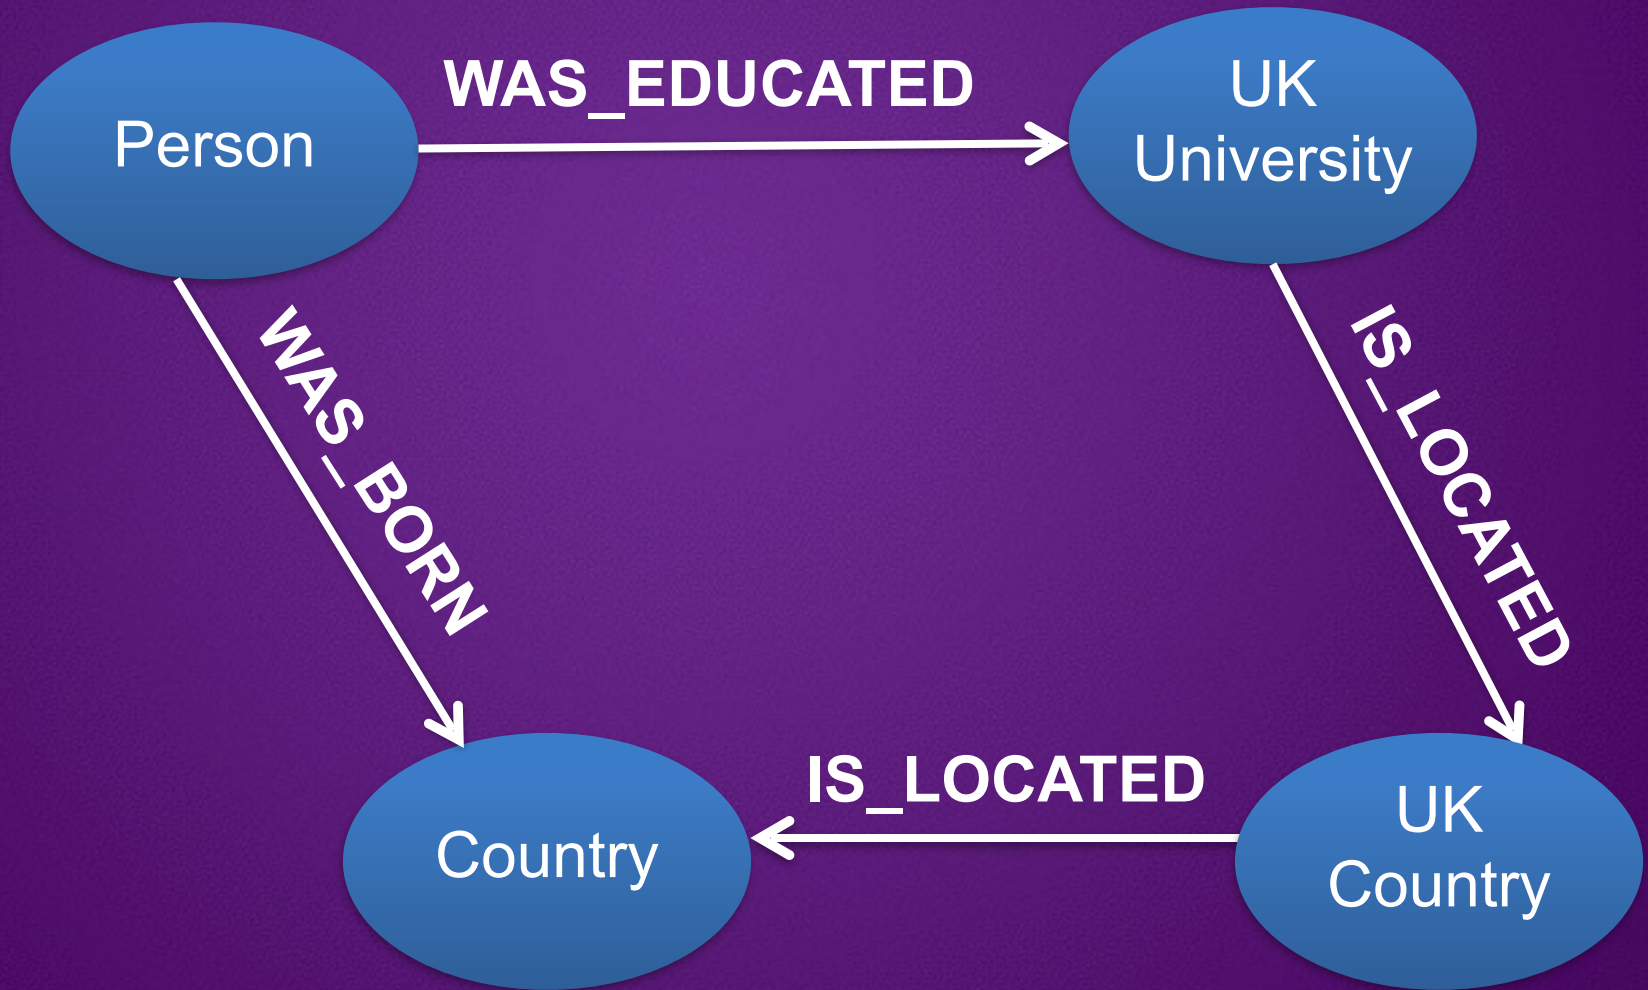
```

1 SELECT
2     p.name, c.name as country, u.name as university
3 FROM
4     person p
5     LEFT JOIN country c ON c.ID = p.country
6     LEFT JOIN university u ON p.university = u.ID
7 WHERE
8     u.name = 'The University of Manchester'
    
```



name	country	university
Basil	UK	The University of Manchester

Relational DB vs Graph DB



Cypher: declarative query language



```
1 MATCH
2 (p1)-[r1:WAS_EDUCATED]-(u),
3 (p2)-[r2:WAS_BORN]→(c)
4 WHERE u.name = 'The University of Manchester'
5 RETURN p1.name as name, c.name as country, u.name as university
```

	name	country	university
1	"Basil"	"UK"	"The University of Manchester"

Cypher: declarative query language



Performance

Write queries

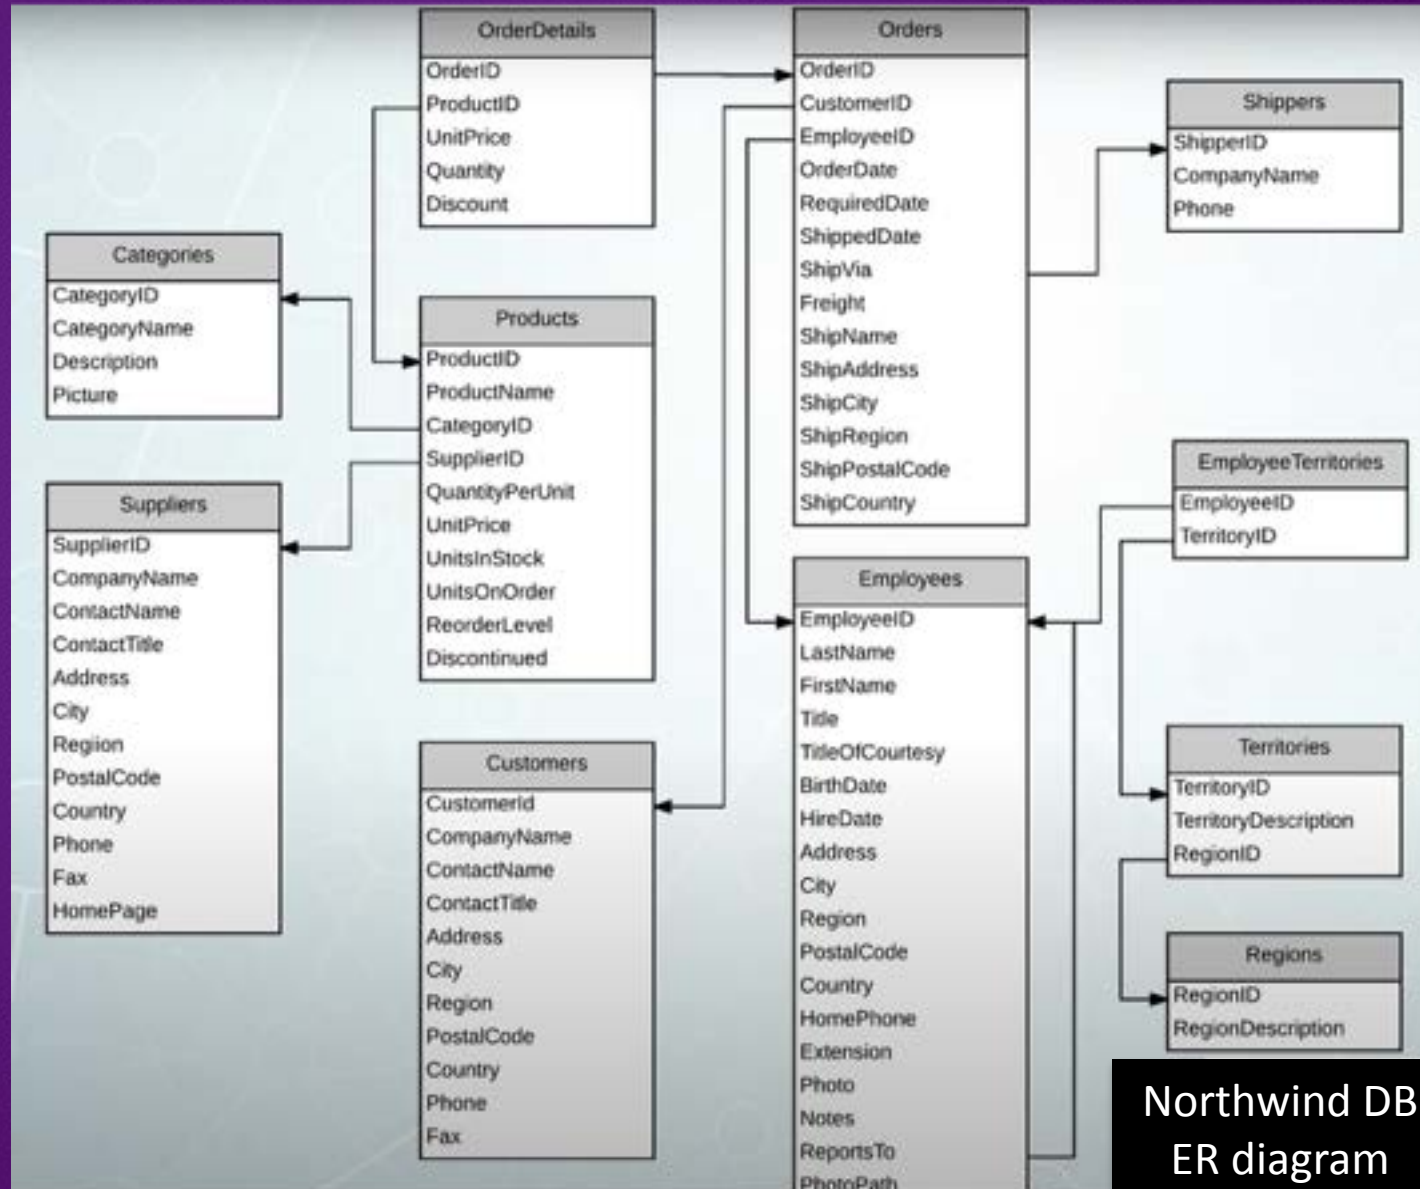
Maintenance

Model and store relationships



Let's see if
I understand
this ER
diagram...
...primary
keys...
foreign
keys...

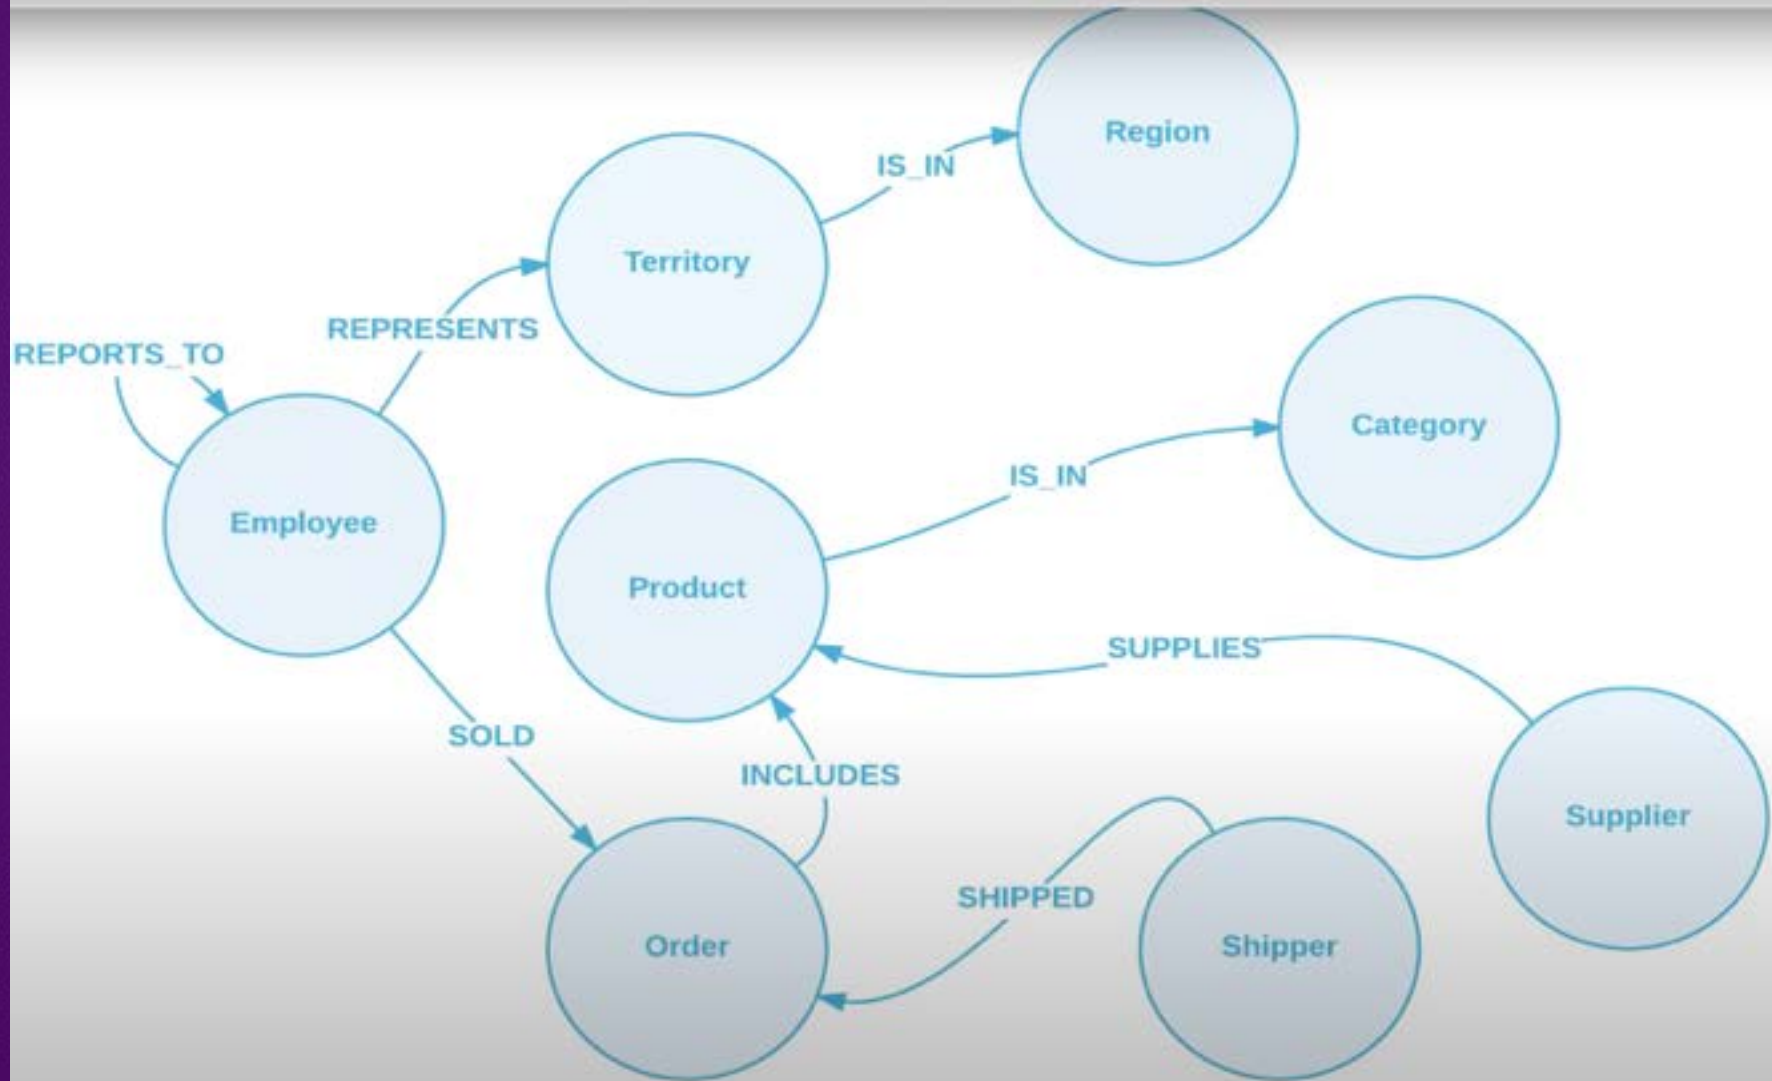
Converting Northwind DB to Graph DB



Northwind DB
ER diagram

Converting Northwind DB to Graph DB

The graph model for Northwind dataset

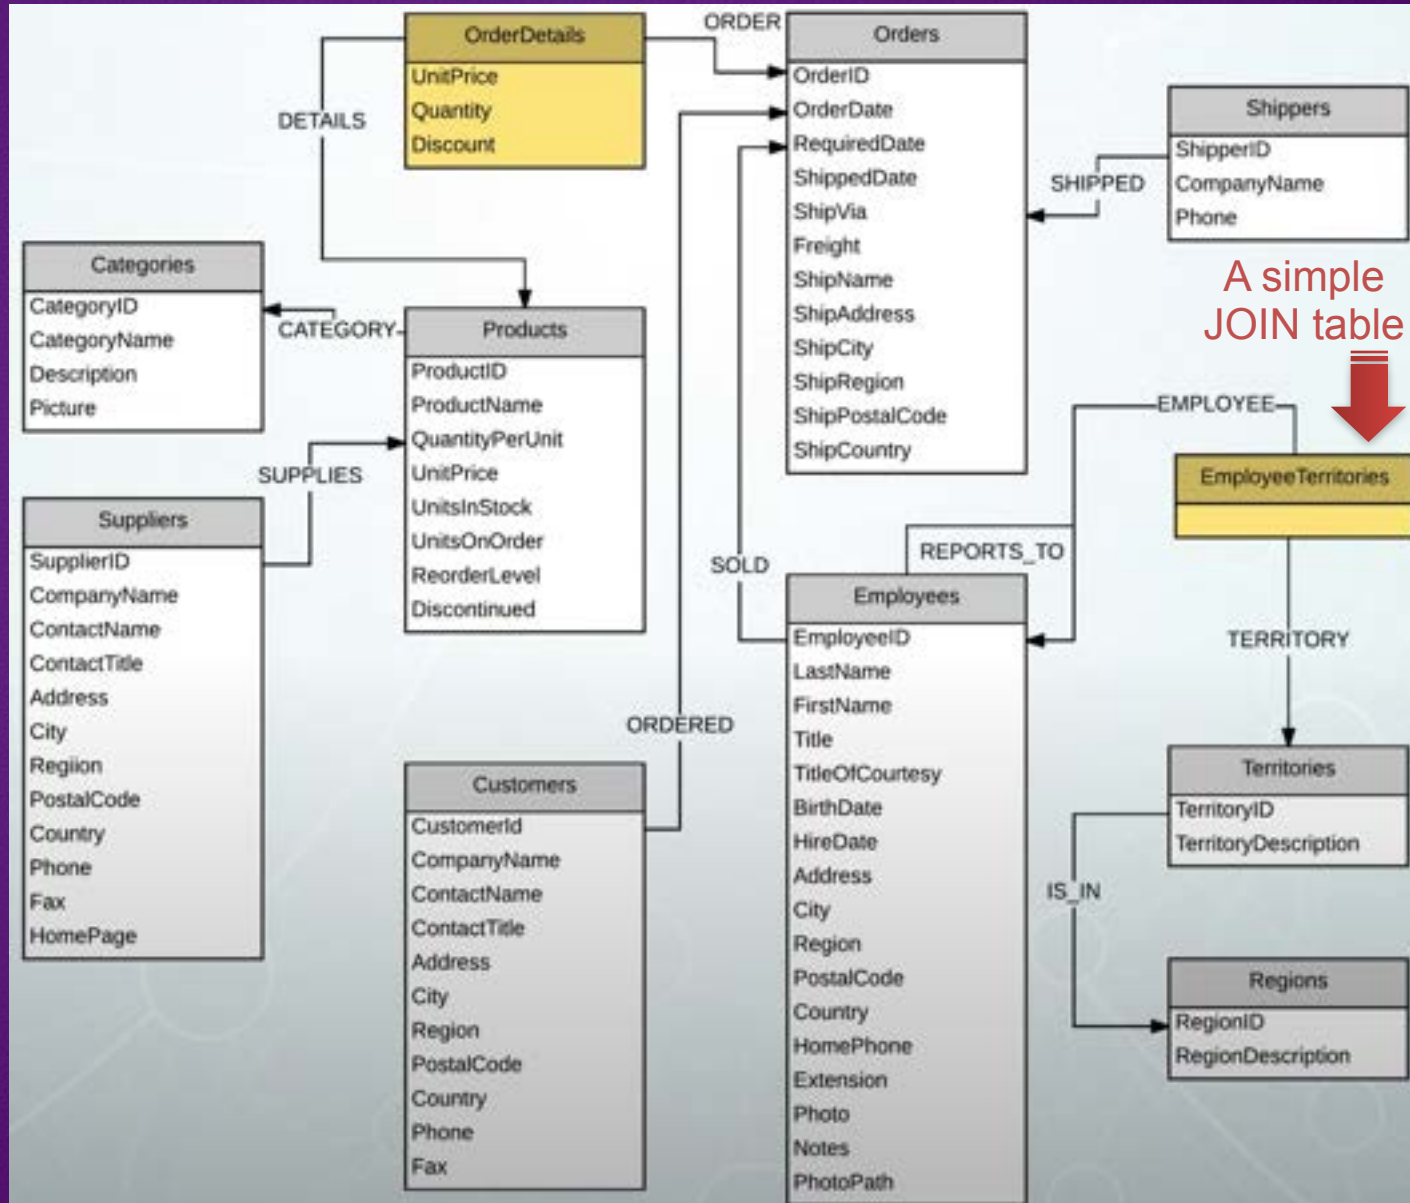


Converting Northwind DB to Graph DB

Step 1: locate foreign keys, delete them, and replace them with relationships

Step 2: locate JOIN tables, and

.....

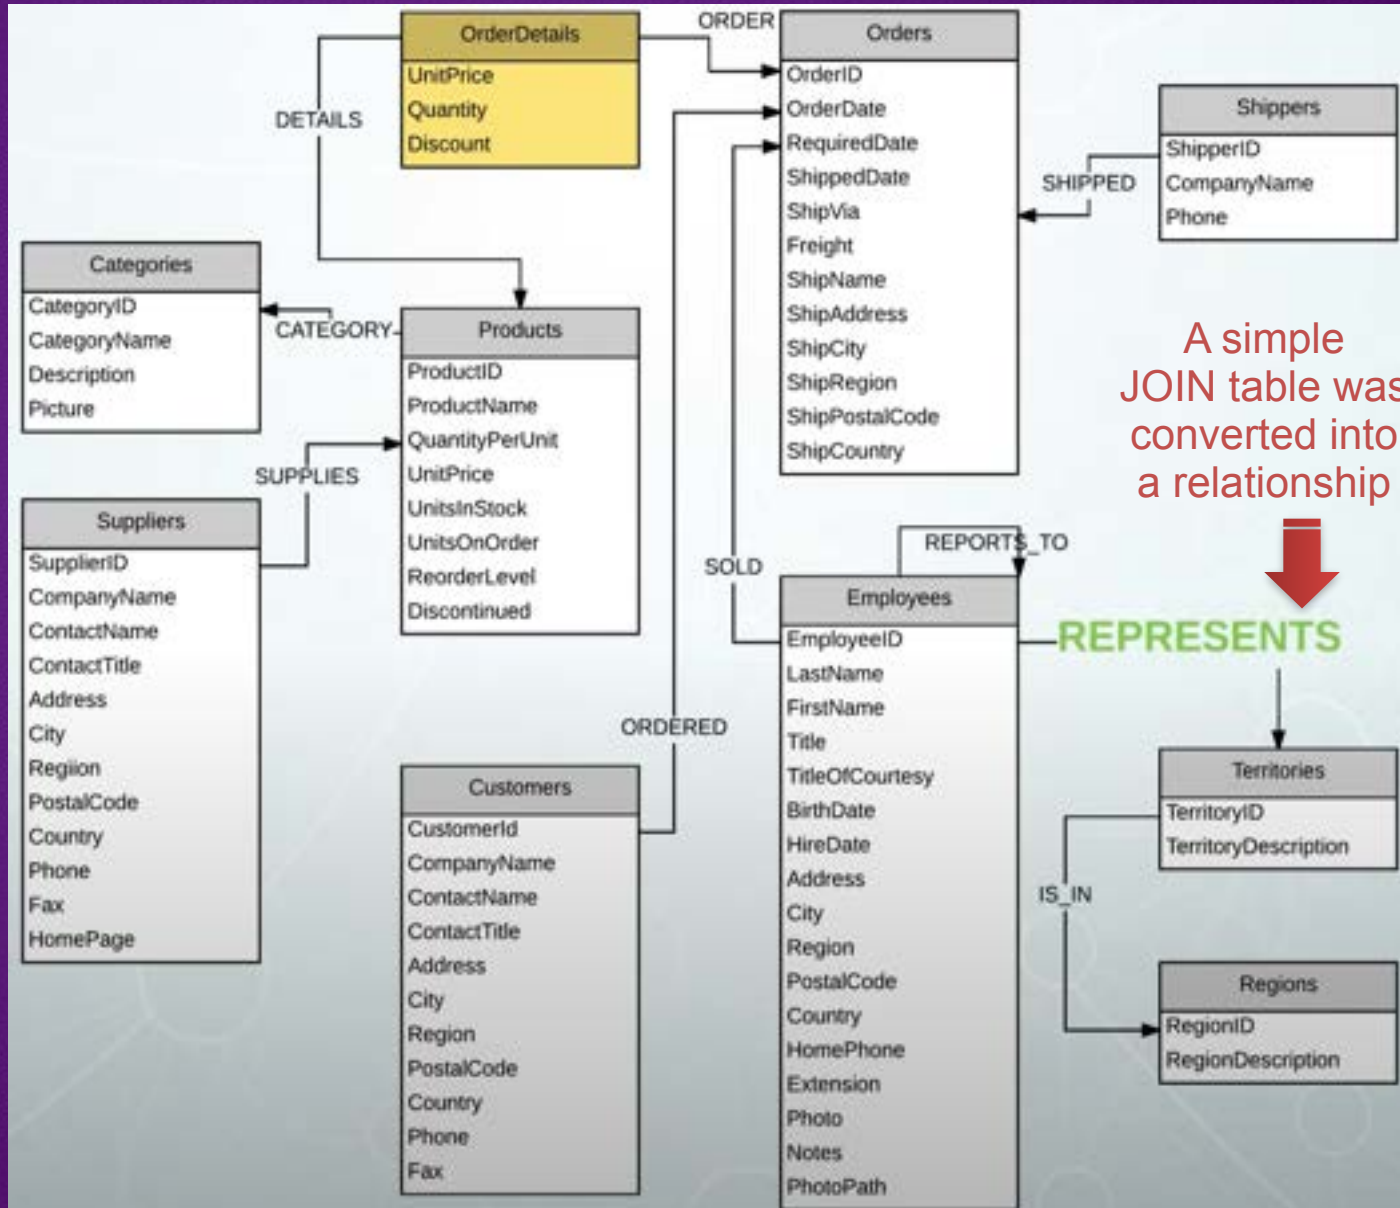


Converting Northwind DB to Graph DB

Step 1: locate foreign keys, delete them, and replace them with relationships

Step 2: locate JOIN tables, and

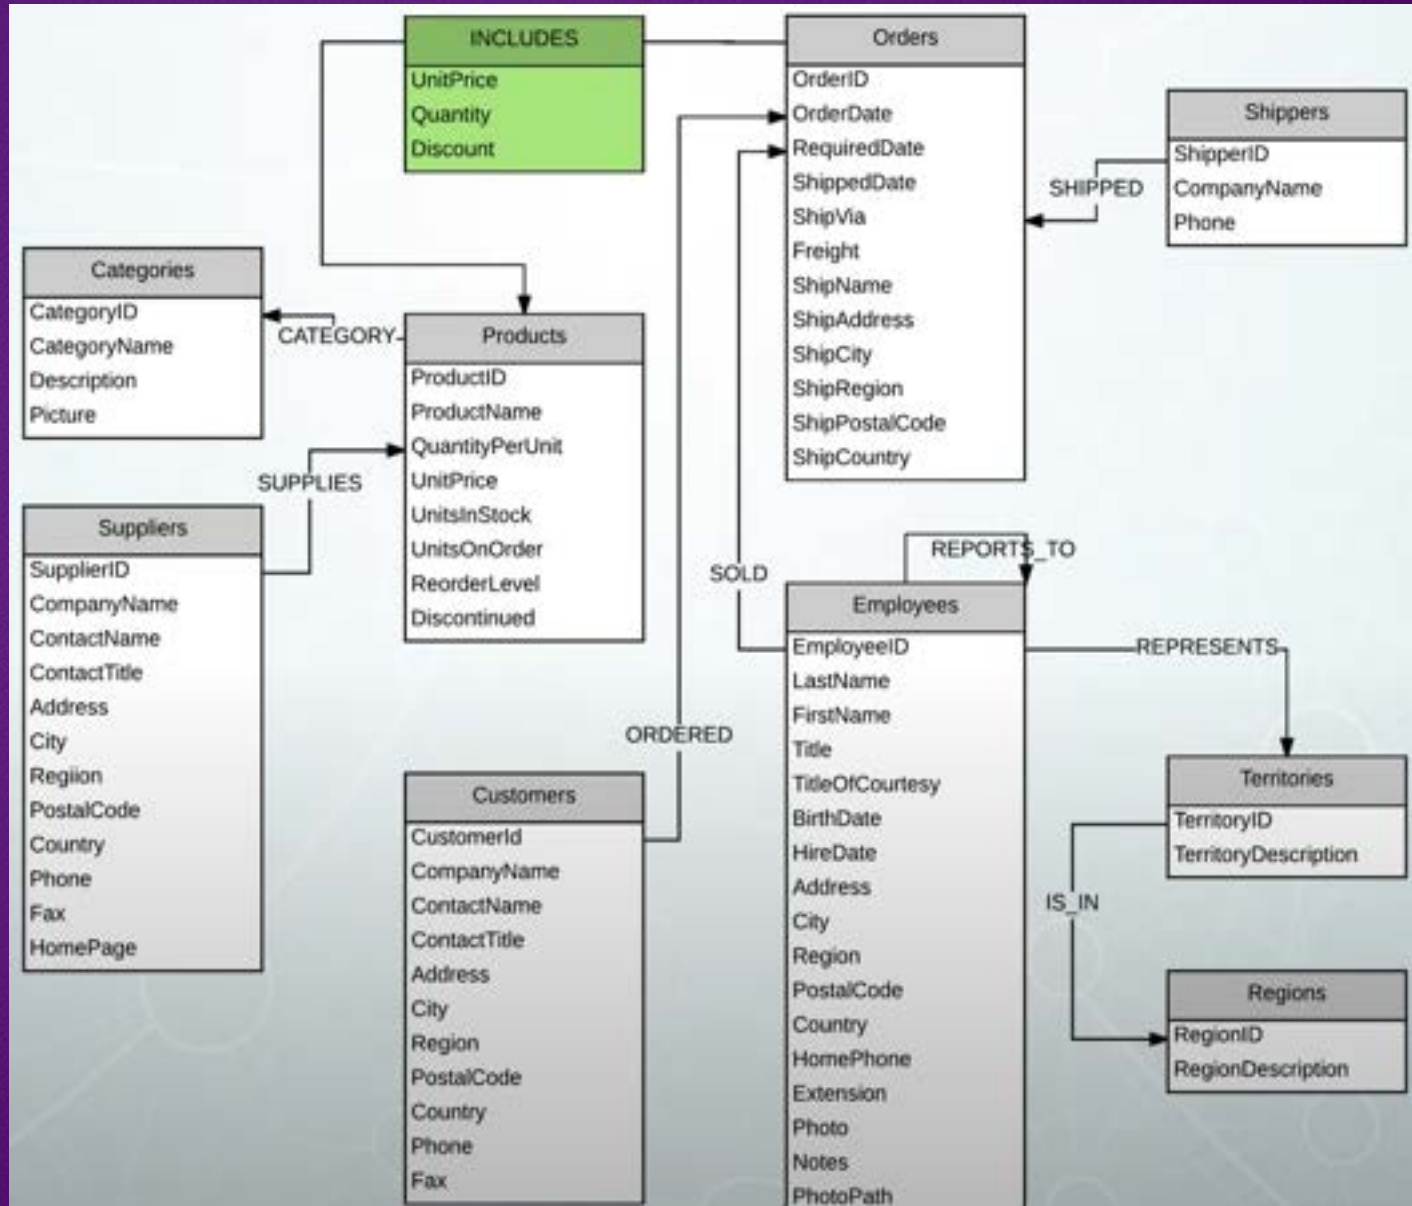
.....

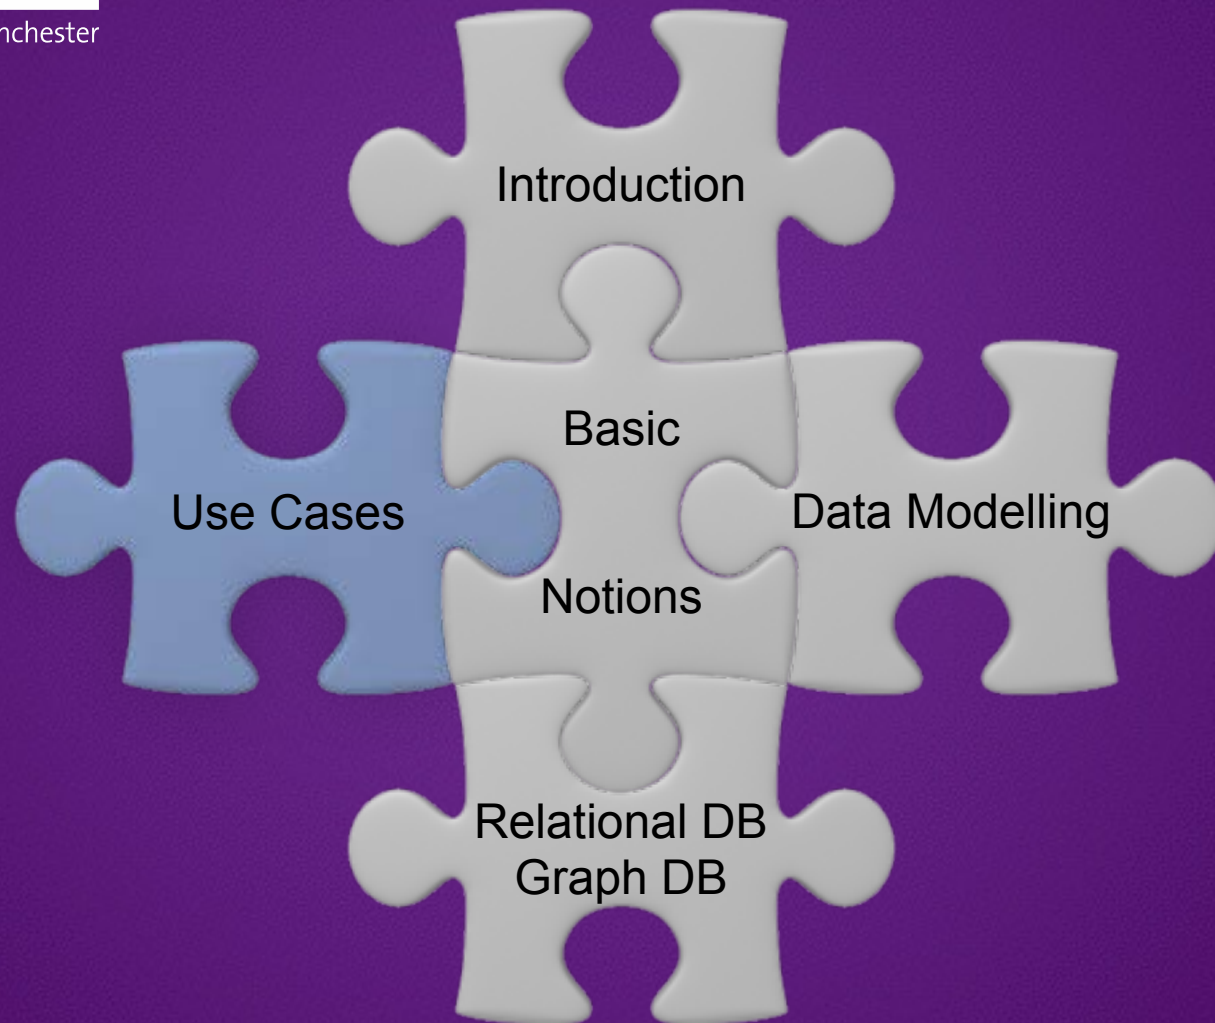


Converting Northwind DB to Graph DB

Step 1: locate foreign keys, delete them, and replace them with relationships

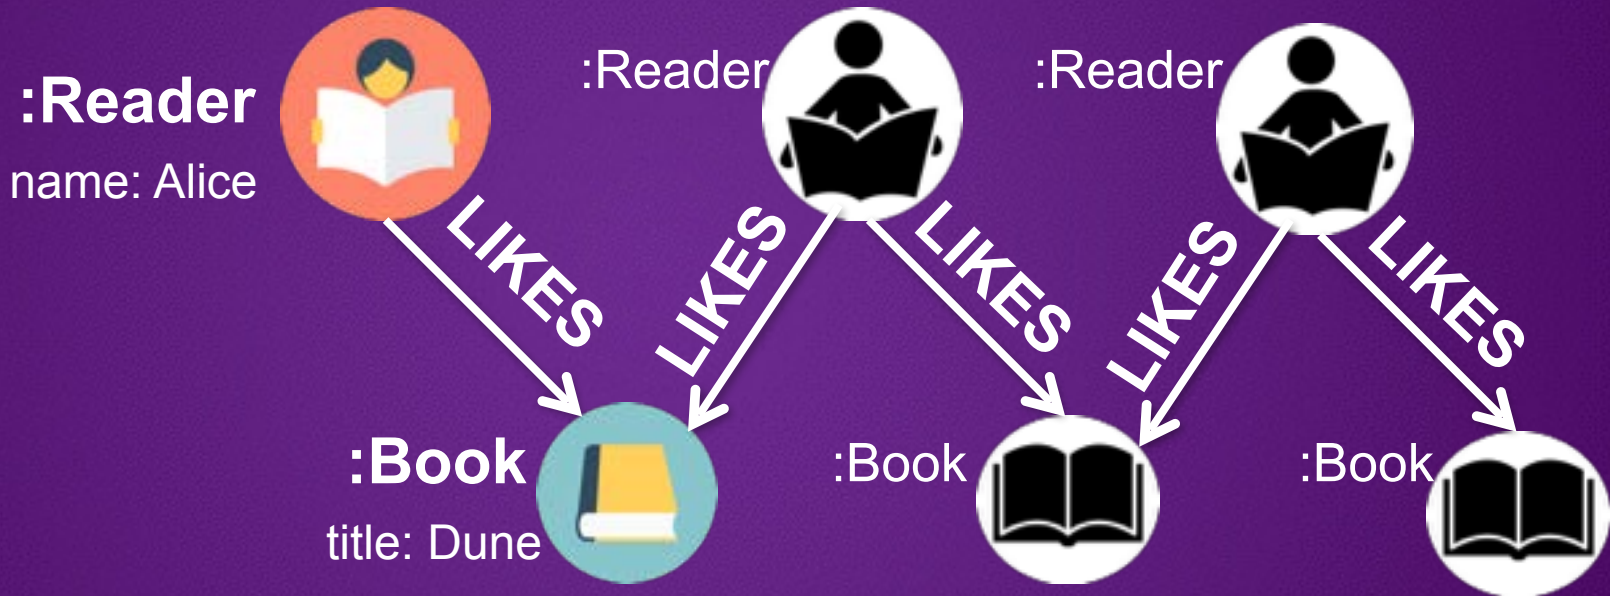
Step 2: locate JOIN tables, and convert them into relationships





Book review data model

Simple data model



... finding books for Alice according to other readers ...

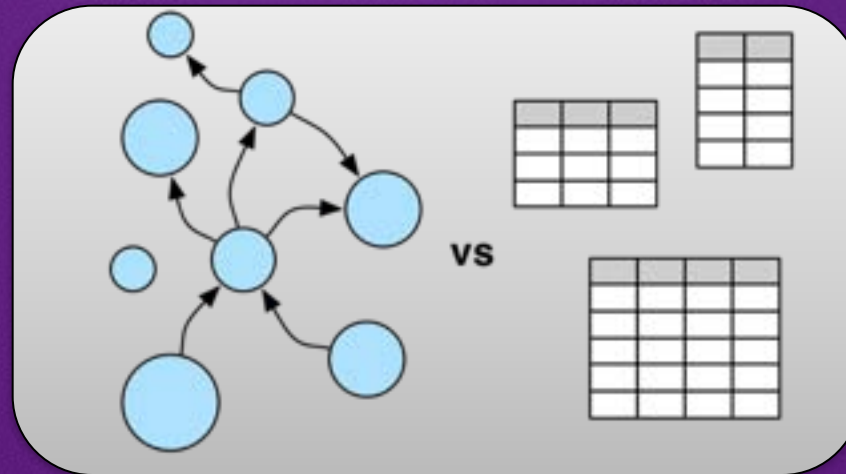
```
MATCH (:Reader {name:'Alice'})-[:LIKES]->(:Book {title:'Dune'})
      <-[:LIKES]-(:Reader)-[:LIKES]->(books:Book)
RETURN books.title
```

Common graph technology use cases: Fraud Detection



Fraud
Detection

Banks and insurance companies lose billions of dollars every year to fraud.



Graph databases offer new methods of uncovering fraud by looking at the connections that link individual data points

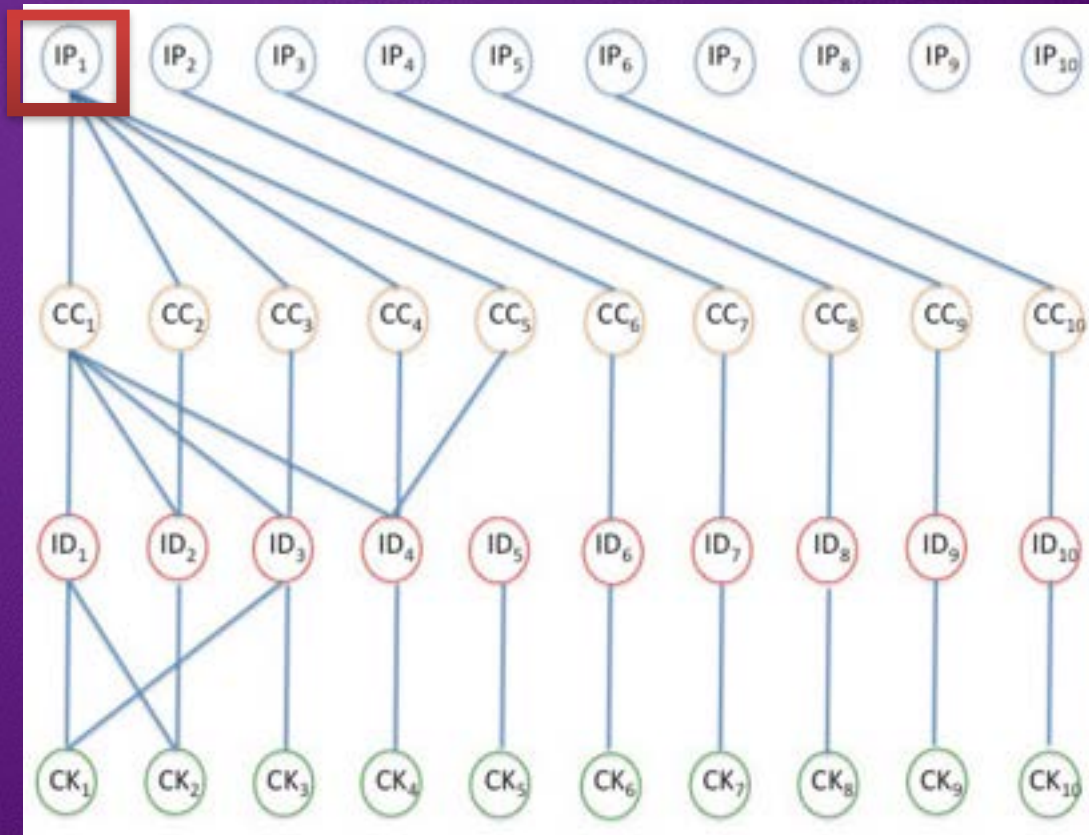
Common graph technology use cases: Fraud Detection



Fraud
Detection


Example of
fraud-detection
with link analysis

Look for multiple transactions from the same
IP with different credit cards.



Google: Knowledge Graph entities

Google Knowledge Graph has millions of entries that describe real-world entities, such as people and places.



Some types of entities found in the **Google Knowledge Graph**

Book

BookSeries

EducationalOrganization

Event

GovernmentOrganization

LocalBusiness

Movie

MovieSeries

MusicAlbum

MusicGroup

MusicRecording

Organization

Periodical

Person

Place

SportsTeam

TVEpisode

TVSeries

VideoGame

VideoGameSeries

WebSite

Google: Knowledge Graph entities

Movie is a type of entity in the **Google Knowledge Graph**

Movie

A Schema.org Type

Thing > CreativeWork > Movie


[more...]

A movie.

Property	Expected Type	Description
Properties from Movie		
actor	Person	An actor, e.g. in TV, radio, movie, video games etc., or in an event. Actors can be associated with individual items or with a series, episode, clip. Supersedes actors .
countryOfOrigin	Country	<p>The country of origin of something, including products as well as creative works such as movie and TV content.</p> <p>In the case of TV and movie, this would be the country of the principle offices of the production company or individual responsible for the movie. For other kinds of CreativeWork it is difficult to provide fully general guidance, and properties such as contentLocation and locationCreated may be more applicable.</p> <p>In the case of products, the country of origin of the product. The exact interpretation of this may vary by context and product type, and cannot be fully enumerated here.</p>
director	Person	A director of e.g. TV, radio, movie, video gaming etc. content, or of an event. Directors can be associated with individual items or with a series, episode, clip. Supersedes directors .
duration	Duration	The duration of the item (movie, audio recording, event, etc.) in ISO 8601 date format .
musicBy	MusicGroup or Person	The composer of the soundtrack.

Common graph technology use cases: Knowledge Graph

Knowledge graphs are a type of graph.



Organisations are using knowledge graphs to improve the reasoning skills

Knowledge graphs can reason (e.g. using a description logic reasoner) about the underlying data.

A key differentiator:

Resource Description Framework (RDF)

SPARQL is a RDF query language

Neo4j supports SPARQL

Amazon Neptune is compatible with SPARQL 1.1.

Common graph technology use cases: Knowledge Graph

The results of SPARQL queries
can be RDF graphs.



Knowledge
Graph

An RDF graph is a set of RDF triples.

An RDF triple



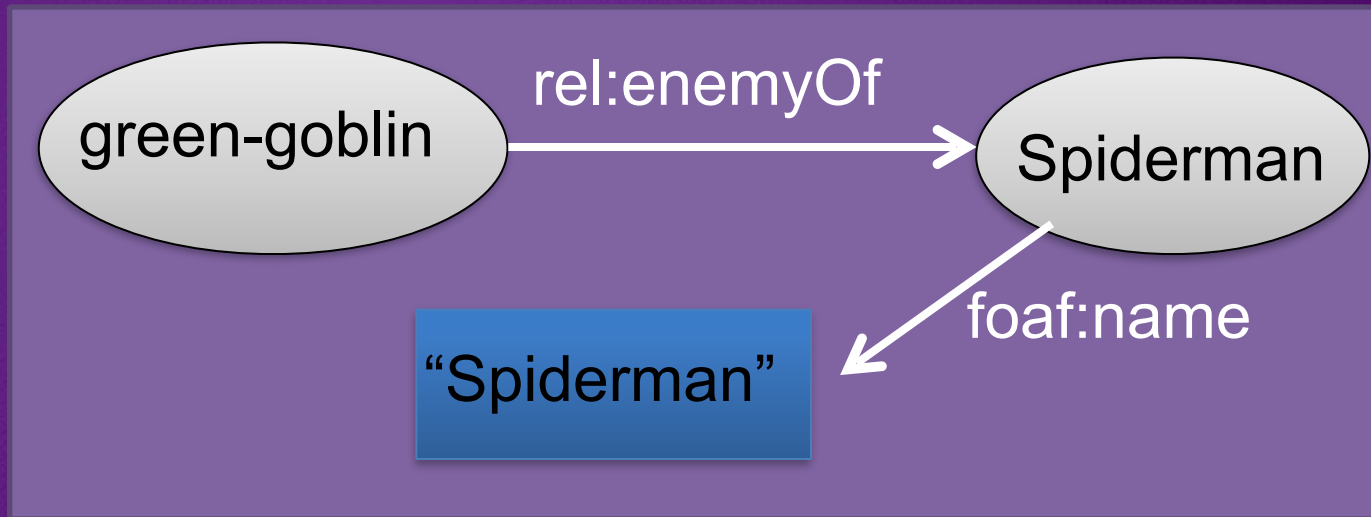
Common graph technology use cases: Knowledge Graph



Knowledge
Graph

Example of an RDF Graph

<https://www.w3.org/TR/turtle/>



RDF Graph uses
Universal Resource Identifiers (URIs)

```
<http://example.org/#green-goblin>  
<http://www.perceive.net/schemas/relationship/enemyOf>  
<http://example.org/#spiderman> .
```

```
<http://example.org/#spiderman>  
<http://xmlns.com/foaf/0.1/name>  
"Spiderman" .
```

Using Amazon Neptune to build an Enterprise Knowledge Graph

<https://aws.amazon.com/neptune/knowledge-graphs-on-aws/>



Knowledge
Graph



*“Amazon Neptune is a key part of the toolkit we use to continually expand **Alexa’s knowledge graph** for our tens of millions of Alexa customers”*

David Hardcastle,
Director of Amazon Alexa.

Why do you need a knowledge graph?

Insights with Machine Learning: use machine learning with knowledge graphs for better decision making and knowledge discovery.

Build virtual assistants, chatbots or question-answering systems: build context-aware systems that can derive an answer based on queries and a vast knowledge base.



Images
shown may
have a
copyright



Questions

