

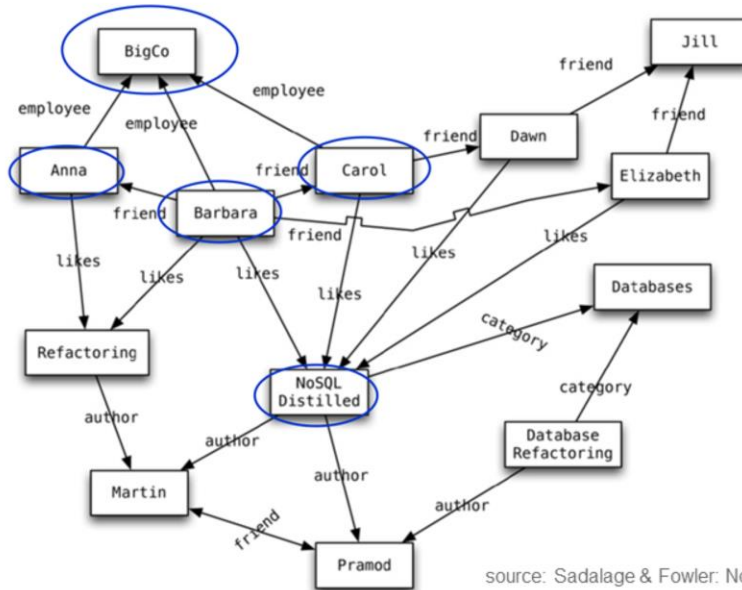
Agenda



- Graph Databases
- Neo4j
 - Basic information
 - Data model
 - Cypher query language
 - Structure and examples
 - Other interfaces: Experience with Web UI
 - Java API (embedded database)
 - Traversal of the graph
 - Traversal framework
 - Examples

Graph Databases: Example

Neo4j
CREATE
nodes
graph
relationships
depth
node
RETURN
name
java
Person
get
Property
Cypher
Found
Andres
traversal
MAYO




source: Sadalage & Fowler: NoSQL Distilled, 2012

Neo4j: Basic Info



- **Open source** graph database
- Initial release: 2007
 - Current version 4.2
- Written in: **Java**
- OS: cross-platform
- Full **transactions** (ACID)
- Partitioning: supported by queries
 - [since 4.0](#), by Neo4j **Fabric**
- **Replication**: Master-slave
 - Eventual consistency

```
USE demo_graph( userShardMapping( $userid ) )
// User Defined Mapping Function (UDMF)
// For implementing sharding schemes,
// i.e. how to locate the shard containing
// specific piece of data
```



Source: neo4j.com

Neo4j 4.0 now allows for [sharding](#) – a result of careful engineering (and at least one PhD in parallel computing) – which **distributes and parallelizes queries and aggregations over multiple databases.**

Neo4j Fabric is Neo4j's solution to graph sharding by allowing users to break a larger graph down into individual, smaller graphs and store them in separate databases. For graphs that are highly-connected, this means some level of data redundancy to maintain the relationships between entities.

Sharding is possible because of new capabilities like **federated queries.**

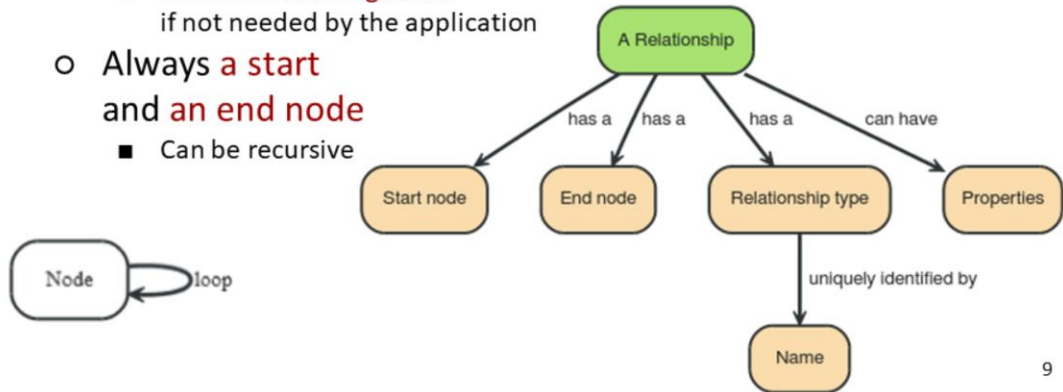
Queries against disjointed graphs

Neo4j not only handles queries against large, sharded graphs that share the same data model; it also supports queries across disjointed graphs. In effect, this makes all the data in an enterprise's graph databases searchable with a single query.



Data Model: Relationships

- **Directed relationships (edges)**
 - Incoming and outgoing **edge**
 - Equally **efficient traversal** in both directions
 - Direction **can be ignored** if not needed by the application
 - Always **a start** and **an end node**
 - Can be recursive



Use of Neo4j



- **Two** ways to **use** Neo4j:
 - **Embedded**: Used directly within a Java application
 - **Self-standing** server + connections

- **Various types of connections**
 - Neo4j Cypher Shell
 - HTTP API
 - uses Cypher query language
 - **Web GUI**
 - using Cypher **query language**
 - Standard **Java API**
 - Gremlin graph traversal language (plugin), etc.

Neo4j Command-line Querying



- **Cypher** shell

- `./bin/cypher-shell`
- can also be installed separately, but shipped with the server

HTTP API



- Query/update **operations** using **HTTP** protocol
 - GET, POST methods
 - data sent/received in JSON
- Fully **transactional** in the latest version
- Example: create node with “name” property

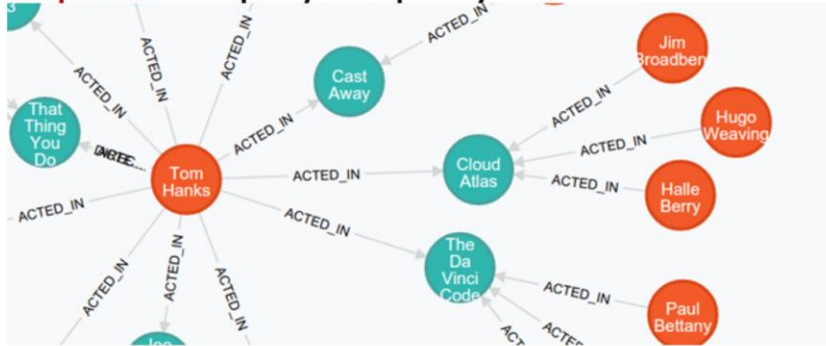
```
curl -i -X POST http://localhost:7474/db/neo4j/tx/commit  
-H "Content-Type: application/json; charset=UTF-8" --user  
"neo4j" -d '{ "statements": [ { "statement": "CREATE (n  
$props) RETURN n", "parameters": { "props": { "name":  
"John Doe" } } } ] }'
```

<https://neo4j.com/docs/http-api/current/>

Neo4j Web Interface



- By default, running on <http://localhost:7474/>
 - default credentials: neo4j/neo4j
- Online **interpreter** of **Cypher**
- **Graphical** display of query **results**



Cypher Language



- Neo4j graph **query language**
 - For querying and updating
- **Declarative** – we say **what** we want
 - **Not how** to get it
 - **Not** necessary to express **traversals**
- **Human-readable**
- Inspired by SQL and SPARQL
- Still growing = syntax changes are often



Cypher: Clauses

- **MATCH:** The graph **pattern** to match
- **WHERE:** **Filtering** criteria
- **RETURN:** What to return

- **WITH:** Divides a query into multiple parts
- **CREATE:** Creates nodes and relationships.
- **DELETE:** Remove nodes, relationships, properties
- **SET:** Set values to **properties**

<https://neo4j.com/developer/cypher/>

Cypher: Creating Nodes (Examples)



```
CREATE (n);
```

(create a node, assign to var n)

```
Created 1 node, returned 0 rows
```

```
CREATE (a: Person {name : 'Jan'}) RETURN a;
```

(create a node with label 'Person' and 'name' property Jan')

```
Created 1 node, set 1 property, returned 1 row
```

Task 1: Add people with names: John, Jack, Andres

Cypher: Creating Paths



```
CREATE p = (andres: Person {name: 'Andres'})
```

```
-[:WORKS_AT]->
```

```
(neo)
```

```
<-[:WORKS_AT]-
```

```
(michael: Person {name:'Michael'})
```

```
RETURN p ;
```

(all parts of the pattern are created)

```
P [Node[4] {name:"Andres"}, :WORKS_AT[2]
 {}, Node[5] {}, :WORKS_AT[3] {}, Node[6] {name:"Michael"}]
1 row
Nodes created: 3
Relationships created: 2
Properties set: 2
```

To create just a relationship, use
MATCH and WHERE

Task 3: Movies Database



- Go over the “Movies” demo prepared by Neo4j

localhost

- Download the data from the course page ([movies-insert.cypher](#))

- Copy the file to Stratus VM

stratus

- Import by cypher shell

```
# cd neo4j-community-3.1.4
# bin/cypher-shell -u neo4j -p <pwd> <movies-insert.cypher
Added 171 nodes, Created 253 relationships, Set 564
properties, Added 171 labels
```

“<” sends the file to
stdin of cypher shell

Neo4j as Embedded Database



- either use **.jar** packages from the **distribution**
- ...or download packages from **Maven** repository
 - package `org.neo4j:neo4j:3.0.0`
 - **dependencies** automatically **loaded**
 - newest versions available in **repository**

localhost

<http://repo.maven.apache.org/maven2/>

- ...or download **project** from the course web

```
$ unzip neo4j-exercise.zip
```

```
$ module add idea-2019.2
```

```
$ idea.sh
```


Traversal Framework



- A **traversal** is influenced by
 - **Starting node(s)** where the traversal will begin
 - **Expanders** – define what to traverse
 - i.e., relationship direction and type
 - **Order** – depth-first / breadth-first
 - **Uniqueness** – visit nodes (relationships, paths) only once
 - **Evaluator** – what to return
and whether to stop or continue beyond current position

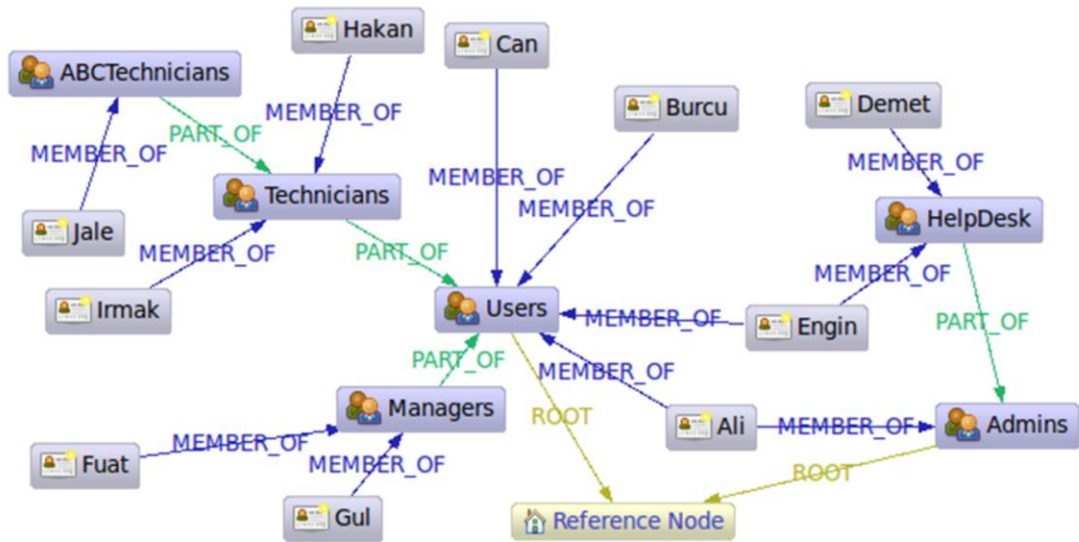
Traversal = TraversalDescription + **starting** node(s)

Traversal Framework – Java API (2)



- `org.neo4j...Evaluator`
 - Used for deciding at each node: **should** the traversal **continue**, and should the node be included in the result
 - `INCLUDE_AND_CONTINUE`: Include this node in the result and continue the traversal
 - `INCLUDE_AND_PRUNE`: Include this node, do not continue traversal
 - `EXCLUDE_AND_CONTINUE`: Exclude this node, but continue traversal
 - `EXCLUDE_AND_PRUNE`: Exclude this node and do not continue
 - **Pre-defined evaluators:**
 - `Evaluators.toDepth(int depth) / Evaluators.fromDepth(int depth),`
 - `Evaluators.excludeStartPosition()`
 - ...

Sample Data



Query: Find All “Admins”



```
Node admins = getNodeByName( "Admins" );
TraversalDescription desc = graphDb.traversalDescription()
    .breadthFirst()
    .evaluator( Evaluators.excludeStartPosition() )
    .relationships(RoleRels.PART_OF, Direction.INCOMING)
    .relationships(RoleRels.MEMBER_OF, Direction.INCOMING);
Traverser traverser = desc.traverse(admins);
StringBuilder output = new StringBuilder();

for ( Node node : traverser.nodes() ) {
    output.append("Found: ")
        .append(node.getProperty(NAME))
        .append(" at depth: ")
        .append(path.length()) .append("\n");
}
```

```
Found: HelpDesk at depth: 1
Found: Ali at depth: 1
Found: Engin at depth: 2
Found: Demet at depth: 2
```

Query: Get Group Membership of a User

```
Node jale = getNodeByName( "Jale" );
desc = graphDb.traversalDescription()
    .depthFirst()
    .evaluator( Evaluators.excludeStartPosition() )
    .relationships(RoleRels.MEMBER_OF, Direction.OUTGOING)
    .relationships(RoleRels.PART_OF, Direction.OUTGOING);

traverser = traversalDescription.traverse( jale );
```

```
Found: ABCTechnicians at depth: 1
Found: Technicians at depth: 2
Found: Users at depth: 3
```

Query: Get All Groups

```
Node referenceNode = getNodeByName( "Reference_Node" ) ;
desc = graphDb.traversalDescription()
    .breadthFirst()
    .evaluator( Evaluators.excludeStartPosition() )
    .relationships(RoleRels.ROOT, Direction.INCOMING )
    .relationships(RoleRels.PART_OF, Direction.INCOMING);

traverser = desc.traverse( referenceNode );
```

```
Found: Admins at depth: 1
Found: Users at depth: 1
Found: HelpDesk at depth: 2
Found: Managers at depth: 2
Found: Technicians at depth: 2
Found: ABCTechnicians at depth: 3
```

Query: Get All Members in the Database

```
Node referenceNode = getNodeByName( "Reference_Node" ) ;
desc = graphDb.traversalDescription()
    .breadthFirst()
    .evaluator(Evaluators.includeWhereLastRelationshipTypeIs
        (RoleRels.MEMBER_OF ) );

traverser =
desc.traverse( referenceNode );
```

```
Found: Ali at depth: 2
Found: Engin at depth: 2
Found: Burcu at depth: 2
Found: Can at depth: 2
Found: Demet at depth: 3
Found: Gul at depth: 3
Found: Fuat at depth: 3
Found: Hakan at depth: 3
Found: Irmak at depth: 3
Found: Jale at depth: 4
```


Task 4: Movies in Embedded Mode



localhost

- Use the **Movie** database in the **embedded** mode
 - download the Java Maven [project](#) from course page
 - insert the Movie database using Cypher
 - The code is prepared in `MoviesBuild.java`
 - source data in `movies-insert.cypher`

Task 5: Query Movies in Embedded Mode



- Find all **actors** who played in a movie with **Keanu Reeves**.
- Find all **directors** of movies where acted Tom Hanks.

