




# 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







# Neo4j: Basic Info





- **Open source** graph database
- Initial release: 2007
  - Current version 3.3
- Written in: **Java**
- OS: cross-platform
- Full **transactions** (ACID)
- Partitioning: None
- **Replication**: Master-slave
  - Eventual consistency


# Data Model: Nodes



- Fundamental unit: **node**
- Nodes have **properties**
  - **Key-value** pairs
  - **null** is **not** a **valid** property value
    - nulls can be modelled by the absence of a key
- Nodes have **labels**
  - labels typically express "type of node"



# Data Model: Properties



Type	Description
boolean	true/false
byte	8-bit integer
short	16-bit integer
int	32-bit integer
long	64-bit integer
float	32-bit IEEE 754 floating-point number
double	64-bit IEEE 754 floating-point number
char	16-bit unsigned integers representing Unicode characters
String	sequence of Unicode characters

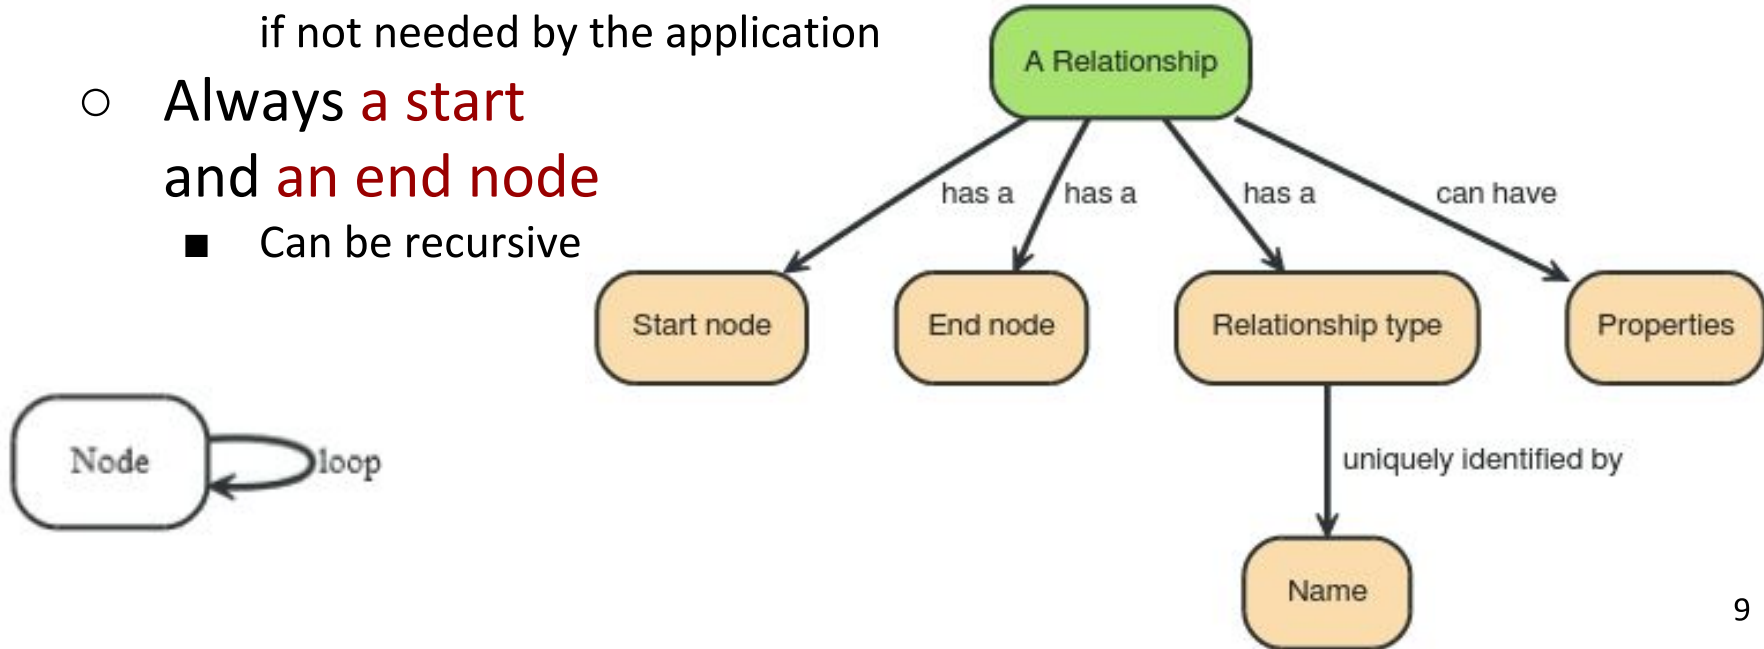


# Data Model: Relationships




- 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 Shell
  - HTTP REST API
    - also using Cypher query language
  - **Web GUI**
    - using Cypher **query language**
  - Standard **Java API**
  - Gremlin graph traversal language (plugin), etc.

# Neo4j in Server mode



- Virtual machine <http://stratus.fi.muni.cz>


- Template “PA195 - Neo4j”

```
$ ssh root@... -L 7474:localhost:7474  
                -L 7687:localhost:7687  
# neo4j-community-3.1.4/bin/neo4j start
```

- or Install on your own:


- **download** from <https://neo4j.com/> to /var/tmp
- tar xvzf neo4j-community-\*.tar.gz
  - module add jdk
- ./bin/neo4j start

# Neo4j Shell



- **command-line** shell (deprecated)
  - `./bin/neo4j-shell`
  - <http://neo4j.com/docs/stable/shell.html>
  - 1. Get **information** about the database
    - and some administration commands
  - 2. Running **Cypher** queries
  - 3. **Browsing** the graph like in Unix shell file system
    - commands like `cd`, `ls` and `pwd`
- **Cypher** shell
  - `./bin/cypher-shell`

# REST API




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

```
curl -i -X POST http://localhost:7474/db/data/node -H  
"Content-Type: application/json; charset=UTF-8" --user  
"neo4j" -d '{ "name": "Jan" }'
```




# 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
- **START**: Starting **points** in the graph
  - by explicit index lookups or by node IDs (both **deprecated**)
- **WITH**: Divides a query into multiple parts
- **CREATE**: Creates nodes and relationships.
- **DELETE**: Remove nodes, relationships, properties
- **SET**: Set values to **properties**



# 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
```

# Cypher: Creating Relationships



```
MATCH (a {name:'John'}), (b {name:'Jack'})
```

```
CREATE a-[r:Friend]->b
```

```
RETURN r ;
```

*(create a relation Friend between John and Jack)*

Created 1 relationship, returned 1 row

```
START a=node(1), b=node(2)
```


```
CREATE (a)-[r:RELTYPE {name : a.name + '->' + b.name }]->(b)
```

```
RETURN r
```

*(set property 'name' of the relationship)*

Created 1 node, set 1 property, returned 1 row

# Cypher: Creating Paths




```
CREATE p = (andres: Person {name:
'Andres'})-[:WORKS_AT]->(neo)-[:WORKS_AT]-(michael: Person
{name:'Michael'})
RETURN p ;
```

To create just a relationship, use  
MATCH and WHERE

*(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
```

# Cypher: Changing Properties



```
MATCH (n: Person {name: 'Andres'})
```

```
SET n.surname = 'Taylor'
```

```
RETURN n
```

*(find a node with name 'Andres' and set its surname 'Taylor')*


```
n
```

```
Node[0] {name:"Andres", surname:"Taylor"}
```

```
1 row
```

```
Properties set: 1
```

# Task 1: Update Queries



**Modify** the nodes so that these queries return something:

```
MATCH (p: Person)
WHERE p.age > 18 AND p.age < 30
RETURN p.name
```


*(return names of all adult people under 30)*

```
MATCH (user: Person {name: 'Andres'})-[:FRIEND]->(follower)
RETURN user.name, follower.name
```

*(find all 'friends' of 'Andres')*



# Cypher: Delete



```
MATCH (n: Person {name: 'Andres'})
```

```
DELETE n
```

*(delete all Persons with name 'Andres')*

```
Cannot delete node<3>, because it still has relationships.
```

```
MATCH (n: Person {name: 'Andres'}), ((n)-[r]-())
```

```
DELETE r,n
```

*(first, we must delete all relationships of node with name 'Andres')*


```
Nodes deleted: 1
```

```
Relationships deleted: 1
```






# Task 2: Query Movies



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



# Neo4j: “Hello World” – Java API




```
String PATH="some_directory";  
GraphDatabaseService graphDb;
```

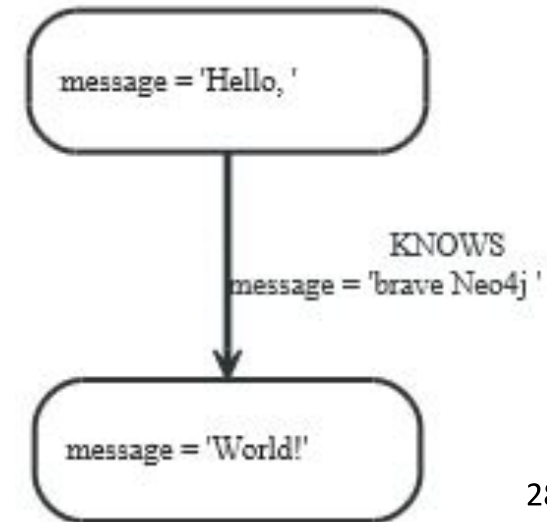
```
// starting a database  
graphDb = new GraphDatabaseFactory().newEmbeddedDatabase(new  
File(PATH));
```

```
Node firstNode, secondNode;  
Relationship relationship;
```


# Neo4j: "Hello World" (2)



```
// create a small graph:  
firstNode = graphDb.createNode();  
firstNode.setProperty( "message", "Hello, " );  
secondNode = graphDb.createNode();  
secondNode.setProperty( "message", "World!" );  
  
relationship = firstNode.createRelationshipTo  
    (secondNode,  
     RelationshipType.withName("KNOWS"));  
  
relationship.setProperty  
    ("message", "brave Neo4j ");
```



# Neo4j: Transactions



```
// all writes (creating, deleting and updating any data)
// have to be performed in a transaction:
try (Transaction tx = graphDb.beginTx()) {

    (...)

    // print the result:
    System.out.print(firstNode.getProperty("message"));
    System.out.print(relationship.getProperty("message"));
    System.out.println(secondNode.getProperty("message"));

    // transaction operations
    tx.success();
}
```


# Data Model: Path & Traversal



- **Path** = specific nodes + connecting relationships
  - Path can be a **result** of a query or a traversal

- **Traversing a graph** = visiting its nodes, following relationships according to some **rules**

- Typically, a subgraph is visited
- Neo4j: Traversal framework in Java API, Cypher, Gremlin



# 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



- `org.neo4j...TraversalDescription`
  - The main **interface** for defining **traversals**
    - Can specify branch ordering `breadthFirst()` / `depthFirst()`
- `.relationships()`
  - Specify the **relationship types** to traverse
    - e.g., traverse only edge types: `FRIEND`, `RELATIVE`
    - Empty (default) = traverse all relationships
  - Can also specify **direction**
    - `Direction.BOTH`
    - `Direction.INCOMING`
    - `Direction.OUTGOING`



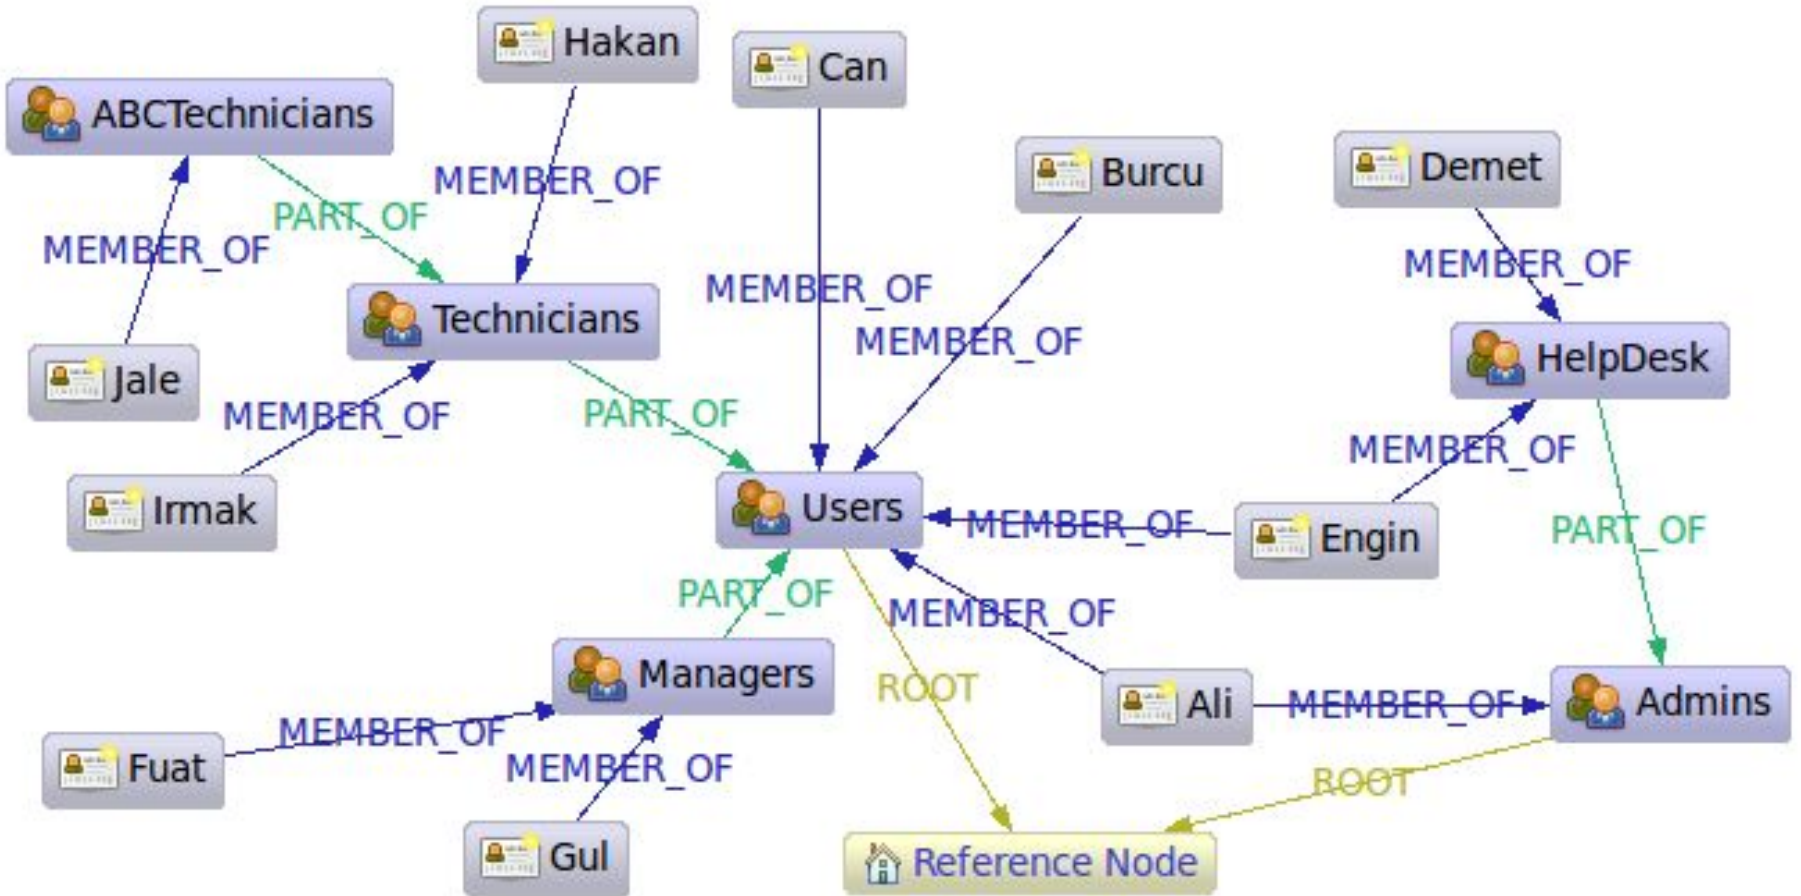
# Traversal Framework – Java API (2)




- `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
```

# Access to Nodes




- How to **get to the starting** node(s) before traversal
  1. Using **internal identifiers** (unique generated IDs)
    - **not recommended** because Neo4j does reuse freed IDs
  2. Using specified **properties**
    - one of the properties is typically “ID” (natural user-specified ID)
    - recommended, properties can be **indexed**
      - automatic indexes
  3. Using “**labels**”
    - group nodes into “**subsets**” (named graph)
    - a node can have **more** than one label
      - belong to more subsets

```
Node ali =  
graphDb.findNode(Label.label("Person"), "name", "Ali");
```



# Task 3: 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 3: 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.



