

Querying Graphs with Neo4j

TABLE OF CONTENTS

Preface	2
Introduction to Neo4j	2
What is a Graph Database	2
Cypher	3
Getting Started	3
Installing Neo4j	3
Accessing Neo4j Browser	3
Creating a new graph database	3
Basic Data Retrieval	4
Retrieving nodes	4
Retrieving relationships	4
Combining node and relationship retrieval	4
Filtering and Sorting	4
Using WHERE to filter nodes and relationships	4
Applying multiple conditions	4
Sorting query results	4
Aggregation and Grouping	4
Using COUNT, SUM, AVG, MIN, MAX	4
GROUP BY clause	4
Filtering aggregated results with HAVING	5
Advanced Relationship Traversal	5
Traversing multiple relationships	5
Variable-length relationships	5
Controlling traversal direction	5
Pattern Matching with MATCH	5
Matching specific patterns	5
Optional match with OPTIONAL MATCH	5
Using patterns as placeholders	5
Working with Path Results	5
Returning paths in queries	5
Filtering paths based on conditions	6
Limiting the number of paths	6

Modifying Data with CREATE, UPDATE, DELETE	6
Creating nodes and relationships	6
Updating property values	6
Deleting nodes, relationships, and properties	6
Indexes and Constraints	6
Creating indexes for faster querying	6
Adding uniqueness constraints	6
Dropping indexes and constraints	6
Combining Cypher Queries	6
Using WITH for result pipelining	6
Chaining multiple queries	7
Using subqueries	7
Importing Data into Neo4j	7
Using Cypher's LOAD CSV for CSV imports	7
Integrating with ETL tools	7
Data Modeling Considerations	7
Performance Tuning	7
Profiling queries for optimization	7
Understanding query execution plans	7
Working with Dates and Times	7
Storing and querying date/time values	7
Performing date calculations	8
Handling time zones	8
User-Defined Procedures and Functions	8
Creating custom procedures and functions	8
Loading and using APOC library	8
Extending Query Capabilities	8
Exporting Query Results	8
Exporting to CSV	8
JSON and other formats	8
Additional Resources	8

PREFACE

This cheatsheet is your guide to effectively querying graphs using Neo4j. Whether you're a seasoned database professional looking to expand your skills or a curious enthusiast eager to dive into the world of graph data, this resource is designed to provide you with quick and concise information to get you started.

INTRODUCTION TO NEO4J

Neo4j is a leading graph database management system that enables efficient storage and Querying Graphs of connected data. It's designed to work with highly interconnected data models, making it suitable for applications such as social networks, recommendation systems, fraud detection, and more.

Key Concepts:

Concept	Description
Nodes	Fundamental building blocks representing entities in the graph database's domain.
Relationships	Connections between nodes that convey meaningful information and context between the connected nodes.
Properties	Key-value pairs attached to nodes and relationships for storing additional data or attributes.
Graph Modeling	Graph databases employ a schema-less model that offers flexibility to adapt to changing data structures.

WHAT IS A GRAPH DATABASE

A graph database is a specialized type of database designed to store and manage data using graph structures. In a graph database, data is modeled as nodes, relationships, and properties. It's a way to represent and store complex relationships and connections between various entities in a more

intuitive and efficient manner compared to traditional relational databases.

Graph databases offer several advantages, especially when dealing with highly interconnected data:

Advantages Of Graph Databases	Description
Efficient Relationship Handling	Graph databases excel at traversing relationships efficiently. This makes them ideal for scenarios where understanding connections and relationships is a critical part of the data.
Flexible Data Modeling	Graph databases adopt a schema-less approach, enabling easy adaptation of data models as requirements evolve. This adaptability is particularly beneficial for dynamic data structures.
Complex Queries	Graph databases excel at handling complex queries involving relationships and patterns. They can uncover hidden relationships and insights that might be challenging for traditional databases.
Use Cases	Graph databases are well-suited for applications like social networks, recommendation systems, fraud detection, and knowledge graphs. They shine where relationships are as important as data.

Advantages Of Graph Databases	Description
Query Performance	Graph databases generally offer superior query performance when retrieving related data, thanks to their optimized traversal mechanisms.
Natural Representation	Graph databases provide a more natural way to model and represent real-world scenarios, aligning well with how humans perceive and understand relationships.

However, it's important to note that while graph databases excel in certain use cases, they might not be the optimal choice for every type of application. Choosing the right database technology depends on the specific needs of your project, including data structure, query patterns, and performance requirements.

CYPHER

Neo4j uses its own language for Querying Graphs called **Cypher**. Cypher is specifically designed for querying and manipulating graph data in the Neo4j database. It provides a powerful and expressive way to interact with the graph database, making it easier to work with nodes, relationships, and their properties.

Cypher is designed to be human-readable and closely resembles patterns in natural language when describing graph patterns. It allows you to express complex queries in a concise and intuitive manner. Cypher queries are written using ASCII art-like syntax to represent nodes, relationships, and patterns within the graph.

For example, a simple Cypher query to retrieve all nodes labeled as "Person" and their names might look like:

```
MATCH (p:Person)
```

```
RETURN p.name
```

In this query, **MATCH** is used to specify the pattern you're looking for, **(p:Person)** defines a node labeled as "Person," and **RETURN** specifies what information to retrieve.

Cypher also supports a wide range of functionalities beyond basic querying, including creating nodes and relationships, filtering, sorting, aggregating data, and more. It's a central tool for interacting with Neo4j databases effectively and efficiently.

It's important to note that while Cypher is specific to Neo4j, other graph databases might have their own query languages or might support other query languages like GraphQL, SPARQL, etc., depending on the database technology being used.

GETTING STARTED

To begin using Neo4j for Querying Graphs, follow these steps:

INSTALLING NEO4J

Download and install Neo4j from the [official website](#). Choose the appropriate version based on your operating system. Follow the installation instructions for a smooth setup.

ACCESSING NEO4J BROWSER

Neo4j Browser is a web-based interface that allows you to interact with your graph database using Cypher queries. After installing Neo4j, you can access the browser by navigating to <https://localhost:7474> in your web browser.

CREATING A NEW GRAPH DATABASE

Once you're in Neo4j Browser, you can create a new graph database using Cypher. For example, to create a node with a "Person" label and a "name" property, run:

```
CREATE (:Person {name: 'John'})
```

BASIC DATA RETRIEVAL

To retrieve data from your Neo4j database, you can use the **MATCH** clause along with patterns to specify what you're looking for.

RETRIEVING NODES

To retrieve all nodes with a specific label, use the **MATCH** clause followed by the label:

```
MATCH (p:Person)
RETURN p
```

RETRIEVING RELATIONSHIPS

To retrieve specific relationships between nodes, use the **MATCH** clause with the desired pattern:

```
MATCH (p1:Person)-[:FRIENDS_WITH]-
>(p2:Person)
RETURN p1, p2
```

COMBINING NODE AND RELATIONSHIP RETRIEVAL

You can retrieve both nodes and relationships in a single query:

```
MATCH (p1:Person)-[r:FRIENDS_WITH]-
>(p2:Person)
RETURN p1, r, p2
```

FILTERING AND SORTING

Use the **WHERE** clause to filter query results based on specific conditions.

USING WHERE TO FILTER NODES AND RELATIONSHIPS

Filter nodes based on property values:

```
MATCH (p:Person)
WHERE p.age > 30
```

```
RETURN p
```

APPLYING MULTIPLE CONDITIONS

Combine conditions using logical operators

```
MATCH (p:Person)
WHERE p.age > 30 AND p.location =
'New York'
RETURN p
```

SORTING QUERY RESULTS

Use the **ORDER BY** clause to sort results:

```
MATCH (p:Person)
RETURN p.name
ORDER BY p.age DESC
```

AGGREGATION AND GROUPING

Aggregation functions allow you to summarize and analyze data.

USING COUNT, SUM, AVG, MIN, MAX

Aggregate functions work on numeric properties:

```
MATCH (p:Person)
RETURN COUNT(p) AS totalPeople,
AVG(p.age) AS avgAge
```

GROUP BY CLAUSE

Group data based on specific properties:

```
MATCH (p:Person)
RETURN p.location, AVG(p.age) AS
avgAge
GROUP BY p.location
```

FILTERING AGGREGATED RESULTS WITH HAVING

Filter groups using the **HAVING** clause

```
MATCH (p:Person)
RETURN p.location, AVG(p.age) AS
avgAge
GROUP BY p.location
HAVING avgAge > 30
```

ADVANCED RELATIONSHIP TRAVERSAL

Neo4j power of Querying Graphs lies in traversing complex relationships.

TRAVERSING MULTIPLE RELATIONSHIPS

Navigate through multiple relationships:

```
MATCH (p:Person)-[:FRIENDS_WITH]-
>(:Person)-[:LIKES]->(m:Movie)
RETURN p, m
```

VARIABLE-LENGTH RELATIONSHIPS

Use the asterisk (*) syntax for variable-length paths:

```
MATCH (p:Person)-
[:FRIENDS_WITH*1..2]-
>(friend:Person)
RETURN p, friend
```

CONTROLLING TRAVERSAL DIRECTION

Specify traversal direction with arrow notation:

```
MATCH (p:Person)-[:FRIENDS_WITH]-
>(friend:Person)
RETURN p, friend
```

PATTERN MATCHING WITH MATCH

Patterns allow you to specify the structure of your data.

MATCHING SPECIFIC PATTERNS

Match nodes and relationships based on patterns:

```
MATCH (p:Person)-[:FRIENDS_WITH]-
>(friend:Person)
WHERE p.name = 'Alice'
RETURN friend
```

OPTIONAL MATCH WITH OPTIONAL MATCH

Include optional relationships in the pattern:

```
MATCH (p:Person)
OPTIONAL MATCH (p)-[:LIKES]-
>(m:Movie)
RETURN p, m
```

USING PATTERNS AS PLACEHOLDERS

Use variables to match patterns conditionally:

```
MATCH (p:Person)-[:FRIENDS_WITH]-
>(friend:Person)
WITH friend, size((friend)-[:LIKES]-
>()) AS numLikes
WHERE numLikes > 2
RETURN friend
```

WORKING WITH PATH RESULTS

Paths represent sequences of nodes and relationships.

RETURNING PATHS IN QUERIES

Use the **MATCH** clause to return paths:

```
MATCH path = (p:Person)-
[:FRIENDS_WITH]->(:Person)-[:LIKES]-
```

```
>(m:Movie)
RETURN path
```

FILTERING PATHS BASED ON CONDITIONS

Filter paths based on specific criteria:

```
MATCH path = (p:Person)-
[:FRIENDS_WITH]->(friend:Person)
WHERE size((friend)-[:LIKES]->()) >
2
RETURN path
```

LIMITING THE NUMBER OF PATHS

Use the **LIMIT** clause to restrict results:

```
MATCH path = (p:Person)-
[:FRIENDS_WITH]->(:Person)-[:LIKES]-
>(m:Movie)
RETURN path
LIMIT 5
```

MODIFYING DATA WITH CREATE, UPDATE, DELETE

Cypher allows you to create, update, and delete data.

CREATING NODES AND RELATIONSHIPS

Use the **CREATE** clause to add nodes and relationships:

```
CREATE (p:Person {name: 'Eve', age:
28})
```

UPDATING PROPERTY VALUES

Use the **SET** clause to update properties:

```
MATCH (p:Person {name: 'Eve'})
SET p.age = 29
```

DELETING NODES, RELATIONSHIPS, AND PROPERTIES

Use the **DELETE** clause to remove data:

```
MATCH (p:Person {name: 'Eve'})
DELETE p
```

INDEXES AND CONSTRAINTS

Indexes and constraints enhance query performance and data integrity.

CREATING INDEXES FOR FASTER QUERYING

Create an index on a property for faster retrieval:

```
CREATE INDEX ON :Person(name)
```

ADDING UNIQUENESS CONSTRAINTS

Enforce uniqueness on properties:

```
CREATE CONSTRAINT ON (p:Person)
ASSERT p.email IS UNIQUE
```

DROPPING INDEXES AND CONSTRAINTS

Remove indexes and constraints as needed:

```
DROP INDEX ON :Person(name)
DROP CONSTRAINT ON (p:Person) ASSERT
p.email IS UNIQUE
```

COMBINING CYPHER QUERIES

Combine multiple queries for more complex operations.

USING WITH FOR RESULT PIPELINING

Pass query results to the next part of the query:

```
MATCH (p:Person)
```

```
WITH p
MATCH (p)-[:FRIENDS_WITH]-
>(friend:Person)
RETURN p, friend
```

CHAINING MULTIPLE QUERIES

Chain queries together using semicolons:

```
MATCH (p:Person)
RETURN p.name;
MATCH (m:Movie)
RETURN m.title
```

USING SUBQUERIES

Embed subqueries within larger queries:

```
MATCH (p:Person)
WHERE p.age > (SELECT AVG(age) FROM
Person)
RETURN p
```

IMPORTING DATA INTO NEO4J

Importing external data into Neo4j as graphs for querying is a common task.

USING CYPHER'S LOAD CSV FOR CSV IMPORTS

Load data from CSV files into the graph:

```
LOAD CSV WITH HEADERS FROM
'file:///people.csv' AS row
CREATE (:Person {name: row.name,
age: toInteger(row.age)})
```

INTEGRATING WITH ETL TOOLS

Use ETL (Extract, Transform, Load) tools like Neo4j ETL or third-party tools to automate data imports.

DATA MODELING CONSIDERATIONS

Plan your graph model and relationships before importing data to ensure optimal performance and queryability.

PERFORMANCE TUNING

Optimize your queries for better performance.

PROFILING QUERIES FOR OPTIMIZATION

Use the **PROFILE** keyword to analyze query execution:

```
PROFILE MATCH (p:Person)-
[:FRIENDS_WITH]->(:Person)
RETURN p
```

UNDERSTANDING QUERY EXECUTION PLANS

Analyze query plans to identify bottlenecks and optimizations.

Tips for improving query performance:

- Use indexes for property-based filtering.
- Avoid unnecessary traversals by using specific patterns.
- Profile and analyze slow queries to identify improvements.

WORKING WITH DATES AND TIMES

Store, query, and manipulate date and time values.

STORING AND QUERYING DATE/TIME VALUES

Store date/time properties and query them using comparisons:

```
MATCH (p:Person)
WHERE p.birthdate > date('1990-01-
01')
RETURN p
```


PERFORMING DATE CALCULATIONS

Perform calculations on date properties:

```
MATCH (p:Person)
SET p.age = date().year -
p.birthdate.year
RETURN p
```

HANDLING TIME ZONES

Use the `datetime()` function to work with time zones:

```
MATCH (m:Movie)
SET m.releaseDate = datetime('2023-
07-01T00:00:00Z')
RETURN m
```

USER-DEFINED PROCEDURES AND FUNCTIONS

Extend Cypher's capabilities with user-defined procedures and functions.

CREATING CUSTOM PROCEDURES AND FUNCTIONS

Write custom procedures using Java and integrate them into your Cypher queries.

LOADING AND USING APOC LIBRARY

APOC (Awesome Procedures on Cypher) is a popular library of procedures and functions:

```
CALL apoc.date.parse('2023-07-01',
's', 'yyyy-MM-dd') YIELD value
RETURN value.year AS year
```

EXTENDING QUERY CAPABILITIES

User-defined functions allow you to encapsulate logic and reuse it in queries.

EXPORTING QUERY RESULTS

Export query results for further analysis.

EXPORTING TO CSV

Use the `EXPORT CSV` clause to export data to a CSV file:

```
MATCH (p:Person)
RETURN p.name, p.age
EXPORT CSV WITH HEADERS FROM
'file:///people.csv'
```

JSON AND OTHER FORMATS

For JSON export, use the APOC library:

```
CALL
apoc.export.json.query('MATCH
(p:Person) RETURN p';,
'people.json', {})
```

ADDITIONAL RESOURCES

Title	Description
Neo4j Documentation	Official documentation for Neo4j, including guides, tutorials, and reference materials.
Neo4j Community Forum	An online community forum where you can ask questions, share knowledge, and engage with other Neo4j users.
Cypher Query Language Manual	In-depth guide to the Cypher query language, explaining its syntax, functions, and usage.
Graph Databases for Beginners	A beginner-friendly guide to graph databases, their benefits, and how they compare to other database types.

Title	Description
Neo4j Online Training	Paid and free online courses provided by Neo4j to learn about graph databases and how to work with Neo4j effectively.
YouTube: Neo4j Channel	Neo4j's official YouTube channel with video tutorials, webinars, and talks about graph databases and Neo4j features.
GitHub: Neo4j Examples	Repository containing sample code and examples for various use cases, helping you understand practical applications of Neo4j.



JCG delivers over 1 million pages each month to more than 700K software developers, architects and decision makers. JCG offers something for everyone, including news, tutorials, cheat sheets, research guides, feature articles, source code and more.

Copyright © 2014 Exelixis Media P.C. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

CHEATSHEET FEEDBACK
 WELCOME
support@javacodegeeks.com

SPONSORSHIP
 OPPORTUNITIES
sales@javacodegeeks.com