


The Neo4j Manual



The Neo4j Manual v1.9.7

The Neo4j Team neo4j.org <<http://neo4j.org/>>
www.neotechnology.com <<http://www.neotechnology.com/>>

The Neo4j Manual v1.9.7

by The Neo4j Team neo4j.org <<http://neo4j.org/>> www.neotechnology.com <<http://www.neotechnology.com/>>

Publication date 2014-04-09 15:22:11

Copyright © 2014 Neo Technology

Starting points

- [What is a graph database?](#)
- [Cypher Query Language](#)
- [Using Neo4j embedded in Java applications](#)
- [Languages / Remote Client Libraries](#)
- [Neo4j Server](#)
- [REST API](#)
- [Server Security](#)

License: Creative Commons 3.0

This book is presented in open source and licensed through Creative Commons 3.0. You are free to copy, distribute, transmit, and/or adapt the work. This license is based upon the following conditions:

Attribution. You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).

Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under the same, similar or a compatible license.

Any of the above conditions can be waived if you get permission from the copyright holder.

In no way are any of the following rights affected by the license:

- Your fair dealing or fair use rights
- The author's moral rights
- Rights other persons may have either in the work itself or in how the work is used, such as publicity or privacy rights



Note

For any reuse or distribution, you must make clear to the others the license terms of this work. The best way to do this is with a direct link to this page: <http://creativecommons.org/licenses/by-sa/3.0/> <<http://creativecommons.org/licenses/by-sa/3.0/>>

Table of Contents

Preface	v
I. Introduction	1
1. Neo4j Highlights	2
2. Graph Database Concepts	3
3. The Neo4j Graph Database	11
II. Tutorials	20
4. Using Neo4j embedded in Java applications	21
5. Neo4j Remote Client Libraries	47
6. The Traversal Framework	53
7. Data Modeling Examples	62
8. Languages	99
9. Extending the Neo4j Server	100
III. Cypher Query Language	107
10. Introduction	108
11. Syntax	117
12. Reading Clauses	126
13. Writing Clauses	161
14. Functions	173
15. From SQL to Cypher	192
IV. Reference	199
16. Capabilities	200
17. Transaction Management	207
18. Data Import	215
19. Indexing	219
20. Graph Algorithms	238
21. Neo4j Server	240
22. REST API	251
23. Deprecations	365
V. Operations	366
24. Installation & Deployment	367
25. Configuration & Performance	377
26. High Availability	408
27. Backup	431
28. Security	436
29. Monitoring	442
VI. Tools	458
30. Web Interface	459
31. Neo4j Shell	460
VII. Community	476
32. Community Support	477
33. Contributing to Neo4j	478
A. Manpages	502
neo4j	503
neo4j-shell	505
neo4j-backup	506
neo4j-arbiter	508

Preface

This is the reference manual for Neo4j version 1.9.7, authored by the Neo4j Team.

The main parts of the manual are:

- [Part I, “Introduction”](#) — introducing graph database concepts and Neo4j.
- [Part II, “Tutorials”](#) — learn how to use Neo4j.
- [Part IV, “Reference”](#) — detailed information on Neo4j.
- [Part V, “Operations”](#) — how to install and maintain Neo4j.
- [Part VI, “Tools”](#) — guides on tools.
- [Part VII, “Community”](#) — getting help from, contributing to.
- [Appendix A, *Manpages*](#) — command line documentation.

The material is practical, technical, and focused on answering specific questions. It addresses how things work, what to do and what to avoid to successfully run Neo4j in a production environment.

The goal is to be thumb-through and rule-of-thumb friendly.

Each section should stand on its own, so you can hop right to whatever interests you. When possible, the sections distill “rules of thumb” which you can keep in mind whenever you wander out of the house without this manual in your back pocket.

The included code examples are executed when Neo4j is built and tested. Also, the REST API request and response examples are captured from real interaction with a Neo4j server. Thus, the examples are always in sync with how Neo4j actually works.

Who should read this?

The topics should be relevant to architects, administrators, developers and operations personnel.

Part I. Introduction

This part gives a bird's eye view of what a graph database is, and then outlines some specifics of Neo4j.

Chapter 1. Neo4j Highlights

As a robust, scalable and high-performance database, Neo4j is suitable for full enterprise deployment or a subset of the full server can be used in lightweight projects.

It features:

- true ACID transactions,
- high availability,
- scales to billions of nodes and relationships,
- high speed querying through traversals,
- declarative graph query language.

Proper ACID behavior is the foundation of data reliability. Neo4j enforces that all operations that modify data occur within a transaction, guaranteeing consistent data. This robustness extends from single instance embedded graphs to multi-server high availability installations. For details, see [Chapter 17, *Transaction Management*](#).

Reliable graph storage can easily be added to any application. A graph can scale in size and complexity as the application evolves, with little impact on performance. Whether starting new development, or augmenting existing functionality, Neo4j is only limited by physical hardware.

A single server instance can handle a graph of billions of nodes and relationships. When data throughput is insufficient, the graph database can be distributed among multiple servers in a high availability configuration. See [Chapter 26, *High Availability*](#) to learn more.

The graph database storage shines when storing richly-connected data. Querying is performed through traversals, which can perform millions of traversal steps per second. A traversal step resembles a *join* in a RDBMS.

Chapter 2. Graph Database Concepts

This chapter contains an introduction to the graph data model and also compares it to other data models used when persisting data.


2.1. What is a Graph Database?

A graph database stores data in a graph, the most generic of data structures, capable of elegantly representing any kind of data in a highly accessible way. Let's follow along some graphs, using them to express graph concepts. We'll "read" a graph by following arrows around the diagram to form sentences.

2.1.1. A Graph contains Nodes and Relationships

"A Graph —records data in→ Nodes —which have→ Properties"

The simplest possible graph is a single Node, a record that has named values referred to as Properties. A Node could start with a single Property and grow to a few million Properties, though that can get a little awkward. At some point it makes sense to distribute the data into multiple nodes, organized with explicit Relationships.



2.1.2. Relationships organize the Graph


"Nodes —are organized by→ Relationships —which also have→ Properties"

Relationships organize Nodes into arbitrary structures, allowing a Graph to resemble a List, a Tree, a Map, or a compound Entity – any of which can be combined into yet more complex, richly interconnected structures.

2.1.3. Query a Graph with a Traversal

"A Traversal —navigates→ a Graph; it —identifies→ Paths —which order→ Nodes"

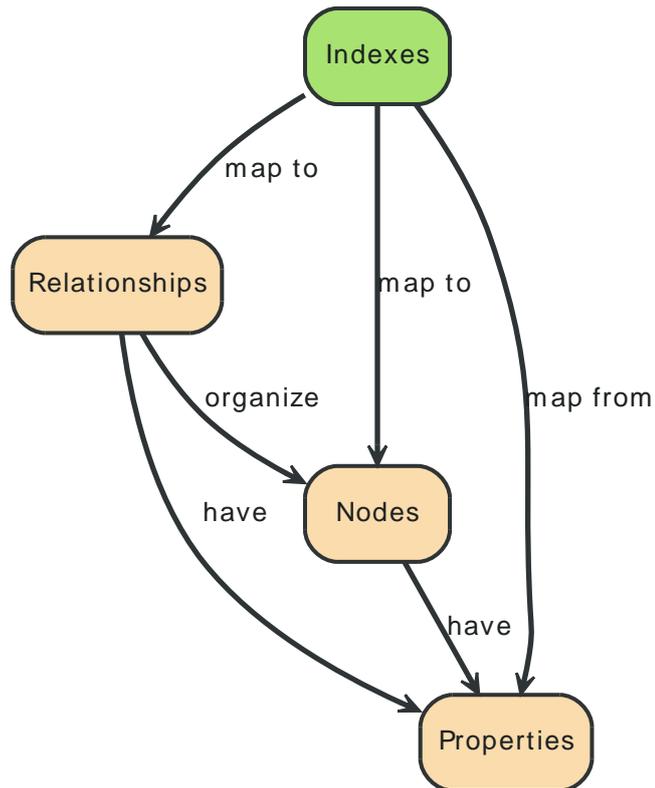
A Traversal is how you query a Graph, navigating from starting Nodes to related Nodes according to an algorithm, finding answers to questions like "what music do my friends like that I don't yet own," or "if this power supply goes down, what web services are affected?"



2.1.4. Indexes look-up Nodes or Relationships

“An Index —maps from→ Properties —to either→ Nodes or Relationships”


Often, you want to find a specific Node or Relationship according to a Property it has. Rather than traversing the entire graph, use an Index to perform a look-up, for questions like “find the Account for username master-of-graphs.”



2.1.5. Neo4j is a Graph Database

“A Graph Database —manages a→ Graph and —also manages related→ Indexes”

Neo4j is a commercially supported open-source graph database. It was designed and built from the ground-up to be a reliable database, optimized for graph structures instead of tables. Working with Neo4j, your application gets all the expressiveness of a graph, with all the dependability you expect out of a database.



2.2. Comparing Database Models

A Graph Database stores data structured in the Nodes and Relationships of a graph. How does this compare to other persistence models? Because a graph is a generic structure, let's compare how a few models would look in a graph.

2.2.1. A Graph Database transforms a RDBMS

Topple the stacks of records in a relational database while keeping all the relationships, and you'll see a graph. Where an RDBMS is optimized for aggregated data, Neo4j is optimized for highly connected data.

Figure 2.1. RDBMS

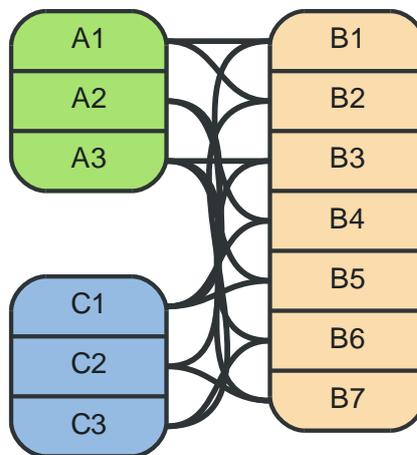
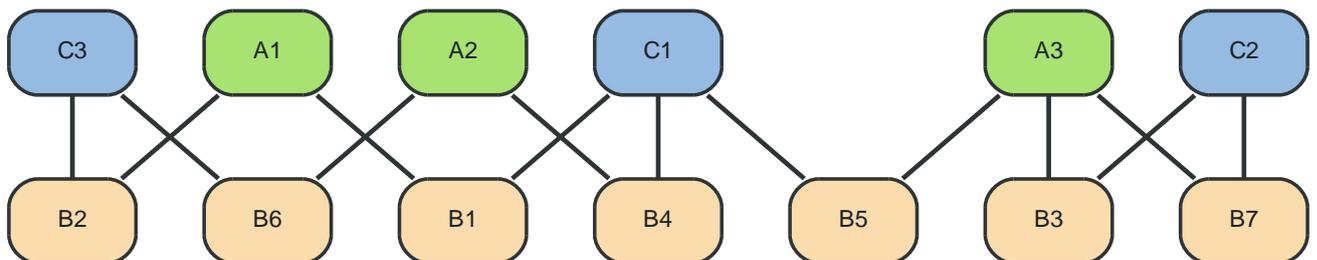


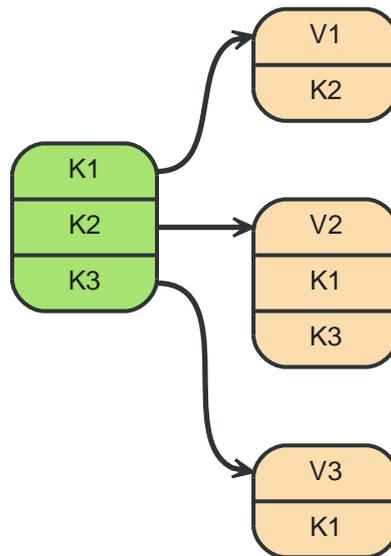
Figure 2.2. Graph Database as RDBMS



2.2.2. A Graph Database elaborates a Key-Value Store

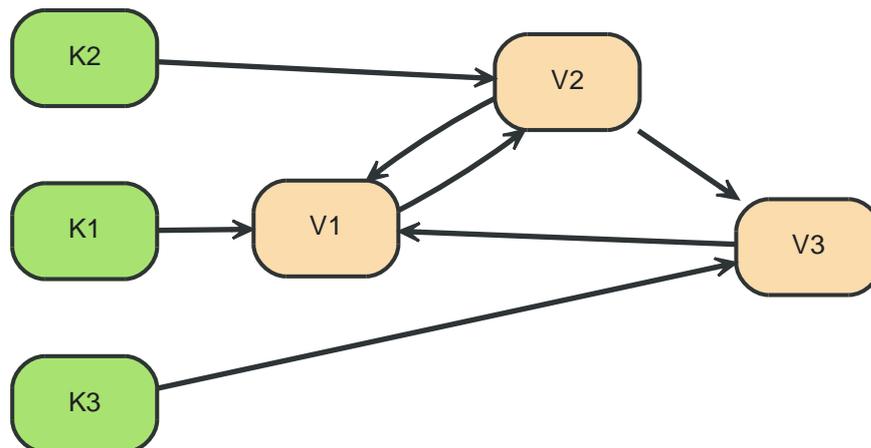
A Key-Value model is great for lookups of simple values or lists. When the values are themselves interconnected, you've got a graph. Neo4j lets you elaborate the simple data structures into more complex, interconnected data.

Figure 2.3. Key-Value Store



k^* represents a key, v^* a value. Note that some keys point to other keys as well as plain values.

Figure 2.4. Graph Database as Key-Value Store




2.2.3. A Graph Database relates Column-Family

Column Family (BigTable-style) databases are an evolution of key-value, using "families" to allow grouping of rows. Stored in a graph, the families could become hierarchical, and the relationships among data becomes explicit.

2.2.4. A Graph Database navigates a Document Store


The container hierarchy of a document database accommodates nice, schema-free data that can easily be represented as a tree. Which is of course a graph. Refer to other documents (or document elements) within that tree and you have a more expressive representation of the same data. When in Neo4j, those relationships are easily navigable.

Figure 2.5. Document Store



D=Document, s=Subdocument, v=Value, D2/S2 = reference to subdocument in (other) document.

Figure 2.6. Graph Database as Document Store




Chapter 3. The Neo4j Graph Database

This chapter goes into more detail on the data model and behavior of Neo4j.


3.1. Nodes

The fundamental units that form a graph are nodes and relationships. In Neo4j, both nodes and relationships can contain [properties](#).

Nodes are often used to represent *entities*, but depending on the domain relationships may be used for that purpose as well.




Let's start out with a really simple graph, containing only a single node with one property:



3.2. Relationships


Relationships between nodes are a key part of a graph database. They allow for finding related data. Just like nodes, relationships can have [properties](#).



A relationship connects two nodes, and is guaranteed to have valid start and end nodes.



As relationships are always directed, they can be viewed as outgoing or incoming relative to a node, which is useful when traversing the graph:



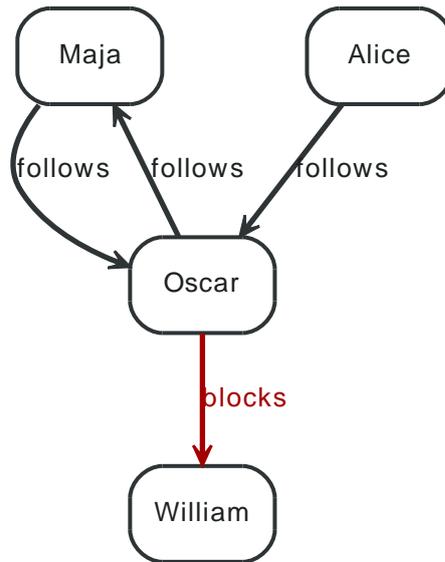
Relationships are equally well traversed in either direction. This means that there is no need to add duplicate relationships in the opposite direction (with regard to traversal or performance).

While relationships always have a direction, you can ignore the direction where it is not useful in your application.

Note that a node can have relationships to itself as well:



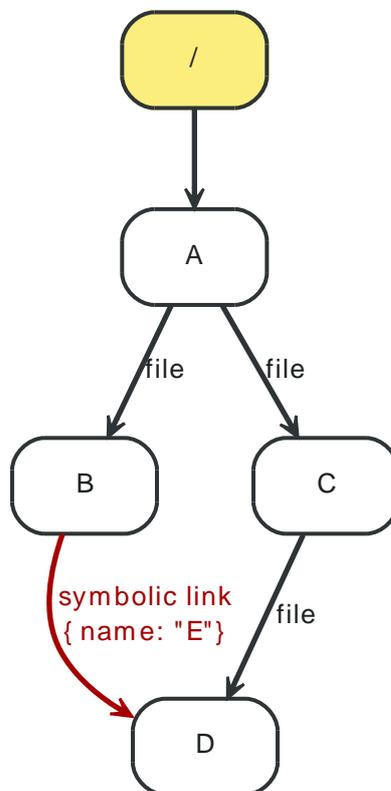
To further enhance graph traversal all relationships have a relationship type. Note that the word *type* might be misleading here, you could rather think of it as a *label*. The following example shows a simple social network with two relationship types.



Using relationship direction and type

What	How
get who a person follows	outgoing follows relationships, depth one
get the followers of a person	incoming follows relationships, depth one
get who a person blocks	outgoing blocks relationships, depth one
get who a person is blocked by	incoming blocks relationships, depth one

This example is a simple model of a file system, which includes symbolic links:



Depending on what you are looking for, you will use the direction and type of relationships during traversal.

What	How
get the full path of a file	incoming file relationships
get all paths for a file	incoming file and symbolic link relationships
get all files in a directory	outgoing file and symbolic link relationships, depth one
get all files in a directory, excluding symbolic links	outgoing file relationships, depth one
get all files in a directory, recursively	outgoing file and symbolic link relationships

3.3. Properties


Both nodes and relationships can have properties.

Properties are key-value pairs where the key is a string. Property values can be either a primitive or an array of one primitive type. For example `String`, `int` and `int[]` values are valid for properties.



Note

`null` is not a valid property value. Nulls can instead be modeled by the absence of a key.



Property value types

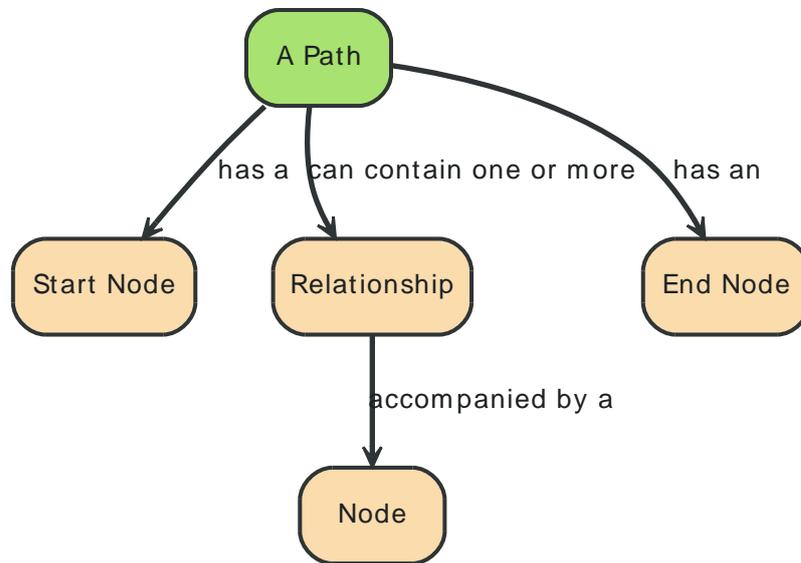
Type	Description	Value range
boolean		true/false
byte	8-bit integer	-128 to 127, inclusive
short	16-bit integer	-32768 to 32767, inclusive
int	32-bit integer	-2147483648 to 2147483647, inclusive
long	64-bit integer	-9223372036854775808 to 9223372036854775807, inclusive
float	32-bit IEEE 754 floating-point number	
double	64-bit IEEE 754 floating-point number	

Type	Description	Value range
char	16-bit unsigned integers representing Unicode characters	u0000 to uffff (0 to 65535)
String	sequence of Unicode characters	

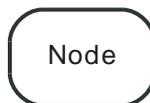
For further details on float/double values, see [Java Language Specification](http://docs.oracle.com/javase/specs/jls/se5.0/html/typesValues.html#4.2.3) <<http://docs.oracle.com/javase/specs/jls/se5.0/html/typesValues.html#4.2.3>>.

3.4. Paths


A path is one or more nodes with connecting relationships, typically retrieved as a query or traversal result.



The shortest possible path has length zero and looks like this:



A path of length one:



Another path of length one:



3.5. Traversal

Traversing a graph means visiting its nodes, following relationships according to some rules. In most cases only a subgraph is visited, as you already know where in the graph the interesting nodes and relationships are found.

Neo4j comes with a callback based traversal API which lets you specify the traversal rules. At a basic level there's a choice between traversing breadth- or depth-first.

For an in-depth introduction to the traversal framework, see [Chapter 6, *The Traversal Framework*](#). For Java code examples see [Section 4.5, "Traversal"](#).

Other options to traverse or query graphs in Neo4j are [Cypher](#) and [Gremlin](#).

Part II. Tutorials

The tutorial part describes how to set up your environment, and write programs using Neo4j. It takes you from Hello World to advanced usage of graphs.

Chapter 4. Using Neo4j embedded in Java applications

It's easy to use Neo4j embedded in Java applications. In this chapter you will find everything needed — from setting up the environment to doing something useful with your data.

4.1. Include Neo4j in your project

After selecting the appropriate [edition](#) for your platform, embed Neo4j in your Java application by including the Neo4j library jars in your build. The following sections will show how to do this by either altering the build path directly or by using dependency management.

4.1.1. Add Neo4j to the build path

Get the Neo4j libraries from one of these sources:

- Extract a Neo4j [download](http://neo4j.org/download/) zip/tarball, and use the *jar* files found in the *lib/* directory.
- Use the *jar* files available from [Maven Central Repository](http://search.maven.org/#search|gal|g%3A%22org.neo4j%22)

Add the jar files to your project:

JDK tools

Append to `-classpath`

Eclipse

- Right-click on the project and then go *Build Path* → *Configure Build Path*. In the dialog, choose *Add External JARs*, browse to the Neo4j *lib/* directory and select all of the jar files.
- Another option is to use [User Libraries](http://help.eclipse.org/indigo/index.jsp?topic=/org.eclipse.jdt.doc.user/reference/preferences/java/buildpath/ref-preferences-user-libraries.htm)

IntelliJ IDEA

See [Libraries, Global Libraries, and the Configure Library dialog](http://www.jetbrains.com/idea/webhelp/configuring-project-and-global-libraries.html)

NetBeans

- Right-click on the *Libraries* node of the project, choose *Add JAR/Folder*, browse to the Neo4j *lib/* directory and select all of the jar files.
- You can also handle libraries from the project node, see [Managing a Project's Classpath](http://netbeans.org/kb/docs/java/project-setup.html#projects-classpath)

4.1.2. Add Neo4j as a dependency

For an overview of the main Neo4j artifacts, see [Neo4j editions](#). The artifacts listed there are top-level artifacts that will transitively include the actual Neo4j implementation. You can either go with the top-level artifact or include the individual components directly. The examples included here use the top-level artifact approach.

Maven

Maven dependency.

```
<project>
...
<dependencies>
<dependency>
<groupId>org.neo4j</groupId>
<artifactId>neo4j</artifactId>
<version>1.9.7</version>
</dependency>
...
</dependencies>
```

```
...
</project>
```

Where the `artifactId` is found in [Neo4j editions](#).

Eclipse and Maven

For development in [Eclipse](http://www.eclipse.org) <<http://www.eclipse.org>>, it is recommended to install the [m2e plugin](http://www.eclipse.org/m2e/) <<http://www.eclipse.org/m2e/>> and let Maven manage the project build classpath instead, see above. This also adds the possibility to build your project both via the command line with Maven and have a working Eclipse setup for development.

Ivy

Make sure to resolve dependencies from Maven Central, for example using this configuration in your `ivysettings.xml` file:

```
<ivysettings>
  <settings defaultResolver="main"/>
  <resolvers>
    <chain name="main">
      <filesystem name="local">
        <artifact pattern="${ivy.settings.dir}/repository/[artifact]-[revision].[ext]" />
      </filesystem>
      <ibiblio name="maven_central" root="http://repo1.maven.org/maven2/" m2compatible="true"/>
    </chain>
  </resolvers>
</ivysettings>
```

With that in place you can add Neo4j to the mix by having something along these lines to your `ivy.xml` file:

```
..
<dependencies>
  ..
  <dependency org="org.neo4j" name="neo4j" rev="1.9.7"/>
  ..
</dependencies>
..
```

Where the `name` is found in [Neo4j editions](#).

Gradle

The example below shows an example gradle build script for including the Neo4j libraries.

```
def neo4jVersion = "1.9.7"
apply plugin: 'java'
repositories {
  mavenCentral()
}
dependencies {
  compile "org.neo4j:neo4j:${neo4jVersion}"
}
```

Where the `coordinates` (`org.neo4j:neo4j` in the example) are found in [Neo4j editions](#).

4.1.3. Starting and stopping

To create a new database or open an existing one you instantiate an [EmbeddedGraphDatabase](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/EmbeddedGraphDatabase.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/EmbeddedGraphDatabase.html>>.

```
graphDb = new GraphDatabaseFactory().newEmbeddedDatabase( DB_PATH );
registerShutdownHook( graphDb );
```

**Note**

The `EmbeddedGraphDatabase` instance can be shared among multiple threads. Note however that you can't create multiple instances pointing to the same database.

To stop the database, call the `shutdown()` method:

```
graphDb.shutdown();
```

To make sure Neo4j is shut down properly you can add a shutdown hook:

```
private static void registerShutdownHook( final GraphDatabaseService graphDb )
{
    // Registers a shutdown hook for the Neo4j instance so that it
    // shuts down nicely when the VM exits (even if you "Ctrl-C" the
    // running application).
    Runtime.getRuntime().addShutdownHook( new Thread()
    {
        @Override
        public void run()
        {
            graphDb.shutdown();
        }
    } );
}
```

Starting an embedded database with configuration settings

To start Neo4j with configuration settings, a Neo4j properties file can be loaded like this:

```
GraphDatabaseService graphDb = new GraphDatabaseFactory()
    .newEmbeddedDatabaseBuilder( "target/database/location" )
    .loadPropertiesFromFile( pathToConfig + "neo4j.properties" )
    .newGraphDatabase();
```

Or you could of course create you own `Map<String, String>` programmatically and use that instead, like so:

```
Map<String, String> config = new HashMap<String, String>();
config.put( "neostore.nodestore.db.mapped_memory", "10M" );
config.put( "string_block_size", "60" );
config.put( "array_block_size", "300" );
GraphDatabaseService graphDb = new GraphDatabaseFactory()
    .newEmbeddedDatabaseBuilder( "target/database/location" )
    .setConfig( config )
    .newGraphDatabase();
```

For configuration settings, see [Chapter 25, Configuration & Performance](#).

Starting an embedded read-only instance

If you want a *read-only view* of the database, create an instance this way:

```
Map<String, String> config = new HashMap<String, String>();
config.put( "read_only", "true" );
graphDb = new GraphDatabaseFactory().newEmbeddedDatabaseBuilder(
    "target/read-only-db/location" )
    .setConfig( config )
    .newGraphDatabase();
```

Obviously the database has to already exist in this case.

**Note**

Concurrent access to the same database files by multiple (read-only or write) instances is not supported.

4.2. Hello World

Learn how to create and access nodes and relationships. For information on project setup, see [Section 4.1, “Include Neo4j in your project”](#).

Remember, from [Section 2.1, “What is a Graph Database?”](#), that a Neo4j graph consist of:

- Nodes that are connected by
- Relationships, with
- Properties on both nodes and relationships.

All relationships have a type. For example, if the graph represents a social network, a relationship type could be `KNOWS`. If a relationship of the type `KNOWS` connects two nodes, that probably represents two people that know each other. A lot of the semantics (that is the meaning) of a graph is encoded in the relationship types of the application. And although relationships are directed they are equally well traversed regardless of which direction they are traversed.



Tip

The source code of this example is found here: [EmbeddedNeo4j.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/EmbeddedNeo4j.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/EmbeddedNeo4j.java>>

4.2.1. Prepare the database

Relationship types can be created by using an enum. In this example we only need a single relationship type. This is how to define it:

```
private static enum RelTypes implements RelationshipType
{
    KNOWS
}
```

We also prepare some variables to use:

```
GraphDatabaseService graphDb;
Node firstNode;
Node secondNode;
Relationship relationship;
```

The next step is to start the database server. Note that if the directory given for the database doesn't already exist, it will be created.

```
graphDb = new GraphDatabaseFactory().newEmbeddedDatabase( DB_PATH );
registerShutdownHook( graphDb );
```

Note that starting a database server is an expensive operation, so don't start up a new instance every time you need to interact with the database! The instance can be shared by multiple threads. Transactions are thread confined.

As seen, we register a shutdown hook that will make sure the database shuts down when the JVM exits. Now it's time to interact with the database.

4.2.2. Wrap writes in a transaction

All writes (creating, deleting and updating any data) have to be performed in a transaction. This is a conscious design decision, since we believe transaction demarcation to be an important part of working with a real enterprise database. Now, transaction handling in Neo4j is very easy:

```
Transaction tx = graphDb.beginTx();
```

```

try
{
    // Updating operations go here
    tx.success();
}
finally
{
    tx.finish();
}

```

For more information on transactions, see [Chapter 17, Transaction Management](#) and [Java API for Transaction](#) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Transaction.html>>.

4.2.3. Create a small graph

Now, let's create a few nodes. The API is very intuitive. Feel free to have a look at the JavaDocs at <http://components.neo4j.org/neo4j/1.9.7/apidocs/>. They're included in the distribution, as well. Here's how to create a small graph consisting of two nodes, connected with one relationship and some properties:

```


firstNode = graphDb.createNode();
firstNode.setProperty( "message", "Hello, " );
secondNode = graphDb.createNode();
secondNode.setProperty( "message", "World!" );

relationship = firstNode.createRelationshipTo( secondNode, RelTypes.KNOWS );
relationship.setProperty( "message", "brave Neo4j " );

```

We now have a graph that looks like this:

Figure 4.1. Hello World Graph



4.2.4. Print the result

After we've created our graph, let's read from it and print the result.

```

System.out.print( firstNode.getProperty( "message" ) );
System.out.print( relationship.getProperty( "message" ) );
System.out.print( secondNode.getProperty( "message" ) );

```

Which will output:

Hello, brave Neo4j World!

4.2.5. Remove the data

In this case we'll remove the data before committing:

```

// let's remove the data
firstNode.getSingleRelationship( RelTypes.KNOWS, Direction.OUTGOING ).delete();

```

```
firstNode.delete();  
secondNode.delete();
```

Note that deleting a node which still has relationships when the transaction commits will fail. This is to make sure relationships always have a start node and an end node.

4.2.6. Shut down the database server

Finally, shut down the database server *when the application finishes*:

```
graphDb.shutdown();
```

4.3. User database with index

You have a user database, and want to retrieve users by name.



Tip

The source code used in this example is found here: [EmbeddedNeo4jWithIndexing.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/EmbeddedNeo4jWithIndexing.java)
<<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/EmbeddedNeo4jWithIndexing.java>>

We have created two helper methods to handle user names and adding users to the database:

```
private static String idToUserName( final int id )
{
    return "user" + id + "@neo4j.org";
}

private static Node createAndIndexUser( final String username )
{
    Node node = graphDb.createNode();
    node.setProperty( USERNAME_KEY, username );
    nodeIndex.add( node, USERNAME_KEY, username );
    return node;
}
```

The next step is to start the database server:

```
graphDb = new GraphDatabaseFactory().newEmbeddedDatabase( DB_PATH );
nodeIndex = graphDb.index().forNodes( "nodes" );
registerShutdownHook();
```

It's time to add the users:

```
Transaction tx = graphDb.beginTx();
try
{
    // Create some users and index their names with the IndexService
    for ( int id = 0; id < 100; id++ )
    {
        Node userNode = createAndIndexUser( idToUserName( id ) );
    }
}
```

And here's how to find a user by Id:

```
int idToFind = 45;
String userName = idToUserName( idToFind );
Node foundUser = nodeIndex.get( USERNAME_KEY, userName ).getSingle();
System.out.println( "The username of user " + idToFind + " is "
    + foundUser.getProperty( USERNAME_KEY ) );
```

4.4. Basic unit testing

The basic pattern of unit testing with Neo4j is illustrated by the following example.

To access the Neo4j testing facilities you should have the `neo4j-kernel tests.jar` on the classpath during tests. You can download it from Maven Central: [org.neo4j:neo4j-kernel](http://search.maven.org/#search|g|a%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-kernel%22) <<http://search.maven.org/#search|g|a%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-kernel%22>>.

Using Maven as a dependency manager you would typically add this dependency together with JUnit and Hamcrest like so:

Maven dependency.

```
<project>
...
<dependencies>
  <dependency>
    <groupId>org.neo4j</groupId>
    <artifactId>neo4j-kernel</artifactId>
    <version>1.9.7</version>
    <type>test-jar</type>
    <scope>test</scope>
  </dependency>
  <dependency>
    <groupId>junit</groupId>
    <artifactId>junit-dep</artifactId>
    <version>4.11</version>
    <scope>test</scope>
  </dependency>
  <dependency>
    <groupId>org.hamcrest</groupId>
    <artifactId>hamcrest-all</artifactId>
    <version>1.3</version>
    <scope>test</scope>
  </dependency>
  ...
</dependencies>
...
</project>
```

Observe that the `<type>test-jar</type>` is crucial. Without it you would get the common `neo4j-kernel jar`, not the one containing the testing facilities.

With that in place, we're ready to code our tests.



Tip

For the full source code of this example see: [Neo4jBasicDocTest.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/Neo4jBasicDocTest.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/Neo4jBasicDocTest.java>>

Before each test, create a fresh database:

```
@Before
public void prepareTestDatabase()
{
  graphDb = new TestGraphDatabaseFactory().newImpermanentDatabase();
}
```

After the test has executed, the database should be shut down:

```
@After
public void destroyTestDatabase()
{
```

```
graphDb.shutdown();
}
```

During a test, create nodes and check to see that they are there, while enclosing write operations in a transaction.

```
Transaction tx = graphDb.beginTx();

Node n = null;
try
{
    n = graphDb.createNode();
    n.setProperty( "name", "Nancy" );
    tx.success();
}
catch ( Exception e )
{
    tx.failure();
}
finally
{
    tx.finish();
}

// The node should have an id greater than 0, which is the id of the
// reference node.
assertThat( n.getId(), is( greaterThan( 0L ) ) );

// Retrieve a node by using the id of the created node. The id's and
// property should match.
Node foundNode = graphDb.getNodeById( n.getId() );
assertThat( foundNode.getId(), is( n.getId() ) );
assertThat( (String) foundNode.getProperty( "name" ), is( "Nancy" ) );
```

If you want to set configuration parameters at database creation, it's done like this:

```
Map<String, String> config = new HashMap<String, String>();
config.put( "neostore.nodestore.db.mapped_memory", "10M" );
config.put( "string_block_size", "60" );
config.put( "array_block_size", "300" );
GraphDatabaseService db = new TestGraphDatabaseFactory()
    .newImpermanentDatabaseBuilder()
    .setConfig( config )
    .newGraphDatabase();
```

4.5. Traversal


For reading about traversals, see [Chapter 6, The Traversal Framework](#).

For more examples of traversals, see [Chapter 7, Data Modeling Examples](#).

4.5.1. The Matrix

This is the first graph we want to traverse into:

Figure 4.2. Matrix node space view



Tip

The source code of this example is found here: [NewMatrix.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/NewMatrix.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/NewMatrix.java>>

Friends and friends of friends.

```
private static Traverser getFriends(
    final Node person )
{
    TraversalDescription td = Traversal.description()
        .breadthFirst()
        .relationships( RelTypes.KNOWS, Direction.OUTGOING )
        .evaluator( Evaluators.excludeStartPosition() );
    return td.traverse( person );
}
```

Let's perform the actual traversal and print the results:

```
int numberOfFriends = 0;
String output = neoNode.getProperty( "name" ) + "'s friends:\n";
Traverser friendsTraverser = getFriends( neoNode );
for ( Path friendPath : friendsTraverser )
{
    output += "At depth " + friendPath.length() + " => "
        + friendPath.endNode()
        .getProperty( "name" ) + "\n";
    numberOfFriends++;
}
```

```
output += "Number of friends found: " + numberOfFriends + "\n";
```

Which will give us the following output:

```
Thomas Anderson's friends:
At depth 1 => Trinity
At depth 1 => Morpheus
At depth 2 => Cypher
At depth 3 => Agent Smith
Number of friends found: 4
```

Who coded the Matrix?

```
private static Traverser findHackers( final Node startNode )
{
    TraversalDescription td = Traversal.description()
        .breadthFirst()
        .relationships( RelTypes.CODED_BY, Direction.OUTGOING )
        .relationships( RelTypes.KNOWS, Direction.OUTGOING )
        .evaluator(
            Evaluators.includeWhereLastRelationshipTypeIs( RelTypes.CODED_BY ) );
    return td.traverse( startNode );
}
```

Print out the result:

```
String output = "Hackers:\n";
int numberOfHackers = 0;
Traverser traverser = findHackers( getNeoNode() );
for ( Path hackerPath : traverser )
{
    output += "At depth " + hackerPath.length() + " => "
        + hackerPath.endNode()
        .getProperty( "name" ) + "\n";
    numberOfHackers++;
}
output += "Number of hackers found: " + numberOfHackers + "\n";
```

Now we know who coded the Matrix:

```
Hackers:
At depth 4 => The Architect
Number of hackers found: 1
```

Walking an ordered path

This example shows how to use a path context holding a representation of a path.




Tip

The source code of this example is found here: [OrderedPath.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/orderedpath/OrderedPath.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/orderedpath/OrderedPath.java>>

Create a toy graph.

```
Node A = db.createNode();
Node B = db.createNode();
Node C = db.createNode();
Node D = db.createNode();

A.createRelationshipTo( C, REL2 );
C.createRelationshipTo( D, REL3 );
A.createRelationshipTo( B, REL1 );
B.createRelationshipTo( C, REL2 );
```



Now, the order of relationships (REL1 → REL2 → REL3) is stored in an `ArrayList`. Upon traversal, the Evaluator can check against it to ensure that only paths are included and returned that have the predefined order of relationships:

Define how to walk the path.

```

final ArrayList<RelationshipType> orderedPathContext = new ArrayList<RelationshipType>();
orderedPathContext.add( REL1 );
orderedPathContext.add( withName( "REL2" ) );
orderedPathContext.add( withName( "REL3" ) );
TraversalDescription td = Traversal.description()
    .evaluator( new Evaluator()
    {
        @Override
        public Evaluation evaluate( final Path path )
        {
            if ( path.length() == 0 )
            {
                return Evaluation.EXCLUDE_AND_CONTINUE;
            }
            RelationshipType expectedType = orderedPathContext.get( path.length() - 1 );
            boolean isExpectedType = path.lastRelationship()
                .isType( expectedType );
            boolean included = path.length() == orderedPathContext.size() && isExpectedType;
            boolean continued = path.length() < orderedPathContext.size() && isExpectedType;
            return Evaluation.of( included, continued );
        }
    } )
    .uniqueness( Uniqueness.NODE_PATH );
  
```

Note that we set the uniqueness to `Uniqueness.NODE_PATH` <http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Uniqueness.html#NODE_PATH> as we want to be able to revisit the same node during the traversal, but not the same path.

Perform the traversal and print the result.

```

Traverser traverser = td.traverse( A );
PathPrinter pathPrinter = new PathPrinter( "name" );
for ( Path path : traverser )
{
    output += Traversal.pathToString( path, pathPrinter );
}

```

Which will output:

```
(A)--[REL1]-->(B)--[REL2]-->(C)--[REL3]-->(D)
```

In this case we use a custom class to format the path output. This is how it's done:

```

static class PathPrinter implements Traversal.PathDescriptor<Path>
{
    private final String nodePropertyKey;

    public PathPrinter( String nodePropertyKey )
    {
        this.nodePropertyKey = nodePropertyKey;
    }

    @Override
    public String nodeRepresentation( Path path, Node node )
    {
        return "(" + node.getProperty( nodePropertyKey, "" ) + ")";
    }

    @Override
    public String relationshipRepresentation( Path path, Node from, Relationship relationship )
    {
        String prefix = "--", suffix = "--";
        if ( from.equals( relationship.getEndNode() ) )
        {
            prefix = "<--";
        }
        else
        {
            suffix = "-->";
        }
        return prefix + "[" + relationship.getType().name() + "]" + suffix;
    }
}


```

For options regarding output of a Path, see the [Traversal](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Traversal.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Traversal.html>> class.

4.5.2. Uniqueness of Paths in traversals

This example is demonstrating the use of node uniqueness. Below an imaginary domain graph with Principals that own pets that are descendant to other pets.

Figure 4.3. Descendants Example Graph



In order to return all descendants of `Pet0` which have the relation `owns` to `Principal1` (`Pet1` and `Pet3`), the Uniqueness of the traversal needs to be set to `NODE_PATH` rather than the default `NODE_GLOBAL` so that nodes can be traversed more than once, and paths that have different nodes but can have some nodes in common (like the start and end node) can be returned.

```
final Node target = data.get().get( "Principal1" );
TraversalDescription td = Traversal.description()
    .uniqueness( Uniqueness.NODE_PATH )
    .evaluator( new Evaluator()
{
    @Override
    public Evaluation evaluate( Path path )
    {
        if ( path.endNode().equals( target ) )
        {
            return Evaluation.INCLUDE_AND_PRUNE;
        }
        return Evaluation.EXCLUDE_AND_CONTINUE;
    }
} );

Traverser results = td.traverse( start );
```

This will return the following paths:

```
(3)--[descendant,0]-->(1)<--[owns,3]--(5)
(3)--[descendant,2]-->(4)<--[owns,5]--(5)
```

In the default `path.toString()` implementation, `(1)--[knows,2]-->(4)` denotes a node with ID=1 having a relationship with ID 2 or type `knows` to a node with ID=4.

Let's create a new `TraversalDescription` from the old one, having `NODE_GLOBAL` uniqueness to see the difference.



Tip

The `TraversalDescription` object is immutable, so we have to use the new instance returned with the new uniqueness setting.

```
TraversalDescription nodeGlobalTd = td.uniqueness( Uniqueness.NODE_GLOBAL );
results = nodeGlobalTd.traverse( start );
```

Now only one path is returned:

```
(3)--[descendant,0]-->(1)<--[owns,3]--(5)
```

4.5.3. Social network



Note

The following example uses the new enhanced traversal API.

Social networks (known as social graphs out on the web) are natural to model with a graph. This example shows a very simple social model that connects friends and keeps track of status updates.




Tip

The source code of the example is found here: [socnet](https://github.com/neo4j/neo4j/tree/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/socnet) <<https://github.com/neo4j/neo4j/tree/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/socnet>>

Simple social model

Figure 4.4. Social network data model




The data model for a social network is pretty simple: Persons with names and StatusUpdates with timestamped text. These entities are then connected by specific relationships.

- Person
 - friend: relates two distinct Person instances (no self-reference)
 - status: connects to the most recent StatusUpdate
- StatusUpdate
 - next: points to the next StatusUpdate in the chain, which was posted before the current one

Status graph instance

The StatusUpdate list for a Person is a linked list. The head of the list (the most recent status) is found by following status. Each subsequent StatusUpdate is connected by next.

Here's an example where Andreas Kollegger micro-blogged his way to work in the morning:



To read the status updates, we can create a traversal, like so:

```

TraversalDescription traversal = Traversal.description().
    depthFirst().
    relationships( NEXT );
  
```

This gives us a traverser that will start at one `StatusUpdate`, and will follow the chain of updates until they run out. Traversers are lazy loading, so it's performant even when dealing with thousands of statuses — they are not loaded until we actually consume them.

Activity stream

Once we have friends, and they have status messages, we might want to read our friends status' messages, in reverse time order — latest first. To do this, we go through these steps:

1. Gather all friend's status update iterators in a list — latest date first.
2. Sort the list.
3. Return the first item in the list.
4. If the first iterator is exhausted, remove it from the list. Otherwise, get the next item in that iterator.
5. Go to step 2 until there are no iterators left in the list.

Animated, the sequence looks like [this](http://www.slideshare.net/systay/pattern-activity-stream) <<http://www.slideshare.net/systay/pattern-activity-stream>>.

The code looks like:

```

PositionedIterator<StatusUpdate> first = statuses.get(0);
StatusUpdate returnVal = first.current();

if ( !first.hasNext() )
{
    statuses.remove( 0 );
}
else
  
```

```
{  
    first.next();  
    sort();  
}  
return returnVal;
```

4.6. Domain entities

This page demonstrates one way to handle domain entities when using Neo4j. The principle at use is to wrap the entities around a node (the same approach can be used with relationships as well).



Tip

The source code of the examples is found here: [Person.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/socnet/Person.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/socnet/Person.java>>

First off, store the node and make it accessible inside the package:

```
private final Node underlyingNode;

Person( Node personNode )
{
    this.underlyingNode = personNode;
}

protected Node getUnderlyingNode()
{
    return underlyingNode;
}
```

Delegate attributes to the node:

```
public String getName()
{
    return (String)underlyingNode.getProperty( NAME );
}
```

Make sure to override these methods:

```
@Override
public int hashCode()
{
    return underlyingNode.hashCode();
}

@Override
public boolean equals( Object o )
{
    return o instanceof Person &&
        underlyingNode.equals( ( (Person)o ).getUnderlyingNode() );
}

@Override
public String toString()
{
    return "Person[" + getName() + "]";
}
```

4.7. Graph Algorithm examples



Tip

The source code used in the example is found here: [PathFindingDocTest.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/PathFindingDocTest.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/PathFindingDocTest.java>>

Calculating the shortest path (least number of relationships) between two nodes:

```
Node startNode = graphDb.createNode();
Node middleNode1 = graphDb.createNode();
Node middleNode2 = graphDb.createNode();
Node middleNode3 = graphDb.createNode();
Node endNode = graphDb.createNode();
createRelationshipsBetween( startNode, middleNode1, endNode );
createRelationshipsBetween( startNode, middleNode2, middleNode3, endNode );

// Will find the shortest path between startNode and endNode via
// "MY_TYPE" relationships (in OUTGOING direction), like f.ex:
//
// (startNode)-->(middleNode1)-->(endNode)
//
PathFinder<Path> finder = GraphAlgoFactory.shortestPath(
    Traversal.expanderForTypes( ExampleTypes.MY_TYPE, Direction.OUTGOING ), 15 );
Iterable<Path> paths = finder.findAllPaths( startNode, endNode );
```


Using [Dijkstra's algorithm](http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm) <http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm> to calculate cheapest path between node A and B where each relationship can have a weight (i.e. cost) and the path(s) with least cost are found.

```
PathFinder<WeightedPath> finder = GraphAlgoFactory.dijkstra(
    Traversal.expanderForTypes( ExampleTypes.MY_TYPE, Direction.BOTH ), "cost" );

WeightedPath path = finder.findSinglePath( nodeA, nodeB );

// Get the weight for the found path
path.weight();
```

Using [A*](http://en.wikipedia.org/wiki/A*_search_algorithm) <http://en.wikipedia.org/wiki/A*_search_algorithm> to calculate the cheapest path between node A and B, where cheapest is for example the path in a network of roads which has the shortest length between node A and B. Here's our example graph:



```
Node nodeA = createNode( "name", "A", "x", 0d, "y", 0d );
Node nodeB = createNode( "name", "B", "x", 7d, "y", 0d );
Node nodeC = createNode( "name", "C", "x", 2d, "y", 1d );
Relationship relAB = createRelationship( nodeA, nodeC, "length", 2d );
Relationship relBC = createRelationship( nodeC, nodeB, "length", 3d );
Relationship relAC = createRelationship( nodeA, nodeB, "length", 10d );

EstimateEvaluator<Double> estimateEvaluator = new EstimateEvaluator<Double>()
```

```
{
    public Double getCost( final Node node, final Node goal )
    {
        double dx = (Double) node.getProperty( "x" ) - (Double) goal.getProperty( "x" );
        double dy = (Double) node.getProperty( "y" ) - (Double) goal.getProperty( "y" );
        double result = Math.sqrt( Math.pow( dx, 2 ) + Math.pow( dy, 2 ) );
        return result;
    }
};
PathFinder<WeightedPath> astar = GraphAlgoFactory.aStar(
    Traversal.expanderForAllTypes(),
    CommonEvaluators.doubleCostEvaluator( "length" ), estimateEvaluator );
WeightedPath path = astar.findSinglePath( nodeA, nodeB );
```

4.8. Reading a management attribute

The [JmxUtils](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/jmx/JmxUtils.html) <http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/jmx/JmxUtils.html> class includes methods to access Neo4j management beans. The common JMX service can be used as well, but from your code you probably rather want to use the approach outlined here.



Tip

The source code of the example is found here: [JmxDocTest.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/JmxDocTest.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/JmxDocTest.java>

This example shows how to get the start time of a database:

```
private static Date getStartTimeFromManagementBean(
    GraphDatabaseService graphDbService )
{
    ObjectName objectName = JmxUtils.getObjectNames( graphDbService, "Kernel" );
    Date date = JmxUtils.getAttribute( objectName, "KernelStartTime" );
    return date;
}
```

Depending on which Neo4j edition you are using different sets of management beans are available.

- For all editions, see the [org.neo4j.jmx](http://components.neo4j.org/neo4j-jmx/1.9.7/apidocs/org/neo4j/jmx/package-summary.html) <http://components.neo4j.org/neo4j-jmx/1.9.7/apidocs/org/neo4j/jmx/package-summary.html> package.
- For the Advanced and Enterprise editions, see the [org.neo4j.management](http://components.neo4j.org/neo4j-management/1.9.7/apidocs/org/neo4j/management/package-summary.html) <http://components.neo4j.org/neo4j-management/1.9.7/apidocs/org/neo4j/management/package-summary.html> package as well.

4.9. How to create unique nodes

This section is about how to ensure uniqueness of a property when creating nodes. For an overview of the topic, see [Section 17.6, “Creating unique nodes”](#).

4.9.1. Get or create unique node

By using `put-if-absent` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#putIfAbsent%28T,%20java.lang.String,%20java.lang.Object%29>> functionality, entity uniqueness can be guaranteed using an index. Here the index acts as the lock and will only lock the smallest part needed to guarantee uniqueness across threads and transactions. To get the more high-level `get-or-create` functionality make use of `UniqueFactory` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/UniqueFactory.html>> as seen in the example below.

Create a factory for unique nodes at application start.

```
Transaction transaction = graphDb.beginTx();
try
{
    UniqueFactory<Node> factory = new UniqueFactory.UniqueNodeFactory( graphDb, "users" )
    {
        @Override
        protected void initialize( Node created, Map<String, Object> properties )
        {
            created.setProperty( "name", properties.get( "name" ) );
        }
    };
    return factory;
}
finally
{
    transaction.finish();
}
```

Use the unique node factory to get or create a node.

```
Transaction transaction = graphDb.beginTx();
try
{
    Node node = factory.getOrCreate( "name", username );
    transaction.success();
    return node;
}
finally
{
    transaction.finish();
}
```

4.9.2. Pessimistic locking for node creation



Important

While this is a working solution, please consider using the preferred [Section 4.9.1, “Get or create unique node”](#) instead.

One might be tempted to use Java synchronization for pessimistic locking, but this is dangerous. By mixing locks in Neo4j and in the Java runtime, it is easy to produce deadlocks that are not detectable by Neo4j. As long as all locking is done by Neo4j, all deadlocks will be detected and avoided. Also, a solution using manual synchronization doesn't ensure uniqueness in an HA environment.

This example uses a single “lock node” for locking. We create it only as a place to put locks, nothing else.

Create a lock node at application start.

```
Transaction tx = graphdb().beginTx();
try
{
    final Node lockNode = graphdb().createNode();
    tx.success();
    return lockNode;
}
finally
{
    tx.finish();
}
```

Use the lock node to ensure nodes are not created concurrently.

```
Transaction tx = graphDb.beginTx();
try
{
    Index<Node> usersIndex = graphDb.index().forNodes( "users" );
    Node userNode = usersIndex.get( "name", username ).getSingle();
    if ( userNode != null ) return userNode;

    tx.acquireWriteLock( lockNode );
    userNode = usersIndex.get( "name", username ).getSingle();
    if ( userNode == null )
    {
        userNode = graphDb.createNode();
        userNode.setProperty( "name", username );
        usersIndex.add( userNode, "name", username );
    }
    tx.success();
    return userNode;
}
finally
{
    tx.finish();
}
```

Note that finishing the transaction will release the lock on the lock node.

4.10. Execute Cypher Queries from Java



Tip

The full source code of the example: [JavaQuery.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/cypher/src/test/java/org/neo4j/cypher/javacompat/JavaQuery.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/cypher/src/test/java/org/neo4j/cypher/javacompat/JavaQuery.java>>

In Java, you can use the [Cypher query language](#) as per the example below. First, let's add some data.

```
GraphDatabaseService db = new GraphDatabaseFactory().newEmbeddedDatabase( DB_PATH );
Transaction tx = db.beginTx();
try
{
    Node myNode = db.createNode();
    myNode.setProperty( "name", "my node" );
    tx.success();
}
finally
{
    tx.finish();
}
```

Execute a query:

```
ExecutionEngine engine = new ExecutionEngine( db );
ExecutionResult result = engine.execute( "start n=node(*) where n.name! = 'my node' return n, n.name" );
```



Note

Keep the ExecutionEngine around, don't create a new one for each query!

The result will be:

```
+-----+
| n          | n.name |
+-----+
| Node[1]{name:"my node"} | "my node" |
+-----+
1 row
```



Caution

The classes used here are from the `org.neo4j.cypher.javacompat` package, *not* `org.neo4j.cypher`, see link to the Java API below.

You can get a list of the columns in the result:

```
List<String> columns = result.columns();
```

This contains:

```
[n, n.name]
```

To fetch the result items from a single column, do like this:

```
Iterator<Node> n_column = result.columnAs( "n" );
for ( Node node : IteratorUtil.asIterable( n_column ) )
{
    // note: we're grabbing the name property from the node,
    // not from the n.name in this case.
    nodeResult = node + ": " + node.getProperty( "name" );
}
```

In this case there's only one node in the result:

```
Node[1]: my node
```

To get all columns, do like this instead:

```
for ( Map<String, Object> row : result )
{
    for ( Entry<String, Object> column : row.entrySet() )
    {
        rows += column.getKey() + ": " + column.getValue() + "; ";
    }
    rows += "\n";
}
```

This outputs:

```
n.name: my node; n: Node[1];
```



Caution

`dumpToString()`, `columnAs()` and `iterator()` cannot be called more than once on the same `ExecutionResult` object. You should instead use only one and if you need the facilities of the other methods on the same query result instead create a new `ExecutionResult`.



Caution

When using an `ExecutionResult`, you'll need to exhaust it by using any of the iterating methods (`columnAs()` and `iterator()`) on it. Failing to do so will not properly clean up resources used by the `ExecutionResult`, leading to unwanted behaviour, such as leaking transactions.

For more information on the Java interface to Cypher, see the [Java API](http://components.neo4j.org/neo4j-cypher/1.9.7/apidocs/index.html) <<http://components.neo4j.org/neo4j-cypher/1.9.7/apidocs/index.html>>.

For more information and examples for Cypher, see [Part III, “Cypher Query Language”](#) and [Chapter 7, *Data Modeling Examples*](#).

Chapter 5. Neo4j Remote Client Libraries

The included Java example shows a “low-level” approach to using the Neo4j REST API from Java. For other options, see below.

Neo4j REST clients contributed by the community. Please see <http://www.neo4j.org/drivers> for the current set of REST clients!

5.1. How to use the REST API from Java

5.1.1. Creating a graph through the REST API from Java

The REST API uses HTTP and JSON, so that it can be used from many languages and platforms. Still, when getting started it's useful to see some patterns that can be re-used. In this brief overview, we'll show you how to create and manipulate a simple graph through the REST API and also how to query it.

For these examples, we've chosen the [Jersey](http://jersey.java.net/) <http://jersey.java.net/> client components, which are easily [downloaded](http://jersey.java.net/nonav/documentation/latest/user-guide.html#chapter_deps) <http://jersey.java.net/nonav/documentation/latest/user-guide.html#chapter_deps> via Maven.

5.1.2. Start the server

Before we can perform any actions on the server, we need to start it as per [Section 21.1, "Server Installation"](#).

```
WebResource resource = Client.create()
    .resource( SERVER_ROOT_URI );
ClientResponse response = resource.get( ClientResponse.class );

System.out.println( String.format( "GET on [%s], status code [%d]",
    SERVER_ROOT_URI, response.getStatus() ) );
response.close();
```

If the status of the response is 200 OK, then we know the server is running fine and we can continue. If the code fails to connect to the server, then please have a look at [Chapter 21, Neo4j Server](#).



Note

If you get any other response than 200 OK (particularly 4xx or 5xx responses) then please check your configuration and look in the log files in the *data/log* directory.

5.1.3. Creating a node

The REST API uses POST to create nodes. Encapsulating that in Java is straightforward using the Jersey client:

```
final String nodeEntryPointUri = SERVER_ROOT_URI + "node";
// http://localhost:7474/db/data/node

WebResource resource = Client.create()
    .resource( nodeEntryPointUri );
// POST {} to the node entry point URI
ClientResponse response = resource.accept( MediaType.APPLICATION_JSON )
    .type( MediaType.APPLICATION_JSON )
    .entity( "{}" )
    .post( ClientResponse.class );

final URI location = response.getLocation();
System.out.println( String.format(
    "POST to [%s], status code [%d], location header [%s]",
    nodeEntryPointUri, response.getStatus(), location.toString() ) );
response.close();

return location;
```

If the call completes successfully, under the covers it will have sent a HTTP request containing a JSON payload to the server. The server will then have created a new node in the database and responded with a 201 Created response and a Location header with the URI of the newly created node.

In our example, we call this functionality twice to create two nodes in our database.

5.1.4. Adding properties

Once we have nodes in our database, we can use them to store useful data. In this case, we're going to store information about music in our database. Let's start by looking at the code that we use to create nodes and add properties. Here we've added nodes to represent "Joe Strummer" and a band called "The Clash".

```
URI firstNode = createNode();
addProperty( firstNode, "name", "Joe Strummer" );
URI secondNode = createNode();
addProperty( secondNode, "band", "The Clash" );
```

Inside the `addProperty` method we determine the resource that represents properties for the node and decide on a name for that property. We then proceed to `PUT` the value of that property to the server.

```
String propertyUri = nodeUri.toString() + "/properties/" + propertyName;
// http://localhost:7474/db/data/node/{node_id}/properties/{property_name}

WebResource resource = Client.create()
    .resource( propertyUri );
ClientResponse response = resource.accept( MediaType.APPLICATION_JSON )
    .type( MediaType.APPLICATION_JSON )
    .entity( "\"" + propertyValue + "\"" )
    .put( ClientResponse.class );

System.out.println( String.format( "PUT to [%s], status code [%d]",
    propertyUri, response.getStatus() ) );
response.close();
```

If everything goes well, we'll get a `204 No Content` back indicating that the server processed the request but didn't echo back the property value.

5.1.5. Adding relationships

Now that we have nodes to represent Joe Strummer and The Clash, we can relate them. The REST API supports this through a `POST` of a relationship representation to the start node of the relationship. Correspondingly in Java we `POST` some JSON to the URI of our node that represents Joe Strummer, to establish a relationship between that node and the node representing The Clash.

```
URI relationshipUri = addRelationship( firstNode, secondNode, "singer",
    "{ \"from\" : \"1976\", \"until\" : \"1986\" }" );
```

Inside the `addRelationship` method, we determine the URI of the Joe Strummer node's relationships, and then `POST` a JSON description of our intended relationship. This description contains the destination node, a label for the relationship type, and any attributes for the relation as a JSON collection.

```
private static URI addRelationship( URI startNode, URI endNode,
    String relationshipType, String jsonAttributes )
    throws URISyntaxException
{
    URI fromUri = new URI( startNode.toString() + "/relationships" );
    String relationshipJson = generateJsonRelationship( endNode,
        relationshipType, jsonAttributes );

    WebResource resource = Client.create()
        .resource( fromUri );
    // POST JSON to the relationships URI
    ClientResponse response = resource.accept( MediaType.APPLICATION_JSON )
```

```

        .type( MediaType.APPLICATION_JSON )
        .entity( relationshipJson )
        .post( ClientResponse.class );

    final URI location = response.getLocation();
    System.out.println( String.format(
        "POST to [%s], status code [%d], location header [%s]",
        fromUri, response.getStatus(), location.toString() ) );

    response.close();
    return location;
}

```

If all goes well, we receive a 201 Created status code and a Location header which contains a URI of the newly created relation.

5.1.6. Add properties to a relationship

Like nodes, relationships can have properties. Since we're big fans of both Joe Strummer and the Clash, we'll add a rating to the relationship so that others can see he's a 5-star singer with the band.

```
addMetadataToProperty( relationshipUri, "stars", "5" );
```

Inside the `addMetadataToProperty` method, we determine the URI of the properties of the relationship and PUT our new values (since it's PUT it will always overwrite existing values, so be careful).

```

private static void addMetadataToProperty( URI relationshipUri,
    String name, String value ) throws URISyntaxException
{
    URI propertyUri = new URI( relationshipUri.toString() + "/properties" );
    String entity = toJsonNameValuePairCollection( name, value );
    WebResource resource = Client.create()
        .resource( propertyUri );
    ClientResponse response = resource.accept( MediaType.APPLICATION_JSON )
        .type( MediaType.APPLICATION_JSON )
        .entity( entity )
        .put( ClientResponse.class );

    System.out.println( String.format(
        "PUT [%s] to [%s], status code [%d]", entity, propertyUri,
        response.getStatus() ) );
    response.close();
}

```

Assuming all goes well, we'll get a 204 OK response back from the server (which we can check by calling `ClientResponse.getStatus()`) and we've now established a very small graph that we can query.

5.1.7. Querying graphs

As with the embedded version of the database, the Neo4j server uses graph traversals to look for data in graphs. Currently the Neo4j server expects a JSON payload describing the traversal to be POST-ed at the starting node for the traversal (though this is *likely to change* in time to a GET-based approach).

To start this process, we use a simple class that can turn itself into the equivalent JSON, ready for POST-ing to the server, and in this case we've hardcoded the traverser to look for all nodes with outgoing relationships with the type "singer".

```

// TraversalDefinition turns into JSON to send to the Server
TraversalDefinition t = new TraversalDefinition();
t.setOrder( TraversalDefinition.DEPTH_FIRST );
t.setUniqueness( TraversalDefinition.NODE );

```

```
t.setMaxDepth( 10 );
t.setReturnFilter( TraversalDefinition.ALL );
t.setRelationships( new Relation( "singer", Relation.OUT ) );
```

Once we have defined the parameters of our traversal, we just need to transfer it. We do this by determining the URI of the traversers for the start node, and then POST-ing the JSON representation of the traverser to it.

```
URI traverserUri = new URI( startNode.toString() + "/traverse/node" );
WebResource resource = Client.create()
    .resource( traverserUri );
String jsonTraverserPayload = t.toJson();
ClientResponse response = resource.accept( MediaType.APPLICATION_JSON )
    .type( MediaType.APPLICATION_JSON )
    .entity( jsonTraverserPayload )
    .post( ClientResponse.class );

System.out.println( String.format(
    "POST [%s] to [%s], status code [%d], returned data: "
    + System.getProperty( "line.separator" ) + "%s",
    jsonTraverserPayload, traverserUri, response.getStatus(),
    response.getEntity( String.class ) ) );
response.close();
```

Once that request has completed, we get back our dataset of singers and the bands they belong to:

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/82/relationships/out",
  "data" : {
    "band" : "The Clash",
    "name" : "Joe Strummer"
  },
  "traverse" : "http://localhost:7474/db/data/node/82/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/82/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/82/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/82/relationships/all",
  "self" : "http://localhost:7474/db/data/node/82",
  "properties" : "http://localhost:7474/db/data/node/82/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/82/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/82/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/82/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/82/relationships"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/83/relationships/out",
  "data" : {
  },
  "traverse" : "http://localhost:7474/db/data/node/83/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/83/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/83/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/83/relationships/all",
  "self" : "http://localhost:7474/db/data/node/83",
  "properties" : "http://localhost:7474/db/data/node/83/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/83/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/83/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/83/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/83/relationships"
} ]
```

5.1.8. Phew, is that it?

That's a flavor of what we can do with the REST API. Naturally any of the HTTP idioms we provide on the server can be easily wrapped, including removing nodes and relationships through DELETE. Still if you've gotten this far, then switching `.post()` for `.delete()` in the Jersey client code should be straightforward.

5.1.9. What's next?

The HTTP API provides a good basis for implementers of client libraries, it's also great for HTTP and REST folks. In the future though we expect that idiomatic language bindings will appear to take advantage of the REST API while providing comfortable language-level constructs for developers to use, much as there are similar bindings for the embedded database.

5.1.10. Appendix: the code

- [CreateSimpleGraph.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/CreateSimpleGraph.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/CreateSimpleGraph.java>
- [Relation.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/Relation.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/Relation.java>
- [TraversalDefinition.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/TraversalDefinition.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/TraversalDefinition.java>

Chapter 6. The Traversal Framework


The [Neo4j Traversal API](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/package-summary.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/package-summary.html>> is a callback based, lazily executed way of specifying desired movements through a graph in Java. Some traversal examples are collected under [Section 4.5, “Traversal”](#).

Other options to traverse or query graphs in Neo4j are [Cypher](#) and [Gremlin](#).

6.1. Main concepts

Here follows a short explanation of all different methods that can modify or add to a traversal description.

- *Expanders* — define what to traverse, typically in terms of relationship direction and type.
- *Order* — for example depth-first or breadth-first.
- *Uniqueness* — visit nodes (relationships, paths) only once.
- *Evaluator* — decide what to return and whether to stop or continue traversal beyond the current position.
- *Starting nodes* where the traversal will begin.



See [Section 6.2, “Traversal Framework Java API”](#) for more details.

6.2. Traversal Framework Java API

The traversal framework consists of a few main interfaces in addition to `Node` and `Relationship`: `TraversalDescription`, `Evaluator`, `Traverser` and `Uniqueness` are the main ones. The `Path` interface also has a special purpose in traversals, since it is used to represent a position in the graph when evaluating that position. Furthermore the `PathExpander` (replacing `RelationshipExpander`) and `Expander` interfaces are central to traversals, but users of the API rarely need to implement them. There are also a set of interfaces for advanced use, when explicit control over the traversal order is required: `BranchSelector`, `BranchOrderingPolicy` and `TraversalBranch`.

6.2.1. TraversalDescription

The `TraversalDescription` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html>> is the main interface used for defining and initializing traversals. It is not meant to be implemented by users of the traversal framework, but rather to be provided by the implementation of the traversal framework as a way for the user to describe traversals. `TraversalDescription` instances are immutable and its methods returns a new `TraversalDescription` that is modified compared to the object the method was invoked on with the arguments of the method.

Relationships

Adds a relationship type to the list of relationship types to traverse. By default that list is empty and it means that it will traverse *all relationships*, irregardless of type. If one or more relationships are added to this list *only the added* types will be traversed. There are two methods, one [including direction](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#relationships) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#relationships>> and another one [excluding direction](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#relationships) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#relationships>>, where the latter traverses relationships in [both directions](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Direction.html#BOTH) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Direction.html#BOTH>>.

6.2.2. Evaluator

`Evaluator` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Evaluator.html>>s are used for deciding, at each position (represented as a `Path`): should the traversal continue, and/or should the node be included in the result. Given a `Path`, it asks for one of four actions for that branch of the traversal:

- `Evaluation.INCLUDE_AND_CONTINUE`: Include this node in the result and continue the traversal
- `Evaluation.INCLUDE_AND_PRUNE`: Include this node in the result, but don't continue the traversal
- `Evaluation.EXCLUDE_AND_CONTINUE`: Exclude this node from the result, but continue the traversal
- `Evaluation.EXCLUDE_AND_PRUNE`: Exclude this node from the result and don't continue the traversal

More than one evaluator can be added. Note that evaluators will be called for all positions the traverser encounters, even for the start node.

6.2.3. Traverser

The `Traverser` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html>> object is the result of invoking `traverse()` <[http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#traverse\(org.neo4j.graphdb.Node\)](http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#traverse(org.neo4j.graphdb.Node))> of a `TraversalDescription` object. It represents a traversal positioned in the graph, and a specification of the format of the result. The actual traversal is performed lazily each time the `next()`-method of the iterator of the `Traverser` is invoked.

6.2.4. Uniqueness

Sets the rules for how positions can be revisited during a traversal as stated in [Uniqueness](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Uniqueness.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Uniqueness.html>>. Default if not set is `NODE_GLOBAL` <http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Uniqueness.html#NODE_GLOBAL>.

A Uniqueness can be supplied to the TraversalDescription to dictate under what circumstances a traversal may revisit the same position in the graph. The various uniqueness levels that can be used in Neo4j are:

- `NONE`: Any position in the graph may be revisited.
- `NODE_GLOBAL` uniqueness: No node in the entire graph may be visited more than once. This could potentially consume a lot of memory since it requires keeping an in-memory data structure remembering all the visited nodes.
- `RELATIONSHIP_GLOBAL` uniqueness: no relationship in the entire graph may be visited more than once. For the same reasons as `NODE_GLOBAL` uniqueness, this could use up a lot of memory. But since graphs typically have a larger number of relationships than nodes, the memory overhead of this uniqueness level could grow even quicker.
- `NODE_PATH` uniqueness: A node may not occur previously in the path reaching up to it.
- `RELATIONSHIP_PATH` uniqueness: A relationship may not occur previously in the path reaching up to it.
- `NODE_RECENT` uniqueness: Similar to `NODE_GLOBAL` uniqueness in that there is a global collection of visited nodes each position is checked against. This uniqueness level does however have a cap on how much memory it may consume in the form of a collection that only contains the most recently visited nodes. The size of this collection can be specified by providing a number as the second argument to the TraversalDescription.uniqueness()-method along with the uniqueness level.
- `RELATIONSHIP_RECENT` uniqueness: Works like `NODE_RECENT` uniqueness, but with relationships instead of nodes.

Depth First / Breadth First

These are convenience methods for setting preorder [depth-first](http://en.wikipedia.org/wiki/Depth-first_search) <http://en.wikipedia.org/wiki/Depth-first_search>/ [breadth-first](http://en.wikipedia.org/wiki/Breadth-first_search) <http://en.wikipedia.org/wiki/Breadth-first_search> BranchSelector/ordering policies. The same result can be achieved by calling the `order` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#order>> method with ordering policies from the [Traversal](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Traversal.html#preorderDepthFirst) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Traversal.html#preorderDepthFirst>> [factory](http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/kernel/Traversal.html#preorderBreadthFirst) <<http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/kernel/Traversal.html#preorderBreadthFirst>>, or to write your own BranchSelector/BranchOrderingPolicy and pass in.

6.2.5. Order — How to move through branches?

A more generic version of depthFirst/breadthFirst methods in that it allows an arbitrary [BranchOrderingPolicy](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/BranchOrderingPolicy.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/BranchOrderingPolicy.html>> to be injected into the description.

6.2.6. BranchSelector

A BranchSelector is used for selecting which branch of the traversal to attempt next. This is used for implementing traversal orderings. The traversal framework provides a few basic ordering implementations:

- `Traversal.preorderDepthFirst()`: Traversing depth first, visiting each node before visiting its child nodes.
- `Traversal.postorderDepthFirst()`: Traversing depth first, visiting each node after visiting its child nodes.
- `Traversal.preorderBreadthFirst()`: Traversing breadth first, visiting each node before visiting its child nodes.
- `Traversal.postorderBreadthFirst()`: Traversing breadth first, visiting each node after visiting its child nodes.



Note

Please note that breadth first traversals have a higher memory overhead than depth first traversals.

`BranchSelector` carries state and hence needs to be uniquely instantiated for each traversal. Therefore it is supplied to the `TraversalDescription` through a `BranchOrderingPolicy` interface, which is a factory of `BranchSelector` instances.

A user of the Traversal framework rarely needs to implement his own `BranchSelector` or `BranchOrderingPolicy`, it is provided to let graph algorithm implementors provide their own traversal orders. The Neo4j Graph Algorithms package contains for example a `BestFirst` order `BranchSelector`/`BranchOrderingPolicy` that is used in `BestFirst` search algorithms such as A* and Dijkstra.

BranchOrderingPolicy

A factory for creating `BranchSelector`s to decide in what order branches are returned (where a branch's position is represented as a `Path` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Path.html>> from the start node to the current node). Common policies are `depth-first` <[http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#depthFirst\(\)](http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#depthFirst())> and `breadth-first` <[http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#breadthFirst\(\)](http://components.neo4j.org/neo4j-kernel/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#breadthFirst())> and that's why there are convenience methods for those. For example, calling `TraversalDescription#depthFirst()` <[http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#depthFirst\(\)](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html#depthFirst())> is equivalent to:

```
description.order( Traversal.preorderDepthFirst() );
```

TraversalBranch

An object used by the `BranchSelector` to get more branches from a certain branch. In essence these are a composite of a `Path` and a `RelationshipExpander` that can be used to get new `TraversalBranch` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalBranch.html>>es from the current one.

6.2.7. Path

A `Path` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Path.html>> is a general interface that is part of the Neo4j API. In the traversal API of Neo4j the use of Paths are twofold. Traversers can return their results in the form of the Paths of the visited positions in the graph that are marked for being returned. Path objects are also used in the evaluation of positions in the graph, for determining if the traversal should continue from a certain point or not, and whether a certain position should be included in the result set or not.

6.2.8. PathExpander/RelationshipExpander

The traversal framework use PathExpanders (replacing RelationshipExpander) to discover the relationships that should be followed from a particular path to further branches in the traversal.

6.2.9. Expander

A more generic version of relationships where a RelationshipExpander is injected, defining all relationships to be traversed for any given node. By default (and when using relationships) a [default expander](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Traversal.html#emptyExpander) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Traversal.html#emptyExpander>> is used, where any particular order of relationships isn't guaranteed. There's another implementation which guarantees that relationships are traversed in [order of relationship type](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/OrderByTypeExpander.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/OrderByTypeExpander.html>>, where types are iterated in the order they were added.


The Expander interface is an extension of the RelationshipExpander interface that makes it possible to build customized versions of an Expander. The implementation of TraversalDescription uses this to provide methods for defining which relationship types to traverse, this is the usual way a user of the API would define a RelationshipExpander — by building it internally in the TraversalDescription.

All the RelationshipExpanders provided by the Neo4j traversal framework also implement the Expander interface. For a user of the traversal API it is easier to implement the PathExpander/RelationshipExpander interface, since it only contains one method — the method for getting the relationships from a path/node, the methods that the Expander interface adds are just for building new Expanders.

6.2.10. How to use the Traversal framework

A [traversal description](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html>> is built using a fluent interface and such a description can then spawn [traversers](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html>>.

Figure 6.1. Traversal Example Graph



With the definition of the RelationshipTypes as

```
private enum Rels implements RelationshipType
{
    LIKES, KNOWS
}
```

The graph can be traversed with for example the following traverser, starting at the “Joe” node:

```
for ( Path position : Traversal.description()
    .depthFirst()
    .relationships( Rels.KNOWS )
    .relationships( Rels.LIKES, Direction.INCOMING )
    .evaluator( Evaluators.toDepth( 5 ) )
    .traverse( node ) )
{
    output += position + "\n";
}
```

The traversal will output:

```
(7)
(7)<--[LIKES,1]--(4)
(7)<--[LIKES,1]--(4)--[KNOWS,6]-->(1)
(7)<--[LIKES,1]--(4)--[KNOWS,6]-->(1)--[KNOWS,4]-->(6)
(7)<--[LIKES,1]--(4)--[KNOWS,6]-->(1)--[KNOWS,4]-->(6)--[KNOWS,3]-->(5)
```

```
(7)<--[LIKES,1]--(4)--[KNOWS,6]-->(1)--[KNOWS,4]-->(6)--[KNOWS,3]-->(5)--[KNOWS,2]-->(2)
(7)<--[LIKES,1]--(4)--[KNOWS,6]-->(1)<--[KNOWS,5]--(3)
```

Since [TraversalDescription](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/TraversalDescription.html) are immutable it is also useful to create template descriptions which holds common settings shared by different traversals. For example, let's start with this traverser:

```
final TraversalDescription FRIENDS_TRAVERSAL = Traversal.description()
    .depthFirst()
    .relationships( Rels.KNOWS )
    .uniqueness( Uniqueness.RELATIONSHIP_GLOBAL );
```

This traverser would yield the following output (we will keep starting from the “Joe” node):

```
(7)
(7)--[KNOWS,0]-->(2)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)<--[KNOWS,4]--(1)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)<--[KNOWS,4]--(1)<--[KNOWS,5]--(3)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)<--[KNOWS,4]--(1)<--[KNOWS,6]--(4)
```

Now let's create a new traverser from it, restricting depth to three:

```
for ( Path path : FRIENDS_TRAVERSAL
    .evaluator( Evaluators.toDepth( 3 ) )
    .traverse( node ) )
{
    output += path + "\n";
}
```

This will give us the following result:

```
(7)
(7)--[KNOWS,0]-->(2)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)
```

Or how about from depth two to four? That's done like this:

```
for ( Path path : FRIENDS_TRAVERSAL
    .evaluator( Evaluators.fromDepth( 2 ) )
    .evaluator( Evaluators.toDepth( 4 ) )
    .traverse( node ) )
{
    output += path + "\n";
}
```

This traversal gives us:

```
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)
(7)--[KNOWS,0]-->(2)<--[KNOWS,2]--(5)<--[KNOWS,3]--(6)<--[KNOWS,4]--(1)
```

For various useful evaluators, see the [Evaluators](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Evaluators.html) Java API or simply implement the [Evaluator](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Evaluator.html) interface yourself.

If you're not interested in the [Path](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Path.html)s, but the [Node](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Node.html)s you can transform the traverser into an iterable of [nodes](#)

[components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html#nodes\(\)>](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html#nodes()>) like this:

```
for ( Node currentNode : FRIENDS_TRAVERSAL
    .traverse( node )
    .nodes() )
{
    output += currentNode.getProperty( "name" ) + "\n";
}
```

In this case we use it to retrieve the names:

```
Joe
Sara
Peter
Dirk
Lars
Ed
Lisa
```

Relationships <[http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html#relationships\(\)>](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/traversal/Traverser.html#relationships()>) are fine as well, here's how to get them:

```
for ( Relationship relationship : FRIENDS_TRAVERSAL
    .traverse( node )
    .relationships() )
{
    output += relationship.getType() + "\n";
}
```

Here the relationship types are written, and we get:

```
KNOWS
KNOWS
KNOWS
KNOWS
KNOWS
KNOWS
```



Tip

The source code for the traversers in this example is available at: [TraversalExample.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/TraversalExample.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/main/java/org/neo4j/examples/TraversalExample.java>>


Chapter 7. Data Modeling Examples

The following chapters contain simplified examples of how different domains can be modeled using Neo4j. The aim is not to give full examples, but to suggest possible ways to think using nodes, relationships, graph patterns and data locality in traversals.

The examples use Cypher queries a lot, read [Part III, “Cypher Query Language”](#) for more information.

7.1. User roles in graphs

This is an example showing a hierarchy of roles. What's interesting is that a tree is not sufficient for storing this kind of structure, as elaborated below.



This is an implementation of an example found in the article [A Model to Represent Directed Acyclic Graphs \(DAG\) on SQL Databases](http://www.codeproject.com/Articles/22824/A-Model-to-Represent-Directed-Acyclic-Graphs-DAG-o) <<http://www.codeproject.com/Articles/22824/A-Model-to-Represent-Directed-Acyclic-Graphs-DAG-o>> by [Kemal Erdogan](http://www.codeproject.com/script/Articles/MemberArticles.aspx?amid=274518) <<http://www.codeproject.com/script/Articles/MemberArticles.aspx?amid=274518>>. The article discusses how to store [directed acyclic graphs](http://en.wikipedia.org/wiki/Directed_acyclic_graph) (DAGs) in SQL based DBs. DAGs are almost trees, but with a twist: it may be possible to reach the same node through different paths. Trees are restricted from this possibility, which makes them much easier to handle. In our case it is “Ali” and “Engin”, as they are both admins and users and thus reachable through these group nodes. Reality often looks this way and can't be captured by tree structures.

In the article an SQL Stored Procedure solution is provided. The main idea, that also have some support from scientists, is to pre-calculate all possible (transitive) paths. Pros and cons of this approach:

- decent performance on read
- low performance on insert
- wastes *lots* of space
- relies on stored procedures

In Neo4j storing the roles is trivial. In this case we use `PART_OF` (green edges) relationships to model the group hierarchy and `MEMBER_OF` (blue edges) to model membership in groups. We also connect the top level groups to the reference node by `ROOT` relationships. This gives us a useful partitioning of the graph. Neo4j has no predefined relationship types, you are free to create any relationship types and give them the semantics you want.

Lets now have a look at how to retrieve information from the graph. The Java code is using the Neo4j Traversal API (see [Section 6.2, “Traversal Framework Java API”](#)), the queries are done using [Cypher](#).

7.1.1. Get the admins

```
Node admins = getNodeByName( "Admins" );
TraversalDescription traversalDescription = Traversal.description()
    .breadthFirst()
    .evaluator( Evaluators.excludeStartPosition() )
    .relationships( RoleRels.PART_OF, Direction.INCOMING )
    .relationships( RoleRels.MEMBER_OF, Direction.INCOMING );
Traverser traverser = traversalDescription.traverse( admins );
```

resulting in the output

```
Found: HelpDesk at depth: 0
Found: Ali at depth: 0
Found: Engin at depth: 1
Found: Demet at depth: 1
```

The result is collected from the traverser using this code:

```
String output = "";
for ( Path path : traverser )
{
    Node node = path.endNode();
    output += "Found: " + node.getProperty( NAME ) + " at depth: "
        + ( path.length() - 1 ) + "\n";
}
}
```

In Cypher, a similar query would be:

```
START admins=node(14)
MATCH admins<-[:PART_OF*0..]-group<-[:MEMBER_OF]-user
RETURN user.name, group.name
```

resulting in:

user.name	group.name
"Ali"	"Admins"
"Engin"	"HelpDesk"
"Demet"	"HelpDesk"
3 rows	

7.1.2. Get the group memberships of a user

Using the Neo4j Java Traversal API, this query looks like:

```
Node jale = getNodeByName( "Jale" );
TraversalDescription traversalDescription = Traversal.description()
    .depthFirst()
    .evaluator( Evaluators.excludeStartPosition() )
    .relationships( RoleRels.MEMBER_OF, Direction.OUTGOING )
    .relationships( RoleRels.PART_OF, Direction.OUTGOING );
traverser = traversalDescription.traverse( jale );
```

resulting in:

```
Found: ABCTechnicians at depth: 0
Found: Technicians at depth: 1
Found: Users at depth: 2
```

In Cypher:

```
START jale=node(10)
```

```
MATCH jale-[:MEMBER_OF]->()-[:PART_OF*0..]->group
RETURN group.name
```

group.name
"ABCTechnicians"
"Technicians"
"Users"
3 rows

7.1.3. Get all groups

In Java:

```
Node referenceNode = getNodeByName( "Reference_Node" ) ;
traversalDescription = Traversal.description()
    .breadthFirst()
    .evaluator( Evaluators.excludeStartPosition() )
    .relationships( RoleReIs.ROOT, Direction.INCOMING )
    .relationships( RoleReIs.PART_OF, Direction.INCOMING );
traverser = traversalDescription.traverse( referenceNode );
```

resulting in:

```
Found: Admins at depth: 0
Found: Users at depth: 0
Found: HelpDesk at depth: 1
Found: Managers at depth: 1
Found: Technicians at depth: 1
Found: ABCTechnicians at depth: 2
```

In Cypher:

```
START refNode=node(16)
MATCH refNode<-[:ROOT]->()-[:PART_OF*0..]-group
RETURN group.name
```

group.name
"Admins"
"HelpDesk"
"Users"
"Managers"
"Technicians"
"ABCTechnicians"
6 rows

7.1.4. Get all members of all groups

Now, let's try to find all users in the system being part of any group.

in Java:

```
traversalDescription = Traversal.description()
    .breadthFirst()
    .evaluator(
        Evaluators.includeWhereLastRelationshipTypeIs( RoleReIs.MEMBER_OF ) );
```

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

```
Found: Ali at depth: 1  
Found: Engin at depth: 1  
Found: Burcu at depth: 1  
Found: Can at depth: 1  
Found: Demet at depth: 2  
Found: Gul at depth: 2  
Found: Fuat at depth: 2  
Found: Hakan at depth: 2  
Found: Irmak at depth: 2  
Found: Jale at depth: 3
```

In Cypher, this looks like:

```
START refNode=node(16)  
MATCH refNode<-[:ROOT]->root, p=root<-[:PART_OF*0..]-()-[:MEMBER_OF]-user  
RETURN user.name, min(length(p))  
ORDER BY min(length(p)), user.name
```

and results in the following output:

user.name	min(length(p))
"Ali"	1
"Burcu"	1
"Can"	1
"Engin"	1
"Demet"	2
"Fuat"	2
"Gul"	2
"Hakan"	2
"Irmak"	2
"Jale"	3
10 rows	

As seen above, querying even more complex scenarios can be done using comparatively short constructs in Java and other query mechanisms.


7.2. ACL structures in graphs

This example gives a generic overview of an approach to handling Access Control Lists (ACLs) in graphs, and a simplified example with concrete queries.

7.2.1. Generic approach

In many scenarios, an application needs to handle security on some form of managed objects. This example describes one pattern to handle this through the use of a graph structure and traversers that build a full permissions-structure for any managed object with exclude and include overriding possibilities. This results in a dynamic construction of ACLs based on the position and context of the managed object.

The result is a complex security scheme that can easily be implemented in a graph structure, supporting permissions overriding, principal and content composition, without duplicating data anywhere.



Technique

As seen in the example graph layout, there are some key concepts in this domain model:

- The managed content (folders and files) that are connected by **HAS_CHILD_CONTENT** relationships
- The Principal subtree pointing out principals that can act as ACL members, pointed out by the **PRINCIPAL** relationships.
- The aggregation of principals into groups, connected by the **IS_MEMBER_OF** relationship. One principal (user or group) can be part of many groups at the same time.
- The **SECURITY** — relationships, connecting the content composite structure to the principal composite structure, containing a addition/removal modifier property ("RW").

Constructing the ACL

The calculation of the effective permissions (e.g. Read, Write, Execute) for a principal for any given ACL-managed node (content) follows a number of rules that will be encoded into the permissions-traversal:

Top-down-Traversal

This approach will let you define a generic permission pattern on the root content, and then refine that for specific sub-content nodes and specific principals.

1. Start at the content node in question traverse upwards to the content root node to determine the path to it.
2. Start with a effective optimistic permissions list of "all permitted" (111 in a bit encoded ReadWriteExecute case) or 000 if you like pessimistic security handling (everything is forbidden unless explicitly allowed).
3. Beginning from the topmost content node, look for any SECURITY relationships on it.
4. If found, look if the principal in question is part of the end-principal of the SECURITY relationship.
5. If yes, add the "+" permission modifiers to the existing permission pattern, revoke the "-" permission modifiers from the pattern.
6. If two principal nodes link to the same content node, first apply the more generic principals modifiers.
7. Repeat the security modifier search all the way down to the target content node, thus overriding more generic permissions with the set on nodes closer to the target node.

The same algorithm is applicable for the bottom-up approach, basically just traversing from the target content node upwards and applying the security modifiers dynamically as the traverser goes up.


Example

Now, to get the resulting access rights for e.g. "user 1" on the "My File.pdf" in a Top-Down approach on the model in the graph above would go like:

1. Traveling upward, we start with "Root folder", and set the permissions to 11 initially (only considering Read, Write).
2. There are two SECURITY relationships to that folder. User 1 is contained in both of them, but "root" is more generic, so apply it first then "All principals" +W +R → 11.
3. "Home" has no SECURITY instructions, continue.
4. "user1 Home" has SECURITY. First apply "Regular Users" (-R -W) → 00, Then "user 1" (+R +W) → 11.
5. The target node "My File.pdf" has no SECURITY modifiers on it, so the effective permissions for "User 1" on "My File.pdf" are ReadWrite → 11.

7.2.2. Read-permission example

In this example, we are going to examine a tree structure of directories and files. Also, there are users that own files and roles that can be assigned to users. Roles can have permissions on directory or files structures (here we model only canRead, as opposed to full rwx Unix permissions) and be nested. A more thorough example of modeling ACL structures can be found at [How to Build Role-Based Access Control in SQL](http://www.xaprb.com/blog/2006/08/16/how-to-build-role-based-access-control-in-sql/) <http://www.xaprb.com/blog/2006/08/16/how-to-build-role-based-access-control-in-sql/>.



Find all files in the directory structure

In order to find all files contained in this structure, we need a variable length query that follows all contains relationships and retrieves the nodes at the other end of the leaf relationships.

```
START root=node:node_auto_index(name = 'FileRoot')
MATCH root-[:contains*0..]->(parentDir)-[:leaf]->file
RETURN file
```

resulting in:

file
Node[11]{name:"File1"}
Node[10]{name:"File2"}
2 rows

What files are owned by whom?

If we introduce the concept of ownership on files, we then can ask for the owners of the files we find — connected via owns relationships to file nodes.

```
START root=node:node_auto_index(name = 'FileRoot')
MATCH root-[:contains*0..]->()-[:leaf]->file<-[:owns]-user
RETURN file, user
```

Returning the owners of all files below the FileRoot node.

file	user
Node[11]{name:"File1"}	Node[8]{name:"User1"}
2 rows	

file	user
Node[10]{name:"File2"}	Node[7]{name:"User2"}
2 rows	

Who has access to a File?

If we now want to check what users have read access to all Files, and define our ACL as

- The root directory has no access granted.
- Any user having a role that has been granted `canRead` access to one of the parent folders of a File has read access.

In order to find users that can read any part of the parent folder hierarchy above the files, Cypher provides optional variable length path.

```
START file=node:node_auto_index('name:File*')
MATCH file<-[:leaf]-()-[:contains*0..]-dir<-[:canRead]-role-[:member]->readUser
RETURN file.name, dir.name, role.name, readUser.name
```

This will return the `file`, and the directory where the user has the `canRead` permission along with the user and their role.

file.name	dir.name	role.name	readUser.name
"File2"	"Desktop"	<null>	<null>
"File2"	"HomeU2"	<null>	<null>
"File2"	"Home"	<null>	<null>
"File2"	"FileRoot"	"SUDOers"	"Admin1"
"File2"	"FileRoot"	"SUDOers"	"Admin2"
"File1"	"HomeU1"	<null>	<null>
"File1"	"Home"	<null>	<null>
"File1"	"FileRoot"	"SUDOers"	"Admin1"
"File1"	"FileRoot"	"SUDOers"	"Admin2"
9 rows			

The results listed above contain `null` values for optional path segments, which can be mitigated by either asking several queries or returning just the really needed values.


7.3. Linked Lists

A powerful feature of using a graph database, is that you can create your own in-graph data structures — like a linked list.

This datastructure uses a single node as the list reference. The reference has an outgoing relationship to the head of the list, and an incoming relationship from the last element of the list. If the list is empty, the reference will point to it self.

Something like this:

Figure 7.1. Graph



To initialize an empty linked list, we simply create an empty node, and make it link to itself.

Query.

```
CREATE root-[:LINK]->root // no 'value' property assigned to root
RETURN root
```

Adding values is done by finding the relationship where the new value should be placed in, and replacing it with a new node, and two relationships to it.

Query.

```
START root=node:node_auto_index(name = "ROOT")
MATCH root-[:LINK*0..]->before, // before could be same as root
after-[:LINK*0..]->root, // after could be same as root
before-[old:LINK]->after
WHERE before.value? < 25 // This is the value, which would normally
AND 25 < after.value? // be supplied through a parameter.
CREATE before-[:LINK]->({value:25})-[:LINK]->after
DELETE old
```

Deleting a value, conversely, is done by finding the node with the value, and the two relationships going in and out from it, and replacing with a new value.


Query.

```
START root=node:node_auto_index(name = "ROOT")
MATCH root-[:LINK*0..]->before,
before-[delBefore:LINK]->del-[delAfter:LINK]->after,
after-[:LINK*0..]->root
WHERE del.value! = 10
CREATE before-[:LINK]->after
DELETE del, delBefore, delAfter
```

7.4. Hyperedges

Imagine a user being part of different groups. A group can have different roles, and a user can be part of different groups. He also can have different roles in different groups apart from the membership. The association of a User, a Group and a Role can be referred to as a *HyperEdge*. However, it can be easily modeled in a property graph as a node that captures this n-ary relationship, as depicted below in the U1G2R1 node.

Figure 7.2. Graph



7.4.1. Find Groups

To find out in what roles a user is for a particular groups (here *Group2*), the following query can traverse this HyperEdge node and provide answers.

Query.

```
START n=node:node_auto_index(name = 'User1')
MATCH n-[:hasRoleInGroup]->hyperEdge-[:hasGroup]->group, hyperEdge-[:hasRole]->role
WHERE group.name = "Group2"
RETURN role.name
```

The role of User1 is returned:

Result

role.name
"Role1"
1 row

7.4.2. Find all groups and roles for a user

Here, find all groups and the roles a user has, sorted by the name of the role.

Query.

```
START n=node:node_auto_index(name = "User1")
MATCH n-[:hasRoleInGroup]->hyperEdge-[:hasGroup]->group, hyperEdge-[:hasRole]->role
RETURN role.name, group.name
ORDER BY role.name asc
```

The groups and roles of User1 are returned:

Result

role.name	group.name
"Role1"	"Group2"
"Role2"	"Group1"
2 rows	


7.4.3. Find common groups based on shared roles

Assume a more complicated graph:

1. Two user nodes User1, User2.
2. User1 is in Group1, Group2, Group3.
3. User1 has Role1, Role2 in Group1; Role2, Role3 in Group2; Role3, Role4 in Group3 (hyper edges).
4. User2 is in Group1, Group2, Group3.
5. User2 has Role2, Role5 in Group1; Role3, Role4 in Group2; Role5, Role6 in Group3 (hyper edges).

The graph for this looks like the following (nodes like U1G2R23 representing the HyperEdges):

Figure 7.3. Graph



To return Group1 and Group2 as User1 and User2 share at least one common role in these two groups, the query looks like this:

Query.

```
START u1=node:node_auto_index(name = 'User1'), u2=node:node_auto_index(name = 'User2')
MATCH u1-[:hasRoleInGroup]->hyperEdge1-[:hasGroup]->group,
      hyperEdge1-[:hasRole]->role,
      u2-[:hasRoleInGroup]->hyperEdge2-[:hasGroup]->group,
```

```
hyperEdge2-[:hasRole]->role
RETURN group.name, count(role)
ORDER BY group.name ASC
```

The groups where User1 and User2 share at least one common role:


Result

group.name	count(role)
"Group1"	1
"Group2"	1
2 rows	

7.5. Basic friend finding based on social neighborhood

Imagine an example graph like the following one:

Figure 7.4. Graph



To find out the friends of Joe's friends that are not already his friends, the query looks like this:

Query.

```

START joe=node:node_auto_index(name = "Joe")
MATCH joe-[:knows*2..2]-friend_of_friend
WHERE not(joe-[:knows]-friend_of_friend)
RETURN friend_of_friend.name, COUNT(*)
ORDER BY COUNT(*) DESC, friend_of_friend.name
  
```


This returns a list of friends-of-friends ordered by the number of connections to them, and secondly by their name.

Result

friend_of_friend.name	COUNT(*)
"Ian"	2
"Derrick"	1
"Jill"	1
3 rows	

7.6. Co-favorited places

Figure 7.5. Graph



7.6.1. Co-favorited places — users who like x also like y

Find places that people also like who favorite this place:

- Determine who has favorited place x.
- What else have they favorited that is not place x.

Query.

```
START place=node:node_auto_index(name = "CoffeeShop1")
MATCH place<-[:favorite]-person-[:favorite]->stuff
RETURN stuff.name, count(*)
ORDER BY count(*) DESC, stuff.name
```

The list of places that are favorited by people that favorited the start place.

Result

stuff.name	count(*)
"MelsPlace"	2
"CoffeShop2"	1
"SaunaX"	1
3 rows	

7.6.2. Co-Tagged places — places related through tags

Find places that are tagged with the same tags:

- Determine the tags for place x.
- What else is tagged the same as x that is not x.

Query.

```
START place=node:node_auto_index(name = "CoffeeShop1")
MATCH place-[:tagged]->tag<-[:tagged]-otherPlace
RETURN otherPlace.name, collect(tag.name)
ORDER BY length(collect(tag.name)) DESC, otherPlace.name
```


This query returns other places than CoffeeShop1 which share the same tags; they are ranked by the number of tags.

Result

otherPlace.name	collect(tag.name)
"MelsPlace"	["Cool", "Cosy"]
"CoffeeShop2"	["Cool"]
"CoffeeShop3"	["Cosy"]
3 rows	

7.7. Find people based on similar favorites

Figure 7.6. Graph



To find out the possible new friends based on them liking similar things as the asking person, use a query like this:

Query.

```

START me=node:node_auto_index(name = "Joe")
MATCH me-[:favorite]->stuff<-[:favorite]-person
WHERE NOT(me-[:friend]-person)
RETURN person.name, count(stuff)
ORDER BY count(stuff) DESC
  
```


The list of possible friends ranked by them liking similar stuff that are not yet friends is returned.

Result

person.name	count(stuff)
"Derrick"	2
"Jill"	1
2 rows	

7.8. Find people based on mutual friends and groups

Figure 7.7. Graph



In this scenario, the problem is to determine mutual friends and groups, if any, between persons. If no mutual groups or friends are found, there should be a \emptyset returned.

Query.

```

START me=node(5), other=node(4, 3)
MATCH pGroups=me-[?:member_of_group]->mg<-[?:member_of_group]-other,
pMutualFriends=me-[?:knows]->mf<-[?:knows]-other
RETURN other.name as name,
count(distinct pGroups) AS mutualGroups,
count(distinct pMutualFriends) AS mutualFriends
ORDER BY mutualFriends DESC
  
```


The question we are asking is — how many unique paths are there between me and Jill, the paths being common group memberships, and common friends. If the paths are mandatory, no results will be returned if me and Bob lack any common friends, and we don't want that. To make a path optional, you have to make at least one of it's relationships optional. That makes the whole path optional.

Result

name	mutualGroups	mutualFriends
"Jill"	1	1
"Bob"	1	\emptyset
2 rows		

7.9. Find friends based on similar tagging

Figure 7.8. Graph



To find people similar to me based on the taggings of their favorited items, one approach could be:

- Determine the tags associated with what I favorite.
- What else is tagged with those tags?
- Who favorites items tagged with the same tags?
- Sort the result by how many of the same things these people like.

Query.

```

START me=node:node_auto_index(name = "Joe")
MATCH me-[:favorite]->myFavorites-[:tagged]->tag<-[:tagged]-theirFavorites<-[:favorite]-people
WHERE NOT(me=people)
RETURN people.name as name, count(*) as similar_favs
ORDER BY similar_favs DESC
  
```


The query returns the list of possible friends ranked by them liking similar stuff that are not yet friends.

Result

name	similar_favs
"Sara"	2
"Derrick"	1
2 rows	

7.10. Multirelational (social) graphs

Figure 7.9. Graph



This example shows a multi-relational network between persons and things they like. A multi-relational graph is a graph with more than one kind of relationship between nodes.

Query.

```

START me=node:node_auto_index(name = 'Joe')
MATCH me-[r1:FOLLOWS|LOVES]->other-[r2]->me
WHERE type(r1)=type(r2)
RETURN other.name, type(r1)


```

The query returns people that FOLLOWS or LOVES Joe back.

Result

other.name	type(r1)
"Sara"	"FOLLOWS"
"Maria"	"FOLLOWS"
"Maria"	"LOVES"
3 rows	

7.11. Implementing newsfeeds in a graph



Implementation of newsfeed or timeline feature is a frequent requirement for social applications. The following examples are inspired by [Newsfeed feature powered by Neo4j Graph Database](http://techfin.in/2012/10/newsfeed-feature-powered-by-neo4j-graph-database/) <http://techfin.in/2012/10/newsfeed-feature-powered-by-neo4j-graph-database/>. The query asked here is:

Starting at me, retrieve the time-ordered status feed of the status updates of me and all friends that are connected via a CONFIRMED FRIEND relationship to me.

Query.

```

START me=node:node_auto_index(name='Joe')
MATCH me-[rels:FRIEND*0..1]-myfriend
WHERE ALL(r in rels
WHERE r.status = 'CONFIRMED')
WITH myfriend
  
```

```
MATCH myfriend-[:STATUS]-latestupdate-[:NEXT*0..1]-statusupdates
RETURN myfriend.name as name, statusupdates.date as date, statusupdates.text as text
ORDER BY statusupdates.date DESC
LIMIT 3
```

To understand the strategy, let's divide the query into five steps:

1. First Get the list of all my friends (along with me) through FRIEND relationship (MATCH me-[rels:FRIEND*0..1]-myfriend). Also, the WHERE predicate can be added to check whether the friend request is pending or confirmed.
2. Get the latest status update of my friends through Status relationship (MATCH myfriend-[:STATUS]-latestupdate).
3. Get subsequent status updates (along with the latest one) of my friends through NEXT relationships (MATCH myfriend-[:STATUS]-latestupdate-[:NEXT*0..1]-statusupdates) which will give you the latest and one additional statusupdate, adjust 0..1 to whatever suits your case.
4. Sort the status updates by posted date (ORDER BY statusupdates.date DESC).
5. LIMIT the number of updates you need in every query (LIMIT x SKIP x*y).

Result

name	date	text
"Joe"	6	"Joe status2"
"Bob"	4	"bobs status2"
"Joe"	3	"Joe status1"
3 rows		

Here, the example shows how to add a new status update into the existing data for a user.

Query.


```
START me=node:node_auto_index(name='Bob')
MATCH me-[r?:STATUS]-secondlatestupdate
DELETE r
WITH me, secondlatestupdate
CREATE me-[:STATUS]->(latest_update{text:'Status',date:123})
WITH latest_update,secondlatestupdate
CREATE latest_update-[:NEXT]-secondlatestupdate
WHERE secondlatestupdate <> null
RETURN latest_update.text as new_status
```

Dividing the query into steps, this query resembles adding new item in middle of a doubly linked list:

1. Get the latest update (if it exists) of the user through the STATUS relationship (MATCH me-[r?:STATUS]-secondlatestupdate).
2. Delete the STATUS relationship between user and secondlatestupdate (if it exists), as this would become the second latest update now and only the latest update would be added through a STATUS relationship, all earlier updates would be connected to their subsequent updates through a NEXT relationship. (DELETE r).
3. Now, create the new statusupdate node (with text and date as properties) and connect this with the user through a STATUS relationship (CREATE me-[:STATUS]->(latest_update{text:'Status',date:123})).
4. Now, create a NEXT relationship between the latest status update and the second latest status update (if it exists) (CREATE latest_update-[:NEXT]-secondlatestupdate WHERE secondlatestupdate <> null).


Result

new_status
"Status"
1 row
Nodes created: 1
Relationships created: 2
Properties set: 2
Relationships deleted: 1



7.12. Boosting recommendation results

Figure 7.10. Graph



This query finds the recommended friends for the origin that are working at the same place as the origin, or know a person that the origin knows, also, the origin should not already know the target. This recommendation is weighted for the weight of the relationship r_2 , and boosted with a factor of 2, if there is an activity-property on that relationship

Query.

```
START origin=node:node_auto_index(name = "Clark Kent")
MATCH origin-[r1:KNOWS|WORKSAT]-(c)-[r2:KNOWS|WORKSAT]-candidate
WHERE type(r1)=type(r2) AND (NOT (origin-[:KNOWS]-candidate))
RETURN origin.name as origin, candidate.name as candidate,
SUM(ROUND(r2.weight + (COALESCE(r2.activity?, 0) * 2))) as boost
ORDER BY boost desc
LIMIT 10
```


This returns the recommended friends for the origin nodes and their recommendation score.

Result

origin	candidate	boost
"Clark Kent"	"Perry White"	22
"Clark Kent"	"Anderson Cooper"	4
2 rows		

7.13. Calculating the clustering coefficient of a network

Figure 7.11. Graph



In this example, adapted from [Niko Gamulins blog post on Neo4j for Social Network Analysis](http://mypetprojects.blogspot.se/2012/06/social-network-analysis-with-neo4j.html) <<http://mypetprojects.blogspot.se/2012/06/social-network-analysis-with-neo4j.html>>, the graph in question is showing the 2-hop relationships of a sample person as nodes with KNOWS relationships.

The **clustering coefficient** <http://en.wikipedia.org/wiki/Clustering_coefficient> of a selected node is defined as the probability that two randomly selected neighbors are connected to each other. With the number of neighbors as n and the number of mutual connections between the neighbors r the calculation is:

The number of possible connections between two neighbors is $n! / (2!(n-2)!) = 4! / (2!(4-2)!) = 24/4 = 6$, where n is the number of neighbors $n = 4$ and the actual number r of connections is 1. Therefore the clustering coefficient of node 1 is $1/6$.

n and r are quite simple to retrieve via the following query:

Query.

```

START a = node(*)
MATCH (a)--(b)
WITH a, count(distinct b) as n
MATCH (a)--()-[r]-(a)
WHERE a.name! = "startnode"
RETURN n, count(distinct r) as r
  
```

This returns n and r for the above calculations.

Result

n	r
4	1
1 row	

7.14. Pretty graphs

This section is showing how to create some of the [named pretty graphs on Wikipedia](http://en.wikipedia.org/wiki/Gallery_of_named_graphs) <http://en.wikipedia.org/wiki/Gallery_of_named_graphs>.

7.14.1. Star graph

The graph is created by first creating a center node, and then once per element in the range, creates a leaf node and connects it to the center.

Query.


```
CREATE center
foreach( x in range(1,6) :
CREATE leaf, center-[:X]->leaf
)
RETURN id(center) as id;
```

The query returns the id of the center node.

Result

id
8
1 row
Nodes created: 7
Relationships created: 6

Figure 7.12. Graph



7.14.2. Wheel graph

This graph is created in a number of steps:

- Create a center node.
- Once per element in the range, create a leaf and connect it to the center.
- Select 2 leafs from the center node and connect them.
- Find the minimum and maximum leaf and connect these.

- Return the id of the center node.

Query.

```

CREATE center
foreach( x in range(1,6) :
CREATE leaf={count:x}, center-[:X]->leaf
)
WITH center
MATCH large_leaf<--center-->small_leaf
WHERE large_leaf.count = small_leaf.count + 1
CREATE small_leaf-[:X]->large_leaf

WITH center, min(small_leaf.count) as min, max(large_leaf.count) as max
MATCH first_leaf<--center-->last_leaf
WHERE first_leaf.count = min AND last_leaf.count = max
CREATE last_leaf-[:X]->first_leaf


RETURN id(center) as id
    
```

The query returns the id of the center node.

Result

id
8
1 row
Nodes created: 7
Relationships created: 12
Properties set: 6

Figure 7.13. Graph



7.14.3. Complete graph

For this graph, a root node is created, and used to hang a number of nodes from. Then, two nodes are selected, hanging from the center, with the requirement that the id of the first is less than the id of the next. This is to prevent double relationships and self relationships. Using said match, relationships between all these nodes are created. Lastly, the center node and all relationships connected to it are removed.

Query.


```
CREATE center
foreach( x in range(1,6) :
CREATE leaf={count : x}, center-[:X]->leaf
)
==== center ====
MATCH leaf1<--center-->leaf2
WHERE id(leaf1)<id(leaf2)
CREATE leaf1-[:X]->leaf2
==== center ====
MATCH center-[r]->( )
DELETE center,r;
```

Nothing is returned by this query.

Result

(empty result)
Nodes created: 7
Relationships created: 21
Properties set: 6
Nodes deleted: 1
Relationships deleted: 6

Figure 7.14. Graph



7.14.4. Friendship graph

This query first creates a center node, and then once per element in the range, creates a cycle graph and connects it to the center

Query.

```
CREATE center
```


```
foreach( x in range(1,3) :
CREATE leaf1, leaf2, center-[:X]->leaf1, center-[:X]->leaf2, leaf1-[:X]->leaf2
)
RETURN ID(center) as id
```

The id of the center node is returned by the query.

Result

id
8
1 row
Nodes created: 7
Relationships created: 9

Figure 7.15. Graph



7.15. A multilevel indexing structure (path tree)


In this example, a multi-level tree structure is used to index event nodes (here Event1, Event2 and Event3, in this case with a YEAR-MONTH-DAY granularity, making this a timeline indexing structure. However, this approach should work for a wide range of multi-level ranges.

The structure follows a couple of rules:

- Events can be indexed multiple times by connecting the indexing structure leafs with the events via a VALUE relationship.
- The querying is done in a path-range fashion. That is, the start- and end path from the indexing root to the start and end leafs in the tree are calculated
- Using Cypher, the queries following different strategies can be expressed as path sections and put together using one single query.

The graph below depicts a structure with 3 Events being attached to an index structure at different leafs.


Figure 7.16. Graph



7.15.1. Return zero range

Here, only the events indexed under one leaf (2010-12-31) are returned. The query only needs one path segment rootPath (color Green) through the index.

Figure 7.17. Graph



Query.

```

START root=node:node_auto_index(name = 'Root')
MATCH rootPath=root-[:`2010`]->()-[:`12`]->()-[:`31`]->leaf, leaf-[:VALUE]->event
RETURN event.name
ORDER BY event.name ASC
    
```

Returning all events on the date 2010-12-31, in this case Event1 and Event2


Result

event.name
"Event1"
"Event2"
2 rows

7.15.2. Return the full range

In this case, the range goes from the first to the last leaf of the index tree. Here, startPath (color Greenyellow) and endPath (color Green) span up the range, valuePath (color Blue) is then connecting the leafs, and the values can be read from the middle node, hanging off the values (color Red) path.

Figure 7.18. Graph



Query.

```

START root=node:node_auto_index(name = 'Root')
MATCH startPath=root-[:'2010']->()-[:'12']->()-[:'31']->startLeaf,
      endPath=root-[:'2011']->()-[:'01']->()-[:'03']->endLeaf,
      valuePath=startLeaf-[:NEXT*0..]->middle-[:NEXT*0..]->endLeaf,
      values=middle-[:VALUE]->event
RETURN event.name
ORDER BY event.name ASC
    
```

Returning all events between 2010-12-31 and 2011-01-03, in this case all events.


Result

event.name
"Event1"
"Event2"
"Event2"
"Event3"
4 rows

7.15.3. Return partly shared path ranges

Here, the query range results in partly shared paths when querying the index, making the introduction of and common path segment `commonPath` (color Black) necessary, before spanning up `startPath` (color Greenyellow) and `endPath` (color Darkgreen). After that, `valuePath` (color Blue) connects the leafs and the indexed values are returned off `values` (color Red) path.

Figure 7.19. Graph



Query.

```

START root=node:node_auto_index(name = 'Root')
MATCH commonPath=root-[:`2011`]->()-[:`01`]->commonRootEnd,
      startPath=commonRootEnd-[:`01`]->startLeaf,
      endPath=commonRootEnd-[:`03`]->endLeaf,
      valuePath=startLeaf-[:NEXT*0..]->middle-[:NEXT*0..]->endLeaf,
      values=middle-[:VALUE]->event
RETURN event.name
ORDER BY event.name ASC
  
```

Returning all events between 2011-01-01 and 2011-01-03, in this case Event2 and Event3.

Result

event.name
"Event2"
2 rows

event.name
"Event3"
2 rows

7.16. Complex similarity computations

7.16.1. Calculate similarities by complex calculations

Here, a similarity between two players in a game is calculated by the number of times they have eaten the same food.

Query.


```
START me=node:node_auto_index(name = "me")
MATCH me-[r1:ATE]->food<-[r2:ATE]-you
==== me,count(distinct r1) as H1,count(distinct r2) as H2,you ====
MATCH me-[r1:ATE]->food<-[r2:ATE]-you
RETURN sum((1-ABS(r1.times/H1-r2.times/H2))*(r1.times+r2.times)/(H1+H2)) as similarity
```

The two players and their similarity measure.

Result

similarity
-30.0
1 row

Figure 7.20. Graph



7.17. The Graphity activity stream model

7.17.1. Find Activity Streams in a network without scaling penalty

This is an approach for scaling the retrieval of activity streams in a friend graph put forward by Rene Pickard as [Graphity](http://www.rene-pickhardt.de/graphity-an-efficient-graph-model-for-retrieving-the-top-k-news-feeds-for-users-in-social-networks/) <http://www.rene-pickhardt.de/graphity-an-efficient-graph-model-for-retrieving-the-top-k-news-feeds-for-users-in-social-networks/>. In short, a linked list is created for every persons friends in the order that the last activities of these friends have occurred. When new activities occur for a friend, all the ordered friend lists that this friend is part of are reordered, transferring computing load to the time of new event updates instead of activity stream reads.



Tip

This approach of course makes excessive use of relationship types. This needs to be taken into consideration when designing a production system with this approach. See [Section 16.5, “Capacity”](#) for the maximum number of relationship types.

To find the activity stream for a person, just follow the linked list of the friend list, and retrieve the needed amount of activities form the respective activity list of the friends.

Query.


```
START me=node:node_auto_index(name = "Jane")
MATCH p=me-[:jane_knows*]->friend, friend-[:has]->status
RETURN me.name, friend.name, status.name, length(p)
ORDER BY length(p)
```

The returns the activity stream for Jane.

Result

me.name	friend.name	status.name	length(p)
"Jane"	"Bill"	"Bill_s1"	1
"Jane"	"Joe"	"Joe_s1"	2
"Jane"	"Bob"	"Bob_s1"	3
3 rows			

Figure 7.21. Graph



Chapter 8. Languages

Please see <http://www.neo4j.org/drivers> for the current set of drivers!

Chapter 9. Extending the Neo4j Server

The Neo4j Server can be extended by either plugins or unmanaged extensions. For more information on the server, see [Chapter 21, *Neo4j Server*](#).

9.1. Server Plugins

Quick info

- The server's functionality can be extended by adding plugins.
- Plugins are user-specified code which extend the capabilities of the database, nodes, or relationships.
- The neo4j server will then advertise the plugin functionality within representations as clients interact via HTTP.

Plugins provide an easy way to extend the Neo4j REST API with new functionality, without the need to invent your own API. Think of plugins as server-side scripts that can add functions for retrieving and manipulating nodes, relationships, paths, properties or indices.



Tip

If you want to have full control over your API, and are willing to put in the effort, and understand the risks, then Neo4j server also provides hooks for [unmanaged extensions](#) based on JAX-RS.

The needed classes reside in the [org.neo4j:server-api](http://search.maven.org/#search|gav|1lg%3A%22org.neo4j%22%20AND%20a%3A%22server-api%22) <<http://search.maven.org/#search|gav|1lg%3A%22org.neo4j%22%20AND%20a%3A%22server-api%22>> jar file. See the linked page for downloads and instructions on how to include it using dependency management. For Maven projects, add the Server API dependencies in your pom.xml like this:

```
<dependency>
  <groupId>org.neo4j</groupId>
  <artifactId>server-api</artifactId>
  <version>${neo4j-version}</version>
</dependency>
```

Where *\${neo4j-version}* is the intended version.

To create a plugin, your code must inherit from the [ServerPlugin](http://components.neo4j.org/server-api/1.9.7/apidocs/org/neo4j/server/plugins/ServerPlugin.html) <<http://components.neo4j.org/server-api/1.9.7/apidocs/org/neo4j/server/plugins/ServerPlugin.html>> class. Your plugin should also:

- ensure that it can produce an (Iterable of) Node, Relationship or Path, any Java primitive or String or an instance of a `org.neo4j.server.rest.repr.Representation`
- specify parameters,
- specify a point of extension and of course
- contain the application logic.
- make sure that the discovery point type in the `@PluginTarget` and the `@Source` parameter are of the same type.

An example of a plugin which augments the database (as opposed to nodes or relationships) follows:

Get all nodes or relationships plugin.

```
@Description( "An extension to the Neo4j Server for getting all nodes or relationships" )
public class GetAll extends ServerPlugin
{
  @Name( "get_all_nodes" )
  @Description( "Get all nodes from the Neo4j graph database" )
  @PluginTarget( GraphDatabaseService.class )
  public Iterable<Node> getAllNodes( @Source GraphDatabaseService graphDb )
  {
```

```

        return GlobalGraphOperations.at( graphDb ).getAllNodes();
    }

    @Description( "Get all relationships from the Neo4j graph database" )
    @PluginTarget( GraphDatabaseService.class )
    public Iterable<Relationship> getAllRelationships( @Source GraphDatabaseService graphDb )
    {
        return GlobalGraphOperations.at( graphDb ).getAllRelationships();
    }
}

```

The full source code is found here: [GetAll.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/plugins/GetAll.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/plugins/GetAll.java>

Find the shortest path between two nodes plugin.

```

public class ShortestPath extends ServerPlugin
{
    @Description( "Find the shortest path between two nodes." )
    @PluginTarget( Node.class )
    public Iterable<Path> shortestPath(
        @Source Node source,
        @Description( "The node to find the shortest path to." )
        @Parameter( name = "target" ) Node target,
        @Description( "The relationship types to follow when searching for the shortest path(s). " +
            "Order is insignificant, if omitted all types are followed." )
        @Parameter( name = "types", optional = true ) String[] types,
        @Description( "The maximum path length to search for, default value (if omitted) is 4." )
        @Parameter( name = "depth", optional = true ) Integer depth )
    {
        Expander expander;
        if ( types == null )
        {
            expander = Traversal.expanderForAllTypes();
        }
        else
        {
            expander = Traversal.emptyExpander();
            for ( int i = 0; i < types.length; i++ )
            {
                expander = expander.add( DynamicRelationshipType.withName( types[i] ) );
            }
        }
        PathFinder<Path> shortestPath = GraphAlgoFactory.shortestPath(
            expander, depth == null ? 4 : depth.intValue() );
        return shortestPath.findAllPaths( source, target );
    }
}

```

The full source code is found here: [ShortestPath.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/plugins/ShortestPath.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/plugins/ShortestPath.java>

To deploy the code, simply compile it into a .jar file and place it onto the server classpath (which by convention is the plugins directory under the Neo4j server home directory).



Tip

Make sure the directories listings are retained in the jarfile by either building with default Maven, or with `jar -cvf myext.jar *`, making sure to jar directories instead of specifying single files.

The .jar file must include the file `META-INF/services/org.neo4j.server.plugins.ServerPlugin` with the fully qualified name of the implementation class. This is an example with multiple entries, each on a separate line:

```
org.neo4j.examples.server.plugins.GetAll
org.neo4j.examples.server.plugins.DepthTwo
org.neo4j.examples.server.plugins.ShortestPath
```

The code above makes an extension visible in the database representation (via the `@PluginTarget` annotation) whenever it is served from the Neo4j Server. Simply changing the `@PluginTarget` parameter to `Node.class` or `Relationship.class` allows us to target those parts of the data model should we wish. The functionality extensions provided by the plugin are automatically advertised in representations on the wire. For example, clients can discover the extension implemented by the above plugin easily by examining the representations they receive as responses from the server, e.g. by performing a GET on the default database URI:

```
curl -v http://localhost:7474/db/data/
```

The response to the GET request will contain (by default) a JSON container that itself contains a container called "extensions" where the available plugins are listed. In the following case, we only have the `GetAll` plugin registered with the server, so only its extension functionality is available. Extension names will be automatically assigned, based on method names, if not specifically specified using the `@Name` annotation.

```
{
  "extensions-info" : "http://localhost:7474/db/data/ext",
  "node" : "http://localhost:7474/db/data/node",
  "node_index" : "http://localhost:7474/db/data/index/node",
  "relationship_index" : "http://localhost:7474/db/data/index/relationship",
  "reference_node" : "http://localhost:7474/db/data/node/0",
  "extensions_info" : "http://localhost:7474/db/data/ext",
  "extensions" : {
    "GetAll" : {
      "get_all_nodes" : "http://localhost:7474/db/data/ext/GetAll/graphdb/get_all_nodes",
      "get_all_relationships" : "http://localhost:7474/db/data/ext/GetAll/graphdb/getAllRelationships"
    }
  }
}
```

Performing a GET on one of the two extension URIs gives back the meta information about the service:

```
curl http://localhost:7474/db/data/ext/GetAll/graphdb/get_all_nodes
```

```
{
  "extends" : "graphdb",
  "description" : "Get all nodes from the Neo4j graph database",
  "name" : "get_all_nodes",
  "parameters" : [ ]
}
```

To use it, just POST to this URL, with parameters as specified in the description and encoded as JSON data content. F.ex for calling the shortest path extension (URI gotten from a GET to <http://localhost:7474/db/data/node/123>):

```
curl -X POST http://localhost:7474/db/data/ext/GetAll/node/123/shortestPath \
-H "Content-Type: application/json" \
-d '{"target":"http://localhost:7474/db/data/node/456&depth=5"}'
```

If everything is OK a response code `200` and a list of zero or more items will be returned. If nothing is returned (null returned from extension) an empty result and response code `204` will be returned. If the extension throws an exception response code `500` and a detailed error message is returned.

Extensions that do any kind of write operation will have to manage their own transactions, i.e. transactions aren't managed automatically.

Through this model, any plugin can naturally fit into the general hypermedia scheme that Neo4j espouses — meaning that clients can still take advantage of abstractions like Nodes, Relationships

and Paths with a straightforward upgrade path as servers are enriched with plugins (old clients don't break).

9.2. Unmanaged Extensions

Quick info

- **Danger: Men at Work!** The unmanaged extensions are a way of deploying arbitrary JAX-RS code into the Neo4j server.
- The unmanaged extensions are exactly that: unmanaged. If you drop poorly tested code into the server, it's highly likely you'll degrade its performance, so be careful.

Some projects want extremely fine control over their server-side code. For this we've introduced an unmanaged extension API.



Warning

This is a sharp tool, allowing users to deploy arbitrary [JAX-RS](http://en.wikipedia.org/wiki/JAX-RS) classes to the server and so you should be careful when thinking about using this. In particular you should understand that it's easy to consume lots of heap space on the server and hinder performance if you're not careful.

Still, if you understand the disclaimer, then you load your JAX-RS classes into the Neo4j server simply by adding adding a `@Context` annotation to your code, compiling against the JAX-RS jar and any Neo4j jars you're making use of. Then add your classes to the runtime classpath (just drop it in the lib directory of the Neo4j server). In return you get access to the hosted environment of the Neo4j server like logging through the `org.neo4j.server.logging.Logger`.

In your code, you get access to the underlying `GraphDatabaseService` through the `@Context` annotation like so:

```
public MyCoolService( @Context GraphDatabaseService database )
{
    // Have fun here, but be safe!
}
```

Remember, the unmanaged API is a very sharp tool. It's all too easy to compromise the server by deploying code this way, so think first and see if you can't use the managed extensions in preference. However, a number of context parameters can be automatically provided for you, like the reference to the database.

In order to specify the mount point of your extension, a full class looks like this:

Unmanaged extension example.

```
@Path( "/helloworld" )
public class HelloWorldResource
{
    private final GraphDatabaseService database;

    public HelloWorldResource( @Context GraphDatabaseService database )
    {
        this.database = database;
    }

    @GET
    @Produces( MediaType.TEXT_PLAIN )
    @Path(("/{nodeId}") )
    public Response hello( @PathParam( "nodeId" ) long nodeId )
    {
```

```
// Do stuff with the database
return Response.status( Status.OK ).entity(
    ("Hello World, nodeId=" + nodeId).getBytes( Charset.forName("UTF-8") ) ).build();
}
}
```

The full source code is found here: [HelloWorldResource.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/unmanaged/HelloWorldResource.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/server-examples/src/main/java/org/neo4j/examples/server/unmanaged/HelloWorldResource.java>

Build this code, and place the resulting jar file (and any custom dependencies) into the `$NEO4J_SERVER_HOME/plugins` directory, and include this class in the `neo4j-server.properties` file, like so:



Tip

Make sure the directories listings are retained in the jarfile by either building with default Maven, or with `jar -cvf myext.jar *`, making sure to jar directories instead of specifying single files.

```
#Comma separated list of JAXRS packages containing JAXRS Resource, one package name for each mountpoint.
org.neo4j.server.thirdparty_jaxrs_classes=org.neo4j.examples.server.unmanaged=/examples/unmanaged
```

Which binds the hello method to respond to GET requests at the URI: `http://{neo4j_server}:{neo4j_port}/examples/unmanaged/helloworld/{nodeId}`

```
curl http://localhost:7474/examples/unmanaged/helloworld/123
```

which results in

```
Hello World, nodeId=123
```

Part III. Cypher Query Language

The Cypher part is the authoritative source for details on the Cypher Query Language. For an introduction, see [Section 10.1, “What is Cypher?”](#).

Chapter 10. Introduction

To get an overview of Cypher, continue reading [Section 10.1, “What is Cypher?”](#). The rest of this chapter deals with the context of Cypher statements, like for example transaction management and how to use parameters. For the Cypher language reference itself see other chapters at [Part III, “Cypher Query Language”](#).

10.1. What is Cypher?

Cypher is a declarative graph query language that allows for expressive and efficient querying and updating of the graph store without having to write traversals through the graph structure in code. Cypher is still growing and maturing, and that means that there probably will be breaking syntax changes. It also means that it has not undergone the same rigorous performance testing as other Neo4j components.

Cypher is designed to be a humane query language, suitable for both developers and (importantly, we think) operations professionals who want to make ad-hoc queries on the database. Our guiding goal is to make the simple things simple, and the complex things possible. Its constructs are based on English prose and neat iconography, which helps to make it (somewhat) self-explanatory.

Cypher is inspired by a number of different approaches and builds upon established practices for expressive querying. Most of the keywords like `WHERE` and `ORDER BY` are inspired by [SQL](http://en.wikipedia.org/wiki/SQL) <<http://en.wikipedia.org/wiki/SQL>>. Pattern matching borrows expression approaches from [SPARQL](http://en.wikipedia.org/wiki/SPARQL) <<http://en.wikipedia.org/wiki/SPARQL>>.

Being a declarative language, Cypher focuses on the clarity of expressing *what* to retrieve from a graph, not *how* to do it, in contrast to imperative languages like Java, and scripting languages like [Gremlin](http://gremlin.tinkerpop.com) <<http://gremlin.tinkerpop.com>> (supported via the [Section 22.18, “Gremlin Plugin”](#)) and [the JRuby Neo4j bindings](http://neo4j.rubyforge.org/) <<http://neo4j.rubyforge.org/>>. This makes the concern query optimization an implementation detail which is not forced upon to the user.


The query language is comprised of several distinct clauses.

- `START`: Starting points in the graph, obtained via index lookups or by element IDs.
- `MATCH`: The graph pattern to match, bound to the starting points in `START`.
- `WHERE`: Filtering criteria.
- `RETURN`: What to return.
- `CREATE`: Creates nodes and relationships.
- `DELETE`: Removes nodes, relationships and properties.
- `SET`: Set values to properties.
- `FOREACH`: Performs updating actions once per element in a list.
- `WITH`: Divides a query into multiple, distinct parts.

Let's see three of them in action.

Imagine an example graph like the following one:

Figure 10.1. Example Graph



For example, here is a query which finds a user called John in an index and then traverses the graph looking for friends of John's friends (though not his direct friends) before returning both John and any friends-of-friends that are found.

```

START john=node:node_auto_index(name = 'John')
MATCH john-[:friend]->()-[:friend]->fof
RETURN john, fof
  
```

Resulting in:

john	fof
Node[4]{name: "John"}	Node[2]{name: "Maria"}
Node[4]{name: "John"}	Node[3]{name: "Steve"}
2 rows	

Next up we will add filtering to set more parts in motion:

In this next example, we take a list of users (by node ID) and traverse the graph looking for those other users that have an outgoing `friend` relationship, returning only those followed users who have a name property starting with `s`.

```

START user=node(5,4,1,2,3)
MATCH user-[:friend]->follower
WHERE follower.name =~ 'S.*'
RETURN user, follower.name
  
```

Resulting in:

user	follower.name
Node[5]{name: "Joe"}	"Steve"
Node[4]{name: "John"}	"Sara"
2 rows	

To use Cypher from Java, see [Section 4.10, "Execute Cypher Queries from Java"](#). For more Cypher examples, see [Chapter 7, Data Modeling Examples](#) as well.

10.2. Updating the graph

Cypher can be used for both querying and updating your graph.

10.2.1. The Structure of Updating Queries

Quick info

- A Cypher query part can't both match and update the graph at the same time.
- Every part can either read and match on the graph, or make updates on it.

If you read from the graph, and then update the graph, your query implicitly has two parts — the reading is the first part, and the writing is the second. If your query is read-only, Cypher will be lazy, and not actually pattern match until you ask for the results. Here, the semantics are that *all* the reading will be done before any writing actually happens. This is very important — without this it's easy to find cases where the pattern matcher runs into data that is being created by the very same query, and all bets are off. That road leads to Heisenbugs, Brownian motion and cats that are dead and alive at the same time.

First reading, and then writing, is the only pattern where the query parts are implicit — any other order and you have to be explicit about your query parts. The parts are separated using the `WITH` statement. `WITH` is like the event horizon — it's a barrier between a plan and the finished execution of that plan.

When you want to filter using aggregated data, you have to chain together two reading query parts — the first one does the aggregating, and the second query filters on the results coming from the first one.

```
START n=node(...)
MATCH n-[:friend]-friend
WITH n, count(friend) as friendsCount
WHERE friendsCount > 3
RETURN n, friendsCount
```

Using `WITH`, you specify how you want the aggregation to happen, and that the aggregation has to be finished before Cypher can start filtering.

You can chain together as many query parts as you have JVM heap for.

10.2.2. Returning data

Any query can return data. If your query only reads, it has to return data — it serves no purpose if it doesn't, and it is not a valid Cypher query. Queries that update the graph don't have to return anything, but they can.

After all the parts of the query comes one final `RETURN` statement. `RETURN` is not part of any query part — it is a period symbol after an eloquent statement. When `RETURN` is legal, it's also legal to use `SKIP/LIMIT` and `ORDER BY`.

If you return graph elements from a query that has just deleted them — beware, you are holding a pointer that is no longer valid. Operations on that node might fail mysteriously and unpredictably.

10.3. Transactions

Any query that updates the graph will run in a transaction. An updating query will always either fully succeed, or not succeed at all.

Cypher will either create a new transaction or run inside an existing one:

- If no transaction exists in the running context Cypher will create one and commit it once the query finishes.
- In case there already exists a transaction in the running context, the query will run inside it, and nothing will be persisted to disk until that transaction is successfully committed.

This can be used to have multiple queries be committed as a single transaction:

1. Open a transaction,
2. run multiple updating Cypher queries,
3. and commit all of them in one go.

Note that a query will hold the changes in memory until the whole query has finished executing. A large query will consequently need a JVM with lots of heap space.

All iterators returned from an execution result should be exhausted fully to ensure that resources bound to them will be properly closed. Resources include transactions started by the query, so failing to do so may, for example, lead to deadlocks or other weird behaviour.

10.4. Parameters

Cypher supports querying with parameters. This allows developers to not to have to do string building to create a query, and it also makes caching of execution plans much easier for Cypher.

Parameters can be used for literals and expressions in the `WHERE` clause, for the index key and index value in the `START` clause, index queries, and finally for node/relationship ids. Parameters can not be used as for property names, relationship types and labels, since since these patterns are part of the query structure that is compiled into a query plan.

Accepted names for parameter are letters and number, and any combination of these.

Here follows a few examples of how you can use parameters from Java.

Parameter for node id.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "id", 0 );
ExecutionResult result = engine.execute( "start n=node({id}) return n.name", params );
```

Parameter for node object.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "node", andreasNode );
ExecutionResult result = engine.execute( "start n=node({node}) return n.name", params );
```

Parameter for multiple node ids.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "id", Arrays.asList( 0, 1, 2 ) );
ExecutionResult result = engine.execute( "start n=node({id}) return n.name", params );
```

Parameter for string literal.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "name", "Johan" );
ExecutionResult result =
    engine.execute( "start n=node(0,1,2) where n.name = {name} return n", params );
```

Parameter for index key and value.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "key", "name" );
params.put( "value", "Michaela" );
ExecutionResult result =
    engine.execute( "start n=node:people({key} = {value}) return n", params );
```



Caution

Using a parameter for index keys has been deprecated.

Parameter for index query.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "query", "name:Andreas" );
ExecutionResult result = engine.execute( "start n=node:people({query}) return n", params );
```

Numeric parameters for `SKIP` and `LIMIT`.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "s", 1 );
params.put( "l", 1 );
ExecutionResult result =
    engine.execute( "start n=node(0,1,2) return n.name skip {s} limit {l}", params );
```

Parameter for regular expression.

```
Map<String, Object> params = new HashMap<String, Object>();
params.put( "regex", ".*h.*" );
ExecutionResult result =
    engine.execute( "start n=node(0,1,2) where n.name =~ {regex} return n.name", params );
```

Parameter setting properties on node.

```
Map<String, Object> n1 = new HashMap<String, Object>();
n1.put( "name", "Andres" );
n1.put( "position", "Developer" );

Map<String, Object> params = new HashMap<String, Object>();
params.put( "props", n1 );
engine.execute( "START n=node(0) SET n = {props}", params );
```

10.5. Compatibility

Cypher is still changing rather rapidly. Parts of the changes are internal — we add new pattern matchers, aggregators and other optimizations, which hopefully makes your queries run faster.

Other changes are directly visible to our users — the syntax is still changing. New concepts are being added and old ones changed to fit into new possibilities. To guard you from having to keep up with our syntax changes, Cypher allows you to use an older parser, but still gain the speed from new optimizations.

There are two ways you can select which parser to use. You can configure your database with the configuration parameter `cypher_parser_version`, and enter which parser you'd like to use (1.7, 1.8 and 1.9 are supported now). Any Cypher query that doesn't explicitly say anything else, will get the parser you have configured.

The other way is on a query by query basis. By simply pre-pending your query with "CYPHER 1.7", that particular query will be parsed with the 1.7 version of the parser. Example:

```
CYPHER 1.7 START n=node(0)
WHERE n.foo = "bar"
RETURN n
```

10.6. Query Performance

Cypher works very hard to execute queries as fast as possible.

However, when optimizing for maximum query execution performance, it may be helpful to rephrase queries using knowledge about the domain and the application.

The overall goal of manual query performance optimization is to ensure that only necessary data is retrieved from the graph. At least data should get filtered out as early as possible in order to reduce the amount of work that has to be done at later stages of query execution. This also goes for what gets returned: avoid returning whole nodes and relationships — instead, pick the data you need and return only that. You should also make sure to set an upper limit on variable length patterns, so they don't cover larger portions of the dataset than needed.

Each Cypher query gets optimized and transformed into an execution plan by the Cypher execution engine. To minimize the resources used for this, make sure to use parameters instead of literals when possible. This allows Cypher to re-use your queries instead of having to parse and build new execution plans.

Chapter 11. Syntax

The nitty-gritty details of Cypher syntax.

11.1. Operators

11.1.1. Mathematical operators

The mathematical operators are +, -, *, / and %, ^.

11.1.2. Comparison operators

The comparison operators are =, <>, <, >, <=, >=.

11.1.3. Boolean operators

The boolean operators are AND, `OR, NOT.

11.1.4. String operators

Strings can be concatenated using the + operator.

11.1.5. Collection operators

Collections can be concatenated using the + operator.

11.1.6. Property operators

Since Neo4j is a schema-free graph database, Cypher has two special operators: ? and !.

They are used on properties, and are used to deal with missing values. A comparison on a property that does not exist would normally cause an error. Instead of having to always check if the property exists before comparing its value with something else, the special property operators can be used.

The question mark makes the comparison always return true if the property is missing, and the exclamation mark makes the comparator return false.

This predicate will evaluate to true if n.prop is missing.

```
WHERE n.prop? = "foo"
```

This predicate will evaluate to false if n.prop is missing.

```
WHERE n.prop! = "foo"
```



Warning

Mixing the two in the same comparison will lead to unpredictable results.

This is really syntactic sugar that expands to this:

```
WHERE n.prop? = "foo" ⇒ WHERE (not(has(n.prop)) OR n.prop = "foo")
```

```
WHERE n.prop! = "foo" ⇒ WHERE (has(n.prop) AND n.prop = "foo")
```

11.2. Expressions

11.2.1. Expressions in general

An expression in Cypher can be:

- A numeric literal (integer or double): 13, 40000, 3.14.
- A string literal: "Hello", 'World'.
- A boolean literal: true, false, TRUE, FALSE.
- An identifier: n, x, rel, myFancyIdentifier, `A name with weird stuff in it[]!`.
- A property: n.prop, x.prop, rel.thisProperty, myFancyIdentifier.`(weird property name)`.
- A nullable property: it's a property, with a question mark or exclamation mark — n.prop?, rel.thisProperty!.
- A parameter: {param}, {0}
- A collection of expressions: ["a", "b"], [1,2,3], ["a", 2, n.property, {param}], [].
- A function call: length(p), nodes(p).
- An aggregate function: avg(x.prop), count(*).
- Relationship types: :REL_TYPE, :`REL TYPE`, :REL1|REL2.
- A path-pattern: a-->()<--b.
- A predicate expression is an expression that returns true or false: a.prop = "Hello", length(p) > 10, has(a.name)

11.2.2. Note on string literals

String literals can contain these escape sequences.

Escape sequence	Character
\t	Tab
\b	Backspace
\n	Newline
\r	Carriage return
\f	Form feed
\'	Single quote
\"	Double quote
\\	Backslash

11.3. Identifiers

When you reference parts of the pattern, you do so by naming them. The names you give the different parts are called identifiers.

In this example:

```
START n=node(1) MATCH n-->b RETURN b
```

The identifiers are `n` and `b`.

Identifier names are case sensitive, and can contain underscores and alphanumeric characters (a-z, 0-9), but must start with a letter. If other characters are needed, you can quote the identifier using backquote (```) signs.

The same rules apply to property names.

11.4. Comments

To add comments to your queries, use double slash. Examples:

```
START n=node(1) RETURN n //This is an end of line comment
```

```
START n=node(1)
//This is a whole line comment
RETURN n
```

```
START n=node(1) WHERE n.property = "//This is NOT a comment" RETURN n
```

11.5. Patterns

Patterns and pattern-matching are at the very heart of Cypher, and so it's very important to understand patterns, to be able to be effective with Cypher.

For instance, using patterns, you describe the shape of the data you're looking for in the `MATCH` clause. You describe the pattern, and Cypher will figure out how to get that data for you. The idea is for you to draw your query on a whiteboard, naming the interesting parts of the pattern, so you can then use values from these parts to create the result set you are looking for.

Patterns have bound points, or starting points. They are the parts of the pattern that are already "bound" to a set of graph nodes or relationships. All parts of the pattern must be directly or indirectly connected to a starting point — a pattern where parts of the pattern are not reachable from any starting point will be rejected.

Clause	Optional	Multiple rel. types	Varlength	Paths	Maps
Match	Yes	Yes	Yes	Yes	-
Create	-	-	-	Yes	Yes
Create Unique	-	-	-	Yes	Yes
Expressions	-	Yes	Yes	-	-

11.5.1. Patterns for related nodes

The description of the pattern is made up of one or more paths, separated by commas. A path is a sequence of nodes and relationships that always start and end in nodes. An example path would be:

```
(a)-->(b)
```

This is a path starting from the pattern node `a`, with an outgoing relationship from it to pattern node `b`.

Paths can be of arbitrary length, and the same node may appear in multiple places in the path.

Path patterns are expressions, and since these expressions are collections, they can also be used as predicates (where a non-empty collection signifies true).

Node identifiers can be used with or without surrounding parenthesis. The following match is semantically identical to the one we saw above — the difference is purely aesthetic.

```
a-->b
```

If you don't care about a node, you don't need to name it. Empty parenthesis are used for these nodes, like so:

```
a-->()<--b
```

11.5.2. Working with relationships

If you need to work with the relationship between two nodes, you can name it.

```
a-[r]->b
```

If you don't care about the direction of the relationship, you can omit the arrow at either end of the relationship, like this:

```
a--b
```

Relationships have types. When you are only interested in a specific relationship type, you can specify this like so:

```
a-[ :REL_TYPE ]->b
```

If multiple relationship types are acceptable, you can list them, separating them with the pipe symbol `|` like this:

```
a-[r:TYPE1|TYPE2]->b
```

This pattern matches a relationship of type `TYPE1` or `TYPE2`, going from `a` to `b`. The relationship is named `r`. Multiple relationship types can not be used with `CREATE` or `CREATE UNIQUE`.

11.5.3. Optional relationships

An optional relationship is matched when it is found, but replaced by a `null` otherwise. Normally, if no matching relationship is found, that sub-graph is not matched. Optional relationships could be considered the Cypher equivalent of the outer join in SQL and can only be used in a `MATCH` clause.

Optional relationships are marked with a question mark. They allow you to write queries like this one:

Query.

```
START me=node(*)
MATCH me-->friend-[?]->friend_of_friend
RETURN friend, friend_of_friend
```

The query above says “for every person, give me all their friends, and their friends friends, if they have any.”

Optionality is transitive — if a part of the pattern can only be reached from a bound point through an optional relationship, that part is also optional. In the pattern above, the only bound point in the pattern is `me`. Since the relationship between `friend` and `children` is optional, `children` is an optional part of the graph.

Also, named paths that contain optional parts are also optional — if any part of the path is `null`, the whole path is `null`.

In the following examples, `b` and `p` are all optional and can contain `null`:

Query.

```
START a=node(4)
MATCH p = a-[?]->b
RETURN b
```

Query.

```
START a=node(4)
MATCH p = a-[?*]->b
RETURN b
```

Query.

```
START a=node(4)
MATCH p = a-[?]->x-->b
RETURN b
```

Query.

```
START a=node(4), x=node(3)
MATCH p = shortestPath( a-[?*]->x )
RETURN p
```

11.5.4. Controlling depth

A pattern relationship can span multiple graph relationships. These are called variable length relationships, and are marked as such using an asterisk (`*`):

(a)-[*]->(b)

This signifies a path starting on the pattern node a, following only outgoing relationships, until it reaches pattern node b. Any number of relationships can be followed searching for a path to b, so this can be a very expensive query, depending on what your graph looks like.

You can set a minimum set of steps that can be taken, and/or the maximum number of steps:

(a)-[*3..5]->(b)

This is a variable length relationship containing at least three graph relationships, and at most five.

Variable length relationships can not be used with CREATE and CREATE UNIQUE.

As a simple example, let's take the query below:

Query.

```
START me=node(3)
MATCH me-[:KNOWS*1..2]-remote_friend
RETURN remote_friend
```

Result

remote_friend
(empty result)
0 row

This query starts from one node, and follows KNOWS relationships two or three steps out, and then stops.

11.5.5. Assigning to path identifiers

In a graph database, a path is a very important concept. A path is a collection of nodes and relationships, that describe a path in the graph. To assign a path to a path identifier, you simply assign a path pattern to an identifier, like so:

p = (a)-[*3..5]->(b)

You can do this in MATCH, CREATE and CREATE UNIQUE, but not when using patterns as expressions.

Example of the three in a single query:

Query.

```
START me=node(3)
MATCH p1 = me-[*2]-friendOfFriend
CREATE p2 = me-[:MARRIED_TO]-(:wife {name:"Gunhild"})
CREATE UNIQUE p3 = wife-[:KNOWS]-friendOfFriend
RETURN p1,p2,p3
```

11.5.6. Setting properties

Nodes and relationships are important, but Neo4j uses properties on both of these to allow for far richer graph models.

Properties are expressed in patterns using a map-construct, simply curly brackets surrounding a number of key-expression pairs, separated by commas, e.g. { name: "Andres", sport: "Brazilian Ju-Jitsu" }. If the map is supplied through a parameter, the normal parameter expression is used: { paramName }.

Patterns are also used to mutate the graph with CREATE and CREATE UNIQUE. Maps are only used by CREATE and CREATE UNIQUE. In CREATE they are used to set the properties on the newly created nodes and relationships. When used with CREATE UNIQUE, they are used to try to match a pattern element with the

corresponding graph element. The match is successful if the properties on the pattern element can be matched exactly against properties on the graph elements. The graph element can have additional properties, and they do not affect the match. If Neo4j fails to find matching graph elements, the maps is used to set the properties on the newly created elements.


Chapter 12. Reading Clauses

Read data from the database.

12.1. Start

Every query describes a pattern, and in that pattern one can have multiple starting points. A starting point is a relationship or a node where a pattern is anchored. You can either introduce starting points by id, or by index lookups. Note that trying to use an index that doesn't exist will throw an exception.

Figure 12.1. Graph



12.1.1. Node by id

Binding a node as a starting point is done with the `node(*)` function.



Note

Neo4j reuses its internal ids when nodes and relationships are deleted, which means it's bad practice to refer to them this way. Instead, use application generated ids.

Query.

```
START n=node(1)
RETURN n
```

The corresponding node is returned.

Result

n
Node[1]{name:"A"}
1 row

12.1.2. Relationship by id

Binding a relationship as a starting point is done with the `relationship(*)` function, which can also be abbreviated `rel(*)`. See [Section 12.1.1, "Node by id"](#) for more information on Neo4j ids.

Query.

```
START r=relationship(0)
RETURN r
```

The relationship with id 0 is returned.

Result

r
:KNOWS[0] {}
1 row

12.1.3. Multiple nodes by id

Multiple nodes are selected by listing them separated by commas.

Query.

```
START n=node(1, 2, 3)
RETURN n
```

This returns the nodes listed in the START statement.

Result

n
Node[1]{name: "A"}
Node[2]{name: "B"}
Node[3]{name: "C"}
3 rows

12.1.4. All nodes

To get all the nodes, use an asterisk. This can be done with relationships as well.

Query.

```
START n=node(*)
RETURN n
```

This query returns all the nodes in the graph.

Result

n
Node[1]{name: "A"}
Node[2]{name: "B"}
Node[3]{name: "C"}
3 rows

12.1.5. Node by index lookup

When the starting point can be found by using index lookups, it can be done like this: `node: index-name(key = "value")`. In this example, there exists a node index named `nodes`.

Query.

```
START n=node:nodes(name = "A")
RETURN n
```

The query returns the node indexed with the name "A".

Result

n
Node[1]{name: "A"}
1 row

12.1.6. Relationship by index lookup

When the starting point can be found by using index lookups, it can be done like this:

relationship:index-name(key = "value").

Query.

```
START r=relationship:rels(name = "Andrés")
RETURN r
```

The relationship indexed with the name property set to "Andrés" is returned by the query.

Result

r
:KNOWS[0] {name:"Andrés"}
1 row

12.1.7. Node by index query

When the starting point can be found by more complex Lucene queries, this is the syntax to use:

node:index-name("query").This allows you to write more advanced index queries.

Query.

```
START n=node:nodes("name:A")
RETURN n
```

The node indexed with name "A" is returned by the query.

Result

n
Node[1]{name:"A"}
1 row

12.1.8. Multiple starting points

Sometimes you want to bind multiple starting points. Just list them separated by commas.

Query.

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

Both the nodes A and the B are returned.

Result

a	b
Node[1]{name:"A"}	Node[2]{name:"B"}
1 row	

12.2. Match

12.2.1. Introduction




Tip

In the MATCH clause, patterns are used a lot. Read [Section 11.5, “Patterns”](#) for an introduction.

The following graph is used for the examples below:

Figure 12.2. Graph



12.2.2. Related nodes

The symbol -- means *related to*, without regard to type or direction.

Query.

```

START n=node(3)
MATCH (n)--(x)
RETURN x
    
```

All nodes related to A (Anders) are returned by the query.

Result

x
Node[4]{name: "Bossman"}
Node[1]{name: "David"}
Node[5]{name: "Cesar"}
3 rows

12.2.3. Outgoing relationships

When the direction of a relationship is interesting, it is shown by using --> or <--, like this:

Query.

```
START n=node(3)
MATCH (n)-->(x)
RETURN x
```

All nodes that A has outgoing relationships to are returned.

Result

x
Node[4]{name: "Bossman"}
Node[5]{name: "Cesar"}
2 rows

12.2.4. Directed relationships and identifier

If an identifier is needed, either for filtering on properties of the relationship, or to return the relationship, this is how you introduce the identifier.

Query.

```
START n=node(3)
MATCH (n)-[r]->()
RETURN r
```

The query returns all outgoing relationships from node A.

Result

r
:KNOWS[0] {}
:BLOCKS[1] {}
2 rows

12.2.5. Match by relationship type

When you know the relationship type you want to match on, you can specify it by using a colon together with the relationship type.

Query.

```
START n=node(3)
MATCH (n)-[:BLOCKS]->(x)
RETURN x
```

All nodes that are BLOCKed by A are returned by this query.

Result

x
Node[5]{name: "Cesar"}
1 row

12.2.6. Match by multiple relationship types

To match on one of multiple types, you can specify this by chaining them together with the pipe symbol |.

Query.

```
START n=node(3)
MATCH (n)-[:BLOCKS|KNOWS]->(x)
RETURN x
```

All nodes with a BLOCK or KNOWS relationship to A are returned.

Result

x
Node[5]{name: "Cesar"}
Node[4]{name: "Bossman"}
2 rows

12.2.7. Match by relationship type and use an identifier

If you both want to introduce an identifier to hold the relationship, and specify the relationship type you want, just add them both, like this.

Query.

```
START n=node(3)
MATCH (n)-[r:BLOCKS]->(c)
RETURN r
```

All BLOCKS relationships going out from A are returned.

Result

r
:BLOCKS[1] {}
1 row

12.2.8. Relationship types with uncommon characters

Sometime your database will have types with non-letter characters, or with spaces in them. Use ` (backtick) to quote these.

Query.

```
START n=node(3)
MATCH (n)-[r:`TYPE THAT HAS SPACE IN IT`]->(c)
RETURN r
```

This query returns a relationship of a type with spaces in it.

Result

r
:TYPE THAT HAS SPACE IN IT[6] {}
1 row

12.2.9. Multiple relationships

Relationships can be expressed by using multiple statements in the form of `()--()`, or they can be strung together, like this:

Query.

```
START a=node(3)
MATCH (a)-[:KNOWS]->(b)-[:KNOWS]->(c)
RETURN a,b,c
```

The three nodes in the path are returned by the query.

Result

a	b	c
Node[3]{name:"Anders"}	Node[4]{name:"Bossman"}	Node[2]{name:"Emil"}
1 row		

12.2.10. Variable length relationships

Nodes that are a variable number of relationship→node hops away can be found using the following syntax: `-[:TYPE*minHops..maxHops]->`. `minHops` and `maxHops` are optional and default to 1 and infinity respectively. When no bounds are given the dots may be omitted.

Query.

```
START a=node(3), x=node(2, 4)
MATCH a-[:KNOWS*1..3]->x
RETURN a,x
```

This query returns the start and end point, if there is a path between 1 and 3 relationships away.

Result

a	x
Node[3]{name:"Anders"}	Node[2]{name:"Emil"}
Node[3]{name:"Anders"}	Node[4]{name:"Bossman"}
2 rows	

12.2.11. Relationship identifier in variable length relationships

When the connection between two nodes is of variable length, a relationship identifier becomes an collection of relationships.

Query.

```
START a=node(3), x=node(2, 4)
MATCH a-[r:KNOWS*1..3]->x
RETURN r
```

The query returns the relationships, if there is a path between 1 and 3 relationships away.

Result

r
[:KNOWS[0] {}, :KNOWS[3] {}]
[:KNOWS[0] {}]
2 rows

12.2.12. Zero length paths

Using variable length paths that have the lower bound zero means that two identifiers can point to the same node. If the distance between two nodes is zero, they are by definition the same node. Note that when matching zero length paths the result may contain a match even when matching on a relationship type not in use.

Query.

```
START a=node(3)
MATCH p1=a-[:KNOWS*0..1]->b, p2=b-[:BLOCKS*0..1]->c
RETURN a,b,c, length(p1), length(p2)
```

This query will return four paths, some of which have length zero.

Result

a	b	c	length(p1)	length(p2)
Node[3] {name: "Anders"}	Node[3] {name: "Anders"}	Node[3] {name: "Anders"}	0	0
Node[3] {name: "Anders"}	Node[3] {name: "Anders"}	Node[5] {name: "Cesar"}	0	1
Node[3] {name: "Anders"}	Node[4] {name: "Bossman"}	Node[4] {name: "Bossman"}	1	0
Node[3] {name: "Anders"}	Node[4] {name: "Bossman"}	Node[1] {name: "David"}	1	1
4 rows				

12.2.13. Optional relationship

If a relationship is optional, it can be marked with a question mark. This is similar to how a SQL outer join works. If the relationship is there, it is returned. If it's not, `null` is returned in its place. Remember that anything hanging off an optional relationship, is in turn optional, unless it is connected with a bound node through some other path.

Query.

```
START a=node(2)
MATCH a-[:?]->x
RETURN a,x
```

A node, and `null` are returned, since the node has no outgoing relationships.

Result

a	x
Node[2]{name: "Emil"}	<null>
1 row	

12.2.14. Optional typed and named relationship

Just as with a normal relationship, you can decide which identifier it goes into, and what relationship type you need.

Query.

```
START a=node(3)
```

```
MATCH a-[r?:LOVES]->()
RETURN a,r
```

This returns a node, and null, since the node has no outgoing LOVES relationships.

Result

a	r
Node[3]{name:"Anders"}	<null>
1 row	

12.2.15. Properties on optional elements

Returning a property from an optional element that is null will also return null.

Query.

```
START a=node(2)
MATCH a-[?]->x
RETURN x, x.name
```

This returns the element x (null in this query), and null as it's name.

Result

x	x.name
<null>	<null>
1 row	

12.2.16. Complex matching

Using Cypher, you can also express more complex patterns to match on, like a diamond shape pattern.

Query.

```
START a=node(3)
MATCH (a)-[:KNOWS]->(b)-[:KNOWS]->(c), (a)-[:BLOCKS]->(d)-[:KNOWS]->(c)
RETURN a,b,c,d
```

This returns the four nodes in the paths.

Result

a	b	c	d
Node[3]{name:"Anders"}	Node[4]{name:"Bossman"}	Node[2]{name:"Emil"}	Node[5]{name:"Cesar"}
1 row			

12.2.17. Shortest path

Finding a single shortest path between two nodes is as easy as using the shortestPath function. It's done like this:

Query.

```
START d=node(1), e=node(2)
MATCH p = shortestPath( d-[*..15]->e )
RETURN p
```

This means: find a single shortest path between two nodes, as long as the path is max 15 relationships long. Inside of the parenthesis you define a single link of a path — the starting node, the connecting

relationship and the end node. Characteristics describing the relationship like relationship type, max hops and direction are all used when finding the shortest path. You can also mark the path as optional.

Result

p
[Node[1]{name:"David"}, :KNOWS[2] {}, Node[3]{name:"Anders"}, :KNOWS[0] {}, Node[4]{name:"Bossman"}, :KNOWS[3] {}, Node[2]{name:"Emil"}]
1 row

12.2.18. All shortest paths

Finds all the shortest paths between two nodes.

Query.

```
START d=node(1), e=node(2)
MATCH p = allShortestPaths( d-[*..15]->e )
RETURN p
```

This example will find the two directed paths between David and Emil.

Result

p
[Node[1]{name:"David"}, :KNOWS[2] {}, Node[3]{name:"Anders"}, :KNOWS[0] {}, Node[4]{name:"Bossman"}, :KNOWS[3] {}, Node[2]{name:"Emil"}]
[Node[1]{name:"David"}, :KNOWS[2] {}, Node[3]{name:"Anders"}, :BLOCKS[1] {}, Node[5]{name:"Cesar"}, :KNOWS[4] {}, Node[2]{name:"Emil"}]
2 rows

12.2.19. Named path

If you want to return or filter on a path in your pattern graph, you can introduce a named path.

Query.

```
START a=node(3)
MATCH p = a-->b
RETURN p
```

This returns the two paths starting from the first node.

Result

p
[Node[3]{name:"Anders"}, :KNOWS[0] {}, Node[4]{name:"Bossman"}]
[Node[3]{name:"Anders"}, :BLOCKS[1] {}, Node[5]{name:"Cesar"}]
2 rows

12.2.20. Matching on a bound relationship

When your pattern contains a bound relationship, and that relationship pattern doesn't specify direction, Cypher will try to match the relationship where the connected nodes switch sides.

Query.

```
START r=rel(0)
MATCH a-[r]-b
RETURN a,b
```

This returns the two connected nodes, once as the start node, and once as the end node.

Result

a	b
Node[3]{name: "Anders"}	Node[4]{name: "Bossman"}
Node[4]{name: "Bossman"}	Node[3]{name: "Anders"}
2 rows	

12.2.21. Match with OR

Strictly speaking, you can't do OR in your MATCH. It's still possible to form a query that works a lot like OR.

Query.

```
START a=node(3), b=node(2)
MATCH a-[?:KNOWS]-x-[?:KNOWS]-b
RETURN x
```

This query is saying: give me the nodes that are connected to a, or b, or both.


Result

x
Node[4]{name: "Bossman"}
Node[5]{name: "Cesar"}
Node[1]{name: "David"}
3 rows

12.3. Where

If you need filtering apart from the pattern of the data that you are looking for, you can add clauses in the WHERE part of the query.

Figure 12.3. Graph



12.3.1. Boolean operations

You can use the expected boolean operators AND and OR, and also the boolean function NOT().

Query.

```
START n=node(3, 1)
WHERE (n.age < 30 and n.name = "Tobias") or not(n.name = "Tobias")
RETURN n
```

This will return both nodes in the start clause.

Result

n
Node[3]{name: "Andres", age: 36, belt: "white"}
Node[1]{name: "Tobias", age: 25}
2 rows

12.3.2. Filter on node property

To filter on a property, write your clause after the WHERE keyword. Filtering on relationship properties works just the same way.

Query.

```
START n=node(3, 1)
WHERE n.age < 30
RETURN n
```

The "Tobias" node will be returned.

Result

n
Node[1]{name: "Tobias", age: 25}
1 row

12.3.3. Regular expressions

You can match on regular expressions by using `=~ "regexp"`, like this:

Query.

```
START n=node(3, 1)
WHERE n.name =~ 'Tob.*'
RETURN n
```

The "Tobias" node will be returned.

Result

n
Node[1]{name: "Tobias", age: 25}
1 row

12.3.4. Escaping in regular expressions

If you need a forward slash inside of your regular expression, escape it. Remember that back slash needs to be escaped in string literals

Query.

```
START n=node(3, 1)
WHERE n.name =~ 'Some\\/thing'
RETURN n
```

No nodes match this regular expression.

Result

n
(empty result)
0 row

12.3.5. Case insensitive regular expressions

By pre-pending a regular expression with `(?i)`, the whole expression becomes case insensitive.

Query.

```
START n=node(3, 1)
WHERE n.name =~ '(?i)ANDR.*'
RETURN n
```

The node with name "Andres" is returned.

Result

n
Node[3]{name: "Andres", age: 36, belt: "white"}
1 row

12.3.6. Filtering on relationship type

You can put the exact relationship type in the `MATCH` pattern, but sometimes you want to be able to do more advanced filtering on the type. You can use the special property `TYPE` to compare the type with

something else. In this example, the query does a regular expression comparison with the name of the relationship type.

Query.

```
START n=node(3)
MATCH (n)-[r]->()
WHERE type(r) =~ 'K.*'
RETURN r
```

This returns relationships that has a type whose name starts with K.

Result

r
:KNOWS[0] {}
:KNOWS[1] {}
2 rows

12.3.7. Property exists

To only include nodes/relationships that have a property, use the HAS() function and just write out the identifier and the property you expect it to have.

Query.

```
START n=node(3, 1)
WHERE has(n.belt)
RETURN n
```

The node named "Andres" is returned.

Result

n
Node[3]{name:"Andres", age:36, belt:"white"}
1 row

12.3.8. Default true if property is missing

If you want to compare a property on a graph element, but only if it exists, use the nullable property syntax. You can use a question mark if you want missing property to return true, like:

Query.

```
START n=node(3, 1)
WHERE n.belt? = 'white'
RETURN n
```

This returns all nodes, even those without the belt property.

Result

n
Node[3]{name:"Andres", age:36, belt:"white"}
Node[1]{name:"Tobias", age:25}
2 rows

12.3.9. Default false if property is missing

When you need missing property to evaluate to false, use the exclamation mark.

Query.

```
START n=node(3, 1)
WHERE n.belt! = 'white'
RETURN n
```

No nodes without the belt property are returned.

Result

n
Node[3]{name: "Andres", age: 36, belt: "white"}
1 row

12.3.10. Filter on null values

Sometimes you might want to test if a value or an identifier is `null`. This is done just like SQL does it, with `IS NULL`. Also like SQL, the negative is `IS NOT NULL`, although `NOT(IS NULL x)` also works.

Query.

```
START a=node(1), b=node(3, 2)
MATCH a<-[r?]-b
WHERE r is null
RETURN b
```

Nodes that Tobias is not connected to are returned.

Result

b
Node[2]{name: "Peter", age: 34}
1 row

12.3.11. Filter on patterns

Patterns are expressions in Cypher, expressions that return a collection of paths. Collection expressions are also predicates — an empty collection represents `false`, and a non-empty represents `true`.

So, patterns are not only expressions, they are also predicates. The only limitation to your pattern is that you must be able to express it in a single path. You can not use commas between multiple paths like you do in `MATCH`. You can achieve the same effect by combining multiple patterns with `AND`.

Note that you can not introduce new identifiers here. Although it might look very similar to the `MATCH` patterns, the `WHERE` clause is all about eliminating matched subgraphs. `MATCH a-[*]->b` is very different from `WHERE a-[*]->b`; the first will produce a subgraph for every path it can find between `a` and `b`, and the latter will eliminate any matched subgraphs where `a` and `b` do not have a directed relationship chain between them.

Query.

```
START tobias=node(1), others=node(3, 2)
WHERE tobias<--others
RETURN others
```

Nodes that have an outgoing relationship to the "Tobias" node are returned.

Result

others
Node[3]{name:"Andres", age:36, belt:"white"}
1 row

12.3.12. Filter on patterns using NOT

The NOT() function can be used to exclude a pattern.

Query.

```
START persons=node(*), peter=node(2)
WHERE not(persons-->peter)
RETURN persons
```

Nodes that do not have an outgoing relationship to the "Peter" node are returned.

Result

persons
Node[1]{name:"Tobias", age:25}
Node[2]{name:"Peter", age:34}
2 rows

12.3.13. IN operator

To check if an element exists in a collection, you can use the IN operator.

Query.

```
START a=node(3, 1, 2)
WHERE a.name IN ["Peter", "Tobias"]
RETURN a
```

This query shows how to check if a property exists in a literal collection.


Result

a
Node[1]{name:"Tobias", age:25}
Node[2]{name:"Peter", age:34}
2 rows

12.4. Return

In the RETURN part of your query, you define which parts of the pattern you are interested in. It can be nodes, relationships, or properties on these.

Figure 12.4. Graph



12.4.1. Return nodes

To return a node, list it in the RETURN statement.

Query.

```
START n=node(2)
RETURN n
```

The example will return the node.

Result

n
Node[2]{name: "B"}
1 row

12.4.2. Return relationships

To return a relationship, just include it in the RETURN list.

Query.

```
START n=node(1)
MATCH (n)-[r:KNOWS]->(c)
RETURN r
```

The relationship is returned by the example.

Result

r
:KNOWS[0] {}
1 row

12.4.3. Return property

To return a property, use the dot separator, like this:

Query.

```
START n=node(1)
RETURN n.name
```

The value of the property name gets returned.

Result

n.name
"A"
1 row

12.4.4. Return all elements

When you want to return all nodes, relationships and paths found in a query, you can use the * symbol.

Query.

```
START a=node(1)
MATCH p=a-[r]->b
RETURN *
```

This returns the two nodes, the relationship and the path used in the query.

Result

b	a	r	p
Node[2]{name: "B"}	Node[1]{name: "A", happy: "Yes!", age: 55}	:KNOWS[0] {}	[Node[1]{name: "A", happy: "Yes!", age: 55}, :KNOWS[0] {}, Node[2]{name: "B"}]
Node[2]{name: "B"}	Node[1]{name: "A", happy: "Yes!", age: 55}	:BLOCKS[1] {}	[Node[1]{name: "A", happy: "Yes!", age: 55}, :BLOCKS[1] {}, Node[2]{name: "B"}]
2 rows			

12.4.5. Identifier with uncommon characters

To introduce a placeholder that is made up of characters that are outside of the english alphabet, you can use the ` to enclose the identifier, like this:

Query.

```
START `This isn't a common identifier`=node(1)
RETURN `This isn't a common identifier`.happy
```

The node indexed with name "A" is returned

Result

This isn't a common identifier.happy
"Yes!"
1 row

12.4.6. Column alias

If the name of the column should be different from the expression used, you can rename it by using AS <new name>.

Query.

```
START a=node(1)
RETURN a.age AS SomethingTotallyDifferent
```

Returns the age property of a node, but renames the column.

Result

SomethingTotallyDifferent
55
1 row

12.4.7. Optional properties

If a property might or might not be there, you can select it optionally by adding a questionmark to the identifier, like this:

Query.

```
START n=node(1, 2)
RETURN n.age?
```

This example returns the age when the node has that property, or null if the property is not there.

Result

n.age?
55
<null>
2 rows

12.4.8. Other expressions

Any expression can be used as a return item - literals, predicates, properties, functions, and everything else.

Query.

```
START a=node(1)
RETURN a.age > 30, "I'm a literal", length(a-->())
```

Returns a predicate, a literal and function call with a pattern expression parameter.

Result

a.age > 30	"I'm a literal"	length(a-->())
true	"I'm a literal"	2
1 row		

12.4.9. Unique results

DISTINCT retrieves only unique rows depending on the columns that have been selected to output.

Query.

```
START a=node(1)
MATCH (a)-->(b)
RETURN distinct b
```

The node named B is returned by the query, but only once.

Result

b
Node[2]{name: "B"}
1 row

12.5. Aggregation

12.5.1. Introduction

To calculate aggregated data, Cypher offers aggregation, much like SQL's `GROUP BY`.

Aggregate functions take multiple input values and calculate an aggregated value from them. Examples are `AVG` that calculate the average of multiple numeric values, or `MIN` that finds the smallest numeric value in a set of values.

Aggregation can be done over all the matching sub graphs, or it can be further divided by introducing key values. These are non-aggregate expressions, that are used to group the values going into the aggregate functions.

So, if the return statement looks something like this:

```
RETURN n, count(*)
```

We have two return expressions — `n`, and `count(*)`. The first, `n`, is no aggregate function, and so it will be the grouping key. The latter, `count(*)` is an aggregate expression. So the matching subgraphs will be divided into different buckets, depending on the grouping key. The aggregate function will then run on these buckets, calculating the aggregate values.

If you want to use aggregations to sort your result set, the aggregation must be included in the `RETURN` to be used in your `ORDER BY`.

The last piece of the puzzle is the `DISTINCT` keyword. It is used to make all values unique before running them through an aggregate function.

An example might be helpful:

Query.

```
START me=node(1)
MATCH me-->friend-->friend_of_friend
RETURN count(distinct friend_of_friend), count(friend_of_friend)
```


In this example we are trying to find all our friends of friends, and count them. The first aggregate function, `count(distinct friend_of_friend)`, will only see a `friend_of_friend` once — `DISTINCT` removes the duplicates. The latter aggregate function, `count(friend_of_friend)`, might very well see the same `friend_of_friend` multiple times. Since there is no real data in this case, an empty result is returned. See the sections below for real data.

Result

<code>count(distinct friend_of_friend)</code>	<code>count(friend_of_friend)</code>
0	0
1 row	

The following examples are assuming the example graph structure below.

Figure 12.5. Graph



12.5.2. COUNT

COUNT is used to count the number of rows. COUNT can be used in two forms — COUNT(*) which just counts the number of matching rows, and COUNT(<identifier>), which counts the number of non-null values in <identifier>.

12.5.3. Count nodes

To count the number of nodes, for example the number of nodes connected to one node, you can use count(*).

Query.

```
START n=node(2)
MATCH (n)-->(x)
RETURN n, count(*)
```

This returns the start node and the count of related nodes.

Result

n	count(*)
Node[2]{name: "A", property:13}	3
1 row	

12.5.4. Group Count Relationship Types

To count the groups of relationship types, return the types and count them with count(*).

Query.

```
START n=node(2)
MATCH (n)-[r]->()
RETURN type(r), count(*)
```

The relationship types and their group count is returned by the query.

Result

type(r)	count(*)
"KNOWS"	3
1 row	

12.5.5. Count entities

Instead of counting the number of results with `count(*)`, it might be more expressive to include the name of the identifier you care about.

Query.

```
START n=node(2)
MATCH (n)-->(x)
RETURN count(x)
```

The example query returns the number of connected nodes from the start node.

Result

count(x)
3
1 row

12.5.6. Count non-null values

You can count the non-null values by using `count(<identifier>)`.

Query.

```
START n=node(2,3,4,1)
RETURN count(n.property?)
```

The count of related nodes with the property property set is returned by the query.

Result

count(n.property?)
3
1 row

12.5.7. SUM

The SUM aggregation function simply sums all the numeric values it encounters. Nulls are silently dropped. This is an example of how you can use SUM.

Query.

```
START n=node(2,3,4)
RETURN sum(n.property)
```

This returns the sum of all the values in the property property.

Result

sum(n.property)
90
1 row

12.5.8. AVG

AVG calculates the average of a numeric column.

Query.

```
START n=node(2,3,4)
RETURN avg(n.property)
```

The average of all the values in the property property is returned by the example query.

Result

avg(n.property)
30.0
1 row

12.5.9. PERCENTILE_DISC

PERCENTILE_DISC calculates the percentile of a given value over a group, with a percentile from 0.0 to 1.0. It uses a rounding method, returning the nearest value to the percentile. For interpolated values, see PERCENTILE_CONT.

Query.

```
START n=node(2,3,4)
RETURN percentile_disc(n.property, 0.5)
```

The 50th percentile of the values in the property property is returned by the example query. In this case, 0.5 is the median, or 50th percentile.

Result

percentile_disc(n.property, 0.5)
33
1 row

12.5.10. PERCENTILE_CONT

PERCENTILE_CONT calculates the percentile of a given value over a group, with a percentile from 0.0 to 1.0. It uses a linear interpolation method, calculating a weighted average between two values, if the desired percentile lies between them. For nearest values using a rounding method, see PERCENTILE_DISC.

Query.

```
START n=node(2,3,4)
RETURN percentile_cont(n.property, 0.4)
```

The 40th percentile of the values in the property property is returned by the example query, calculated with a weighted average.

Result

percentile_cont(n.property, 0.4)
29.0
1 row

12.5.11. MAX

MAX find the targets value in a numeric column.

Query.

```
START n=node(2,3,4)
RETURN max(n.property)
```

The largest of all the values in the property property is returned.

Result

max(n.property)
44
1 row

12.5.12. MIN

MIN takes a numeric property as input, and returns the smallest value in that column.

Query.

```
START n=node(2,3,4)
RETURN min(n.property)
```

This returns the smallest of all the values in the property property.

Result

min(n.property)
13
1 row

12.5.13. COLLECT

COLLECT collects all the values into a list. It will ignore null values,

Query.

```
START n=node(2,3,4,1)
RETURN collect(n.property?)
```

Returns a single row, with all the values collected.

Result

collect(n.property?)
[13, 33, 44]
1 row

12.5.14. DISTINCT

All aggregation functions also take the DISTINCT modifier, which removes duplicates from the values. So, to count the number of unique eye colors from nodes related to a, this query can be used:

Query.

```
START a=node(2)
MATCH a-->b
RETURN count(distinct b.eyes)
```

Returns the number of eye colors.


Result

count(distinct b.eyes)
2
1 row

12.6. Order by

To sort the output, use the `ORDER BY` clause. Note that you can not sort on nodes or relationships, just on properties on these.

Figure 12.6. Graph



12.6.1. Order nodes by property

`ORDER BY` is used to sort the output.

Query.

```
START n=node(3,1,2)
RETURN n
ORDER BY n.name
```

The nodes are returned, sorted by their name.

Result

n
Node[1]{name: "A", age: 34, length: 170}
Node[2]{name: "B", age: 34}
Node[3]{name: "C", age: 32, length: 185}
3 rows

12.6.2. Order nodes by multiple properties

You can order by multiple properties by stating each identifier in the `ORDER BY` clause. Cypher will sort the result by the first identifier listed, and for equals values, go to the next property in the `ORDER BY` clause, and so on.

Query.

```
START n=node(3,1,2)
```

```
RETURN n
ORDER BY n.age, n.name
```

This returns the nodes, sorted first by their age, and then by their name.

Result

n
Node[3]{name: "C", age: 32, length: 185}
Node[1]{name: "A", age: 34, length: 170}
Node[2]{name: "B", age: 34}
3 rows

12.6.3. Order nodes in descending order

By adding DESC[ENDING] after the identifier to sort on, the sort will be done in reverse order.

Query.

```
START n=node(3,1,2)
RETURN n
ORDER BY n.name DESC
```

The example returns the nodes, sorted by their name reversely.

Result

n
Node[3]{name: "C", age: 32, length: 185}
Node[2]{name: "B", age: 34}
Node[1]{name: "A", age: 34, length: 170}
3 rows

12.6.4. Ordering null

When sorting the result set, null will always come at the end of the result set for ascending sorting, and first when doing descending sort.

Query.

```
START n=node(3,1,2)
RETURN n.length?, n
ORDER BY n.length?
```

The nodes are returned sorted by the length property, with a node without that property last.


Result

n.length?	n
170	Node[1]{name: "A", age: 34, length: 170}
185	Node[3]{name: "C", age: 32, length: 185}
<null>	Node[2]{name: "B", age: 34}
3 rows	

12.7. Limit

LIMIT enables the return of only subsets of the total result.

Figure 12.7. Graph



12.7.1. Return first part

To return a subset of the result, starting from the top, use this syntax:

Query.

```

START n=node(3, 4, 5, 1, 2)
RETURN n
LIMIT 3
  
```

The top three items are returned by the example query.


Result

n
Node[3]{name: "A"}
Node[4]{name: "B"}
Node[5]{name: "C"}
3 rows

12.8. Skip

SKIP enables the return of only subsets of the total result. By using SKIP, the result set will get trimmed from the top. Please note that no guarantees are made on the order of the result unless the query specifies the ORDER BY clause.

Figure 12.8. Graph



12.8.1. Skip first three

To return a subset of the result, starting from the fourth result, use the following syntax:

Query.

```
START n=node(3, 4, 5, 1, 2)
RETURN n
ORDER BY n.name
SKIP 3
```

The first three nodes are skipped, and only the last two are returned in the result.

Result

n
Node[1]{name: "D"}
Node[2]{name: "E"}
2 rows

12.8.2. Return middle two

To return a subset of the result, starting from somewhere in the middle, use this syntax:

Query.

```
START n=node(3, 4, 5, 1, 2)
RETURN n
ORDER BY n.name
SKIP 1
LIMIT 2
```

Two nodes from the middle are returned.

Result

n
Node[4]{name: "B"}
2 rows


n
Node[5]{name: "C"}
2 rows

12.9. With

The ability to chain queries together allows for powerful constructs. In Cypher, the `WITH` clause is used to pipe the result from one query to the next.

`WITH` is also used to separate reading from updating of the graph. Every sub-query of a query must be either read-only or write-only.

Figure 12.9. Graph



12.9.1. Filter on aggregate function results

Aggregated results have to pass through a `WITH` clause to be able to filter on.

Query.

```

START david=node(1)
MATCH david--otherPerson-->()
WITH otherPerson, count(*) as foaf
WHERE foaf > 1
RETURN otherPerson
  
```

The person connected to David with the at least more than one outgoing relationship will be returned by the query.

Result

otherPerson
Node[3]{name: "Anders"}
1 row

12.9.2. Sort results before using collect on them

You can sort your results before passing them to `collect`, thus sorting the resulting collection.

Query.

```
START n=node(*)
WITH n
ORDER BY n.name desc
LIMIT 3
RETURN collect(n.name)
```

A list of the names of people in reverse order, limited to 3, in a collection.

Result

collect(n.name)
["Emil", "David", "Cesar"]
1 row

12.9.3. Limit branching of your path search

You can match paths, limit to a certain number, and then match again using those paths as a base As well as any number of similar limited searches.

Query.

```
START n=node(3)
MATCH n--m
WITH m
ORDER BY m.name desc
LIMIT 1
MATCH m--o
RETURN o.name
```

Starting at Anders, find all matching nodes, order by name descending and get the top result, then find all the nodes connected to that top result, and return their names.

Result

o.name
"Anders"
"Bossman"
2 rows

12.9.4. Alternative syntax of WITH

If you prefer a more visual way of writing your query, you can use equal-signs as delimiters before and after the column list. Use at least three before the column list, and at least three after.



Caution

This syntax has been deprecated.

Query.

```
START david=node(1)
MATCH david--otherPerson-->()
===== otherPerson, count(*) as foaf =====
SET otherPerson.connection_count = foaf
```

For persons connected to David, the connection_count property is set to their number of outgoing relationships.

Result

(empty result)

Properties set: 2

Chapter 13. Writing Clauses

Write data to the database.

13.1. Create

Creating graph elements — nodes and relationships, is done with CREATE.



Tip

In the CREATE clause, patterns are used a lot. Read [Section 11.5, “Patterns”](#) for an introduction.

13.1.1. Create single node

Creating a single node is done by issuing the following query.

Query.

```
CREATE n
```

Nothing is returned from this query, except the count of affected nodes.

Result

(empty result)

Nodes created: 1

13.1.2. Return created node

Creating a single node is done by issuing the following query.

Query.

```
CREATE (a {name : 'Andres'})
RETURN a
```

The newly created node is returned.

Result

a

Node[2]{name:"Andres"}

1 row

Nodes created: 1

Properties set: 1

13.1.3. Create single node and set properties

The values for the properties can be any scalar expressions.



Caution

This syntax has been deprecated, please use the one in [Section 13.1.2, “Return created node”](#).

Query.

```
CREATE n = {name : 'Andres', title : 'Developer'}
```

Nothing is returned from this query.

Result

(empty result)
Nodes created: 1
Properties set: 2

13.1.4. Create a relationship between two nodes

To create a relationship between two nodes, we first get the two nodes. Once the nodes are loaded, we simply create a relationship between them.

Query.

```
START a=node(1), b=node(2)
CREATE a-[r:RELTYPE]->b
RETURN r
```

The created relationship is returned by the query.

Result

r
:RELTYPE[1] {}
1 row
Relationships created: 1

13.1.5. Create a relationship and set properties

Setting properties on relationships is done in a similar manner to how it's done when creating nodes. Note that the values can be any expression.

Query.

```
START a=node(1), b=node(2)
CREATE a-[r:RELTYPE {name : a.name + '<->' + b.name }]->b
RETURN r
```

The newly created relationship is returned by the example query.

Result

r
:RELTYPE[1] {name:"Andres<->Michael"}
1 row
Relationships created: 1
Properties set: 1

13.1.6. Create a full path

When you use `CREATE` and a pattern, all parts of the pattern that are not already in scope at this time will be created.

Query.

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

This query creates three nodes and two relationships in one go, assigns it to a path identifier, and returns it.

Result

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

13.1.7. Create single node from map

You can also create a graph entity from a `Map<String, Object>` map. All the key/value pairs in the map will be set as properties on the created relationship or node.

Query.

```
create ({props})
```

This query can be used in the following fashion:

```
Map<String, Object> props = new HashMap<String, Object>();
props.put( "name", "Andres" );
props.put( "position", "Developer" );

Map<String, Object> params = new HashMap<String, Object>();
params.put( "props", props );
engine.execute( "create ({props})", params );
```

13.1.8. Create multiple nodes from maps

By providing an iterable of maps (`Iterable<Map<String, Object>>`), Cypher will create a node for each map in the iterable. When you do this, you can't create anything else in the same create statement.

Query.

```
create (n {props}) return n
```

This query can be used in the following fashion:

```
Map<String, Object> n1 = new HashMap<String, Object>();
n1.put( "name", "Andres" );
n1.put( "position", "Developer" );

Map<String, Object> n2 = new HashMap<String, Object>();
n2.put( "name", "Michael" );
n2.put( "position", "Developer" );

Map<String, Object> params = new HashMap<String, Object>();
List<Map<String, Object>> maps = Arrays.asList( n1, n2 );
params.put( "props", maps );
engine.execute( "create (n {props}) return n", params );
```

13.2. Create Unique

CREATE UNIQUE is in the middle of MATCH and CREATE — it will match what it can, and create what is missing. CREATE UNIQUE will always make the least change possible to the graph — if it can use parts of the existing graph, it will.


Another difference to MATCH is that CREATE UNIQUE assumes the pattern to be unique. If multiple matching subgraphs are found an exception will be thrown.



Tip

In the CREATE UNIQUE clause, patterns are used a lot. Read [Section 11.5, “Patterns”](#) for an introduction.

The examples start out with the following data set:



13.2.1. Create relationship if it is missing

CREATE UNIQUE is used to describe the pattern that should be found or created.

Query.

```
START left=node(1), right=node(3,4)
CREATE UNIQUE left-[r:KNOWS]->right
RETURN r
```

The left node is matched against the two right nodes. One relationship already exists and can be matched, and the other relationship is created before it is returned.

Result

r
:KNOWS[5] {}
:KNOWS[3] {}
2 rows
Relationships created: 1

13.2.2. Create node if missing

If the pattern described needs a node, and it can't be matched, a new node will be created.

Query.

```
START root=node(2)
CREATE UNIQUE root-[:LOVES]-someone
RETURN someone
```

The root node doesn't have any LOVES relationships, and so a node is created, and also a relationship to that node.

Result

someone
Node[6]{}
1 row
Nodes created: 1
Relationships created: 1

13.2.3. Create nodes with values

The pattern described can also contain values on the node. These are given using the following syntax:
prop : <expression>.

Query.

```
START root=node(2)
CREATE UNIQUE root-[:X]-(leaf {name:'D'})
RETURN leaf
```

No node connected with the root node has the name D, and so a new node is created to match the pattern.

Result

leaf
Node[6]{name:"D"}
1 row
Nodes created: 1
Relationships created: 1
Properties set: 1

13.2.4. Create relationship with values

Relationships to be created can also be matched on values.

Query.

```
START root=node(2)
CREATE UNIQUE root-[r:X {since:'forever'}]-()
RETURN r
```

In this example, we want the relationship to have a value, and since no such relationship can be found, a new node and relationship are created. Note that since we are not interested in the created node, we don't name it.

Result

r
:X[5] {since:"forever"}
1 row
Nodes created: 1
Relationships created: 1
Properties set: 1

13.2.5. Describe complex pattern

The pattern described by CREATE UNIQUE can be separated by commas, just like in MATCH and CREATE.

Query.

```
START root=node(2)
CREATE UNIQUE root-[:FOO]->x, root-[:BAR]->x
RETURN x
```

This example pattern uses two paths, separated by a comma.


Result

x
Node[6]{}
1 row
Nodes created: 1
Relationships created: 2

13.3. Set

Updating properties on nodes and relationships is done with the SET clause. SET can also be used with maps from parameters.

The examples use this graph as a starting point:



13.3.1. Set a property

To set a property on a node or relationship, use SET.

Query.

```
START n = node(2)
SET n.surname = 'Taylor'
RETURN n
```

The newly changed node is returned by the query.

Result

n
Node[2]{name:"Andres", age:36, awesome:true, surname:"Taylor"}
1 row
Properties set: 1

13.3.2. Remove a property

Normally you remove a property by using delete, but it's sometimes handy to do it using the SET command. One example is if the property comes from a parameter.

Query.

```
START n = node(2)
SET n.name = null
RETURN n
```

The node is returned by the query, and the name property is now missing.

Result

n
Node[2]{age:36, awesome:true}
1 row
Properties set: 1

13.3.3. Copying properties between nodes and relationships

You can also use SET to copy all properties from one graph element to another. Remember that doing this will remove all other properties on the receiving graph element.

Query.

```
START at = node(2), pn = node(1)
SET at = pn
RETURN at, pn
```

The Andres node has had all it's properties replaced by the properties in the Peter node.


Result

at	pn
Node[2]{name:"Peter", age:34}	Node[1]{name:"Peter", age:34}
1 row	
Properties set: 3	

13.4. Delete

Removing graph elements — nodes, relationships and properties, is done with DELETE.

The examples start out with the following database:



13.4.1. Delete single node

To remove a node from the graph, you can delete it with the DELETE clause.

Query.

```
START n = node(4)
DELETE n
```

Nothing is returned from this query, except the count of affected nodes.

Result

(empty result)

Nodes deleted: 1

13.4.2. Remove a node and connected relationships

If you are trying to remove a node with relationships on it, you have to remove these as well.

Query.

```
START n = node(3)
MATCH n-[r]-()
DELETE n, r
```

Nothing is returned from this query, except the count of affected nodes.

Result

(empty result)

Nodes deleted: 1

Relationships deleted: 2

13.4.3. Remove a property

Neo4j doesn't allow storing `null` in properties. Instead, if no value exists, the property is just not there. So, to remove a property value on a node or a relationship, is also done with DELETE.

Query.

```
START andres = node(3)
DELETE andres.age
RETURN andres
```

The node is returned, and no property age exists on it.

Result

andres
Node[3]{name:"Andres"}
1 row
Properties set: 1


13.5. Foreach

Collections and paths are key concepts in Cypher. To use them for updating data, you can use the `FOREACH` construct. It allows you to do updating commands on elements in a collection — a path, or a collection created by aggregation.

The identifier context inside of the `foreach` parenthesis is separate from the one outside it. This means that if you `CREATE` a node identifier inside of a `FOREACH`, you will *not* be able to use it outside of the `foreach` statement, unless you match to find it.

Inside of the `FOREACH` parentheses, you can do any of the updating commands — `CREATE`, `CREATE UNIQUE`, `DELETE`, and `FOREACH`.

Figure 13.1. Data for the examples



13.5.1. Mark all nodes along a path

This query will set the property `marked` to true on all nodes along a path.

Query.

```

START begin = node(2), end = node(1)
MATCH p = begin -[*]-> end foreach(n in nodes(p) :
SET n.marked = true)

```

Nothing is returned from this query, but four properties are set.

Result

(empty result)

Properties set: 4

Chapter 14. Functions

This chapter contains information on all functions in Cypher. Note that related information exists in [Section 11.1, “Operators”](#).



Note


Most functions in Cypher will return `null` if the input parameter is `null`.

14.1. Predicates

Predicates are boolean functions that return true or false for a given set of input. They are most commonly used to filter out subgraphs in the WHERE part of a query.

See also [Section 11.1.2, “Comparison operators”](#).

Figure 14.1. Graph



14.1.1. ALL

Tests whether a predicate holds for all element of this collection.

Syntax: ALL(identifier in collection WHERE predicate)

Arguments:

- *collection*: An expression that returns a collection
- *identifier*: This is the identifier that can be used from the predicate.
- *predicate*: A predicate that is tested against all items in the collection.

Query.

```

START a=node(3), b=node(1)
MATCH p=a-[*1..3]->b
WHERE all(x in nodes(p)
WHERE x.age > 30)
RETURN p
  
```

All nodes in the returned paths will have an age property of at least 30.

Result

p
[Node[3]{name:"A", age:38, eyes:"brown"}, :KNOWS[1] {}, Node[5]{name:"C", age:53, eyes:"green"}, :KNOWS[3] {}, Node[1]{name:"D", age:54, eyes:"brown"}]
1 row

14.1.2. ANY

Tests whether a predicate holds for at least one element in the collection.

Syntax: ANY(identifier in collection WHERE predicate)

Arguments:

- *collection*: An expression that returns a collection
- *identifier*: This is the identifier that can be used from the predicate.
- *predicate*: A predicate that is tested against all items in the collection.

Query.

```
START a=node(2)
WHERE any(x in a.array
WHERE x = "one")
RETURN a
```

All nodes in the returned paths has at least one one value set in the array property named array.

Result

a
Node[2]{name:"E", age:41, eyes:"blue", array:["one", "two", "three"]}
1 row

14.1.3. NONE

Returns true if the predicate holds for no element in the collection.

Syntax: NONE(identifier in collection WHERE predicate)

Arguments:

- *collection*: An expression that returns a collection
- *identifier*: This is the identifier that can be used from the predicate.
- *predicate*: A predicate that is tested against all items in the collection.

Query.

```
START n=node(3)
MATCH p=n-[*1..3]->b
WHERE NONE(x in nodes(p)
WHERE x.age = 25)
RETURN p
```

No nodes in the returned paths has a age property set to 25.

Result

p
[Node[3]{name:"A", age:38, eyes:"brown"}, :KNOWS[1] {}, Node[5]{name:"C", age:53, eyes:"green"}]
[Node[3]{name:"A", age:38, eyes:"brown"}, :KNOWS[1] {}, Node[5]{name:"C", age:53, eyes:"green"}, :KNOWS[3] {}, Node[1]{name:"D", age:54, eyes:"brown"}]
2 rows

14.1.4. SINGLE

Returns true if the predicate holds for exactly one of the elements in the collection.

Syntax: SINGLE(identifier in collection WHERE predicate)

Arguments:

- *collection*: An expression that returns a collection
- *identifier*: This is the identifier that can be used from the predicate.
- *predicate*: A predicate that is tested against all items in the collection.

Query.

```
START n=node(3)
MATCH p=n-->b
WHERE SINGLE(var in nodes(p)
WHERE var.eyes = "blue")
RETURN p
```

Exactly one node in every returned path will have the eyes property set to "blue".


Result

p
[Node[3]{name:"A", age:38, eyes:"brown"}, :KNOWS[0] {}, Node[4]{name:"B", age:25, eyes:"blue"}]
1 row

14.2. Scalar functions

Scalar functions return a single value.

Figure 14.2. Graph



14.2.1. LENGTH

To return or filter on the length of a collection, use the LENGTH() function.

Syntax: LENGTH(collection)

Arguments:

- *collection*: An expression that returns a collection

Query.

```

START a=node(3)
MATCH p=a-->b-->c
RETURN length(p)
  
```

The length of the path *p* is returned by the query.

Result

length(p)
2
2
2
3 rows

14.2.2. TYPE

Returns a string representation of the relationship type.

Syntax: TYPE(relationship)

Arguments:

- *relationship*: A relationship.

Query.

```
START n=node(3)
MATCH (n)-[r]->( )
RETURN type(r)
```

The relationship type of r is returned by the query.

Result

type(r)
"KNOWS"
"KNOWS"
2 rows

14.2.3. ID

Returns the id of the relationship or node.

Syntax: ID(property-container)

Arguments:

- *property-container*: A node or a relationship.

Query.

```
START a=node(3, 4, 5)
RETURN ID(a)
```

This returns the node id for three nodes.

Result

ID(a)
3
4
5
3 rows

14.2.4. COALESCE

Returns the first non-null value in the list of expressions passed to it.

Syntax: COALESCE(expression [, expression]*)

Arguments:

- *expression*: The expression that might return null.

Query.

```
START a=node(3)
```

```
RETURN coalesce(a.hairColour?, a.eyes?)
```

Result

coalesce(a.hairColour?, a.eyes?)
"brown"
1 row

14.2.5. HEAD

HEAD returns the first element in a collection.

Syntax: HEAD(expression)

Arguments:

- *expression:* This expression should return a collection of some kind.

Query.

```
START a=node(2)
RETURN a.array, head(a.array)
```

The first node in the path is returned.

Result

a.array	head(a.array)
["one", "two", "three"]	"one"
1 row	

14.2.6. LAST

LAST returns the last element in a collection.

Syntax: LAST(expression)

Arguments:

- *expression:* This expression should return a collection of some kind.

Query.

```
START a=node(2)
RETURN a.array, last(a.array)
```

The last node in the path is returned.

Result

a.array	last(a.array)
["one", "two", "three"]	"three"
1 row	

14.2.7. TIMESTAMP

TIMESTAMP returns the difference, measured in milliseconds, between the current time and midnight, January 1, 1970 UTC. It will return the same value during the whole one query, even if the query is a long running one.

Syntax: `TIMESTAMP()`

Arguments:

Query.

```
START n=node(1)
RETURN timestamp()
```

The time in milliseconds.

Result


timestamp()
1397053501783
1 row

14.3. Collection functions

Collection functions return collections of things — nodes in a path, and so on.

See also [Section 11.1.5, “Collection operators”](#).

Figure 14.3. Graph



14.3.1. NODES

Returns all nodes in a path.

Syntax: NODES(path)

Arguments:

- *path*: A path.

Query.

```

START a=node(3), c=node(2)
MATCH p=a-->b-->c
RETURN NODES(p)
  
```

All the nodes in the path *p* are returned by the example query.

Result

NODES(p)
[Node[3]{name: "A", age: 38, eyes: "brown"}, Node[4]{name: "B", age: 25, eyes: "blue"}, Node[2]{name: "E", age: 41, eyes: "blue", array: ["one", "two", "three"]}]
1 row

14.3.2. RELATIONSHIPS

Returns all relationships in a path.

Syntax: RELATIONSHIPS(path)

Arguments:

- *path*: A path.

Query.

```
START a=node(3), c=node(2)
MATCH p=a-->b-->c
RETURN RELATIONSHIPS(p)
```

All the relationships in the path *p* are returned.

Result

RELATIONSHIPS(p)
[:KNOWS[0] {}, :MARRIED[4] {}]
1 row

14.3.3. EXTRACT

To return a single property, or the value of a function from a collection of nodes or relationships, you can use `EXTRACT`. It will go through a collection, run an expression on every element, and return the results in an collection with these values. It works like the `map` method in functional languages such as Lisp and Scala.

Syntax: EXTRACT(identifier in collection : expression)

Arguments:

- *collection*: An expression that returns a collection
- *identifier*: The closure will have an identifier introduced in it's context. Here you decide which identifier to use.
- *expression*: This expression will run once per value in the collection, and produces the result collection.

Query.

```
START a=node(3), b=node(4), c=node(1)
MATCH p=a-->b-->c
RETURN extract(n in nodes(p) : n.age)
```

The age property of all nodes in the path are returned.

Result

extract(n in nodes(p) : n.age)
[38, 25, 54]
1 row

14.3.4. FILTER

`FILTER` returns all the elements in a collection that comply to a predicate.

Syntax: FILTER(identifier in collection : predicate)

Arguments:

- *collection*: An expression that returns a collection
- *identifier*: This is the identifier that can be used from the predicate.
- *predicate*: A predicate that is tested against all items in the collection.

Query.

```
START a=node(2)
RETURN a.array, filter(x in a.array : length(x) = 3)
```

This returns the property named array and a list of values in it, which have the length 3.

Result

a.array	filter(x in a.array : length(x) = 3)
["one", "two", "three"]	["one", "two"]
1 row	

14.3.5. TAIL

TAIL returns all but the first element in a collection.

Syntax: TAIL(expression)

Arguments:

- *expression*: This expression should return a collection of some kind.

Query.

```
START a=node(2)
RETURN a.array, tail(a.array)
```

This returns the property named array and all elements of that property except the first one.

Result

a.array	tail(a.array)
["one", "two", "three"]	["two", "three"]
1 row	

14.3.6. RANGE

Returns numerical values in a range with a non-zero step value step. Range is inclusive in both ends.

Syntax: RANGE(start, end [, step])

Arguments:

- *start*: A numerical expression.
- *end*: A numerical expression.
- *step*: A numerical expression.

Query.

```
START n=node(1)
RETURN range(0,10), range(2,18,3)
```

Two lists of numbers are returned.

Result

range(0,10)	range(2,18,3)
[0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10]	[2, 5, 8, 11, 14, 17]
1 row	

14.3.7. REDUCE

To run an expression against individual elements of a collection, and store the result of the expression in an accumulator, you can use REDUCE. It will go through a collection, run an expression on every element, storing the partial result in the accumulator. It works like the fold or reduce method in functional languages such as Lisp and Scala.

Syntax: REDUCE(accumulator = initial, identifier in collection : expression)

Arguments:

- *accumulator:* An identifier that will hold the result and the partial results as the collection is iterated
- *initial:* An expression that runs once to give a starting value to the accumulator
- *collection:* An expression that returns a collection
- *identifier:* The closure will have an identifier introduced in it's context. Here you decide which identifier to use.
- *expression:* This expression will run once per value in the collection, and produces the result value.

Query.

```
START a=node(3), b=node(4), c=node(1)
MATCH p=a-->b-->c
RETURN reduce(totalAge = 0, n in nodes(p) : totalAge + n.age)
```

The age property of all nodes in the path are summed and returned as a single value.

Result


reduce(totalAge = 0, n in nodes(p) : totalAge + n.age)
117
1 row

14.4. Mathematical functions

These functions all operate on numerical expressions only, and will return an error if used on any other values.

See also [Section 11.1.1, “Mathematical operators”](#).

Figure 14.4. Graph



14.4.1. ABS

ABS returns the absolute value of a number.

Syntax: ABS(expression)

Arguments:

- *expression*: A numeric expression.

Query.

```
START a=node(3), c=node(2)
RETURN a.age, c.age, abs(a.age - c.age)
```

The absolute value of the age difference is returned.

Result

a.age	c.age	abs(a.age - c.age)
38	41	3.0
1 row		

14.4.2. ROUND

ROUND returns the numerical expression, rounded to the nearest integer.

Syntax: ROUND(expression)

Arguments:

- *expression*: A numerical expression.

Query.

```
START a=node(1)
RETURN round(3.141592)
```

Result

round(3.141592)
3
1 row

14.4.3. SQRT

SQRT returns the square root of a number.

Syntax: SQRT(expression)

Arguments:

- *expression*: A numerical expression

Query.

```
START a=node(1)
RETURN sqrt(256)
```

Result

sqrt(256)
16.0
1 row

14.4.4. SIGN

SIGN returns the signum of a number — zero if the expression is zero, -1 for any negative number, and 1 for any positive number.

Syntax: SIGN(expression)

Arguments:

- *expression*: A numerical expression

Query.

```
START n=node(1)
RETURN sign(-17), sign(0.1)
```

Result


sign(-17)	sign(0.1)
-1.0	1.0
1 row	

14.5. String functions

These functions all operate on string expressions only, and will return an error if used on any other values. Except STR(), which converts to strings.

See also [Section 11.1.4, “String operators”](#).

Figure 14.5. Graph



14.5.1. STR

STR returns a string representation of the expression.

Syntax: STR(expression)

Arguments:

- *expression*: An expression that returns anything

Query.

```
START n=node(1)
RETURN str(1)
```

A string.

Result

str(1)
"1"
1 row

14.5.2. REPLACE

REPLACE returns a string with the search string replaced by the replace string. It replaces all occurrences.

Syntax: REPLACE(original, search, replace)

Arguments:

- *original*: An expression that returns a string
- *search*: An expression that returns a string to search for
- *replace*: An expression that returns the string to replace the search string with

Query.

```
START n=node(1)
RETURN replace("hello", "l", "w")
```

A string.

Result

replace("hello", "l", "w")
"hewwo"
1 row

14.5.3. SUBSTRING

SUBSTRING returns a substring of the original, with a 0-based index start and length. If length is omitted, it returns a substring from start until the end of the string.

Syntax: SUBSTRING(original, start [, length])

Arguments:

- *original*: An expression that returns a string
- *start*: An expression that returns a positive number
- *length*: An expression that returns a positive number

Query.

```
START n=node(1)
RETURN substring("hello", 1, 3), substring("hello", 2)
```

A string.

Result

substring("hello", 1, 3)	substring("hello", 2)
"ell"	"llo"
1 row	

14.5.4. LEFT

LEFT returns a string containing the left n characters of the original string.

Syntax: LEFT(original, length)

Arguments:

- *original*: An expression that returns a string
- *n*: An expression that returns a positive number

Query.

```
START n=node(1)
RETURN left("hello", 3)
```

A String.

Result

left("hello", 3)
"hel"
1 row

14.5.5. RIGHT

RIGHT returns a string containing the right n characters of the original string.

Syntax: RIGHT(original, length)

Arguments:

- *original*: An expression that returns a string
- *n*: An expression that returns a positive number

Query.

```
START n=node(1)
RETURN right("hello", 3)
```

A string.

Result

right("hello", 3)
"llo"
1 row

14.5.6. LTRIM

LTRIM returns the original string with whitespace removed from the left side.

Syntax: LTRIM(original)

Arguments:

- *original*: An expression that returns a string

Query.

```
START n=node(1)
RETURN ltrim(" hello")
```

A string.

Result

ltrim(" hello")
"hello"
1 row

14.5.7. RTRIM

RTRIM returns the original string with whitespace removed from the right side.

Syntax: RTRIM(original)

Arguments:

- *original*: An expression that returns a string

Query.

```
START n=node(1)
RETURN rtrim("hello ")
```

A string.

Result

rtrim("hello ")
"hello"
1 row

14.5.8. TRIM

TRIM returns the original string with whitespace removed from both sides.

Syntax: TRIM(original)

Arguments:

- *original*: An expression that returns a string

Query.

```
START n=node(1)
RETURN trim(" hello ")
```

A string.

Result

trim(" hello ")
"hello"
1 row

14.5.9. LOWER

LOWER returns the original string in lowercase.

Syntax: LOWER(original)

Arguments:

- *original*: An expression that returns a string

Query.

```
START n=node(1)
RETURN lower("HELLO")
```

A string.

Result

lower("HELLO")
"hello"
1 row

14.5.10. UPPER

UPPER returns the original string in uppercase.

Syntax: UPPER(original)

Arguments:

- *original*: An expression that returns a string

Query.

```
START n=node(1)
RETURN upper("hello")
```

A string.

Result

upper("hello")
"HELLO"
1 row

Chapter 15. From SQL to Cypher

This guide is for people who understand SQL. You can use that prior knowledge to quickly get going with Cypher and start exploring Neo4j.

15.1. Start

SQL starts with the result you want — we `SELECT` what we want and then declare how to source it. In Cypher, the `START` clause is quite a different concept which specifies starting points in the graph from which the query will execute.

From a SQL point of view, the identifiers in `START` are like table names that point to a set of nodes or relationships. The set can be listed literally, come via parameters, or as I show in the following example, be defined by an index look-up.

So in fact rather than being `SELECT`-like, the `START` clause is somewhere between the `FROM` and the `WHERE` clause in SQL.

SQL Query.

```
SELECT *
FROM "Person"
WHERE name = 'Anakin'
```

NAME	ID	AGE	HAIR
Anakin	1	20	blonde
1 rows			

Cypher Query.

```
START person=node:Person(name = 'Anakin')
RETURN person
```

person
Node[1]{name:"Anakin", id:1, age:20, hair:"blonde"}
1 row

Cypher allows multiple starting points. This should not be strange from a SQL perspective — every table in the `FROM` clause is another starting point.

15.2. Match

Unlike SQL which operates on sets, Cypher predominantly works on sub-graphs. The relational equivalent is the current set of tuples being evaluated during a SELECT query.

The shape of the sub-graph is specified in the MATCH clause. The MATCH clause is analogous to the JOIN in SQL. A normal $a \rightarrow b$ relationship is an inner join between nodes a and b — both sides have to have at least one match, or nothing is returned.

We'll start with a simple example, where we find all email addresses that are connected to the person "Anakin". This is an ordinary one-to-many relationship.

SQL Query.

```
SELECT "Email".*
FROM "Person"
JOIN "Email" ON "Person".id = "Email".person_id
WHERE "Person".name = 'Anakin'
```

ADDRESS	COMMENT	PERSON_ID
anakin@example.com	home	1
anakin@example.org	work	1

2 rows


Cypher Query.

```
START person=node:Person(name = 'Anakin')
MATCH person-[:email]->email
RETURN email
```

email
Node[7]{address:"anakin@example.com", comment:"home"}
Node[8]{address:"anakin@example.org", comment:"work"}

2 rows

There is no join table here, but if one is necessary the next example will show how to do that, writing the pattern relationship like so: $-[r:\text{belongs_to}]->$ will introduce (the equivalent of) join table available as the variable r . In reality this is a named relationship in Cypher, so we're saying "join Person to Group via belongs_to." To illustrate this, consider this image, comparing the SQL model and Neo4j/Cypher.



And here are example queries:

SQL Query.

```
SELECT "Group".*, "Person_Group".*
FROM "Person"
JOIN "Person_Group" ON "Person".id = "Person_Group".person_id
JOIN "Group" ON "Person_Group".Group_id="Group".id
WHERE "Person".name = 'Bridget'
```

NAME	ID	BELONGS_TO_GROUP_ID	PERSON_ID	GROUP_ID
Admin	4	3	2	4
1 rows				

Cypher Query.

```
START person=node:Person(name = 'Bridget')
MATCH person-[r:belongs_to]->group
RETURN group, r
```

group	r
Node[6]{name:"Admin", id:4}	:belongs_to[0] {}
1 row	

An [outer join](http://www.codinghorror.com/blog/2007/10/a-visual-explanation-of-sql-joins.html) is just as easy. Add a question mark `-[?:KNOWS]->` and it's an optional relationship between nodes — the outer join of Cypher.

Whether it's a left outer join, or a right outer join is defined by which side of the pattern has a starting point. This example is a left outer join, because the bound node is on the left side:

SQL Query.

```
SELECT "Person".name, "Email".address
FROM "Person" LEFT
JOIN "Email" ON "Person".id = "Email".person_id
```

NAME	ADDRESS
Anakin	anakin@example.com
Anakin	anakin@example.org
Bridget	<null>
3 rows	

Cypher Query.

```
START person=node:Person('name: *')
MATCH person-[?:email]->email
RETURN person.name, email.address?
```

person.name	email.address?
"Anakin"	"anakin@example.com"
3 rows	

person.name	email.address?
"Anakin"	"anakin@example.org"
"Bridget"	<null>
3 rows	

Relationships in Neo4j are first class citizens — it's like the SQL tables are pre-joined with each other. So, naturally, Cypher is designed to be able to handle highly connected data easily.

One such domain is tree structures — anyone that has tried storing tree structures in SQL knows that you have to work hard to get around the limitations of the relational model. There are even books on the subject.

To find all the groups and sub-groups that Bridget belongs to, this query is enough in Cypher:

Cypher Query.

```
START person=node:Person('name: Bridget')
MATCH person-[:belongs_to*]->group
RETURN person.name, group.name
```

person.name	group.name
"Bridget"	"Admin"
"Bridget"	"Technician"
"Bridget"	"User"
3 rows	

The * after the relationship type means that there can be multiple hops across belongs_to relationships between group and user. Some SQL dialects have recursive abilities, that allow the expression of queries like this, but you may have a hard time wrapping your head around those. Expressing something like this in SQL is hugely impractical if not practically impossible.

15.3. Where

This is the easiest thing to understand — it's the same animal in both languages. It filters out result sets/subgraphs. Not all predicates have an equivalent in the other language, but the concept is the same.

SQL Query.

```
SELECT *
FROM "Person"
WHERE "Person".age > 35 AND "Person".hair = 'blonde'
```

NAME	ID	AGE	HAIR
Bridget	2	40	blonde
1 rows			

Cypher Query.

```
START person=node:Person('name: *')
WHERE person.age > 35 AND person.hair = 'blonde'
RETURN person
```

person
Node[2]{name:"Bridget", id:2, age:40, hair:"blonde"}
1 row

15.4. Return

This is SQL's `SELECT`. We just put it in the end because it felt better to have it there — you do a lot of matching and filtering, and finally, you return something.

Aggregate queries work just like they do in SQL, apart from the fact that there is no explicit `GROUP BY` clause. Everything in the return clause that is not an aggregate function will be used as the grouping columns.

SQL Query.

```
SELECT "Person".name, count(*)
FROM "Person"
GROUP BY "Person".name
ORDER BY "Person".name
```

NAME	C2
Anakin	1
Bridget	1
2 rows	

Cypher Query.

```
START person=node:Person('name: *')
RETURN person.name, count(*)
ORDER BY person.name
```

person.name	count(*)
"Anakin"	1
"Bridget"	1
2 rows	

Order by is the same in both languages — `ORDER BY expression ASC/DESC`. Nothing weird here.

Part IV. Reference

The reference part is the authoritative source for details on Neo4j usage. It covers details on capabilities, transactions, indexing and queries among other topics.

Chapter 16. Capabilities

16.1. Data Security

Some data may need to be protected from unauthorized access (e.g., theft, modification). Neo4j does not deal with data encryption explicitly, but supports all means built into the Java programming language and the JVM to protect data by encrypting it before storing.

Furthermore, data can be easily secured by running on an encrypted datastore at the file system level. Finally, data protection should be considered in the upper layers of the surrounding system in order to prevent problems with scraping, malicious data insertion, and other threats.

16.2. Data Integrity

In order to keep data consistent, there needs to be mechanisms and structures that guarantee the integrity of all stored data. In Neo4j, data integrity is maintained for the core graph engine together with other data sources - see below.

16.2.1. Core Graph Engine

In Neo4j, the whole data model is stored as a graph on disk and persisted as part of every committed transaction. In the storage layer, Relationships, Nodes, and Properties have direct pointers to each other. This maintains integrity without the need for data duplication between the different backend store files.

16.2.2. Different Data Sources

In a number of scenarios, the core graph engine is combined with other systems in order to achieve optimal performance for non-graph lookups. For example, Apache Lucene is frequently used as an additional index system for text queries that would otherwise be very processing-intensive in the graph layer.

To keep these external systems in synchronization with each other, Neo4j provides full Two Phase Commit transaction management, with rollback support over all data sources. Thus, failed index insertions into Lucene can be transparently rolled back in all data sources and thus keep data up-to-date.

16.3. Data Integration

Most enterprises rely primarily on relational databases to store their data, but this may cause performance limitations. In some of these cases, Neo4j can be used as an extension to supplement search/lookup for faster decision making. However, in any situation where multiple data repositories contain the same data, synchronization can be an issue.

In some applications, it is acceptable for the search platform to be slightly out of sync with the relational database. In others, tight data integrity (eg., between Neo4j and RDBMS) is necessary. Typically, this has to be addressed for data changing in real-time and for bulk data changes happening in the RDBMS.

A few strategies for synchronizing integrated data follows.

16.3.1. Event-based Synchronization

In this scenario, all data stores, both RDBMS and Neo4j, are fed with domain-specific events via an event bus. Thus, the data held in the different backends is not actually synchronized but rather replicated.

16.3.2. Periodic Synchronization

Another viable scenario is the periodic export of the latest changes in the RDBMS to Neo4j via some form of SQL query. This allows a small amount of latency in the synchronization, but has the advantage of using the RDBMS as the master for all data purposes. The same process can be applied with Neo4j as the master data source.

16.3.3. Periodic Full Export/Import of Data

Using the Batch Inserter tools for Neo4j, even large amounts of data can be imported into the database in very short times. Thus, a full export from the RDBMS and import into Neo4j becomes possible. If the propagation lag between the RDBMS and Neo4j is not a big issue, this is a very viable solution.

16.4. Availability and Reliability

Most mission-critical systems require the database subsystem to be accessible at all times. Neo4j ensures availability and reliability through a few different strategies.

16.4.1. Operational Availability

In order not to create a single point of failure, Neo4j supports different approaches which provide transparent fallback and/or recovery from failures.

Online backup (Cold spare)

In this approach, a single instance of the master database is used, with Online Backup enabled. In case of a failure, the backup files can be mounted onto a new Neo4j instance and reintegrated into the application.

Online Backup High Availability (Hot spare)

Here, a Neo4j "backup" instance listens to online transfers of changes from the master. In the event of a failure of the master, the backup is already running and can directly take over the load.

High Availability cluster

This approach uses a cluster of database instances, with one (read/write) master and a number of (read-only) slaves. Failing slaves can simply be restarted and brought back online. Alternatively, a new slave may be added by cloning an existing one. Should the master instance fail, a new master will be elected by the remaining cluster nodes.

16.4.2. Disaster Recovery/ Resiliency

In cases of a breakdown of major part of the IT infrastructure, there need to be mechanisms in place that enable the fast recovery and regrouping of the remaining services and servers. In Neo4j, there are different components that are suitable to be part of a disaster recovery strategy.

Prevention

- Online Backup High Availability to other locations outside the current data center.
- Online Backup to different file system locations: this is a simpler form of backup, applying changes directly to backup files; it is thus more suited for local backup scenarios.
- Neo4j High Availability cluster: a cluster of one write-master Neo4j server and a number of read-slaves, getting transaction logs from the master. Write-master failover is handled by quorum election among the read-slaves for a new master.

Detection

- SNMP and JMX monitoring can be used for the Neo4j database.

Correction

- Online Backup: A new Neo4j server can be started directly on the backed-up files and take over new requests.
- Neo4j High Availability cluster: A broken Neo4j read slave can be reinserted into the cluster, getting the latest updates from the master. Alternatively, a new server can be inserted by copying an existing server and applying the latest updates to it.

16.5. Capacity

16.5.1. File Sizes

Neo4j relies on Java's Non-blocking I/O subsystem for all file handling. Furthermore, while the storage file layout is optimized for interconnected data, Neo4j does not require raw devices. Thus, filesizes are only limited by the underlying operating system's capacity to handle large files. Physically, there is no built-in limit of the file handling capacity in Neo4j.

Neo4j tries to memory-map as much of the underlying store files as possible. If the available RAM is not sufficient to keep all data in RAM, Neo4j will use buffers in some cases, reallocating the memory-mapped high-performance I/O windows to the regions with the most I/O activity dynamically. Thus, ACID speed degrades gracefully as RAM becomes the limiting factor.

16.5.2. Read speed

Enterprises want to optimize the use of hardware to deliver the maximum business value from available resources. Neo4j's approach to reading data provides the best possible usage of all available hardware resources. Neo4j does not block or lock any read operations; thus, there is no danger for deadlocks in read operations and no need for read transactions. With a threaded read access to the database, queries can be run simultaneously on as many processors as may be available. This provides very good scale-up scenarios with bigger servers.

16.5.3. Write speed

Write speed is a consideration for many enterprise applications. However, there are two different scenarios:

1. sustained continuous operation and
2. bulk access (e.g., backup, initial or batch loading).

To support the disparate requirements of these scenarios, Neo4j supports two modes of writing to the storage layer.

In transactional, ACID-compliant normal operation, isolation level is maintained and read operations can occur at the same time as the writing process. At every commit, the data is persisted to disk and can be recovered to a consistent state upon system failures. This requires disk write access and a real flushing of data. Thus, the write speed of Neo4j on a single server in continuous mode is limited by the I/O capacity of the hardware. Consequently, the use of fast SSDs is highly recommended for production scenarios.

Neo4j has a Batch Inserter that operates directly on the store files. This mode does not provide transactional security, so it can only be used when there is a single write thread. Because data is written sequentially, and never flushed to the logical logs, huge performance boosts are achieved. The Batch Inserter is optimized for non-transactional bulk import of large amounts of data.

16.5.4. Data size

In Neo4j, data size is mainly limited by the address space of the primary keys for Nodes, Relationships, Properties and RelationshipTypes. Currently, the address space is as follows:

nodes	2^{35} (~ 34 billion)
relationships	2^{35} (~ 34 billion)

Capabilities

properties	2^{36} to 2^{38} depending on property types (maximum ~ 274 billion, always at least ~ 68 billion)
relationship types	2^{15} (~ 32 000)

Chapter 17. Transaction Management

In order to fully maintain data integrity and ensure good transactional behavior, Neo4j supports the ACID properties:

- atomicity: If any part of a transaction fails, the database state is left unchanged.
- consistency: Any transaction will leave the database in a consistent state.
- isolation: During a transaction, modified data cannot be accessed by other operations.
- durability: The DBMS can always recover the results of a committed transaction.

Specifically:

- All modifications to Neo4j data must be wrapped in transactions.
- The default isolation level is `READ_COMMITTED`.
- Data retrieved by traversals is not protected from modification by other transactions.
- Non-repeatable reads may occur (i.e., only write locks are acquired and held until the end of the transaction).
- One can manually acquire write locks on nodes and relationships to achieve higher level of isolation (`SERIALIZABLE`).
- Locks are acquired at the Node and Relationship level.
- Deadlock detection is built into the core transaction management.

17.1. Interaction cycle

All write operations that work with the graph must be performed in a transaction. Transactions are thread confined and can be nested as “flat nested transactions”. Flat nested transactions means that all nested transactions are added to the scope of the top level transaction. A nested transaction can mark the top level transaction for rollback, meaning the entire transaction will be rolled back. To only rollback changes made in a nested transaction is not possible.

When working with transactions the interaction cycle looks like this:

1. Begin a transaction.
2. Operate on the graph performing write operations.
3. Mark the transaction as successful or not.
4. Finish the transaction.

It is very important to finish each transaction. The transaction will not release the locks or memory it has acquired until it has been finished. The idiomatic use of transactions in Neo4j is to use a try-finally block, starting the transaction and then try to perform the write operations. The last operation in the try block should mark the transaction as successful while the finally block should finish the transaction. Finishing the transaction will perform commit or rollback depending on the success status.



Caution

All modifications performed in a transaction are kept in memory. This means that very large updates have to be split into several top level transactions to avoid running out of memory. It must be a top level transaction since splitting up the work in many nested transactions will just add all the work to the top level transaction.

In an environment that makes use of *thread pooling* other errors may occur when failing to finish a transaction properly. Consider a leaked transaction that did not get finished properly. It will be tied to a thread and when that thread gets scheduled to perform work starting a new (what looks to be a) top level transaction it will actually be a nested transaction. If the leaked transaction state is “marked for rollback” (which will happen if a deadlock was detected) no more work can be performed on that transaction. Trying to do so will result in error on each call to a write operation.

17.2. Isolation levels

By default a read operation will read the last committed value unless a local modification within the current transaction exist. The default isolation level is very similar to `READ_COMMITTED`: reads do not block or take any locks so non-repeatable reads can occur. It is possible to achieve a stronger isolation level (such as `REPEATABLE_READ` and `SERIALIZABLE`) by manually acquiring read and write locks.

17.3. Default locking behavior

- When adding, changing or removing a property on a node or relationship a write lock will be taken on the specific node or relationship.
- When creating or deleting a node a write lock will be taken for the specific node.
- When creating or deleting a relationship a write lock will be taken on the specific relationship and both its nodes.

The locks will be added to the transaction and released when the transaction finishes.

17.4. Deadlocks

Since locks are used it is possible for deadlocks to happen. Neo4j will however detect any deadlock (caused by acquiring a lock) before they happen and throw an exception. Before the exception is thrown the transaction is marked for rollback. All locks acquired by the transaction are still being held but will be released when the transaction is finished (in the finally block as pointed out earlier). Once the locks are released other transactions that were waiting for locks held by the transaction causing the deadlock can proceed. The work performed by the transaction causing the deadlock can then be retried by the user if needed.

Experiencing frequent deadlocks is an indication of concurrent write requests happening in such a way that it is not possible to execute them while at the same time live up to the intended isolation and consistency. The solution is to make sure concurrent updates happen in a reasonable way. For example given two specific nodes (A and B), adding or deleting relationships to both these nodes in random order for each transaction will result in deadlocks when there are two or more transactions doing that concurrently. One solution is to make sure that updates always happens in the same order (first A then B). Another solution is to make sure that each thread/transaction does not have any conflicting writes to a node or relationship as some other concurrent transaction. This can for example be achieved by letting a single thread do all updates of a specific type.



Important

Deadlocks caused by the use of other synchronization than the locks managed by Neo4j can still happen. Since all operations in the Neo4j API are thread safe unless specified otherwise, there is no need for external synchronization. Other code that requires synchronization should be synchronized in such a way that it never performs any Neo4j operation in the synchronized block.

17.5. Delete semantics

When deleting a node or a relationship all properties for that entity will be automatically removed but the relationships of a node will not be removed.



Caution

Neo4j enforces a constraint (upon commit) that all relationships must have a valid start node and end node. In effect this means that trying to delete a node that still has relationships attached to it will throw an exception upon commit. It is however possible to choose in which order to delete the node and the attached relationships as long as no relationships exist when the transaction is committed.

The delete semantics can be summarized in the following bullets:

- All properties of a node or relationship will be removed when it is deleted.
- A deleted node can not have any attached relationships when the transaction commits.
- It is possible to acquire a reference to a deleted relationship or node that has not yet been committed.
- Any write operation on a node or relationship after it has been deleted (but not yet committed) will throw an exception
- After commit trying to acquire a new or work with an old reference to a deleted node or relationship will throw an exception.

17.6. Creating unique nodes

In many use cases, a certain level of uniqueness is desired among entities. You could for instance imagine that only one user with a certain e-mail address may exist in a system. If multiple concurrent threads naively try to create the user, duplicates will be created. There are three main strategies for ensuring uniqueness, and they all work across High Availability and single-instance deployments.

17.6.1. Single thread

By using a single thread, no two threads will even try to create a particular entity simultaneously. On High Availability, an external single-threaded client can perform the operations on the cluster.

17.6.2. Get or create

By using `put-if-absent` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#putIfAbsent%28T,%20java.lang.String,%20java.lang.Object%29>> functionality, entity uniqueness can be guaranteed using an index. Here the index acts as the lock and will only lock the smallest part needed to guaranteed uniqueness across threads and transactions.

See [Section 4.9.1, “Get or create unique node”](#) for how to do this using the core Java API. When using the REST API, see [Section 22.11, “Unique Indexes”](#).

17.6.3. Pessimistic locking



Important

While this is a working solution, please consider using the preferred [Section 17.6.2, “Get or create”](#) instead.

By using explicit, pessimistic locking, unique creation of entities can be achieved in a multi-threaded environment. It is most commonly done by locking on a single or a set of common nodes.

See [Section 4.9.2, “Pessimistic locking for node creation”](#) for how to do this using the core Java API.

17.7. Transaction events

Transaction event handlers can be registered to receive Neo4j Transaction events. Once it has been registered at a `GraphDatabaseService` instance it will receive events about what has happened in each transaction which is about to be committed. Handlers won't get notified about transactions which haven't performed any write operation or won't be committed (either if `Transaction#success()` hasn't been called or the transaction has been marked as failed `Transaction#failure()`). Right before a transaction is about to be committed the `beforeCommit` method is called with the entire diff of modifications made in the transaction. At this point the transaction is still running so changes can still be made. However there's no guarantee that other handlers will see such changes since the order in which handlers are executed is undefined. This method can also throw an exception and will, in such a case, prevent the transaction from being committed (where a call to `afterRollback` will follow). If `beforeCommit` is successfully executed the transaction will be committed and the `afterCommit` method will be called with the same transaction data as well as the object returned from `beforeCommit`. This assumes that all other handlers (if more were registered) also executed `beforeCommit` successfully.

Chapter 18. Data Import

For high-performance data import, the batch insert facilities described in this chapter are recommended.

Other ways to import data into Neo4j include using Gremlin graph import (see [Section 22.18.2, “Load a sample graph”](#)) or using the Geoff notation (see <http://nigelsmall.com/geoff>).

18.1. Batch Insertion

Neo4j has a batch insertion facility intended for initial imports, which bypasses transactions and other checks in favor of performance. This is useful when you have a big dataset that needs to be loaded once.

Batch insertion is included in the [neo4j-kernel](http://search.maven.org/#search|gal|1|neo4j-kernel) component, which is part of all Neo4j distributions and editions.

Be aware of the following points when using batch insertion:

- The intended use is for initial import of data.
- Batch insertion is *not thread safe*.
- Batch insertion is *non-transactional*.
- Unless shutdown is successfully invoked at the end of the import, the database files *will* be corrupt.



Warning

Always perform batch insertion in a *single thread* (or use synchronization to make only one thread at a time access the batch inserter) and invoke `shutdown` when finished.

18.1.1. Batch Inserter Examples

Creating a batch inserter is similar to how you normally create data in the database, but in this case the low-level [BatchInserter](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserter.html) interface is used. As we have already pointed out, you can't have multiple threads using the batch inserter concurrently without external synchronization.



Tip

The source code of the examples is found here: [BatchInsertDocTest.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/kernel/src/test/java/examples/BatchInsertDocTest.java)

To get hold of a `BatchInserter`, use [BatchInserters](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserters.html) and then go from there:

```
BatchInserter inserter = BatchInserters.inserter( "target/batchinserter-example", fileSystem );
Map<String, Object> properties = new HashMap<String, Object>();
properties.put( "name", "Mattias" );
long mattiasNode = inserter.createNode( properties );
properties.put( "name", "Chris" );
long chrisNode = inserter.createNode( properties );
RelationshipType knows = DynamicRelationshipType.withName( "KNOWS" );
// To set properties on the relationship, use a properties map
// instead of null as the last parameter.
inserter.createRelationship( mattiasNode, chrisNode, knows, null );
inserter.shutdown();
```

To gain good performance you probably want to set some configuration settings for the batch inserter. Read [Section 25.9.2, “Batch insert example”](#) for information on configuring a batch inserter. This is how to start a batch inserter with configuration options:

```
Map<String, String> config = new HashMap<String, String>();
config.put( "neostore.nodestore.db.mapped_memory", "90M" );
BatchInserter inserter = BatchInserters.inserter(
    "target/batchinserter-example-config", fileSystem, config );
// Insert data here ... and then shut down:
inserter.shutdown();
```

In case you have stored the configuration in a file, you can load it like this:

```
InputStream input = fileSystem.openAsInputStream(
    new File( "target/batchinsert-config" ) );
Map<String, String> config = MapUtil.load( input );
BatchInserter inserter = BatchInserters.inserter(
    "target/batchinserter-example-config", fileSystem, config );
// Insert data here ... and then shut down:
inserter.shutdown();
```

18.1.2. Batch Graph Database

In case you already have code for data import written against the normal Neo4j API, you could consider using a batch inserter exposing that API.



Note

This will not perform as good as using the `BatchInserter` API directly.

Also be aware of the following:

- Starting a transaction or invoking `Transaction.finish()` or `Transaction.success()` will do nothing.
- Invoking the `Transaction.failure()` method will generate a `NotInTransaction` exception.
- `Node.delete()` and `Node.traverse()` are not supported.
- `Relationship.delete()` is not supported.
- Event handlers and indexes are not supported.
- `GraphDatabaseService.getRelationshipTypes()`, `getAllNodes()` and `getAllRelationships()` are not supported.

With these precautions in mind, this is how to do it:

```
GraphDatabaseService batchDb =
    BatchInserters.batchDatabase( "target/batchdb-example", fileSystem );
Node mattiasNode = batchDb.createNode();
mattiasNode.setProperty( "name", "Mattias" );
Node chrisNode = batchDb.createNode();
chrisNode.setProperty( "name", "Chris" );
RelationshipType knows = DynamicRelationshipType.withName( "KNOWS" );
mattiasNode.createRelationshipTo( chrisNode, knows );
batchDb.shutdown();
```



Tip

The source code of the example is found here: [BatchInsertDocTest.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/kernel/src/test/java/examples/BatchInsertDocTest.java) <<https://github.com/neo4j/neo4j/blob/1.9.7/community/kernel/src/test/java/examples/BatchInsertDocTest.java>>

18.1.3. Index Batch Insertion

For general notes on batch insertion, see [Section 18.1, “Batch Insertion”](#).

Indexing during batch insertion is done using [BatchInserterIndex](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserterIndex.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserterIndex.html>> which are provided via [BatchInserterIndexProvider](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserterIndexProvider.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserterIndexProvider.html>>. An example:

```
BatchInserter inserter = BatchInserters.inserter( "target/neo4jdb-batchinsert" );
BatchInserterIndexProvider indexProvider =
```

```
new LuceneBatchInserterIndexProvider( inserter );
BatchInserterIndex actors =
    indexProvider.nodeIndex( "actors", MapUtil.stringMap( "type", "exact" ) );
actors.setCacheCapacity( "name", 100000 );

Map<String, Object> properties = MapUtil.map( "name", "Keanu Reeves" );
long node = inserter.createNode( properties );
actors.add( node, properties );

//make the changes visible for reading, use this sparsely, requires IO!
actors.flush();

// Make sure to shut down the index provider as well
indexProvider.shutdown();
inserter.shutdown();
```

The configuration parameters are the same as mentioned in [Section 19.10, “Configuration and fulltext indexes”](#).

Best practices

Here are some pointers to get the most performance out of BatchInserterIndex:

- Try to avoid [flushing](#) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserterIndex.html#flush%28%29>> too often because each flush will result in all additions (since last flush) to be visible to the querying methods, and publishing those changes can be a performance penalty.
- Have (as big as possible) phases where one phase is either only writes or only reads, and don't forget to flush after a write phase so that those changes becomes visible to the querying methods.
- Enable [caching](#) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/unsafe/batchinsert/BatchInserterIndex.html#setCacheCapacity%28java.lang.String,%20int%29>> for keys you know you're going to do lookups for later on to increase performance significantly (though insertion performance may degrade slightly).



Note

Changes to the index are available for reading first after they are flushed to disk. Thus, for optimal performance, read and lookup operations should be kept to a minimum during batchinsertion since they involve IO and impact speed negatively.

Chapter 19. Indexing

Indexing in Neo4j can be done in two different ways:

1. The database itself is a *natural index* consisting of its relationships of different types between nodes. For example a tree structure can be layered on top of the data and used for index lookups performed by a traverser.
2. Separate index engines can be used, with [Apache Lucene](http://lucene.apache.org/core/3_6_2/index.html) <http://lucene.apache.org/core/3_6_2/index.html> being the default backend included with Neo4j.

This chapter demonstrate how to use the second type of indexing, focusing on Lucene.

19.1. Introduction

Indexing operations are part of the [Neo4j index API](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/package-summary.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/package-summary.html>>.

Each index is tied to a unique, user-specified name (for example "first_name" or "books") and can index either [nodes](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Node.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Node.html>> or [relationships](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Relationship.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Relationship.html>>.

The default index implementation is provided by the `neo4j-lucene-index` component, which is included in the standard Neo4j download. It can also be downloaded separately from <http://repo1.maven.org/maven2/org/neo4j/neo4j-lucene-index/>. For Maven users, the `neo4j-lucene-index` component has the coordinates `org.neo4j:neo4j-lucene-index` and should be used with the same version of `org.neo4j:neo4j-kernel`. Different versions of the index and kernel components are not compatible in the general case. Both components are included transitively by the `org.neo4j:neo4j:pom` artifact which makes it simple to keep the versions in sync.

For initial import of data using indexes, see [Section 18.1.3, “Index Batch Insertion”](#).



Note

All modifying index operations must be performed inside a transaction, as with any modifying operation in Neo4j.

19.2. Create

An index is created if it doesn't exist when you ask for it. Unless you give it a custom configuration, it will be created with default configuration and backend.

To set the stage for our examples, let's create some indexes to begin with:

```
IndexManager index = graphDb.index();
Index<Node> actors = index.forNodes( "actors" );
Index<Node> movies = index.forNodes( "movies" );
RelationshipIndex roles = index.forRelationships( "roles" );
```

This will create two node indexes and one relationship index with default configuration. See [Section 19.8, “Relationship indexes”](#) for more information specific to relationship indexes.

See [Section 19.10, “Configuration and fulltext indexes”](#) for how to create *fulltext* indexes.

You can also check if an index exists like this:

```
IndexManager index = graphDb.index();
boolean indexExists = index.existsForNodes( "actors" );
```

19.3. Delete

Indexes can be deleted. When deleting, the entire contents of the index will be removed as well as its associated configuration. A new index can be created with the same name at a later point in time.

```
IndexManager index = graphDb.index();
Index<Node> actors = index.forNodes( "actors" );
actors.delete();
```

Note that the actual deletion of the index is made during the commit of *the surrounding transaction*. Calls made to such an index instance after `delete()` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#delete%28%29>> has been called are invalid inside that transaction as well as outside (if the transaction is successful), but will become valid again if the transaction is rolled back.

19.4. Add

Each index supports associating any number of key-value pairs with any number of entities (nodes or relationships), where each association between entity and key-value pair is performed individually. To begin with, let's add a few nodes to the indexes:

```
// Actors
Node reeves = graphDb.createNode();
reeves.setProperty( "name", "Keanu Reeves" );
actors.add( reeves, "name", reeves.getProperty( "name" ) );
Node bellucci = graphDb.createNode();
bellucci.setProperty( "name", "Monica Bellucci" );
actors.add( bellucci, "name", bellucci.getProperty( "name" ) );
// multiple values for a field, in this case for search only
// and not stored as a property.
actors.add( bellucci, "name", "La Bellucci" );
// Movies
Node theMatrix = graphDb.createNode();
theMatrix.setProperty( "title", "The Matrix" );
theMatrix.setProperty( "year", 1999 );
movies.add( theMatrix, "title", theMatrix.getProperty( "title" ) );
movies.add( theMatrix, "year", theMatrix.getProperty( "year" ) );
Node theMatrixReloaded = graphDb.createNode();
theMatrixReloaded.setProperty( "title", "The Matrix Reloaded" );
theMatrixReloaded.setProperty( "year", 2003 );
movies.add( theMatrixReloaded, "title", theMatrixReloaded.getProperty( "title" ) );
movies.add( theMatrixReloaded, "year", 2003 );
Node malena = graphDb.createNode();
malena.setProperty( "title", "Malèna" );
malena.setProperty( "year", 2000 );
movies.add( malena, "title", malena.getProperty( "title" ) );
movies.add( malena, "year", malena.getProperty( "year" ) );
```


Note that there can be multiple values associated with the same entity and key.

Next up, we'll create relationships and index them as well:

```
// we need a relationship type
DynamicRelationshipType ACTS_IN = DynamicRelationshipType.withName( "ACTS_IN" );
// create relationships
Relationship role1 = reeves.createRelationshipTo( theMatrix, ACTS_IN );
role1.setProperty( "name", "Neo" );
roles.add( role1, "name", role1.getProperty( "name" ) );
Relationship role2 = reeves.createRelationshipTo( theMatrixReloaded, ACTS_IN );
role2.setProperty( "name", "Neo" );
roles.add( role2, "name", role2.getProperty( "name" ) );
Relationship role3 = bellucci.createRelationshipTo( theMatrixReloaded, ACTS_IN );
role3.setProperty( "name", "Persephone" );
roles.add( role3, "name", role3.getProperty( "name" ) );
Relationship role4 = bellucci.createRelationshipTo( malena, ACTS_IN );
role4.setProperty( "name", "Malèna Scordia" );
roles.add( role4, "name", role4.getProperty( "name" ) );
```

After these operations, our example graph looks like this:

Figure 19.1. Movie and Actor Graph



19.5. Remove

Removing <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#remove%28T,%20java.lang.String,%20java.lang.Object%29>> from an index is similar to adding, but can be done by supplying one of the following combinations of arguments:

- entity
- entity, key
- entity, key, value

```
// completely remove bellucci from the actors index
actors.remove( bellucci );
// remove any "name" entry of bellucci from the actors index
actors.remove( bellucci, "name" );
// remove the "name" -> "La Bellucci" entry of bellucci
actors.remove( bellucci, "name", "La Bellucci" );
```

19.6. Update



Important

To update an index entry, the old one must be removed and a new one added. For details on removing index entries, see [Section 19.5, “Remove”](#).

Remember that a node or relationship can be associated with any number of key-value pairs in an index. This means that you can index a node or relationship with many key-value pairs that have the same key. In the case where a property value changes and you’d like to update the index, it’s not enough to just index the new value — you’ll have to remove the old value as well.

Here’s a code example that demonstrates how it’s done:

```
// create a node with a property
// so we have something to update later on
Node fishburn = graphDb.createNode();
fishburn.setProperty( "name", "Fishburn" );
// index it
actors.add( fishburn, "name", fishburn.getProperty( "name" ) );
// update the index entry
// when the property value changes
actors.remove( fishburn, "name", fishburn.getProperty( "name" ) );
fishburn.setProperty( "name", "Laurence Fishburn" );
actors.add( fishburn, "name", fishburn.getProperty( "name" ) );
```

19.7. Search

An index can be searched in two ways, [get](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#get%28java.lang.String,%20java.lang.Object%29) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#get%28java.lang.String,%20java.lang.Object%29>> and [query](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#query%28java.lang.String,%20java.lang.Object%29) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html#query%28java.lang.String,%20java.lang.Object%29>>. The `get` method will return exact matches to the given key-value pair, whereas `query` exposes querying capabilities directly from the backend used by the index. For example the [Lucene query syntax](http://lucene.apache.org/core/3_6_2/queryparsersyntax.html) <http://lucene.apache.org/core/3_6_2/queryparsersyntax.html> can be used directly with the default indexing backend.

19.7.1. Get

This is how to search for a single exact match:

```
IndexHits<Node> hits = actors.get( "name", "Keanu Reeves" );
Node reeves = hits.getSingle();
```

[IndexHits](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html>> is an `Iterable` with some additional useful methods. For example [getSingle\(\)](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html#getSingle%28%29) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html#getSingle%28%29>> returns the first and only item from the result iterator, or `null` if there isn't any hit.

Here's how to get a single relationship by exact matching and retrieve its start and end nodes:

```
Relationship persephone = roles.get( "name", "Persephone" ).getSingle();
Node actor = persephone.getStartNode();
Node movie = persephone.getEndNode();
```

Finally, we can iterate over all exact matches from a relationship index:

```
for ( Relationship role : roles.get( "name", "Neo" ) )
{
    // this will give us Reeves twice
    Node reeves = role.getStartNode();
}
```



Important

In case you don't iterate through all the hits, [IndexHits.close\(\)](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html#close%28%29) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html#close%28%29>> must be called explicitly.

19.7.2. Query

There are two query methods, one which uses a key-value signature where the value represents a query for values with the given key only. The other method is more generic and supports querying for more than one key-value pair in the same query.

Here's an example using the key-query option:

```
for ( Node actor : actors.query( "name", "*ex" ) )
{
    // This will return Reeves and Bellucci
}
```

In the following example the query uses multiple keys:

```
for ( Node movie : movies.query( "title:*Matrix* AND year:1999" ) )
{
    // This will return "The Matrix" from 1999 only.
}
```



Note

Beginning a wildcard search with "*" or "?" is discouraged by Lucene, but will nevertheless work.



Caution

You can't have *any whitespace* in the search term with this syntax. See [Section 19.11.3, "Querying with Lucene Query objects"](#) for how to do that.

19.8. Relationship indexes

An index for relationships is just like an index for nodes, extended by providing support to constrain a search to relationships with a specific start and/or end nodes. These extra methods reside in the [RelationshipIndex](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/RelationshipIndex.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/RelationshipIndex.html>> interface which extends [Index<Relationship>](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html>>.

Example of querying a relationship index:

```
// find relationships filtering on start node
// using exact matches
IndexHits<Relationship> reevesAsNeoHits;
reevesAsNeoHits = roles.get( "name", "Neo", reeves, null );
Relationship reevesAsNeo = reevesAsNeoHits.iterator().next();
reevesAsNeoHits.close();
// find relationships filtering on end node
// using a query
IndexHits<Relationship> matrixNeoHits;
matrixNeoHits = roles.query( "name", "*eo", null, theMatrix );
Relationship matrixNeo = matrixNeoHits.iterator().next();
matrixNeoHits.close();
```

And here's an example for the special case of searching for a specific relationship type:

```
// find relationships filtering on end node
// using a relationship type.
// this is how to add it to the index:
roles.add( reevesAsNeo, "type", reevesAsNeo.getType().name() );
// Note that to use a compound query, we can't combine committed
// and uncommitted index entries, so we'll commit before querying:
tx.success();
tx.finish();
// and now we can search for it:
IndexHits<Relationship> typeHits;
typeHits = roles.query( "type:ACTS_IN AND name:Neo", null, theMatrix );
Relationship typeNeo = typeHits.iterator().next();
typeHits.close();
```

Such an index can be useful if your domain has nodes with a very large number of relationships between them, since it reduces the search time for a relationship between two nodes. A good example where this approach pays dividends is in time series data, where we have readings represented as a relationship per occurrence.

19.9. Scores

The `IndexHits` interface exposes [scoring](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html#currentScore%28%29) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/IndexHits.html#currentScore%28%29>> so that the index can communicate scores for the hits. Note that the result is not sorted by the score unless you explicitly specify that. See [Section 19.11.2, “Sorting”](#) for how to sort by score.

```
IndexHits<Node> hits = movies.query( "title", "The*" );
for ( Node movie : hits )
{
    System.out.println( movie.getProperty( "title" ) + " " + hits.currentScore() );
}
```

19.10. Configuration and fulltext indexes

At the time of creation extra configuration can be specified to control the behavior of the index and which backend to use. For example to create a Lucene fulltext index:

```
IndexManager index = graphDb.index();
Index<Node> fulltextMovies = index.forNodes( "movies-fulltext",
    MapUtil.stringMap( IndexManager.PROVIDER, "lucene", "type", "fulltext" ) );
fulltextMovies.add( theMatrix, "title", "The Matrix" );
fulltextMovies.add( theMatrixReloaded, "title", "The Matrix Reloaded" );
// search in the fulltext index
Node found = fulltextMovies.query( "title", "reloAdEd" ).getSingle();
```

Here's an example of how to create an exact index which is case-insensitive:

```
Index<Node> index = graphDb.index().forNodes( "exact-case-insensitive",
    stringMap( "type", "exact", "to_lower_case", "true" ) );
Node node = graphDb.createNode();
index.add( node, "name", "Thomas Anderson" );
assertContains( index.query( "name", "\"Thomas Anderson\"" ), node );
assertContains( index.query( "name", "\"thoMas ANDerson\"" ), node );
```



Tip

In order to search for tokenized words, the query method has to be used. The get method will only match the full string value, not the tokens.

The configuration of the index is persisted once the index has been created. The provider configuration key is interpreted by Neo4j, but any other configuration is passed onto the backend index (e.g. Lucene) to interpret.

Lucene indexing configuration parameters

Parameter	Possible values	Effect
type	exact, fulltext	exact is the default and uses a Lucene keyword analyzer <http://lucene.apache.org/core/3_6_2/api/core/org/apache/lucene/analysis/KeywordAnalyzer.html>. fulltext uses a white-space tokenizer in its analyzer.
to_lower_case	true, false	This parameter goes together with type: fulltext and converts values to lower case during both additions and querying, making the index case insensitive. Defaults to true.
analyzer	the full class name of an Analyzer <http://lucene.apache.org/core/3_6_2/api/core/org/apache/lucene/analysis/Analyzer.html>	Overrides the type so that a custom analyzer can be used. Note: to_lower_case still affects lowercasing of string queries. If the custom analyzer uppercases the indexed tokens, string queries will not match as expected.

19.11. Extra features for Lucene indexes

19.11.1. Numeric ranges

Lucene supports smart indexing of numbers, querying for ranges and sorting such results, and so does its backend for Neo4j. To mark a value so that it is indexed as a numeric value, we can make use of the [ValueContext](http://components.neo4j.org/neo4j-lucene-index/1.9.7/apidocs/org/neo4j/index/lucene/ValueContext.html) <<http://components.neo4j.org/neo4j-lucene-index/1.9.7/apidocs/org/neo4j/index/lucene/ValueContext.html>> class, like this:

```
movies.add( theMatrix, "year-numeric", new ValueContext( 1999 ).indexNumeric() );
movies.add( theMatrixReloaded, "year-numeric", new ValueContext( 2003 ).indexNumeric() );
movies.add( malena, "year-numeric", new ValueContext( 2000 ).indexNumeric() );

int from = 1997;
int to = 1999;
hits = movies.query( QueryContext.numericRange( "year-numeric", from, to ) );
```



Note

The same type must be used for indexing and querying. That is, you can't index a value as a Long and then query the index using an Integer.

By giving null as from/to argument, an open ended query is created. In the following example we are doing that, and have added sorting to the query as well:

```
hits = movies.query(
    QueryContext.numericRange( "year-numeric", from, null )
    .sortNumeric( "year-numeric", false ) );
```

From/to in the ranges defaults to be *inclusive*, but you can change this behavior by using two extra parameters:

```
movies.add( theMatrix, "score", new ValueContext( 8.7 ).indexNumeric() );
movies.add( theMatrixReloaded, "score", new ValueContext( 7.1 ).indexNumeric() );
movies.add( malena, "score", new ValueContext( 7.4 ).indexNumeric() );

// include 8.0, exclude 9.0
hits = movies.query( QueryContext.numericRange( "score", 8.0, 9.0, true, false ) );
```

19.11.2. Sorting

Lucene performs sorting very well, and that is also exposed in the index backend, through the [QueryContext](http://components.neo4j.org/neo4j-lucene-index/1.9.7/apidocs/org/neo4j/index/lucene/QueryContext.html) <<http://components.neo4j.org/neo4j-lucene-index/1.9.7/apidocs/org/neo4j/index/lucene/QueryContext.html>> class:

```
hits = movies.query( "title", new QueryContext( "*" ).sort( "title" ) );
for ( Node hit : hits )
{
    // all movies with a title in the index, ordered by title
}
// or
hits = movies.query( new QueryContext( "title:*" ).sort( "year", "title" ) );
for ( Node hit : hits )
{
    // all movies with a title in the index, ordered by year, then title
}
```

We sort the results by relevance (score) like this:

```
hits = movies.query( "title", new QueryContext( "The*" ).sortByScore() );
for ( Node movie : hits )
{
```

```
// hits sorted by relevance (score)
}
```

19.11.3. Querying with Lucene Query objects

Instead of passing in Lucene query syntax queries, you can instantiate such queries programmatically and pass in as argument, for example:

```
// a TermQuery will give exact matches
Node actor = actors.query( new TermQuery( new Term( "name", "Keanu Reeves" ) ) ).getSingle();
```

Note that the [TermQuery](http://lucene.apache.org/core/3_6_2/api/core/org/apache/lucene/search/TermQuery.html) <http://lucene.apache.org/core/3_6_2/api/core/org/apache/lucene/search/TermQuery.html> is basically the same thing as using the `get` method on the index.

This is how to perform *wildcard* searches using Lucene Query Objects:

```
hits = movies.query( new WildcardQuery( new Term( "title", "The Matrix*" ) ) );
for ( Node movie : hits )
{
    System.out.println( movie.getProperty( "title" ) );
}
```

Note that this allows for whitespace in the search string.

19.11.4. Compound queries

Lucene supports querying for multiple terms in the same query, like so:

```
hits = movies.query( "title:*Matrix* AND year:1999" );
```



Caution

Compound queries can't search across committed index entries and those who haven't got committed yet at the same time.

19.11.5. Default operator

The default operator (that is whether AND or OR is used in between different terms) in a query is OR. Changing that behavior is also done via the [QueryContext](http://components.neo4j.org/neo4j-lucene-index/1.9.7/apidocs/org/neo4j/index/lucene/QueryContext.html) <<http://components.neo4j.org/neo4j-lucene-index/1.9.7/apidocs/org/neo4j/index/lucene/QueryContext.html>> class:

```
QueryContext query = new QueryContext( "title:*Matrix* year:1999" )
    .defaultOperator( Operator.AND );
hits = movies.query( query );
```

19.11.6. Caching

If your index lookups becomes a performance bottle neck, caching can be enabled for certain keys in certain indexes (key locations) to speed up get requests. The caching is implemented with an [LRU](http://en.wikipedia.org/wiki/Cache_algorithms#Least_Recently_Used) <http://en.wikipedia.org/wiki/Cache_algorithms#Least_Recently_Used> cache so that only the most recently accessed results are cached (with "results" meaning a query result of a get request, not a single entity). You can control the size of the cache (the maximum number of results) per index key.

```
Index<Node> index = graphDb.index().forNodes( "actors" );
((LuceneIndex<Node>) index).setCacheCapacity( "name", 300000 );
```



Caution

This setting is not persisted after shutting down the database. This means: set this value after each startup of the database if you want to keep it.

19.12. Automatic Indexing

Neo4j provides a single index for nodes and one for relationships in each database that automatically follow property values as they are added, deleted and changed on database primitives. This functionality is called *auto indexing* and is controlled both from the database configuration Map and through its own API.

19.12.1. Configuration

By default Auto Indexing is off for both Nodes and Relationships. To configure this in the *neo4j.properties* file, use the configuration keys `node_auto_indexing` and `relationship_auto_indexing`. For embedded mode, use the configuration options `GraphDatabaseSettings.node_auto_indexing` and `GraphDatabaseSettings.relationship_auto_indexing`. In both cases, set the value to `true`. This will enable automatic indexing on startup. Just note that we're not done yet, see below!

To actually auto index something, you have to set which properties should get indexed. You do this by listing the property keys to index on. In the configuration file, use the `node_keys_indexable` and `relationship_keys_indexable` configuration keys. When using embedded mode, use the `GraphDatabaseSettings.node_keys_indexable` and `GraphDatabaseSettings.relationship_keys_indexable` configuration keys. In all cases, the value should be a comma separated list of property keys to index on.

When coding in Java, it's done like this:

```

/*
 * Creating the configuration, adding nodeProp1 and nodeProp2 as
 * auto indexed properties for Nodes and relProp1 and relProp2 as
 * auto indexed properties for Relationships. Only those will be
 * indexed. We also have to enable auto indexing for both these
 * primitives explicitly.
 */
GraphDatabaseService graphDb = new GraphDatabaseFactory().
    newEmbeddedDatabaseBuilder( storeDirectory ).
    setConfig( GraphDatabaseSettings.node_keys_indexable, "nodeProp1,nodeProp2" ).
    setConfig( GraphDatabaseSettings.relationship_keys_indexable, "relProp1,relProp2" ).
    setConfig( GraphDatabaseSettings.node_auto_indexing, "true" ).
    setConfig( GraphDatabaseSettings.relationship_auto_indexing, "true" ).
    newGraphDatabase();

Transaction tx = graphDb.beginTx();
Node node1 = null, node2 = null;
Relationship rel = null;
try
{
    // Create the primitives
    node1 = graphDb.createNode();
    node2 = graphDb.createNode();
    rel = node1.createRelationshipTo( node2,
        DynamicRelationshipType.withName( "DYNAMIC" ) );

    // Add indexable and non-indexable properties
    node1.setProperty( "nodeProp1", "nodeProp1Value" );
    node2.setProperty( "nodeProp2", "nodeProp2Value" );
    node1.setProperty( "nonIndexed", "nodeProp2NonIndexedValue" );
    rel.setProperty( "relProp1", "relProp1Value" );
    rel.setProperty( "relPropNonIndexed", "relPropValueNonIndexed" );

    // Make things persistent
    tx.success();
}
catch ( Exception e )

```

```

{
    tx.failure();
}
finally
{
    tx.finish();
}

```

19.12.2. Search

The usefulness of the auto indexing functionality comes of course from the ability to actually query the index and retrieve results. To that end, you can acquire a `ReadableIndex` object from the `AutoIndexer` that exposes all the query and get methods of a full `Index` <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/Index.html>> with exactly the same functionality. Continuing from the previous example, accessing the index is done like this:

```

// Get the Node auto index
ReadableIndex<Node> autoNodeIndex = graphDb.index()
    .getNodeAutoIndexer()
    .getAutoIndex();
// node1 and node2 both had auto indexed properties, get them
assertEquals( node1,
    autoNodeIndex.get( "nodeProp1", "nodeProp1Value" ).getSingle() );
assertEquals( node2,
    autoNodeIndex.get( "nodeProp2", "nodeProp2Value" ).getSingle() );
// node2 also had a property that should be ignored.
assertFalse( autoNodeIndex.get( "nonIndexed",
    "nodeProp2NonIndexedValue" ).hasNext() );

// Get the relationship auto index
ReadableIndex<Relationship> autoRelIndex = graphDb.index()
    .getRelationshipAutoIndexer()
    .getAutoIndex();
// One property was set for auto indexing
assertEquals( rel,
    autoRelIndex.get( "relProp1", "relProp1Value" ).getSingle() );
// The rest should be ignored
assertFalse( autoRelIndex.get( "relPropNonIndexed",
    "relPropValueNonIndexed" ).hasNext() );

```

19.12.3. Runtime Configuration

The same options that are available during database creation via the configuration can also be set during runtime via the `AutoIndexer` API.

Gaining access to the `AutoIndexer` API and adding two `Node` and one `Relationship` properties to auto index is done like so:

```

// Start without any configuration
GraphDatabaseService graphDb = new GraphDatabaseFactory().
    newEmbeddedDatabase( storeDirectory );

// Get the Node AutoIndexer, set nodeProp1 and nodeProp2 as auto
// indexed.
AutoIndexer<Node> nodeAutoIndexer = graphDb.index()
    .getNodeAutoIndexer();
nodeAutoIndexer.startAutoIndexingProperty( "nodeProp1" );
nodeAutoIndexer.startAutoIndexingProperty( "nodeProp2" );

// Get the Relationship AutoIndexer
AutoIndexer<Relationship> relAutoIndexer = graphDb.index()
    .getRelationshipAutoIndexer();
relAutoIndexer.startAutoIndexingProperty( "relProp1" );

```

```
// None of the AutoIndexers are enabled so far. Do that now
nodeAutoIndexer.setEnabled( true );
relAutoIndexer.setEnabled( true );
```



Note

Parameters to the AutoIndexers passed through the Configuration and settings made through the API are cumulative. So you can set some beforehand known settings, do runtime checks to augment the initial configuration and then enable the desired auto indexers - the final configuration is the same regardless of the method used to reach it.

19.12.4. Updating the Automatic Index

Updates to the auto indexed properties happen of course automatically as you update them. Removal of properties from the auto index happens for two reasons. One is that you actually removed the property. The other is that you stopped autoindexing on a property. When the latter happens, any primitive you touch and it has that property, it is removed from the auto index, regardless of any operations on the property. When you start or stop auto indexing on a property, no auto update operation happens currently. If you need to change the set of auto indexed properties and have them re-indexed, you currently have to do this by hand. An example will illustrate the above better:

```
/*
 * Creating the configuration
 */
GraphDatabaseService graphDb = new GraphDatabaseFactory().
    newEmbeddedDatabaseBuilder( storeDirectory ).
    setConfig( GraphDatabaseSettings.node_keys_indexable, "nodeProp1,nodeProp2" ).
    setConfig( GraphDatabaseSettings.node_auto_indexing, "true" ).
    newGraphDatabase();

Transaction tx = graphDb.beginTx();
Node node1 = null, node2 = null, node3 = null, node4 = null;
try
{
    // Create the primitives
    node1 = graphDb.createNode();
    node2 = graphDb.createNode();
    node3 = graphDb.createNode();
    node4 = graphDb.createNode();

    // Add indexable and non-indexable properties
    node1.setProperty( "nodeProp1", "nodeProp1Value" );
    node2.setProperty( "nodeProp2", "nodeProp2Value" );
    node3.setProperty( "nodeProp1", "nodeProp3Value" );
    node4.setProperty( "nodeProp2", "nodeProp4Value" );

    // Make things persistent
    tx.success();
}
catch ( Exception e )
{
    tx.failure();
}
finally
{
    tx.finish();
}
/*
 * Here both nodes are indexed. To demonstrate removal, we stop
```

```

* autoindexing nodeProp1.
*/
AutoIndexer<Node> nodeAutoIndexer = graphDb.index().getNodeAutoIndexer();
nodeAutoIndexer.stopAutoIndexingProperty( "nodeProp1" );

tx = graphDb.beginTx();
try
{
    /*
    * nodeProp1 is no longer auto indexed. It will be
    * removed regardless. Note that node3 will remain.
    */
    node1.setProperty( "nodeProp1", "nodeProp1Value2" );
    /*
    * node2 will be auto updated
    */
    node2.setProperty( "nodeProp2", "nodeProp2Value2" );
    /*
    * remove node4 property nodeProp2 from index.
    */
    node4.removeProperty( "nodeProp2" );
    // Make things persistent
    tx.success();
}
catch ( Exception e )
{
    tx.failure();
}
finally
{
    tx.finish();
}

// Verify
ReadableIndex<Node> nodeAutoIndex = nodeAutoIndexer.getAutoIndex();
// node1 is completely gone
assertFalse( nodeAutoIndex.get( "nodeProp1", "nodeProp1Value" ).hasNext() );
assertFalse( nodeAutoIndex.get( "nodeProp1", "nodeProp1Value2" ).hasNext() );
// node2 is updated
assertFalse( nodeAutoIndex.get( "nodeProp2", "nodeProp2Value" ).hasNext() );
assertEquals( node2,
    nodeAutoIndex.get( "nodeProp2", "nodeProp2Value2" ).getSingle() );
/*
* node3 is still there, despite its nodeProp1 property not being monitored
* any more because it was not touched, in contrast with node1.
*/
assertEquals( node3,
    nodeAutoIndex.get( "nodeProp1", "nodeProp3Value" ).getSingle() );
// Finally, node4 is removed because the property was removed.
assertFalse( nodeAutoIndex.get( "nodeProp2", "nodeProp4Value" ).hasNext() );

```



Caution

If you start the database with auto indexing enabled but different auto indexed properties than the last run, then already auto-indexed properties will be deleted from the index when a value is written to them (assuming the property isn't present in the new configuration). Make sure that the monitored set is what you want before enabling the functionality.

Chapter 20. Graph Algorithms

Neo4j graph algorithms is a component that contains Neo4j implementations of some common algorithms for graphs. It includes algorithms like:

- Shortest paths,
- all paths,
- all simple paths,
- Dijkstra and
- A*.

20.1. Introduction

The graph algorithms are found in the `neo4j-graph-algo` component, which is included in the standard Neo4j download.

- [Javadocs](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphalgo/package-summary.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphalgo/package-summary.html>>
- [Download](http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-graph-algo%22) <<http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-graph-algo%22>>
- [Source code](https://github.com/neo4j/neo4j/tree/1.9.7/community/graph-algo) <<https://github.com/neo4j/neo4j/tree/1.9.7/community/graph-algo>>

For information on how to use `neo4j-graph-algo` as a dependency with Maven and other dependency management tools, see [org.neo4j:neo4j-graph-algo](http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-graph-algo%22) <<http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-graph-algo%22>>. Note that it should be used with the same version of [org.neo4j:neo4j-kernel](http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-kernel%22) <<http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-kernel%22>>. Different versions of the `graph-algo` and `kernel` components are not compatible in the general case. Both components are included transitively by the [org.neo4j:neo4j](http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j%22) <<http://search.maven.org/#search%7Cgav%7C1%7Cg%3A%22org.neo4j%22%20AND%20a%3A%22neo4j%22>> artifact which makes it simple to keep the versions in sync.

The starting point to find and use graph algorithms is [GraphAlgoFactory](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphalgo/GraphAlgoFactory.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphalgo/GraphAlgoFactory.html>>.

For examples, see [Section 4.7, “Graph Algorithm examples”](#) (embedded database) and [Section 22.15, “Built-in Graph Algorithms”](#) (REST API).

Chapter 21. Neo4j Server

21.1. Server Installation

Neo4j can be installed as a server, running either as a headless application or system service.

1. Download the latest release from <http://neo4j.org/download>
 - select the appropriate version for your platform
2. Extract the contents of the archive
 - refer to the top-level extracted directory as `NEO4J_HOME`
3. Use the scripts in the *bin* directory
 - for Linux/macOS, run `$NEO4J_HOME/bin/neo4j start`
 - for Windows, double-click on `%NEO4J_HOME%\bin\Neo4j.bat`
4. Refer to the packaged information in the *doc* directory for details

For information on High Availability, please refer to [Chapter 26, High Availability](#).

21.1.1. As a Windows service

With administrative rights, Neo4j can be installed as a Windows service.

1. Click Start → All Programs → Accessories
2. Right click Command Prompt → Run as Administrator
3. Provide authorization and/or the Administrator password
4. Navigate to `%NEO4J_HOME%`
5. Run `bin\Neo4j.bat install`

To uninstall, run `bin\Neo4j.bat remove` as Administrator.

To query the status of the service, run `bin\Neo4j.bat status`

To start the service from the command prompt, run `bin\Neo4j.bat start`

To stop the service from the command prompt, run `bin\Neo4j.bat stop`



Note

Some users have reported problems on Windows when using the ZoneAlarm firewall. If you are having problems getting large responses from the server, or if Webadmin does not work, try disabling ZoneAlarm. Contact ZoneAlarm support to get information on how to resolve this.

21.1.2. Linux Service

Neo4j can participate in the normal system startup and shutdown process. The following procedure should work on most popular Linux distributions:

1. `cd $NEO4J_HOME`
2. `sudo ./bin/neo4j install`
if asked, enter your password to gain super-user privileges
3. `service neo4j-service status`
should indicate that the server is not running
4. `service neo4j-service start`
will start the server

During installation you will be given the option to select the user Neo4j will run as. You will be asked to supply a username (defaulting to `neo4j`) and if that user is not present on the system it will be created as a system account and the `$NEO4J_HOME/data` directory will be `chown`'ed to that user.

You are encouraged to create a dedicated user for running the service and for that reason it is suggested that you unpack the distribution package under `/opt` or your site specific optional packages directory.

After installation you may have to do some platform specific configuration and performance tuning. For that, refer to [Section 25.11, “Linux specific notes”](#).

To remove the server from the set of startup services, the proper command is

```
service neo4j-service remove
```

+ which will stop the server, if running, and remove it.

Note that if you chose to create a new user account, on `uninstall` you will be prompted to remove it from the system.

21.1.3. Mac OSX

via Homebrew

Using [Homebrew](http://mxcl.github.com/homebrew/) <<http://mxcl.github.com/homebrew/>>, to install the latest stable version of Neo4j Server, issue the following command:

```
brew install neo4j && neo4j start
```

This will get a Neo4j instance running on <http://localhost:7474>. The installation files will reside in `ls /usr/local/Cellar/neo4j/community-{NEO4J_VERSION}/libexec/` — to tweak settings and symlink the database directory if desired.

as a Service

Neo4j can be installed as a Mac `launchd` job:

1. `cd $NEO4J_HOME`
2. `./bin/neo4j install`
3. `launchctl list | grep neo`
should reveal the `launchd "org.neo4j.server.7474"` job for running the Neo4j Server
4. `./bin/neo4j status`
should indicate that the server is running
5. `launchctl stop org.neo4j.server.7474`
should stop the server.
6. `launchctl start org.neo4j.server.7474`
should start the server again.

To remove the `launchctl` service, issue the command

```
./bin/neo4j remove
```

21.1.4. Multiple Server instances on one machine

Neo4j can be set up to run as several instances on one machine, providing for instance several databases for development.

For how to set this up, see [Section 26.6.3, “Alternative setup: Creating a local cluster for testing”](#). Just use the Neo4j edition of your choice, follow the guide and remember to not set the servers to run in HA mode.

21.2. Server Configuration

Quick info

- The server's primary configuration file is found under *conf/neo4j-server.properties*
- The *conf/log4j.properties* file contains the default server logging configuration
- Low-level performance tuning parameters are found in *conf/neo4j.properties*
- Configuration of the daemonizing wrapper are found in *conf/neo4j-wrapper.properties*
- HTTP logging configuration is found in *conf/neo4j-http-logging.xml*

21.2.1. Important server configurations parameters

The main configuration file for the server can be found at *conf/neo4j-server.properties*. This file contains several important settings, and although the defaults are sensible administrators might choose to make changes (especially to the port settings).

Set the location on disk of the database directory like this:

```
org.neo4j.server.database.location=data/graph.db
```



Note

On Windows systems, absolute locations including drive letters need to read *"c:/data/db"*.

Specify the HTTP server port supporting data, administrative, and UI access:

```
org.neo4j.server.webserver.port=7474
```

Specify the client accept pattern for the webserver (default is *127.0.0.1*, localhost only):

```
#allow any client to connect
org.neo4j.server.webserver.address=0.0.0.0
```

For securing the Neo4j Server, see also [Section 28.1, “Securing access to the Neo4j Server”](#)

Set the location of the round-robin database directory which gathers metrics on the running server instance:

```
org.neo4j.server.webadmin.rrd.location=data/graph.db/./rrd
```

Set the URI path for the REST data API through which the database is accessed. This should be a relative path.

```
org.neo4j.server.webadmin.data.uri=/db/data/
```

Setting the management URI for the administration API that the Webadmin tool uses. This should be a relative path.

```
org.neo4j.server.webadmin.management.uri=/db/manage
```

Force the server to use IPv4 network addresses, in *conf/neo4j-wrapper.conf* under the section *Java Additional Parameters* add a new paramter:

```
wrapper.java.additional.3=-Djava.net.preferIPv4Stack=true
```

Specify the number of threads used by the Neo4j Web server to control the level of concurrent HTTP requests that the server will service.

```
org.neo4j.server.webserver.maxthreads=200
```

**Note**

The default value is 10 times the number of CPUs reported available by the JVM.

Low-level performance tuning parameters can be explicitly set by referring to the following property:

```
org.neo4j.server.db.tuning.properties=neo4j.properties
```

If this property isn't set, the server will look for a file called *neo4j.properties* in the same directory as the *neo4j-server.properties* file.

If this property isn't set, and there is no *neo4j.properties* file in the default configuration directory, then the server will log a warning. Subsequently at runtime the database engine will attempt to tune itself based on the prevailing conditions.

21.2.2. Neo4j Database performance configuration

The fine-tuning of the low-level Neo4j graph database engine is specified in a separate properties file, *conf/neo4j.properties*.

The graph database engine has a range of performance tuning options which are enumerated in [Section 21.4, “Server Performance Tuning”](#). Note that other factors than Neo4j tuning should be considered when performance tuning a server, including general server load, memory and file contention, and even garbage collection penalties on the JVM, though such considerations are beyond the scope of this configuration document.

21.2.3. Server logging configuration

Application events within Neo4j server are processed with `java.util.logging` <<http://download.oracle.com/javase/6/docs/technotes/guides/logging/overview.html>> and configured in the file *conf/logging.properties*.

By default it is setup to print INFO level messages both on screen and in a rolling file in *data/log*. Most deployments will choose to use their own configuration here to meet local standards. During development, much useful information can be found in the logs so some form of logging to disk is well worth keeping. On the other hand, if you want to completely silence the console output, set:

```
java.util.logging.ConsoleHandler.level=OFF
```

By default log files are rotated at approximately 10Mb and named consecutively *neo4j.<id>.<rotation sequence #>.log*. To change the naming scheme, rotation frequency and backlog size modify

```
java.util.logging.FileHandler.pattern
java.util.logging.FileHandler.limit
java.util.logging.FileHandler.count
```

respectively to your needs. Details are available at the Javadoc for `java.util.logging.FileHandler` <<http://download.oracle.com/javase/6/docs/api/java/util/logging/FileHandler.html>>.

Apart from log statements originating from the Neo4j server, other libraries report their messages through various frameworks.

Zookeeper is hardwired to use the log4j logging framework. The bundled *conf/log4j.properties* applies for this use only and uses a rolling appender and outputs logs by default to the *data/log* directory.

21.2.4. HTTP logging configuration

As well as logging events happening within the Neo4j server, it is possible to log the HTTP requests and responses that the server consumes and produces. Configuring HTTP logging requires operators to enable and configure the logger and where it will log; and then to optionally configure the log format.



Warning

By default the HTTP logger uses [Common Log Format](http://en.wikipedia.org/wiki/Common_Log_Format) <http://en.wikipedia.org/wiki/Common_Log_Format> meaning that most Web server tooling can automatically consume such logs. In general users should only enable HTTP logging, select an output directory, and if necessary alter the rollover and retention policies.

To enable HTTP logging, edit the *conf/neo4j-server.properties* file to resemble the following:

```
org.neo4j.server.http.log.enabled=true
org.neo4j.server.http.log.config=conf/neo4j-http-logging.xml
```

org.neo4j.server.http.log.enabled=true tells the server that HTTP logging is enabled. HTTP logging can be totally disabled by setting this property to *false*. *org.neo4j.server.http.log.config=conf/neo4j-http-logging.xml* specifies the logging format and rollover policy file that governs how HTTP log output is presented and archived. The defaults provided with Neo4j server uses an hourly log rotation and [Common Log Format](http://en.wikipedia.org/wiki/Common_Log_Format) <http://en.wikipedia.org/wiki/Common_Log_Format>.

If logging is set up to use log files then the server will check that the log file directory exists and is writable. If this check fails, then the server will not startup and will report the failure another available channel like standard out.

21.2.5. Using X-Forwarded-Proto and X-Forwarded-Host to parameterize the base URI for REST responses

There are occasions, for example when you want to host Neo4j server behind a proxy (e.g. one that handles HTTPS traffic), and still have Neo4j respect the base URI of that externally visible proxy.

Ordinarily Neo4j uses the `HOST` header of the HTTP request to construct URIs in its responses. Where a proxy is involved however, this is often undesirable. Instead Neo4j, uses the `X-Forwarded-Host` and `X-Forwarded-Proto` headers provided by proxies to parameterize the URIs in the responses from the databases's REST API. From the outside it looks as if the proxy generated that payload. If an `X-Forwarded-Host` header value contains more than one address (`X-Forwarded-Host` allows comma-and-space separated lists of addresses), Neo4j picks the first, which represents the client request.

In order to take advantage of this functionality, your proxy server must be configured to transmit these headers to the Neo4j server. Failure to transmit both `X-Forwarded-Host` and `X-Forwarded-Proto` headers will result in the original base URI being used.

21.2.6. Other configuration options

Enabling logging from the garbage collector

To get garbage collection logging output you have to pass the corresponding option to the server JVM executable by setting in *conf/neo4j-wrapper.conf* the value

```
wrapper.java.additional.3=-Xloggc:data/log/neo4j-gc.log
```

This line is already present and needs uncommenting. Note also that logging is not directed to console ; You will find the logging statements in *data/log/neo4j-gc.log* or whatever directory you set at the option.

Disabling console types in Webadmin

You may, for security reasons, want to disable the the Neo4j Shell in Webadmin. Shells allow arbitrary code execution, and so they could constitute a security risk if you do not trust all users of your Neo4j Server.

In the *conf/neo4j-server.properties* file:

```
# To disable all shells:  
org.neo4j.server.manage.console_engines=  
  
# To enable only the Neo4j Shell:  
org.neo4j.server.manage.console_engines=shell
```

21.3. Setup for remote debugging

In order to configure the Neo4j server for remote debugging sessions, the Java debugging parameters need to be passed to the Java process through the configuration. They live in the *conf/neo4j-wrapper.properties* file.

In order to specify the parameters, add a line for the additional Java arguments like this:

```
# Java Additional Parameters
wrapper.java.additional.1=-Dorg.neo4j.server.properties=conf/neo4j-server.properties
wrapper.java.additional.2=-Dlog4j.configuration=file:conf/log4j.properties
wrapper.java.additional.3=-agentlib:jdwp=transport=dt_socket,server=y,suspend=n,address=5005 \
-Xdebug-Xnoagent-Djava.compiler=NONE\
-Xrunjdwp:transport=dt_socket,server=y,suspend=n,address=5005
```

This configuration will start a Neo4j server ready for remote debugging attachment at localhost and port 5005. Use these parameters to attach to the process from Eclipse, IntelliJ or your remote debugger of choice after starting the server.

21.4. Server Performance Tuning

At the heart of the Neo4j server is a regular Neo4j storage engine instance. That engine can be tuned in the same way as the other embedded configurations, using the same file format. The only difference is that the server must be told where to find the fine-tuning configuration.

Quick info

- The `neo4j.properties` file is a standard configuration file that databases load in order to tune their memory use and caching strategies.
- See [Section 25.4, “Caches in Neo4j”](#) for more information.

21.4.1. Specifying Neo4j tuning properties

The `conf/neo4j-server.properties` file in the server distribution, is the main configuration file for the server. In this file we can specify a second properties file that contains the database tuning settings (that is, the `neo4j.properties` file). This is done by setting a single property to point to a valid `neo4j.properties` file:

```
org.neo4j.server.db.tuning.properties={neo4j.properties file}
```

On restarting the server the tuning enhancements specified in the `neo4j.properties` file will be loaded and configured into the underlying database engine.

21.4.2. Specifying JVM tuning properties

Tuning the standalone server is achieved by editing the `neo4j-wrapper.conf` file in the `conf` directory of `NEO4J_HOME`.

Edit the following properties:

neo4j-wrapper.conf JVM tuning properties

Property Name	Meaning
<code>wrapper.java.initmemory</code>	initial heap size (in MB)
<code>wrapper.java.maxmemory</code>	maximum heap size (in MB)
<code>wrapper.java.additional.N</code>	additional literal JVM parameter, where N is a number for each

For more information on the tuning properties, see [Section 25.6, “JVM Settings”](#).

21.5. Server Installation in the Cloud

Neo4j on various cloud services either by a user, or as a managed instance on the Neo Technology cloud fabric. Below are instructions for some of these.

21.5.1. Heroku

For the basic setup, please see the [Heroku Quickstart tutorial](https://devcenter.heroku.com/articles/quickstart) <https://devcenter.heroku.com/articles/quickstart>.

To add Neo4j to your Heroku app, do:

```
heroku addons:add neo4j
```

Chapter 22. REST API

The Neo4j REST API is designed with discoverability in mind, so that you can start with a GET on the [Section 22.1, “Service root”](#) and from there discover URIs to perform other requests. The examples below uses URIs in the examples; they are subject to change in the future, so for future-proofness *discover URIs where possible*, instead of relying on the current layout. The default representation is `json` <<http://www.json.org/>>, both for responses and for data sent with POST/PUT requests.

Below follows a listing of ways to interact with the REST API. For language bindings to the REST API, see [Chapter 5, *Neo4j Remote Client Libraries*](#).

To interact with the JSON interface you must explicitly set the request header `Accept:application/json` for those requests that responds with data. You should also set the header `Content-Type:application/json` if your request sends data, for example when you’re creating a relationship. The examples include the relevant request and response headers.


The server supports streaming results, with better performance and lower memory overhead. See [Section 22.2, “Streaming”](#) for more information.

22.1. Service root

22.1.1. Get service root

The service root is your starting point to discover the REST API. It contains the basic starting points for the database, and some version and extension information. The `reference_node` entry will only be present if there is a reference node set and that node actually exists in the database.

Figure 22.1. Final Graph



Example request

- GET `http://localhost:7474/db/data/`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
{
  "extensions" : {
  },
  "node" : "http://localhost:7474/db/data/node",
  "reference_node" : "http://localhost:7474/db/data/node/184",
  "node_index" : "http://localhost:7474/db/data/index/node",
  "relationship_index" : "http://localhost:7474/db/data/index/relationship",
  "extensions_info" : "http://localhost:7474/db/data/ext",
  "relationship_types" : "http://localhost:7474/db/data/relationship/types",
  "batch" : "http://localhost:7474/db/data/batch",
  "cypher" : "http://localhost:7474/db/data/cypher",
  "neo4j_version" : "1.9.7"
}
```

22.2. Streaming


The whole REST API can be transmitted as JSON streams, resulting in better performance and lower memory overhead on the server side. To use it, adjust the request headers for every call, see the example below for details.



Caution

This feature is new, and you should make yourself comfortable with the streamed response style versus the non-streamed API where results are delivered in a single large response. Expect future releases to have streaming enabled by default since it is a far more efficient mechanism for both client and server.

Figure 22.2. Final Graph



Example request

- GET `http://localhost:7474/db/data/`
- Accept: `application/json`
- X-Stream: `true`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8; stream=true`

```

{
  "extensions" : {
  },
  "node" : "http://localhost:7474/db/data/node",
  "reference_node" : "http://localhost:7474/db/data/node/186",
  "node_index" : "http://localhost:7474/db/data/index/node",
  "relationship_index" : "http://localhost:7474/db/data/index/relationship",
  "extensions_info" : "http://localhost:7474/db/data/ext",
  "relationship_types" : "http://localhost:7474/db/data/relationship/types",
  "batch" : "http://localhost:7474/db/data/batch",
  "cypher" : "http://localhost:7474/db/data/cypher",
  "neo4j_version" : "1.9.7"
}
  
```

22.3. Cypher queries

The Neo4j REST API allows querying with Cypher, see [Part III, “Cypher Query Language”](#). The results are returned as a list of string headers (columns), and a data part, consisting of a list of all rows, every row consisting of a list of REST representations of the field value — Node, Relationship, Path or any simple value like String.



Tip


In order to speed up queries in repeated scenarios, try not to use literals but replace them with parameters wherever possible in order to let the server cache query plans, see [Section 22.3.1, “Send queries with parameters”](#) for details.

22.3.1. Send queries with parameters

Cypher supports queries with parameters which are submitted as a JSON map.

```
START x = node:node_auto_index(name={startName})
MATCH path = (x-[r]-friend)
WHERE friend.name = {name}
RETURN TYPE(r)
```

Figure 22.3. Final Graph



Example request

- POST <http://localhost:7474/db/data/cypher>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "query" : "start x = node:node_auto_index(name={startName}) match path = (x-[r]-friend) where friend.name = {name} return TYPE(r)",
  "params" : {
    "startName" : "I",
    "name" : "you"
  }
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8


```
{
  "columns" : [ "TYPE(r)" ],
  "data" : [ [ "know" ] ]
}
```

22.3.2. Send a Query

A simple query returning all nodes connected to node 1, returning the node and the name property, if it exists, otherwise null:

```
START x = node(144)
MATCH x -[r]-> n
RETURN type(r), n.name?, n.age?
```

Figure 22.4. Final Graph



Example request

- POST `http://localhost:7474/db/data/cypher`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "query" : "start x = node(144) match x -[r]-> n return type(r), n.name?, n.age?",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`


```
{
  "columns" : [ "type(r)", "n.name?", "n.age?" ],
  "data" : [ [ "know", "him", 25 ], [ "know", "you", null ] ]
}
```

22.3.3. Return paths

Paths can be returned together with other return types by just specifying returns.

```
START x = node(148)
MATCH path = (x--friend)
RETURN path, friend.name
```

Figure 22.5. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/cypher`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "query" : "start x = node(148) match path = (x--friend) return path, friend.name",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`


```
{
  "columns" : [ "path", "friend.name" ],
  "data" : [ [ {
    "start" : "http://localhost:7474/db/data/node/148",
    "nodes" : [ "http://localhost:7474/db/data/node/148", "http://localhost:7474/db/data/node/147" ],
    "length" : 1,
    "relationships" : [ "http://localhost:7474/db/data/relationship/75" ],
    "end" : "http://localhost:7474/db/data/node/147"
  }, "you" ] ]
}
```

22.3.4. Nested results

When sending queries that return nested results like list and maps, these will get serialized into nested JSON representations according to their types.

```
START n = node(146,145)
RETURN collect(n.name)
```

Figure 22.6. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/cypher`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "query" : "start n = node(146,145) return collect(n.name)",
  "params" : {
  }
}

```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

{
  "columns" : [ "collect(n.name)" ],
  "data" : [ [ [ "I", "you" ] ] ]
}

```

22.3.5. Server errors


Errors on the server will be reported as a JSON-formatted stacktrace and message.

```

START x = node(129)
RETURN x.dummy

```

Figure 22.7. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/cypher`
- Accept: `application/json; charset=UTF-8`

- Content-Type: application/json

```
{
  "query" : "start x = node(129) return x.dummy",
  "params" : {
  }
}
```

Example response

- 400: Bad Request
- Content-Type: application/json; charset=UTF-8

```
{
  "message": "The property \u0027dummy\u0027 does not exist on Node[129]",
  "exception": "EntityNotFoundException",
  "fullName": "org.neo4j.cypher.EntityNotFoundException",
  "stacktrace": [
    "org.neo4j.cypher.internal.helpers.MapSupport$PropertyContainerMap$$anonfun$apply$1.apply(MapSupport.scala:64)",
    "org.neo4j.cypher.internal.helpers.MapSupport$PropertyContainerMap$$anonfun$apply$1.apply(MapSupport.scala:64)",
    "scala.Option.getOrElse(Option.scala:120)",
    "org.neo4j.cypher.internal.helpers.MapSupport$PropertyContainerMap.apply(MapSupport.scala:64)",
    "org.neo4j.cypher.internal.helpers.MapSupport$PropertyContainerMap.apply(MapSupport.scala:52)",
    "org.neo4j.cypher.internal.commands.expressions.Property.apply(Property.scala:31)",
    "org.neo4j.cypher.internal.pipes.ExtractPipe$$anonfun$internalCreateResults$1$$anonfun$apply$1.apply(ExtractPipe.scala:47)",
    "org.neo4j.cypher.internal.pipes.ExtractPipe$$anonfun$internalCreateResults$1$$anonfun$apply$1.apply(ExtractPipe.scala:45)",
    "scala.collection.immutable.Map$Map1.foreach(Map.scala:109)",
    "org.neo4j.cypher.internal.pipes.ExtractPipe$$anonfun$internalCreateResults$1.apply(ExtractPipe.scala:45)",
    "org.neo4j.cypher.internal.pipes.ExtractPipe$$anonfun$internalCreateResults$1.apply(ExtractPipe.scala:44)",
    "scala.collection.Iterator$$anon$11.next(Iterator.scala:328)",
    "scala.collection.Iterator$$anon$11.next(Iterator.scala:328)",
    "org.neo4j.cypher.internal.ClosingIterator$$anonfun$next$1.apply(ClosingIterator.scala:44)",
    "org.neo4j.cypher.internal.ClosingIterator.failIfThrows(ClosingIterator.scala:86)",
    "org.neo4j.cypher.internal.ClosingIterator.next(ClosingIterator.scala:43)",
    "org.neo4j.cypher.PipeExecutionResult.next(PipeExecutionResult.scala:159)",
    "org.neo4j.cypher.PipeExecutionResult.next(PipeExecutionResult.scala:33)",
    "scala.collection.Iterator$$anon$11.next(Iterator.scala:328)",
    "scala.collection.convert.Wrappers$IteratorWrapper.next(Wrappers.scala:30)",
    "org.neo4j.cypher.PipeExecutionResult$$anon$1.next(PipeExecutionResult.scala:75)",
    "org.neo4j.helpers.collection.ExceptionHandlingIterable$1.next(ExceptionHandlingIterable.java:67)",
    "org.neo4j.helpers.collection.IteratorWrapper.next(IteratorWrapper.java:47)",
    "org.neo4j.server.rest.repr.ListRepresentation.serialize(ListRepresentation.java:64)",
    "org.neo4j.server.rest.repr.Serializer.serialize(Serializer.java:75)",
    "org.neo4j.server.rest.repr.MappingSerializer.putList(MappingSerializer.java:61)",
    "org.neo4j.server.rest.repr.CypherResultRepresentation.serialize(CypherResultRepresentation.java:57)",
    "org.neo4j.server.rest.repr.MappingRepresentation.serialize(MappingRepresentation.java:41)",
    "org.neo4j.server.rest.repr.OutputFormat.assemble(OutputFormat.java:180)",
    "org.neo4j.server.rest.repr.OutputFormat.formatRepresentation(OutputFormat.java:132)",
    "org.neo4j.server.rest.repr.OutputFormat.response(OutputFormat.java:119)",
    "org.neo4j.server.rest.repr.OutputFormat.ok(OutputFormat.java:57)",
    "org.neo4j.server.rest.web.CypherService.cypher(CypherService.java:95)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

22.4. Property values

The REST API allows setting properties on nodes and relationships through direct RESTful operations. However, there are restrictions as to what types of values can be used as property values. Allowed value types are as follows:

- Numbers: Both integer values, with capacity as Java's Long type, and floating points, with capacity as Java's Double.
- Booleans
- Strings
- Arrays: Of the above basic types

22.4.1. Arrays

There are two important points to be made about array values. First, all values in the array must be of the same type. That means either all integers, all floats, all booleans or all strings. Mixing types is not currently supported.

Second, storing empty arrays is only possible given certain preconditions. Because the JSON transfer format does not contain type information for arrays, type is inferred from the values in the array. If the array is empty, the Neo4j Server cannot determine the type. In these cases, it will check if an array is already stored for the given property, and will use the stored arrays type when storing the empty array. If no array exists already, the server will reject the request.

22.5. Nodes

22.5.1. Create node

Figure 22.8. Final Graph



Example request

- POST `http://localhost:7474/db/data/node`
- Accept: `application/json; charset=UTF-8`


Example response

- 201: Created
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/node/75`

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/75/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/75/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/75/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/75/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/75/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/75/relationships/all",
  "self" : "http://localhost:7474/db/data/node/75",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/75/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/75/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/75/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/75/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/75/relationships",
  "data" : {
  }
}
```

22.5.2. Create node with properties

Figure 22.9. Final Graph



Example request

- POST `http://localhost:7474/db/data/node`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "foo" : "bar"
}
```

Example response

- 201: Created
- Content-Length: 1108
- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/node/71

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/71/paged/traverse/{returnType}?pageSize,leaseTime",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/71/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/71/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/71/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/71/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/71/relationships/all",
  "self" : "http://localhost:7474/db/data/node/71",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/71/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/71/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/71/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/71/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/71/relationships",
  "data" : {
    "foo" : "bar"
  }
}
```

22.5.3. Get node

Note that the response contains URI/templates for the available operations for getting properties and relationships.

Figure 22.10. Final Graph*Example request*

- GET http://localhost:7474/db/data/node/6
- Accept: application/json; charset=UTF-8

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/6/paged/traverse/{returnType}?pageSize,leaseTime",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/6/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/6/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/6/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/6/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/6/relationships/all",
  "self" : "http://localhost:7474/db/data/node/6",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/6/relationships/out/{-list|&|types}",
}
```

```

"properties" : "http://localhost:7474/db/data/node/6/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/6/relationships/in",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/6/relationships/in/{-list|&|types}",
"create_relationship" : "http://localhost:7474/db/data/node/6/relationships",
"data" : {
}
}

```

22.5.4. Get non-existent node

Figure 22.11. Final Graph



Example request

- GET `http://localhost:7474/db/data/node/1000000`
- Accept: `application/json; charset=UTF-8`

Example response

- 404: Not Found
- Content-Type: `application/json; charset=UTF-8`

```

{
  "message": "Cannot find node with id [1000000] in database.",
  "exception": "NodeNotFoundException",
  "fullName": "org.neo4j.server.rest.web.NodeNotFoundException",
  "stacktrace": [
    "org.neo4j.server.rest.web.DatabaseActions.node(DatabaseActions.java:156)",
    "org.neo4j.server.rest.web.DatabaseActions.getNode(DatabaseActions.java:212)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.getNode(RestfulGraphDatabase.java:237)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}

```

22.5.5. Delete node

Figure 22.12. Final Graph

Example request


- DELETE `http://localhost:7474/db/data/node/72`
- Accept: `application/json; charset=UTF-8`

Example response

- 204: No Content

22.5.6. Nodes with relationships can not be deleted

The relationships on a node has to be deleted before the node can be deleted.

Figure 22.13. Final Graph*Example request*

- DELETE <http://localhost:7474/db/data/node/79>
- Accept: application/json; charset=UTF-8

Example response

- 409: Conflict
- Content-Type: application/json; charset=UTF-8

```
{
  "message": "The node with id 79 cannot be deleted. Check that the node is orphaned before deletion.",
  "exception": "OperationFailureException",
  "fullname": "org.neo4j.server.rest.web.OperationFailureException",
  "stacktrace": [
    "org.neo4j.server.rest.web.DatabaseActions.deleteNode(DatabaseActions.java:233)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.deleteNode(RestfulGraphDatabase.java:251)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

22.6. Relationships

Relationships are a first class citizen in the Neo4j REST API. They can be accessed either stand-alone or through the nodes they are attached to.


The general pattern to get relationships from a node is:

```
GET http://localhost:7474/db/data/node/123/relationships/{dir}/{-list|&|types}
```

Where `dir` is one of `all`, `in`, `out` and `types` is an ampersand-separated list of types. See the examples below for more information.

22.6.1. Get Relationship by ID

Figure 22.14. Final Graph



Example request

- GET `http://localhost:7474/db/data/relationship/211`
- Accept: `application/json; charset=UTF-8`

Example response


- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/331",
  "property" : "http://localhost:7474/db/data/relationship/211/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/211",
  "properties" : "http://localhost:7474/db/data/relationship/211/properties",
  "type" : "know",
  "end" : "http://localhost:7474/db/data/node/330",
  "data" : {
  }
}
```

22.6.2. Create relationship

Upon successful creation of a relationship, the new relationship is returned.

Figure 22.15. Final Graph

*Example request*

- POST <http://localhost:7474/db/data/node/2/relationships>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "to" : "http://localhost:7474/db/data/node/1",
  "type" : "LOVES"
}
```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: <http://localhost:7474/db/data/relationship/1>

```
{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/2",
  "property" : "http://localhost:7474/db/data/relationship/1/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/1",
  "properties" : "http://localhost:7474/db/data/relationship/1/properties",
  "type" : "LOVES",
  "end" : "http://localhost:7474/db/data/node/1",
  "data" : {
  }
}
```

22.6.3. Create a relationship with properties

Upon successful creation of a relationship, the new relationship is returned.

Figure 22.16. Starting Graph





Figure 22.17. Final Graph

*Example request*

- POST <http://localhost:7474/db/data/node/12/relationships>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "to" : "http://localhost:7474/db/data/node/11",
  "type" : "LOVES",
  "data" : {
    "foo" : "bar"
  }
}
```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: <http://localhost:7474/db/data/relationship/8>

```
{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/12",
  "property" : "http://localhost:7474/db/data/relationship/8/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/8",
  "properties" : "http://localhost:7474/db/data/relationship/8/properties",
  "type" : "LOVES",
  "end" : "http://localhost:7474/db/data/node/11",
  "data" : {
    "foo" : "bar"
  }
}
```

22.6.4. Delete relationship

Figure 22.18. Starting Graph




Figure 22.19. Final Graph



Example request


- DELETE <http://localhost:7474/db/data/relationship/204>
- Accept: application/json; charset=UTF-8

Example response

- 204: No Content

22.6.5. Get all properties on a relationship

Figure 22.20. Final Graph



Example request

- GET <http://localhost:7474/db/data/relationship/207/properties>
- Accept: application/json; charset=UTF-8

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
{
  "since" : "1day",
  "cost" : "high"
}
```

22.6.6. Set all properties on a relationship

Figure 22.21. Starting Graph





Figure 22.22. Final Graph



Example request

- PUT `http://localhost:7474/db/data/relationship/212/properties`
- Accept: application/json; charset=UTF-8
- Content-Type: application/json


```
{
  "happy" : false
}
```

Example response

- 204: No Content

22.6.7. Get single property on a relationship

Figure 22.23. Final Graph



Example request

- GET `http://localhost:7474/db/data/relationship/208/properties/cost`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

"high"

22.6.8. Set single property on a relationship

Figure 22.24. Starting Graph





Figure 22.25. Final Graph




Example request

- PUT <http://localhost:7474/db/data/relationship/206/properties/cost>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
"deadly"
```

Example response

- 204: No Content

22.6.9. Get all relationships*Figure 22.26. Final Graph**Example request*

- GET <http://localhost:7474/db/data/node/55/relationships/all>
- Accept: application/json; charset=UTF-8

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "start" : "http://localhost:7474/db/data/node/55",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/27",
  "property" : "http://localhost:7474/db/data/relationship/27/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/27/properties",
  "type" : "LIKES",
  "extensions" : {
  },
  "end" : "http://localhost:7474/db/data/node/56"
}, {
```


```

"start" : "http://localhost:7474/db/data/node/57",
"data" : {
},
"self" : "http://localhost:7474/db/data/relationship/28",
"property" : "http://localhost:7474/db/data/relationship/28/properties/{key}",
"properties" : "http://localhost:7474/db/data/relationship/28/properties",
"type" : "LIKES",
"extensions" : {
},
"end" : "http://localhost:7474/db/data/node/55"
}, {
"start" : "http://localhost:7474/db/data/node/55",
"data" : {
},
"self" : "http://localhost:7474/db/data/relationship/29",
"property" : "http://localhost:7474/db/data/relationship/29/properties/{key}",
"properties" : "http://localhost:7474/db/data/relationship/29/properties",
"type" : "HATES",
"extensions" : {
},
"end" : "http://localhost:7474/db/data/node/58"
} ]

```

22.6.10. Get incoming relationships

Figure 22.27. Final Graph



Example request

- GET `http://localhost:7474/db/data/node/74/relationships/in`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

[ {
"start" : "http://localhost:7474/db/data/node/76",
"data" : {
},
"self" : "http://localhost:7474/db/data/relationship/39",

```


```

"property" : "http://localhost:7474/db/data/relationship/39/properties/{key}",
"properties" : "http://localhost:7474/db/data/relationship/39/properties",
"type" : "LIKES",
"extensions" : {
},
"end" : "http://localhost:7474/db/data/node/74"
} ]

```

22.6.11. Get outgoing relationships

Figure 22.28. Final Graph



Example request

- GET <http://localhost:7474/db/data/node/99/relationships/out>
- Accept: application/json; charset=UTF-8

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```

[ {
  "start" : "http://localhost:7474/db/data/node/99",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/53",
  "property" : "http://localhost:7474/db/data/relationship/53/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/53/properties",
  "type" : "LIKES",
  "extensions" : {
  },
  "end" : "http://localhost:7474/db/data/node/100"
}, {
  "start" : "http://localhost:7474/db/data/node/99",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/55",
  "property" : "http://localhost:7474/db/data/relationship/55/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/55/properties",
  "type" : "HATES",
  "extensions" : {
  }
} ]

```

```


},
"end" : "http://localhost:7474/db/data/node/102"
} ]

```

22.6.12. Get typed relationships

Note that the "&" needs to be encoded like "%26" for example when using [cURL](http://curl.haxx.se/) <http://curl.haxx.se/> from the terminal.

Figure 22.29. Final Graph



Example request

- GET <http://localhost:7474/db/data/node/30/relationships/all/LIKES&HATES>
- Accept: application/json; charset=UTF-8

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```

[ {
  "start" : "http://localhost:7474/db/data/node/30",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/12",
  "property" : "http://localhost:7474/db/data/relationship/12/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/12/properties",
  "type" : "LIKES",
  "extensions" : {
  },
  "end" : "http://localhost:7474/db/data/node/31"
}, {
  "start" : "http://localhost:7474/db/data/node/32",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/13",
  "property" : "http://localhost:7474/db/data/relationship/13/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/13/properties",
  "type" : "LIKES",
  "extensions" : {
  },

```


```

"end" : "http://localhost:7474/db/data/node/30"
}, {
"start" : "http://localhost:7474/db/data/node/30",
"data" : {
},
"self" : "http://localhost:7474/db/data/relationship/14",
"property" : "http://localhost:7474/db/data/relationship/14/properties/{key}",
"properties" : "http://localhost:7474/db/data/relationship/14/properties",
"type" : "HATES",
"extensions" : {
},
"end" : "http://localhost:7474/db/data/node/33"
} ]

```

22.6.13. Get relationships on a node without relationships

Figure 22.30. Final Graph



Example request

- GET <http://localhost:7474/db/data/node/88/relationships/all>
- Accept: application/json; charset=UTF-8

Example response


- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ ]
```

22.7. Relationship types

22.7.1. Get relationship types

Figure 22.31. Final Graph



Example request

- GET `http://localhost:7474/db/data/relationship/types`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json`


```
[ "knowledge", "foo", "know", "has", "related-to", "knows", "some type", "HATES", "tested-together", "FRIENDS", "likes", "KNOWS", "d
```

22.8. Node properties

22.8.1. Set property on node

Setting different properties will retain the existing ones for this node. Note that a single value are submitted not as a map but just as a value (which is valid JSON) like in the example below.

Figure 22.32. Final Graph



Example request

- PUT `http://localhost:7474/db/data/node/311/properties/foo`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
"bar"
```


Example response

- 204: No Content

22.8.2. Update node properties

This will replace all existing properties on the node with the new set of attributes.

Figure 22.33. Final Graph



Example request

- PUT `http://localhost:7474/db/data/node/309/properties`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "age" : "18"
}
```

Example response

- 204: No Content

22.8.3. Get properties for node

Figure 22.34. Final Graph



Example request

- GET `http://localhost:7474/db/data/node/2/properties`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
{
  "foo" : "bar"
}
```

22.8.4. Property values can not be null

This example shows the response you get when trying to set a property to null.

Example request

- POST `http://localhost:7474/db/data/node`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "foo" : null
}
```

Example response

- 400: Bad Request
- Content-Type: `application/json; charset=UTF-8`

```
{
  "message": "Could not set property \"foo\", unsupported type: null",
  "exception": "PropertyValueException",
  "fullname": "org.neo4j.server.rest.web.PropertyValueException",
  "stacktrace": [
    "org.neo4j.server.rest.domain.PropertySettingStrategy.setProperty(PropertySettingStrategy.java:151)",
    "org.neo4j.server.rest.domain.PropertySettingStrategy.setProperties(PropertySettingStrategy.java:92)",
    "org.neo4j.server.rest.web.DatabaseActions.createNode(DatabaseActions.java:191)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.createNode(RestfulGraphDatabase.java:207)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

22.8.5. Property values can not be nested

Nesting properties is not supported. You could for example store the nested JSON as a string instead.

Example request


- POST http://localhost:7474/db/data/node/
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "foo" : {
    "bar" : "baz"
  }
}
```

Example response

- 400: Bad Request
- Content-Type: application/json; charset=UTF-8

```
{
  "message": "Could not set property \"foo\", unsupported type: {bar\u003dbaz}",
  "exception": "PropertyValueException",
  "fullname": "org.neo4j.server.rest.web.PropertyValueException",
  "stacktrace": [
    "org.neo4j.server.rest.domain.PropertySettingStrategy.setProperty(PropertySettingStrategy.java:151)",
    "org.neo4j.server.rest.domain.PropertySettingStrategy.setProperties(PropertySettingStrategy.java:92)",
    "org.neo4j.server.rest.web.DatabaseActions.createNode(DatabaseActions.java:191)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.createNode(RestfulGraphDatabase.java:207)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

22.8.6. Delete all properties from node*Figure 22.35. Final Graph**Example request*

- DELETE http://localhost:7474/db/data/node/181/properties
- Accept: application/json; charset=UTF-8

Example response

- 204: No Content

22.8.7. Delete a named property from a node

To delete a single property from a node, see the example below.



Figure 22.36. Starting Graph

Figure 22.37. Final Graph



Example request

- DELETE `http://localhost:7474/db/data/node/180/properties/name`
- Accept: `application/json; charset=UTF-8`


Example response

- 204: No Content

22.9. Relationship properties

22.9.1. Update relationship properties

Figure 22.38. Final Graph



Example request

- PUT <http://localhost:7474/db/data/relationship/64/properties>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json


```
{  
  "jim" : "tobias"  
}
```

Example response

- 204: No Content

22.9.2. Remove properties from a relationship

Figure 22.39. Final Graph



Example request

- DELETE <http://localhost:7474/db/data/relationship/202/properties>
- Accept: application/json; charset=UTF-8

Example response

- 204: No Content

22.9.3. Remove property from a relationship

See the example request below.

Figure 22.40. Starting Graph





Figure 22.41. Final Graph



Example request

- DELETE `http://localhost:7474/db/data/relationship/205/properties/cost`
- Accept: `application/json; charset=UTF-8`


Example response

- 204: No Content

22.9.4. Remove non-existent property from a relationship

Attempting to remove a property that doesn't exist results in an error.

Figure 22.42. Final Graph



Example request

- DELETE http://localhost:7474/db/data/relationship/203/properties/non-existent
- Accept: application/json; charset=UTF-8

Example response


- 404: Not Found
- Content-Type: application/json; charset=UTF-8

```
{
  "message": "Relationship[203] does not have a property \"non-existent\"",
  "exception": "NoSuchPropertyException",
  "fullname": "org.neo4j.server.rest.web.NoSuchPropertyException",
  "stacktrace": [
    "org.neo4j.server.rest.web.DatabaseActions.removeRelationshipProperty(DatabaseActions.java:645)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.deleteRelationshipProperty(RestfulGraphDatabase.java:620)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

22.9.5. Remove properties from a non-existing relationship

Attempting to remove all properties from a relationship which doesn't exist results in an error.

Figure 22.43. Final Graph

*Example request*

- DELETE http://localhost:7474/db/data/relationship/1234/properties
- Accept: application/json; charset=UTF-8

Example response

- 404: Not Found
- Content-Type: application/json; charset=UTF-8


```
{
  "exception": "RelationshipNotFoundException",
  "fullname": "org.neo4j.server.rest.web.RelationshipNotFoundException",
  "stacktrace": [
    "org.neo4j.server.rest.web.DatabaseActions.relationship(DatabaseActions.java:170)",
    "org.neo4j.server.rest.web.DatabaseActions.removeAllRelationshipProperties(DatabaseActions.java:633)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.deleteAllRelationshipProperties(RestfulGraphDatabase.java:599)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

```
]
}
```

22.9.6. Remove property from a non-existing relationship

Attempting to remove a property from a relationship which doesn't exist results in an error.

Figure 22.44. Final Graph



Example request

- DELETE `http://localhost:7474/db/data/relationship/1234/properties/cost`
- Accept: `application/json; charset=UTF-8`

Example response

- 404: Not Found
- Content-Type: `application/json; charset=UTF-8`

```
{
  "exception": "RelationshipNotFoundException",
  "fullname": "org.neo4j.server.rest.web.RelationshipNotFoundException",
  "stacktrace": [
    "org.neo4j.server.rest.web.DatabaseActions.relationship(DatabaseActions.java:170)",
    "org.neo4j.server.rest.web.DatabaseActions.removeRelationshipProperty(DatabaseActions.java:639)",
    "org.neo4j.server.rest.web.RestfulGraphDatabase.deleteRelationshipProperty(RestfulGraphDatabase.java:620)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

22.10. Indexes

An index can contain either nodes or relationships.



Note

To create an index with default configuration, simply start using it by adding nodes/relationships to it. It will then be automatically created for you.

What default configuration means depends on how you have configured your database. If you haven't changed any indexing configuration, it means the indexes will be using a Lucene-based backend.

All the examples below show you how to do operations on node indexes, but all of them are just as applicable to relationship indexes. Simple change the "node" part of the URL to "relationship".

If you want to customize the index settings, see [Section 22.10.2, “Create node index with configuration”](#).

22.10.1. Create node index



Note

Instead of creating the index this way, you can simply start to use it, and it will be created automatically with default configuration.

Example request

- POST `http://localhost:7474/db/data/index/node/`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "name" : "favorites"
}
```

Example response

- 201: Created
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/index/node/favorites/`

```
{
  "template" : "http://localhost:7474/db/data/index/node/favorites/{key}/{value}"
}
```

22.10.2. Create node index with configuration

This request is only necessary if you want to customize the index settings. If you are happy with the defaults, you can just start indexing nodes/relationships, as non-existent indexes will automatically be created as you do. See [Section 19.10, “Configuration and fulltext indexes”](#) for more information on index configuration.

Example request

- POST `http://localhost:7474/db/data/index/node/`
- Accept: `application/json; charset=UTF-8`

- Content-Type: application/json

```
{
  "name" : "fulltext",
  "config" : {
    "type" : "fulltext",
    "provider" : "lucene"
  }
}
```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/index/node/fulltext/

```
{
  "template" : "http://localhost:7474/db/data/index/node/fulltext/{key}/{value}",
  "type" : "fulltext",
  "provider" : "lucene"
}
```

22.10.3. Delete node index

Example request

- DELETE http://localhost:7474/db/data/index/node/kvnode
- Accept: application/json; charset=UTF-8

Example response

- 204: No Content

22.10.4. List node indexes

Example request

- GET http://localhost:7474/db/data/index/node/
- Accept: application/json; charset=UTF-8

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
{
  "favorites" : {
    "template" : "http://localhost:7474/db/data/index/node/favorites/{key}/{value}",
    "provider" : "lucene",
    "type" : "exact"
  }
}
```

22.10.5. Add node to index

Associates a node with the given key/value pair in the given index.

**Note**

Spaces in the URI have to be encoded as %20.

**Caution**

This does **not** overwrite previous entries. If you index the same key/value/item combination twice, two index entries are created. To do update-type operations, you need to delete the old entry before adding a new one.

Example request

- POST `http://localhost:7474/db/data/index/node/favorites`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "value" : "some value",
  "uri" : "http://localhost:7474/db/data/node/109",
  "key" : "some-key"
}
```

Example response

- 201: Created
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/index/node/favorites/some-key/some%20value/109`

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/109/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/109/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/109/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/109/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/109/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/109/relationships/all",
  "self" : "http://localhost:7474/db/data/node/109",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/109/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/109/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/109/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/109/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/109/relationships",
  "data" : {
  },
  "indexed" : "http://localhost:7474/db/data/index/node/favorites/some-key/some%20value/109"
}
```

22.10.6. Remove all entries with a given node from an index

Example request

- DELETE `http://localhost:7474/db/data/index/node/kvnode/114`
- Accept: `application/json; charset=UTF-8`

Example response

- 204: No Content

22.10.7. Remove all entries with a given node and key from an index

Example request

- DELETE `http://localhost:7474/db/data/index/node/kvnode/kvkey2/116`
- Accept: `application/json; charset=UTF-8`

Example response

- 204: No Content

22.10.8. Remove all entries with a given node, key and value from an index

Example request

- DELETE `http://localhost:7474/db/data/index/node/kvnode/kvkey1/value1/110`
- Accept: `application/json; charset=UTF-8`

Example response

- 204: No Content

22.10.9. Find node by exact match



Note

Spaces in the URI have to be encoded as %20.

Example request

- GET `http://localhost:7474/db/data/index/node/favorites/key/the%2520value`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "indexed" : "http://localhost:7474/db/data/index/node/favorites/key/the%2520value/123",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/123/relationships/out",
  "data" : {
  },
  "traverse" : "http://localhost:7474/db/data/node/123/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/123/relationships/all/{"-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/123/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/123",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/123/relationships/out/{"-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/123/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/123/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/123/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/123/paged/traverse/{returnType}{?pageSize,leaseTime}",
```

```
"all_relationships" : "http://localhost:7474/db/data/node/123/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/123/relationships/in/{-list|&|types}"
} ]
```

22.10.10. Find node by query

The query language used here depends on what type of index you are querying. The default index type is Lucene, in which case you should use the Lucene query language here. Below an example of a fuzzy search over multiple keys.

See: http://lucene.apache.org/core/3_6_2/queryparsersyntax.html

Getting the results with a predefined ordering requires adding the parameter

order=ordering

where ordering is one of index, relevance or score. In this case an additional field will be added to each result, named score, that holds the float value that is the score reported by the query result.

Example request

- GET `http://localhost:7474/db/data/index/node/bobTheIndex?query=Name:Build~0.1%20AND%20Gender:Male`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/122/relationships/out",
  "data" : {
    "Name" : "Builder"
  },
  "traverse" : "http://localhost:7474/db/data/node/122/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/122/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/122/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/122",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/122/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/122/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/122/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/122/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/122/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/122/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/122/relationships/in/{-list|&|types}"
} ]
```

22.11. Unique Indexes

For uniqueness enforcements, there are two modes:

- URL Parameter `uniqueness=get_or_create`: Create a new node/relationship and index it if no existing one can be found. If an existing node/relationship is found, discard the sent data and return the existing node/relationship.
- URL Parameter `uniqueness=create_or_fail`: Create a new node/relationship if no existing one can be found in the index. If an existing node/relationship is found, return a conflict error.

For more information, see [Section 17.6, “Creating unique nodes”](#).

22.11.1. Get or create unique node (create)

The node is created if it doesn't exist in the unique index already.

Example request

- POST `http://localhost:7474/db/data/index/node/people?uniqueness=get_or_create`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "key" : "name",
  "value" : "Tobias",
  "properties" : {
    "name" : "Tobias",
    "sequence" : 1
  }
}
```

Example response

- 201: Created
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/index/node/people/name/Tobias/121`

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/121/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/121/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/121/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/121/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/121/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/121/relationships/all",
  "self" : "http://localhost:7474/db/data/node/121",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/121/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/121/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/121/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/121/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/121/relationships",
  "data" : {
    "sequence" : 1,
    "name" : "Tobias"
  },
  "indexed" : "http://localhost:7474/db/data/index/node/people/name/Tobias/121"
}
```

22.11.2. Get or create unique node (existing)

Here, a node is not created but the existing unique node returned, since another node is indexed with the same data already. The node data returned is then that of the already existing node.

Example request

- POST `http://localhost:7474/db/data/index/node/people?uniqueness=get_or_create`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "key" : "name",
  "value" : "Peter",
  "properties" : {
    "name" : "Peter",
    "sequence" : 2
  }
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/index/node/people/name/Peter/113`

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/113/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/113/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/113/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/113/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/113/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/113/relationships/all",
  "self" : "http://localhost:7474/db/data/node/113",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/113/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/113/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/113/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/113/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/113/relationships",
  "data" : {
    "sequence" : 1,
    "name" : "Peter"
  },
  "indexed" : "http://localhost:7474/db/data/index/node/people/name/Peter/113"
}
```

22.11.3. Create a unique node or return fail (create)

Here, in case of an already existing node, an error should be returned. In this example, no existing indexed node is found and a new node is created.

Example request

- POST `http://localhost:7474/db/data/index/node/people?uniqueness=create_or_fail`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
```

```

"key" : "name",
"value" : "Tobias",
"properties" : {
  "name" : "Tobias",
  "sequence" : 1
}
}

```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/index/node/people/name/Tobias/120

```

{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/120/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/120/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/120/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/120/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/120/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/120/relationships/all",
  "self" : "http://localhost:7474/db/data/node/120",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/120/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/120/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/120/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/120/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/120/relationships",
  "data" : {
    "sequence" : 1,
    "name" : "Tobias"
  },
  "indexed" : "http://localhost:7474/db/data/index/node/people/name/Tobias/120"
}

```

22.11.4. Create a unique node or return fail (fail)

Here, in case of an already existing node, an error should be returned. In this example, an existing node indexed with the same data is found and an error is returned.

Example request

- POST http://localhost:7474/db/data/index/node/people?uniqueness=create_or_fail
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```

{
  "key" : "name",
  "value" : "Peter",
  "properties" : {
    "name" : "Peter",
    "sequence" : 2
  }
}

```

Example response

- 409: Conflict
- Content-Type: application/json; charset=UTF-8

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/107/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/107/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/107/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/107/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/107/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/107/relationships/all",
  "self" : "http://localhost:7474/db/data/node/107",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/107/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/107/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/107/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/107/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/107/relationships",
  "data" : {
    "sequence" : 1,
    "name" : "Peter"
  },
  "indexed" : "http://localhost:7474/db/data/index/node/people/name/Peter/107"
}
```

22.11.5. Add an existing node to unique index (not indexed)

Associates a node with the given key/value pair in the given unique index.

In this example, we are using `create_or_fail` uniqueness.

Example request

- POST `http://localhost:7474/db/data/index/node/favorites?uniqueness=create_or_fail`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "value" : "some value",
  "uri" : "http://localhost:7474/db/data/node/117",
  "key" : "some-key"
}
```

Example response

- 201: Created
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/index/node/favorites/some-key/some%20value/117`

```
{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/117/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/117/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/117/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/117/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/117/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/117/relationships/all",
  "self" : "http://localhost:7474/db/data/node/117",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/117/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/117/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/117/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/117/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/117/relationships",
  "data" : {
  }
}
```

```

},
"indexed" : "http://localhost:7474/db/data/index/node/favorites/some-key/some%20value/117"
}

```

22.11.6. Add an existing node to unique index (already indexed)

In this case, the node already exists in the index, and thus we get a HTTP 409 status response, as we have set the uniqueness to create_or_fail.

Example request

- POST http://localhost:7474/db/data/index/node/favorites?uniqueness=create_or_fail
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```

{
  "value" : "some value",
  "uri" : "http://localhost:7474/db/data/node/119",
  "key" : "some-key"
}

```

Example response

- 409: Conflict
- Content-Type: application/json; charset=UTF-8

```

{
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/119/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/119/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/119/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/119/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/119/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/119/relationships/all",
  "self" : "http://localhost:7474/db/data/node/119",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/119/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/119/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/119/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/119/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/119/relationships",
  "data" : {
  },
  "indexed" : "http://localhost:7474/db/data/index/node/favorites/some-key/some%20value/119"
}

```

22.11.7. Get or create unique relationship (create)

Create a unique relationship in an index. If a relationship matching the given key and value already exists in the index, it will be returned. If not, a new relationship will be created.



Note

The type and direction of the relationship is not regarded when determining uniqueness.

Example request

- POST http://localhost:7474/db/data/index/relationship/MyIndex/?uniqueness=get_or_create
- Accept: application/json; charset=UTF-8

- Content-Type: application/json

```
{
  "key" : "name",
  "value" : "Tobias",
  "start" : "http://localhost:7474/db/data/node/161",
  "end" : "http://localhost:7474/db/data/node/162",
  "type" : "knowledge"
}
```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/index/relationship/MyIndex/name/Tobias/81

```
{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/161",
  "property" : "http://localhost:7474/db/data/relationship/81/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/81",
  "properties" : "http://localhost:7474/db/data/relationship/81/properties",
  "type" : "knowledge",
  "end" : "http://localhost:7474/db/data/node/162",
  "data" : {
    "name" : "Tobias"
  },
  "indexed" : "http://localhost:7474/db/data/index/relationship/MyIndex/name/Tobias/81"
}
```

22.11.8. Get or create unique relationship (existing)

Here, in case of an already existing relationship, the sent data is ignored and the existing relationship returned.

Example request

- POST http://localhost:7474/db/data/index/relationship/rels?uniqueness=get_or_create
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "key" : "name",
  "value" : "Peter",
  "start" : "http://localhost:7474/db/data/node/165",
  "end" : "http://localhost:7474/db/data/node/166",
  "type" : "KNOWS"
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/163",
  "property" : "http://localhost:7474/db/data/relationship/82/properties/{key}",
}
```

```

"self" : "http://localhost:7474/db/data/relationship/82",
"properties" : "http://localhost:7474/db/data/relationship/82/properties",
"type" : "KNOWS",
"end" : "http://localhost:7474/db/data/node/164",
"data" : {
},
"indexed" : "http://localhost:7474/db/data/index/relationship/rels/name/Peter/82"
}

```

22.11.9. Create a unique relationship or return fail (create)

Here, in case of an already existing relationship, an error should be returned. In this example, no existing relationship is found and a new relationship is created.

Example request

- POST http://localhost:7474/db/data/index/relationship/rels?uniqueness=create_or_fail
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```

{
  "key" : "name",
  "value" : "Tobias",
  "start" : "http://localhost:7474/db/data/node/173",
  "end" : "http://localhost:7474/db/data/node/174",
  "type" : "KNOWS"
}

```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: <http://localhost:7474/db/data/index/relationship/rels/name/Tobias/85>

```

{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/173",
  "property" : "http://localhost:7474/db/data/relationship/85/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/85",
  "properties" : "http://localhost:7474/db/data/relationship/85/properties",
  "type" : "KNOWS",
  "end" : "http://localhost:7474/db/data/node/174",
  "data" : {
    "name" : "Tobias"
  },
  "indexed" : "http://localhost:7474/db/data/index/relationship/rels/name/Tobias/85"
}

```

22.11.10. Create a unique relationship or return fail (fail)

Here, in case of an already existing relationship, an error should be returned. In this example, an existing relationship is found and an error is returned.

Example request

- POST http://localhost:7474/db/data/index/relationship/rels?uniqueness=create_or_fail
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "key" : "name",
  "value" : "Peter",
  "start" : "http://localhost:7474/db/data/node/157",
  "end" : "http://localhost:7474/db/data/node/158",
  "type" : "KNOWS"
}
```

Example response

- 409: Conflict
- Content-Type: application/json; charset=UTF-8

```
{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/155",
  "property" : "http://localhost:7474/db/data/relationship/79/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/79",
  "properties" : "http://localhost:7474/db/data/relationship/79/properties",
  "type" : "KNOWS",
  "end" : "http://localhost:7474/db/data/node/156",
  "data" : {
  },
  "indexed" : "http://localhost:7474/db/data/index/relationship/rels/name/Peter/79"
}
```

22.11.11. Add an existing relationship to a unique index (not indexed)

If a relationship matching the given key and value already exists in the index, it will be returned. If not, an HTTP 409 (conflict) status will be returned in this case, as we are using `create_or_fail`.

It's possible to use `get_or_create` uniqueness as well.

**Note**

The type and direction of the relationship is not regarded when determining uniqueness.

Example request

- POST `http://localhost:7474/db/data/index/relationship/rels?uniqueness=create_or_fail`
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "key" : "name",
  "value" : "Peter",
  "uri" : "http://localhost:7474/db/data/relationship/78"
}
```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: `http://localhost:7474/db/data/index/relationship/rels/name/Peter/78`

```
{
  "extensions" : {
  },

```

```

"start" : "http://localhost:7474/db/data/node/153",
"property" : "http://localhost:7474/db/data/relationship/78/properties/{key}",
"self" : "http://localhost:7474/db/data/relationship/78",
"properties" : "http://localhost:7474/db/data/relationship/78/properties",
"type" : "KNOWS",
"end" : "http://localhost:7474/db/data/node/154",
"data" : {
},
"indexed" : "http://localhost:7474/db/data/index/relationship/rels/name/Peter/78"
}

```

22.11.12. Add an existing relationship to a unique index (already indexed)

Example request

- POST `http://localhost:7474/db/data/index/relationship/rels?uniqueness=create_or_fail`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "key" : "name",
  "value" : "Peter",
  "uri" : "http://localhost:7474/db/data/relationship/80"
}

```

Example response

- 409: Conflict
- Content-Type: `application/json; charset=UTF-8`

```

{
  "extensions" : {
  },
  "start" : "http://localhost:7474/db/data/node/159",
  "property" : "http://localhost:7474/db/data/relationship/80/properties/{key}",
  "self" : "http://localhost:7474/db/data/relationship/80",
  "properties" : "http://localhost:7474/db/data/relationship/80/properties",
  "type" : "KNOWS",
  "end" : "http://localhost:7474/db/data/node/160",
  "data" : {
  },
  "indexed" : "http://localhost:7474/db/data/index/relationship/rels/name/Peter/80"
}

```

22.12. Automatic Indexes

To enable automatic indexes in neo4j, set up the database for that, see [Section 19.12.1, “Configuration”](#). With this feature enabled, you can then index and query nodes in these indexes.

22.12.1. Find node by exact match from an automatic index

Automatic index nodes can be found via exact lookups with normal Index REST syntax.

Example request

- GET `http://localhost:7474/db/data/index/auto/node/name/I`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/133/relationships/out",
  "data" : {
    "name" : "I"
  },
  "traverse" : "http://localhost:7474/db/data/node/133/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/133/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/133/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/133",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/133/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/133/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/133/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/133/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/133/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/133/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/133/relationships/in/{-list|&|types}"
} ]
```

22.12.2. Find node by query from an automatic index

See Find node by query for the actual query syntax.

Example request

- GET `http://localhost:7474/db/data/index/auto/node/?query=name:I`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/124/relationships/out",
  "data" : {
    "name" : "I"
  },
  "traverse" : "http://localhost:7474/db/data/node/124/traverse/{returnType}",
}
```

```
"all_typed_relationships" : "http://localhost:7474/db/data/node/124/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/124/properties/{key}",
"self" : "http://localhost:7474/db/data/node/124",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/124/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/124/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/124/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/124/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/124/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/124/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/124/relationships/in/{-list|&|types}"
} ]
```

22.13. Configurable Automatic Indexing

Out of the box auto-indexing supports exact matches since they are created with the default configuration (see [Section 19.12, “Automatic Indexing”](#)) the first time you access them. However it is possible to intervene in the lifecycle of the server before any auto indexes are created to change their configuration.



Warning

This approach *cannot* be used on databases that already have auto-indexes established. To change the auto-index configuration existing indexes would have to be deleted first, so be careful!



Caution

This technique works, but it is not particularly pleasant. Future versions of Neo4j may remove this loophole in favour of a better structured feature for managing auto-indexing configurations.

Auto-indexing must be enabled through configuration before we can create or configure them. Firstly ensure that you’ve added some config like this into your server’s `neo4j.properties` file:

```
node_auto_indexing=true
relationship_auto_indexing=true
node_keys_indexable=name,phone
relationship_keys_indexable=since
```

The `node_auto_indexing` and `relationship_auto_indexing` settings turn auto-indexing on for nodes and relationships respectively. The `node_keys_indexable` key allows you to specify a comma-separated list of node property keys to be indexed. The `relationship_keys_indexable` does the same for relationship property keys.

Next start the server as usual by invoking the start script as described in [Section 21.1, “Server Installation”](#).

Next we have to pre-empt the creation of an auto-index, by telling the server to create an apparently manual index which has the same name as the node (or relationship) auto-index. For example, in this case we’ll create a node auto index whose name is `node_auto_index`, like so:

22.13.1. Create an auto index for nodes with specific configuration

Example request

- POST `http://localhost:7474/db/data/index/node/`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "name" : "node_auto_index",
  "config" : {
    "type" : "fulltext",
    "provider" : "lucene"
  }
}
```

Example response

- 201: Created

- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/index/node/node_auto_index/

```
{
  "template" : "http://localhost:7474/db/data/index/node/node_auto_index/{key}/{value}",
  "type" : "fulltext",
  "provider" : "lucene"
}
```

If you require configured auto-indexes for relationships, the approach is similar:

22.13.2. Create an auto index for relationships with specific configuration

Example request

- POST http://localhost:7474/db/data/index/relationship/
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "name" : "relationship_auto_index",
  "config" : {
    "type" : "fulltext",
    "provider" : "lucene"
  }
}
```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/index/relationship/relationship_auto_index/

```
{
  "template" : "http://localhost:7474/db/data/index/relationship/relationship_auto_index/{key}/{value}",
  "type" : "fulltext",
  "provider" : "lucene"
}
```

In case you're curious how this works, on the server side it triggers the creation of an index which happens to have the same name as the auto index that the database would create for itself. Now when we interact with the database, the index thinks the index is already created so the state machine skips over that step and just gets on with normal day-to-day auto-indexing.



Caution

You have to do this early in your server lifecycle, before any normal auto indexes are created.

There are a few REST calls providing a REST interface to the [AutoIndexer](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/index/AutoIndexer.html) component. The following REST calls work both, for node and relationship by simply changing the respective part of the URL.

22.13.3. Get current status for autoindexing on nodes

Example request

- GET `http://localhost:7474/db/data/index/auto/node/status`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
false
```

22.13.4. Enable node autoindexing

Example request

- PUT `http://localhost:7474/db/data/index/auto/node/status`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
true
```

Example response

- 204: No Content

22.13.5. Lookup list of properties being autoindexed

Example request

- GET `http://localhost:7474/db/data/index/auto/node/properties`
- Accept: `application/json; charset=UTF-8`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ "some-property" ]
```

22.13.6. Add a property for autoindexing on nodes

Example request

- POST `http://localhost:7474/db/data/index/auto/node/properties`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
myProperty1
```

Example response

- 204: No Content

22.13.7. Remove a property for autoindexing on nodes

Example request

- DELETE `http://localhost:7474/db/data/index/auto/node/properties/myProperty1`
- Accept: `application/json; charset=UTF-8`

Example response

- 204: No Content

22.14. Traversals



Warning

The Traversal REST Endpoint executes arbitrary Groovy code under the hood as part of the evaluators definitions. In hosted and open environments, this can constitute a security risk. In these case, consider using declarative approaches like [Part III, “Cypher Query Language”](#) or write your own server side plugin executing the interesting traversals with the Java API (see [Section 9.1, “Server Plugins”](#)) or secure your server, see [Section 28.1, “Securing access to the Neo4j Server”](#).

Traversals are performed from a start node. The traversal is controlled by the URI and the body sent with the request.

returnType

The kind of objects in the response is determined by *traverse/{returnType}* in the URL. returnType can have one of these values:

- node
- relationship
- path: contains full representations of start and end node, the rest are URIs.
- fullpath: contains full representations of all nodes and relationships.

To decide how the graph should be traversed you can use these parameters in the request body:

order

Decides in which order to visit nodes. Possible values:

- breadth_first: see [Breadth-first search](http://en.wikipedia.org/wiki/Breadth-first_search) <http://en.wikipedia.org/wiki/Breadth-first_search>.
- depth_first: see [Depth-first search](http://en.wikipedia.org/wiki/Depth-first_search) <http://en.wikipedia.org/wiki/Depth-first_search>

relationships

Decides which relationship types and directions should be followed. The direction can be one of:

- all
- in
- out

uniqueness

Decides how uniqueness should be calculated. For details on different uniqueness values see the [Java API on Uniqueness](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Uniqueness.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/Uniqueness.html>>. Possible values:

- node_global
- none
- relationship_global
- node_path
- relationship_path

prune_evaluator

Decides whether the traverser should continue down that path or if it should be pruned so that the traverser won't continue down that path. You can write your own prune evaluator as (see [Section 22.14.1, “Traversal using a return filter”](#) or use the built-in none prune evaluator.

return_filter

Decides whether the current position should be included in the result. You can provide your own code for this (see [Section 22.14.1, “Traversal using a return filter”](#)), or use one of the built-in filters:

- all
- all_but_start_node

max_depth

Is a short-hand way of specifying a prune evaluator which prunes after a certain depth. If not specified a max depth of 1 is used and if a `prune_evaluator` is specified instead of a `max_depth`, no max depth limit is set.


The `position` object in the body of the `return_filter` and `prune_evaluator` is a `Path` [Path](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/Path.html) object representing the path from the start node to the current traversal position.

Out of the box, the REST API supports JavaScript code in filters and evaluators. The script body will be executed in a Java context which has access to the full Neo4j [Java API](http://components.neo4j.org/neo4j/1.9.7/apidocs/). See the examples for the exact syntax of the request.

22.14.1. Traversal using a return filter

In this example, the `none` prune evaluator is used and a return filter is supplied in order to return all names containing "t". The result is to be returned as nodes and the max depth is set to 3.

Figure 22.45. Final Graph

*Example request*

- POST <http://localhost:7474/db/data/node/218/traverse/node>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```

{
  "order" : "breadth_first",
  "return_filter" : {
    "body" : "position.endNode().getProperty('name').toLowerCase().contains('t')",
    "language" : "javascript"
  },
  "prune_evaluator" : {
    "body" : "position.length() > 10",
    "language" : "javascript"
  },
  "uniqueness" : "node_global",
  "relationships" : [ {
    "direction" : "all",
    "type" : "knows"
  }, {
    "direction" : "all",
    "type" : "loves"
  } ],
  "max_depth" : 3
}

```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/218/relationships/out",
  "data" : {
    "name" : "Root"
  },
  "traverse" : "http://localhost:7474/db/data/node/218/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/218/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/218/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/218",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/218/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/218/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/218/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/218/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/218/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/218/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/218/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/221/relationships/out",
  "data" : {
    "name" : "Mattias"
  },
  "traverse" : "http://localhost:7474/db/data/node/221/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/221/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/221/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/221",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/221/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/221/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/221/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/221/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/221/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/221/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/221/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/220/relationships/out",
  "data" : {
    "name" : "Peter"
  },
  "traverse" : "http://localhost:7474/db/data/node/220/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/220/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/220/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/220",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/220/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/220/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/220/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/220/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/220/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/220/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/220/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/219/relationships/out",
  "data" : {
    "name" : "Tobias"
  },
  "traverse" : "http://localhost:7474/db/data/node/219/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/219/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/219/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/219",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/219/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/219/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/219/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/219/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/219/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/219/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/219/relationships/in/{-list|&|types}"
} ]
```


```

"traverse" : "http://localhost:7474/db/data/node/219/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/219/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/219/properties/{key}",
"self" : "http://localhost:7474/db/data/node/219",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/219/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/219/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/219/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/219/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/219/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/219/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/219/relationships/in/{-list|&|types}"
} ]

```

22.14.2. Return relationships from a traversal

Figure 22.46. Final Graph



Example request

- POST `http://localhost:7474/db/data/node/212/traverse/relationship`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "order" : "breadth_first",
  "uniqueness" : "none",
  "return_filter" : {
    "language" : "builtin",
    "name" : "all"
  }
}

```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

[ {
  "start" : "http://localhost:7474/db/data/node/212",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/109",
  "property" : "http://localhost:7474/db/data/relationship/109/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/109/properties",
  "type" : "know",
  "extensions" : {
  },
  "end" : "http://localhost:7474/db/data/node/211"
} ]

```


```

}, {
  "start" : "http://localhost:7474/db/data/node/212",
  "data" : {
  },
  "self" : "http://localhost:7474/db/data/relationship/110",
  "property" : "http://localhost:7474/db/data/relationship/110/properties/{key}",
  "properties" : "http://localhost:7474/db/data/relationship/110/properties",
  "type" : "own",
  "extensions" : {
  },
  "end" : "http://localhost:7474/db/data/node/210"
} ]

```

22.14.3. Return paths from a traversal

Figure 22.47. Final Graph



Example request

- POST `http://localhost:7474/db/data/node/202/traverse/path`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "order" : "breadth_first",
  "uniqueness" : "none",
  "return_filter" : {
    "language" : "builtin",
    "name" : "all"
  }
}

```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

[ {
  "start" : "http://localhost:7474/db/data/node/202",
  "nodes" : [ "http://localhost:7474/db/data/node/202" ],
  "length" : 0,
  "relationships" : [ ],
  "end" : "http://localhost:7474/db/data/node/202"
}, {
  "start" : "http://localhost:7474/db/data/node/202",
  "nodes" : [ "http://localhost:7474/db/data/node/202", "http://localhost:7474/db/data/node/201" ],
  "length" : 1,
  "relationships" : [ "http://localhost:7474/db/data/relationship/101" ],
  "end" : "http://localhost:7474/db/data/node/201"
} ], {

```

```


"start" : "http://localhost:7474/db/data/node/202",
"nodes" : [ "http://localhost:7474/db/data/node/202", "http://localhost:7474/db/data/node/200" ],
"length" : 1,
"relationships" : [ "http://localhost:7474/db/data/relationship/102" ],
"end" : "http://localhost:7474/db/data/node/200"
} ]

```

22.14.4. Traversal returning nodes below a certain depth

Here, all nodes at a traversal depth below 3 are returned.

Figure 22.48. Final Graph



Example request

- POST `http://localhost:7474/db/data/node/206/traverse/node`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "return_filter" : {
    "body" : "position.length()<3;",
    "language" : "javascript"
  },
  "prune_evaluator" : {
    "name" : "none",
    "language" : "builtin"
  }
}

```

}

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/206/relationships/out",
  "data" : {
    "name" : "Root"
  },
  "traverse" : "http://localhost:7474/db/data/node/206/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/206/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/206/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/206",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/206/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/206/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/206/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/206/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/206/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/206/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/206/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/209/relationships/out",
  "data" : {
    "name" : "Mattias"
  },
  "traverse" : "http://localhost:7474/db/data/node/209/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/209/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/209/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/209",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/209/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/209/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/209/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/209/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/209/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/209/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/209/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/204/relationships/out",
  "data" : {
    "name" : "Johan"
  },
  "traverse" : "http://localhost:7474/db/data/node/204/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/204/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/204/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/204",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/204/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/204/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/204/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/204/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/204/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/204/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/204/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/205/relationships/out",
  "data" : {
```

```

    "name" : "Emil"
  },
  "traverse" : "http://localhost:7474/db/data/node/205/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/205/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/205/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/205",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/205/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/205/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/205/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/205/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/205/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/205/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/205/relationships/in/{-list|&|types}"
} ]

```

22.14.5. Creating a paged traverser

Paged traversers are created by POST-ing a traversal description to the link identified by the `paged_traverser` key in a node representation. When creating a paged traverser, the same options apply as for a regular traverser, meaning that `node`, `path`, or `fullpath`, can be targeted.

Example request

- POST `http://localhost:7474/db/data/node/298/paged/traverse/node`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "prune_evaluator" : {
    "language" : "builtin",
    "name" : "none"
  },
  "return_filter" : {
    "language" : "javascript",
    "body" : "position.endNode().getProperty('name').contains('1');"
  },
  "order" : "depth_first",
  "relationships" : {
    "type" : "NEXT",
    "direction" : "out"
  }
}

```

Example response

- 201: Created
- Content-Type: `application/json; charset=UTF-8`
- Location: `http://localhost:7474/db/data/node/298/paged/traverse/node/2de413b7e500482080dafa48014df5d8`

```

[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/299/relationships/out",
  "data" : {
    "name" : "1"
  },
  "traverse" : "http://localhost:7474/db/data/node/299/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/299/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/299/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/299",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/299/relationships/out/{-list|&|types}",
} ]

```

```

"properties" : "http://localhost:7474/db/data/node/299/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/299/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/299/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/299/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/299/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/299/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/308/relationships/out",
"data" : {
  "name" : "10"
},
"traverse" : "http://localhost:7474/db/data/node/308/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/308/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/308/properties/{key}",
"self" : "http://localhost:7474/db/data/node/308",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/308/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/308/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/308/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/308/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/308/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/308/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/308/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/309/relationships/out",
"data" : {
  "name" : "11"
},
"traverse" : "http://localhost:7474/db/data/node/309/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/309/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/309/properties/{key}",
"self" : "http://localhost:7474/db/data/node/309",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/309/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/309/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/309/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/309/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/309/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/309/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/309/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/310/relationships/out",
"data" : {
  "name" : "12"
},
"traverse" : "http://localhost:7474/db/data/node/310/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/310/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/310/properties/{key}",
"self" : "http://localhost:7474/db/data/node/310",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/310/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/310/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/310/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/310/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/310/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/310/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/310/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/311/relationships/out",
"data" : {

```

```

    "name" : "13"
  },
  "traverse" : "http://localhost:7474/db/data/node/311/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/311/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/311/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/311",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/311/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/311/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/311/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/311/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/311/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/311/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/311/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/312/relationships/out",
  "data" : {
    "name" : "14"
  },
  "traverse" : "http://localhost:7474/db/data/node/312/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/312/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/312/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/312",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/312/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/312/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/312/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/312/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/312/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/312/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/312/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/313/relationships/out",
  "data" : {
    "name" : "15"
  },
  "traverse" : "http://localhost:7474/db/data/node/313/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/313/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/313/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/313",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/313/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/313/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/313/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/313/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/313/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/313/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/313/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/314/relationships/out",
  "data" : {
    "name" : "16"
  },
  "traverse" : "http://localhost:7474/db/data/node/314/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/314/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/314/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/314",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/314/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/314/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/314/relationships/in",
  "extensions" : {
  },

```

```
"create_relationship" : "http://localhost:7474/db/data/node/314/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/314/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/314/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/314/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/315/relationships/out",
  "data" : {
    "name" : "17"
  },
  "traverse" : "http://localhost:7474/db/data/node/315/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/315/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/315/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/315",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/315/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/315/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/315/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/315/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/315/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/315/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/315/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/316/relationships/out",
  "data" : {
    "name" : "18"
  },
  "traverse" : "http://localhost:7474/db/data/node/316/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/316/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/316/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/316",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/316/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/316/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/316/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/316/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/316/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/316/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/316/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/317/relationships/out",
  "data" : {
    "name" : "19"
  },
  "traverse" : "http://localhost:7474/db/data/node/317/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/317/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/317/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/317",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/317/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/317/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/317/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/317/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/317/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/317/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/317/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/319/relationships/out",
  "data" : {
    "name" : "21"
  },
  "traverse" : "http://localhost:7474/db/data/node/319/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/319/relationships/all/{-list|&|types}",
```

```

"property" : "http://localhost:7474/db/data/node/319/properties/{key}",
"self" : "http://localhost:7474/db/data/node/319",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/319/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/319/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/319/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/319/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/319/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/319/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/319/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/329/relationships/out",
"data" : {
  "name" : "31"
},
"traverse" : "http://localhost:7474/db/data/node/329/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/329/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/329/properties/{key}",
"self" : "http://localhost:7474/db/data/node/329",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/329/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/329/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/329/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/329/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/329/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/329/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/329/relationships/in/{-list|&|types}"
} ]

```

22.14.6. Paging through the results of a paged traverser

Paged traversers hold state on the server, and allow clients to page through the results of a traversal. To progress to the next page of traversal results, the client issues a HTTP GET request on the paged traversal URI which causes the traversal to fill the next page (or partially fill it if insufficient results are available). Note that if a traverser expires through inactivity it will cause a 404 response on the next GET request. Traversers' leases are renewed on every successful access for the same amount of time as originally specified. When the paged traverser reaches the end of its results, the client can expect a 404 response as the traverser is disposed by the server.

Example request

- GET `http://localhost:7474/db/data/node/331/paged/traverse/node/0fd1d13bf2b04bb1886b54908a74f5a1`
- Accept: `application/json`

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/662/relationships/out",
  "data" : {
    "name" : "331"
  },
  "traverse" : "http://localhost:7474/db/data/node/662/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/662/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/662/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/662",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/662/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/662/properties",

```

```
"incoming_relationships" : "http://localhost:7474/db/data/node/662/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/662/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/662/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/662/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/662/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/672/relationships/out",
"data" : {
  "name" : "341"
},
"traverse" : "http://localhost:7474/db/data/node/672/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/672/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/672/properties/{key}",
"self" : "http://localhost:7474/db/data/node/672",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/672/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/672/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/672/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/672/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/672/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/672/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/672/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/682/relationships/out",
"data" : {
  "name" : "351"
},
"traverse" : "http://localhost:7474/db/data/node/682/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/682/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/682/properties/{key}",
"self" : "http://localhost:7474/db/data/node/682",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/682/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/682/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/682/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/682/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/682/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/682/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/682/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/692/relationships/out",
"data" : {
  "name" : "361"
},
"traverse" : "http://localhost:7474/db/data/node/692/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/692/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/692/properties/{key}",
"self" : "http://localhost:7474/db/data/node/692",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/692/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/692/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/692/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/692/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/692/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/692/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/692/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/702/relationships/out",
"data" : {
  "name" : "371"
```

```

},
"traverse" : "http://localhost:7474/db/data/node/702/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/702/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/702/properties/{key}",
"self" : "http://localhost:7474/db/data/node/702",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/702/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/702/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/702/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/702/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/702/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/702/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/702/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/712/relationships/out",
"data" : {
"name" : "381"
},
"traverse" : "http://localhost:7474/db/data/node/712/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/712/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/712/properties/{key}",
"self" : "http://localhost:7474/db/data/node/712",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/712/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/712/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/712/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/712/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/712/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/712/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/712/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/722/relationships/out",
"data" : {
"name" : "391"
},
"traverse" : "http://localhost:7474/db/data/node/722/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/722/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/722/properties/{key}",
"self" : "http://localhost:7474/db/data/node/722",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/722/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/722/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/722/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/722/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/722/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/722/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/722/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/732/relationships/out",
"data" : {
"name" : "401"
},
"traverse" : "http://localhost:7474/db/data/node/732/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/732/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/732/properties/{key}",
"self" : "http://localhost:7474/db/data/node/732",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/732/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/732/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/732/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/732/relationships",

```

```
"paged_traverse" : "http://localhost:7474/db/data/node/732/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/732/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/732/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/741/relationships/out",
  "data" : {
    "name" : "410"
  },
  "traverse" : "http://localhost:7474/db/data/node/741/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/741/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/741/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/741",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/741/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/741/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/741/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/741/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/741/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/741/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/741/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/742/relationships/out",
  "data" : {
    "name" : "411"
  },
  "traverse" : "http://localhost:7474/db/data/node/742/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/742/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/742/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/742",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/742/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/742/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/742/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/742/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/742/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/742/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/742/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/743/relationships/out",
  "data" : {
    "name" : "412"
  },
  "traverse" : "http://localhost:7474/db/data/node/743/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/743/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/743/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/743",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/743/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/743/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/743/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/743/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/743/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/743/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/743/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/744/relationships/out",
  "data" : {
    "name" : "413"
  },
  "traverse" : "http://localhost:7474/db/data/node/744/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/744/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/744/properties/{key}",
```

```
"self" : "http://localhost:7474/db/data/node/744",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/744/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/744/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/744/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/744/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/744/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/744/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/744/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/745/relationships/out",
"data" : {
  "name" : "414"
},
"traverse" : "http://localhost:7474/db/data/node/745/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/745/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/745/properties/{key}",
"self" : "http://localhost:7474/db/data/node/745",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/745/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/745/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/745/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/745/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/745/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/745/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/745/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/746/relationships/out",
"data" : {
  "name" : "415"
},
"traverse" : "http://localhost:7474/db/data/node/746/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/746/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/746/properties/{key}",
"self" : "http://localhost:7474/db/data/node/746",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/746/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/746/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/746/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/746/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/746/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/746/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/746/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/747/relationships/out",
"data" : {
  "name" : "416"
},
"traverse" : "http://localhost:7474/db/data/node/747/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/747/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/747/properties/{key}",
"self" : "http://localhost:7474/db/data/node/747",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/747/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/747/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/747/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/747/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/747/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/747/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/747/relationships/in/{-list|&|types}"
}, {
```

```

"outgoing_relationships" : "http://localhost:7474/db/data/node/748/relationships/out",
"data" : {
  "name" : "417"
},
"traverse" : "http://localhost:7474/db/data/node/748/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/748/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/748/properties/{key}",
"self" : "http://localhost:7474/db/data/node/748",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/748/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/748/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/748/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/748/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/748/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/748/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/748/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/749/relationships/out",
"data" : {
  "name" : "418"
},
"traverse" : "http://localhost:7474/db/data/node/749/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/749/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/749/properties/{key}",
"self" : "http://localhost:7474/db/data/node/749",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/749/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/749/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/749/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/749/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/749/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/749/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/749/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/750/relationships/out",
"data" : {
  "name" : "419"
},
"traverse" : "http://localhost:7474/db/data/node/750/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/750/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/750/properties/{key}",
"self" : "http://localhost:7474/db/data/node/750",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/750/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/750/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/750/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/750/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/750/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/750/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/750/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/752/relationships/out",
"data" : {
  "name" : "421"
},
"traverse" : "http://localhost:7474/db/data/node/752/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/752/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/752/properties/{key}",
"self" : "http://localhost:7474/db/data/node/752",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/752/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/752/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/752/relationships/in",

```

```

"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/752/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/752/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/752/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/752/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/762/relationships/out",
"data" : {
  "name" : "431"
},
"traverse" : "http://localhost:7474/db/data/node/762/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/762/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/762/properties/{key}",
"self" : "http://localhost:7474/db/data/node/762",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/762/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/762/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/762/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/762/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/762/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/762/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/762/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/772/relationships/out",
"data" : {
  "name" : "441"
},
"traverse" : "http://localhost:7474/db/data/node/772/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/772/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/772/properties/{key}",
"self" : "http://localhost:7474/db/data/node/772",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/772/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/772/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/772/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/772/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/772/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/772/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/772/relationships/in/{-list|&|types}"
} ]

```

22.14.7. Paged traverser page size

The default page size is 50 items, but depending on the application larger or smaller pages sizes might be appropriate. This can be set by adding a `pageSize` query parameter.

Example request

- POST `http://localhost:7474/db/data/node/34/paged/traverse/node?pageSize=1`
- Accept: `application/json`
- Content-Type: `application/json`

```

{
  "prune_evaluator" : {
    "language" : "builtin",
    "name" : "none"
  },
  "return_filter" : {
    "language" : "javascript",
    "body" : "position.endNode().getProperty('name').contains('1');"
  }
}

```

```

},
"order" : "depth_first",
"relationships" : {
  "type" : "NEXT",
  "direction" : "out"
}
}

```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: http://localhost:7474/db/data/node/34/paged/traverse/node/10f8894ab27d4e80b50390c846d7fe35

```

[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/35/relationships/out",
  "data" : {
    "name" : "1"
  },
  "traverse" : "http://localhost:7474/db/data/node/35/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/35/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/35/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/35",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/35/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/35/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/35/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/35/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/35/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/35/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/35/relationships/in/{-list|&|types}"
} ]

```

22.14.8. Paged traverser timeout

The default timeout for a paged traverser is 60 seconds, but depending on the application larger or smaller timeouts might be appropriate. This can be set by adding a `leaseTime` query parameter with the number of seconds the paged traverser should last.

Example request

- POST http://localhost:7474/db/data/node/808/paged/traverse/node?leaseTime=10
- Accept: application/json
- Content-Type: application/json

```

{
  "prune_evaluator" : {
    "language" : "builtin",
    "name" : "none"
  },
  "return_filter" : {
    "language" : "javascript",
    "body" : "position.endNode().getProperty('name').contains('1');"
  },
  "order" : "depth_first",
  "relationships" : {
    "type" : "NEXT",
    "direction" : "out"
  }
}

```

Example response

- 201: Created
- Content-Type: application/json; charset=UTF-8
- Location: <http://localhost:7474/db/data/node/808/paged/traverse/node/03c4fb0ef4bf4fd69d4e0104c693e490>

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/809/relationships/out",
  "data" : {
    "name" : "1"
  },
  "traverse" : "http://localhost:7474/db/data/node/809/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/809/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/809/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/809",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/809/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/809/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/809/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/809/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/809/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/809/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/809/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/818/relationships/out",
  "data" : {
    "name" : "10"
  },
  "traverse" : "http://localhost:7474/db/data/node/818/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/818/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/818/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/818",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/818/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/818/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/818/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/818/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/818/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/818/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/818/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/819/relationships/out",
  "data" : {
    "name" : "11"
  },
  "traverse" : "http://localhost:7474/db/data/node/819/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/819/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/819/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/819",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/819/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/819/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/819/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/819/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/819/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/819/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/819/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/820/relationships/out",
```

```

"data" : {
  "name" : "12"
},
"traverse" : "http://localhost:7474/db/data/node/820/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/820/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/820/properties/{key}",
"self" : "http://localhost:7474/db/data/node/820",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/820/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/820/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/820/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/820/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/820/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/820/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/820/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/821/relationships/out",
"data" : {
  "name" : "13"
},
"traverse" : "http://localhost:7474/db/data/node/821/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/821/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/821/properties/{key}",
"self" : "http://localhost:7474/db/data/node/821",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/821/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/821/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/821/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/821/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/821/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/821/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/821/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/822/relationships/out",
"data" : {
  "name" : "14"
},
"traverse" : "http://localhost:7474/db/data/node/822/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/822/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/822/properties/{key}",
"self" : "http://localhost:7474/db/data/node/822",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/822/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/822/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/822/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/822/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/822/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/822/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/822/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/823/relationships/out",
"data" : {
  "name" : "15"
},
"traverse" : "http://localhost:7474/db/data/node/823/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/823/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/823/properties/{key}",
"self" : "http://localhost:7474/db/data/node/823",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/823/relationships/out/{-list|&|types}",
"properties" : "http://localhost:7474/db/data/node/823/properties",
"incoming_relationships" : "http://localhost:7474/db/data/node/823/relationships/in",
"extensions" : {

```

```

    },
    "create_relationship" : "http://localhost:7474/db/data/node/823/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/823/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/823/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/823/relationships/in/{-list|&|types}"
  }, {
    "outgoing_relationships" : "http://localhost:7474/db/data/node/824/relationships/out",
    "data" : {
      "name" : "16"
    },
    },
    "traverse" : "http://localhost:7474/db/data/node/824/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/824/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/824/properties/{key}",
    "self" : "http://localhost:7474/db/data/node/824",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/824/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/824/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/824/relationships/in",
    "extensions" : {
    },
    },
    "create_relationship" : "http://localhost:7474/db/data/node/824/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/824/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/824/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/824/relationships/in/{-list|&|types}"
  }, {
    "outgoing_relationships" : "http://localhost:7474/db/data/node/825/relationships/out",
    "data" : {
      "name" : "17"
    },
    },
    "traverse" : "http://localhost:7474/db/data/node/825/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/825/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/825/properties/{key}",
    "self" : "http://localhost:7474/db/data/node/825",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/825/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/825/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/825/relationships/in",
    "extensions" : {
    },
    },
    "create_relationship" : "http://localhost:7474/db/data/node/825/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/825/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/825/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/825/relationships/in/{-list|&|types}"
  }, {
    "outgoing_relationships" : "http://localhost:7474/db/data/node/826/relationships/out",
    "data" : {
      "name" : "18"
    },
    },
    "traverse" : "http://localhost:7474/db/data/node/826/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/826/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/826/properties/{key}",
    "self" : "http://localhost:7474/db/data/node/826",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/826/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/826/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/826/relationships/in",
    "extensions" : {
    },
    },
    "create_relationship" : "http://localhost:7474/db/data/node/826/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/826/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/826/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/826/relationships/in/{-list|&|types}"
  }, {
    "outgoing_relationships" : "http://localhost:7474/db/data/node/827/relationships/out",
    "data" : {
      "name" : "19"
    },
    },
    "traverse" : "http://localhost:7474/db/data/node/827/traverse/{returnType}",

```

```

    "all_typed_relationships" : "http://localhost:7474/db/data/node/827/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/827/properties/{key}",
    "self" : "http://localhost:7474/db/data/node/827",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/827/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/827/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/827/relationships/in",
    "extensions" : {
    },
    "create_relationship" : "http://localhost:7474/db/data/node/827/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/827/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/827/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/827/relationships/in/{-list|&|types}"
  }, {
    "outgoing_relationships" : "http://localhost:7474/db/data/node/829/relationships/out",
    "data" : {
      "name" : "21"
    },
    "traverse" : "http://localhost:7474/db/data/node/829/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/829/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/829/properties/{key}",
    "self" : "http://localhost:7474/db/data/node/829",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/829/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/829/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/829/relationships/in",
    "extensions" : {
    },
    "create_relationship" : "http://localhost:7474/db/data/node/829/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/829/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/829/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/829/relationships/in/{-list|&|types}"
  }, {
    "outgoing_relationships" : "http://localhost:7474/db/data/node/839/relationships/out",
    "data" : {
      "name" : "31"
    },
    "traverse" : "http://localhost:7474/db/data/node/839/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/839/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/839/properties/{key}",
    "self" : "http://localhost:7474/db/data/node/839",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/839/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/839/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/839/relationships/in",
    "extensions" : {
    },
    "create_relationship" : "http://localhost:7474/db/data/node/839/relationships",
    "paged_traverse" : "http://localhost:7474/db/data/node/839/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "all_relationships" : "http://localhost:7474/db/data/node/839/relationships/all",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/839/relationships/in/{-list|&|types}"
  }
] ]

```

22.15. Built-in Graph Algorithms

Neo4j comes with a number of built-in graph algorithms. They are performed from a start node. The traversal is controlled by the URI and the body sent with the request.

algorithm

The algorithm to choose. If not set, default is `shortestPath`. `algorithm` can have one of these values:

- `shortestPath`
- `allSimplePaths`
- `allPaths`
- `dijkstra` (optional with `cost_property` and `default_cost` parameters)


max_depth

The maximum depth as an integer for the algorithms like `ShortestPath`, where applicable. Default is 1.

22.15.1. Find all shortest paths

The `shortestPath` algorithm can find multiple paths between the same nodes, like in this example.

Figure 22.49. Final Graph



Example request

- `POST http://localhost:7474/db/data/node/50/paths`

- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "to" : "http://localhost:7474/db/data/node/45",
  "max_depth" : 3,
  "relationships" : {
    "type" : "to",
    "direction" : "out"
  },
  "algorithm" : "shortestPath"
}
```

Example response


- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "start" : "http://localhost:7474/db/data/node/50",
  "nodes" : [ "http://localhost:7474/db/data/node/50", "http://localhost:7474/db/data/node/49", "http://localhost:7474/db/data/node/45" ],
  "length" : 2,
  "relationships" : [ "http://localhost:7474/db/data/relationship/50", "http://localhost:7474/db/data/relationship/59" ],
  "end" : "http://localhost:7474/db/data/node/45"
}, {
  "start" : "http://localhost:7474/db/data/node/50",
  "nodes" : [ "http://localhost:7474/db/data/node/50", "http://localhost:7474/db/data/node/46", "http://localhost:7474/db/data/node/45" ],
  "length" : 2,
  "relationships" : [ "http://localhost:7474/db/data/relationship/51", "http://localhost:7474/db/data/relationship/57" ],
  "end" : "http://localhost:7474/db/data/node/45"
} ]
```

22.15.2. Find one of the shortest paths between nodes

If no path algorithm is specified, a ShortestPath algorithm with a max depth of 1 will be chosen. In this example, the `max_depth` is set to 3 in order to find the shortest path between 3 linked nodes.

Figure 22.50. Final Graph

*Example request*

- POST <http://localhost:7474/db/data/node/34/path>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```

{
  "to" : "http://localhost:7474/db/data/node/29",
  "max_depth" : 3,
  "relationships" : {
    "type" : "to",
    "direction" : "out"
  },
  "algorithm" : "shortestPath"
}

```

Example response


- 200: OK
- Content-Type: application/json; charset=UTF-8

```

{
  "start" : "http://localhost:7474/db/data/node/34",
  "nodes" : [ "http://localhost:7474/db/data/node/34", "http://localhost:7474/db/data/node/30", "http://localhost:7474/db/data/node/29" ],
  "length" : 2,
  "relationships" : [ "http://localhost:7474/db/data/relationship/28", "http://localhost:7474/db/data/relationship/34" ],
}

```

```
"end" : "http://localhost:7474/db/data/node/29"  
}
```

22.15.3. Execute a Dijkstra algorithm with similar weights on relationships*Figure 22.51. Final Graph**Example request*

- POST `http://localhost:7474/db/data/node/40/path`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "to" : "http://localhost:7474/db/data/node/43",
  "cost_property" : "cost",
  "relationships" : {
    "type" : "to",
    "direction" : "out"
  },
  "algorithm" : "dijkstra"
}
```


Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
{
  "weight" : 2.0,
  "start" : "http://localhost:7474/db/data/node/40",
  "nodes" : [ "http://localhost:7474/db/data/node/40", "http://localhost:7474/db/data/node/41", "http://localhost:7474/db/data/node/43" ],
  "length" : 2,
  "relationships" : [ "http://localhost:7474/db/data/relationship/37", "http://localhost:7474/db/data/relationship/38" ],
  "end" : "http://localhost:7474/db/data/node/43"
}
```

22.15.4. Execute a Dijkstra algorithm with weights on relationships

Figure 22.52. Final Graph



Example request

- POST `http://localhost:7474/db/data/node/24/path`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "to" : "http://localhost:7474/db/data/node/27",
  "cost_property" : "cost",
  "relationships" : {
    "type" : "to",
    "direction" : "out"
  },
  "algorithm" : "dijkstra"
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
{
  "weight" : 6.0,
  "start" : "http://localhost:7474/db/data/node/24",
  "nodes" : [ "http://localhost:7474/db/data/node/24", "http://localhost:7474/db/data/node/25", "http://localhost:7474/db/data/node/26", "http://localhost:7474/db/data/node/27" ],
  "length" : 6,
  "relationships" : [ "http://localhost:7474/db/data/relationship/14", "http://localhost:7474/db/data/relationship/16", "http://localhost:7474/db/data/relationship/18" ],
  "end" : "http://localhost:7474/db/data/node/27"
}
```

22.16. Batch operations

22.16.1. Execute multiple operations in batch

This lets you execute multiple API calls through a single HTTP call, significantly improving performance for large insert and update operations.

The batch service expects an array of job descriptions as input, each job description describing an action to be performed via the normal server API.

This service is transactional. If any of the operations performed fails (returns a non-2xx HTTP status code), the transaction will be rolled back and all changes will be undone.

Each job description should contain a `to` attribute, with a value relative to the data API root (so <http://localhost:7474/db/data/node> becomes just `/node`), and a `method` attribute containing HTTP verb to use.

Optionally you may provide a `body` attribute, and an `id` attribute to help you keep track of responses, although responses are guaranteed to be returned in the same order the job descriptions are received.

The following figure outlines the different parts of the job descriptions:





Figure 22.53. Final Graph



Example request

- POST <http://localhost:7474/db/data/batch>
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
[ {
  "method" : "PUT",
  "to" : "/node/119/properties",
  "body" : {
    "age" : 1
  },
  "id" : 0
}, {
  "method" : "GET",
```

```

    "to" : "/node/119",
    "id" : 1
  }, {
    "method" : "POST",
    "to" : "/node",
    "body" : {
      "age" : 1
    },
    "id" : 2
  }, {
    "method" : "POST",
    "to" : "/node",
    "body" : {
      "age" : 1
    },
    "id" : 3
  } ]
} ]

```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```

[ {
  "id" : 0,
  "from" : "/node/119/properties"
}, {
  "id" : 1,
  "body" : {
    "extensions" : {
    },
    "paged_traverse" : "http://localhost:7474/db/data/node/119/paged/traverse/{returnType}?pageSize,leaseTime)",
    "outgoing_relationships" : "http://localhost:7474/db/data/node/119/relationships/out",
    "traverse" : "http://localhost:7474/db/data/node/119/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/119/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/119/properties/{key}",
    "all_relationships" : "http://localhost:7474/db/data/node/119/relationships/all",
    "self" : "http://localhost:7474/db/data/node/119",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/119/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/119/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/119/relationships/in",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/119/relationships/in/{-list|&|types}",
    "create_relationship" : "http://localhost:7474/db/data/node/119/relationships",
    "data" : {
      "age" : 1
    }
  },
  "from" : "/node/119"
}, {
  "id" : 2,
  "location" : "http://localhost:7474/db/data/node/120",
  "body" : {
    "extensions" : {
    },
    "paged_traverse" : "http://localhost:7474/db/data/node/120/paged/traverse/{returnType}?pageSize,leaseTime)",
    "outgoing_relationships" : "http://localhost:7474/db/data/node/120/relationships/out",
    "traverse" : "http://localhost:7474/db/data/node/120/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/120/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/120/properties/{key}",
    "all_relationships" : "http://localhost:7474/db/data/node/120/relationships/all",
    "self" : "http://localhost:7474/db/data/node/120",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/120/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/120/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/120/relationships/in",

```

```

    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/120/relationships/in/{-list|&|types}",
    "create_relationship" : "http://localhost:7474/db/data/node/120/relationships",
    "data" : {
      "age" : 1
    }
  },
  "from" : "/node"
}, {
  "id" : 3,
  "location" : "http://localhost:7474/db/data/node/121",
  "body" : {
    "extensions" : {
    },
    "paged_traverse" : "http://localhost:7474/db/data/node/121/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "outgoing_relationships" : "http://localhost:7474/db/data/node/121/relationships/out",
    "traverse" : "http://localhost:7474/db/data/node/121/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/121/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/121/properties/{key}",
    "all_relationships" : "http://localhost:7474/db/data/node/121/relationships/all",
    "self" : "http://localhost:7474/db/data/node/121",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/121/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/121/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/121/relationships/in",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/121/relationships/in/{-list|&|types}",
    "create_relationship" : "http://localhost:7474/db/data/node/121/relationships",
    "data" : {
      "age" : 1
    }
  }
},
"from" : "/node"
} ]


```

22.16.2. Refer to items created earlier in the same batch job

The batch operation API allows you to refer to the URI returned from a created resource in subsequent job descriptions, within the same batch call.

Use the `{[JOB ID]}` special syntax to inject URIs from created resources into JSON strings in subsequent job descriptions.

Figure 22.54. Final Graph



Example request

- POST `http://localhost:7474/db/data/batch`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

[ {
  "method" : "POST",

```

```

    "to" : "/node",
    "id" : 0,
    "body" : {
      "name" : "bob"
    }
  }, {
    "method" : "POST",
    "to" : "/node",
    "id" : 1,
    "body" : {
      "age" : 12
    }
  }, {
    "method" : "POST",
    "to" : "{0}/relationships",
    "id" : 3,
    "body" : {
      "to" : "{1}",
      "data" : {
        "since" : "2010"
      },
      "type" : "KNOWS"
    }
  }, {
    "method" : "POST",
    "to" : "/index/relationship/my_rels",
    "id" : 4,
    "body" : {
      "key" : "since",
      "value" : "2010",
      "uri" : "{3}"
    }
  }
]

```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```

[ {
  "id" : 0,
  "location" : "http://localhost:7474/db/data/node/112",
  "body" : {
    "extensions" : {
    },
    "paged_traverse" : "http://localhost:7474/db/data/node/112/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "outgoing_relationships" : "http://localhost:7474/db/data/node/112/relationships/out",
    "traverse" : "http://localhost:7474/db/data/node/112/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/112/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/112/properties/{key}",
    "all_relationships" : "http://localhost:7474/db/data/node/112/relationships/all",
    "self" : "http://localhost:7474/db/data/node/112",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/112/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/112/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/112/relationships/in",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/112/relationships/in/{-list|&|types}",
    "create_relationship" : "http://localhost:7474/db/data/node/112/relationships",
    "data" : {
      "name" : "bob"
    }
  },
  "from" : "/node"
}, {
  "id" : 1,

```


```

"location" : "http://localhost:7474/db/data/node/113",
"body" : {
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/113/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/113/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/113/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/113/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/113/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/113/relationships/all",
  "self" : "http://localhost:7474/db/data/node/113",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/113/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/113/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/113/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/113/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/113/relationships",
  "data" : {
    "age" : 12
  }
},
"from" : "/node"
}, {
  "id" : 3,
  "location" : "http://localhost:7474/db/data/relationship/58",
  "body" : {
    "extensions" : {
    },
    "start" : "http://localhost:7474/db/data/node/112",
    "property" : "http://localhost:7474/db/data/relationship/58/properties/{key}",
    "self" : "http://localhost:7474/db/data/relationship/58",
    "properties" : "http://localhost:7474/db/data/relationship/58/properties",
    "type" : "KNOWS",
    "end" : "http://localhost:7474/db/data/node/113",
    "data" : {
      "since" : "2010"
    }
  },
  "from" : "http://localhost:7474/db/data/node/112/relationships"
}, {
  "id" : 4,
  "location" : "http://localhost:7474/db/data/index/relationship/my_rels/since/2010/58",
  "body" : {
    "extensions" : {
    },
    "start" : "http://localhost:7474/db/data/node/112",
    "property" : "http://localhost:7474/db/data/relationship/58/properties/{key}",
    "self" : "http://localhost:7474/db/data/relationship/58",
    "properties" : "http://localhost:7474/db/data/relationship/58/properties",
    "type" : "KNOWS",
    "end" : "http://localhost:7474/db/data/node/113",
    "data" : {
      "since" : "2010"
    }
  },
  "indexed" : "http://localhost:7474/db/data/index/relationship/my_rels/since/2010/58"
},
"from" : "/index/relationship/my_rels"
} ]

```

22.16.3. Execute multiple operations in batch streaming

Figure 22.55. Final Graph



Example request

- POST http://localhost:7474/db/data/batch
- Accept: application/json
- Content-Type: application/json
- X-Stream: true

```
[ {
  "method" : "PUT",
  "to" : "/node/290/properties",
  "body" : {
    "age" : 1
  },
  "id" : 0
}, {
  "method" : "GET",
  "to" : "/node/290",
  "id" : 1
}, {
  "method" : "POST",
  "to" : "/node",
  "body" : {
    "age" : 1
  },
  "id" : 2
}, {
  "method" : "POST",
  "to" : "/node",
  "body" : {
    "age" : 1
  },
  "id" : 3
} ]
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "id" : 0,
  "from" : "/node/290/properties",
  "body" : null,
  "status" : 204
}, {
  "id" : 1,
```

```

"from" : "/node/290",
"body" : {
  "extensions" : {
  },
  "paged_traverse" : "http://localhost:7474/db/data/node/290/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "outgoing_relationships" : "http://localhost:7474/db/data/node/290/relationships/out",
  "traverse" : "http://localhost:7474/db/data/node/290/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/290/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/290/properties/{key}",
  "all_relationships" : "http://localhost:7474/db/data/node/290/relationships/all",
  "self" : "http://localhost:7474/db/data/node/290",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/290/relationships/out/{-list|&|types}",
  "properties" : "http://localhost:7474/db/data/node/290/properties",
  "incoming_relationships" : "http://localhost:7474/db/data/node/290/relationships/in",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/290/relationships/in/{-list|&|types}",
  "create_relationship" : "http://localhost:7474/db/data/node/290/relationships",
  "data" : {
    "age" : 1
  }
},
"status" : 200
}, {
  "id" : 2,
  "from" : "/node",
  "body" : {
    "extensions" : {
    },
    "paged_traverse" : "http://localhost:7474/db/data/node/291/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "outgoing_relationships" : "http://localhost:7474/db/data/node/291/relationships/out",
    "traverse" : "http://localhost:7474/db/data/node/291/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/291/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/291/properties/{key}",
    "all_relationships" : "http://localhost:7474/db/data/node/291/relationships/all",
    "self" : "http://localhost:7474/db/data/node/291",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/291/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/291/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/291/relationships/in",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/291/relationships/in/{-list|&|types}",
    "create_relationship" : "http://localhost:7474/db/data/node/291/relationships",
    "data" : {
      "age" : 1
    }
  }
},
"location" : "http://localhost:7474/db/data/node/291",
"status" : 201
}, {
  "id" : 3,
  "from" : "/node",
  "body" : {
    "extensions" : {
    },
    "paged_traverse" : "http://localhost:7474/db/data/node/292/paged/traverse/{returnType}{?pageSize,leaseTime}",
    "outgoing_relationships" : "http://localhost:7474/db/data/node/292/relationships/out",
    "traverse" : "http://localhost:7474/db/data/node/292/traverse/{returnType}",
    "all_typed_relationships" : "http://localhost:7474/db/data/node/292/relationships/all/{-list|&|types}",
    "property" : "http://localhost:7474/db/data/node/292/properties/{key}",
    "all_relationships" : "http://localhost:7474/db/data/node/292/relationships/all",
    "self" : "http://localhost:7474/db/data/node/292",
    "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/292/relationships/out/{-list|&|types}",
    "properties" : "http://localhost:7474/db/data/node/292/properties",
    "incoming_relationships" : "http://localhost:7474/db/data/node/292/relationships/in",
    "incoming_typed_relationships" : "http://localhost:7474/db/data/node/292/relationships/in/{-list|&|types}",
    "create_relationship" : "http://localhost:7474/db/data/node/292/relationships",
    "data" : {
      "age" : 1
    }
  }
}

```

```
    }  
  },  
  "location" : "http://localhost:7474/db/data/node/292",  
  "status" : 201  
} ]
```

22.17. WADL Support

The Neo4j REST API is a truly RESTful interface relying on hypermedia controls (links) to advertise permissible actions to users. Hypermedia is a dynamic interface style where declarative constructs (semantic markup) are used to inform clients of their next legal choices just in time.



Caution

RESTful APIs cannot be modelled by static interface description languages like WSDL or WADL.

However for some use cases, developers may wish to expose WADL descriptions of the Neo4j REST API, particularly when using tooling that expects such.

In those cases WADL generation may be enabled by adding to your server's `neo4j.properties` file:

```
unsupported_wadl_generation_enabled=true
```



Caution

WADL is not an officially supported part of the Neo4j server API because WADL is insufficiently expressive to capture the set of potential interactions a client can drive with Neo4j server. Expect the WADL description to be incomplete, and in some cases contradictory to the real API. In any cases where the WADL description disagrees with the REST API, the REST API should be considered authoritative. WADL generation may be withdrawn at any point in the Neo4j release cycle.

22.18. Gremlin Plugin



Gremlin <<http://gremlin.tinkerpop.com>> is a Groovy based Graph Traversal Language. It provides a very expressive way of explicitly scripting traversals through a Neo4j graph.

The Neo4j Gremlin Plugin provides an endpoint to send Gremlin scripts to the Neo4j Server. The scripts are executed on the server database and the results are returned as Neo4j Node and Relationship representations. This keeps the types throughout the REST API consistent. The results are quite verbose when returning Neo4j Node, Relationship or Graph representations. On the other hand, just return properties like in the [Section 22.18.4, “Send a Gremlin Script - JSON encoded with table results”](#) example for responses tailored to specific needs.



Warning

The Gremlin plugin lets you execute arbitrary Groovy code under the hood. In hosted and open environments, this can constitute a security risk. In these case, consider using declarative approaches like [Part III, “Cypher Query Language”](#) or write your own server side plugin executing the interesting Gremlin or Java routines, see [Section 9.1, “Server Plugins”](#) or secure your server, see [Section 28.1, “Securing access to the Neo4j Server”](#).



Tip

When returning results from pipes like `g.v(0).in()`, make sure to iterate through the results in order not to return the pipe object but its content, like `g.v(0).in().iterate()`. For more caveats, see [Gremlin Troubleshooting](https://github.com/tinkerpop/gremlin/wiki/Troubleshooting) <<https://github.com/tinkerpop/gremlin/wiki/Troubleshooting>>


22.18.1. Send a Gremlin Script - URL encoded

Scripts can be sent as URL-encoded In this example, the graph has been autoindexed by Neo4j, so we can look up the name property on nodes.

Raw script source

```
g.idx('node_auto_index')[[name:'I']].out
```

Figure 22.56. Final Graph



Example request

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/x-www-form-urlencoded`

```
script=g.idx%28%27node_auto_index%27%29%5B%5Bname%3A%27I%27%5D%5D.out
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/13/relationships/out",
  "data" : {
    "name" : "you"
  },
  "traverse" : "http://localhost:7474/db/data/node/13/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/13/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/13/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/13",
  "properties" : "http://localhost:7474/db/data/node/13/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/13/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/13/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/13/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/13/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/13/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/13/relationships/in/{-list|&|types}"
} ]
```


22.18.2. Load a sample graph

Importing a graph from a [GraphML](http://graphml.graphdrawing.org/) <<http://graphml.graphdrawing.org/>> file can be achieved through the Gremlin GraphMLReader. The following script imports a small GraphML file from an URL into Neo4j, resulting in the depicted graph. The underlying database is auto-indexed, see [Section 19.12, “Automatic Indexing”](#) so the script can return the imported node by index lookup.

Raw script source

```
g.clear()
g.loadGraphML('file:/var/lib/jenkins/jobs/neo4j-1.9-release-candidate/workspace/target/gremlin-plugin/target/test-classes/graphml.xml')
g.idx('node_auto_index')[[name:'you']]
```

Figure 22.57. Final Graph



Example request

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`

- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "g.clear();g.loadGraphML('file:/var/lib/jenkins/jobs/neo4j-1.9-release-candidate/workspace/target/gremlin-plugin/target.
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/45/relationships/out",
  "data" : {
    "name" : "you"
  },
  "traverse" : "http://localhost:7474/db/data/node/45/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/45/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/45/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/45",
  "properties" : "http://localhost:7474/db/data/node/45/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/45/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/45/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/45/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/45/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/45/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/45/relationships/in/{-list|&|types}"
} ]
```


22.18.3. Sort a result using raw Groovy operations

The following script returns a sorted list of all nodes connected via outgoing relationships to node 1, sorted by their name-property.

Raw script source

```
g.idx('node_auto_index')[[name:'I']].out.sort{it.name}
```

Figure 22.58. Final Graph



Example request

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "g.idx('node_auto_index')[[name:'I']].out.sort{it.name}",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/66/relationships/out",
  "data" : {
    "name" : "him"
  },
  "traverse" : "http://localhost:7474/db/data/node/66/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/66/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/66/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/66",
  "properties" : "http://localhost:7474/db/data/node/66/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/66/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/66/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/66/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/66/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/66/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/66/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/65/relationships/out",
  "data" : {
    "name" : "you"
  },
  "traverse" : "http://localhost:7474/db/data/node/65/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/65/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/65/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/65",
  "properties" : "http://localhost:7474/db/data/node/65/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/65/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/65/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/65/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/65/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/65/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/65/relationships/in/{-list|&|types}"
} ]
```


22.18.4. Send a Gremlin Script - JSON encoded with table results

To send a Script JSON encoded, set the payload Content-Type Header. In this example, find all the things that my friends like, and return a table listing my friends by their name, and the names of the things they like in a table with two columns, ignoring the third named step variable I. Remember that everything in Gremlin is an iterator - in order to populate the result table t, iterate through the pipes with iterate().

Raw script source

```
t= new Table()
g.v(49).as('I').out('know').as('friend').out('like').as('likes').table(t,['friend','likes']){it.name}{it.name}.iterate()
t
```

Figure 22.59. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "script" : "t= new Table();g.v(49).as('I').out('know').as('friend').out('like').as('likes').table(t,['friend','likes']){it.name}{it.name}.cap",
  "params" : {
  }
}

```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```


{
  "columns" : [ "friend", "likes" ],
  "data" : [ [ "Joe", "cats" ], [ "Joe", "dogs" ] ]
}

```

22.18.5. Returning nested pipes*Raw script source*

```
g.v(23).as('I').out('know').as('friend').out('like').as('likes').table(new Table()){it.name}{it.name}.cap
```

Figure 22.60. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "script" : "g.v(23).as('I').out('know').as('friend').out('like').as('likes').table(new Table()){it.name}{it.name}.cap;",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "data" : [ [ "I", "Joe", "cats" ], [ "I", "Joe", "dogs" ] ],
  "columns" : [ "I", "friend", "likes" ]
} ]
```

22.18.6. Set script variables

To set variables in the bindings for the Gremlin Script Engine on the server, you can include a `params` parameter with a String representing a JSON map of variables to set to initial values. These can then be accessed as normal variables within the script.

Raw script source

```
meaning_of_life
```

Figure 22.61. Final Graph

Example request

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "script" : "meaning_of_life",
  "params" : {
    "meaning_of_life" : 42.0
  }
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
42.0
```


22.18.7. Send a Gremlin Script with variables in a JSON Map

Send a Gremlin Script, as JSON payload and additional parameters

Raw script source

```
g.v(me).out
```

Figure 22.62. Final Graph



Example request

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "g.v(me).out",
  "params" : {
    "me" : "32"
  }
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/31/relationships/out",
  "data" : {
    "name" : "you"
  },
  "traverse" : "http://localhost:7474/db/data/node/31/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/31/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/31/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/31",
}
```

```

"properties" : "http://localhost:7474/db/data/node/31/properties",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/31/relationships/out/{-list|&|types}",
"incoming_relationships" : "http://localhost:7474/db/data/node/31/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/31/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/31/paged/traverse/{returnType}{?pageSize,leaseTime}",
"all_relationships" : "http://localhost:7474/db/data/node/31/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/31/relationships/in/{-list|&|types}"
} ]

```


22.18.8. Return paths from a Gremlin script

The following script returns paths. Paths in Gremlin consist of the pipes that make up the path from the starting pipes. The server is returning JSON representations of their content as a nested list.

Raw script source

```
g.v(3).out.name.paths
```

Figure 22.63. Final Graph



Example request

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "script" : "g.v(3).out.name.paths",
  "params" : {
  }
}

```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

[ [ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/3/relationships/out",
  "data" : {
    "name" : "I"
  },
  "traverse" : "http://localhost:7474/db/data/node/3/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/3/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/3/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/3",
  "properties" : "http://localhost:7474/db/data/node/3/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/3/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/3/relationships/in",

```

```

"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/3/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/3/paged/traverse/{returnType}?pageSize,leaseTime",
"all_relationships" : "http://localhost:7474/db/data/node/3/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/3/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/1/relationships/out",
"data" : {
  "name" : "you"
},
"traverse" : "http://localhost:7474/db/data/node/1/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/1/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/1/properties/{key}",
"self" : "http://localhost:7474/db/data/node/1",
"properties" : "http://localhost:7474/db/data/node/1/properties",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/1/relationships/out/{-list|&|types}",
"incoming_relationships" : "http://localhost:7474/db/data/node/1/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/1/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/1/paged/traverse/{returnType}?pageSize,leaseTime",
"all_relationships" : "http://localhost:7474/db/data/node/1/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/1/relationships/in/{-list|&|types}"
}, "you" ], [ {
"outgoing_relationships" : "http://localhost:7474/db/data/node/3/relationships/out",
"data" : {
  "name" : "I"
},
"traverse" : "http://localhost:7474/db/data/node/3/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/3/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/3/properties/{key}",
"self" : "http://localhost:7474/db/data/node/3",
"properties" : "http://localhost:7474/db/data/node/3/properties",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/3/relationships/out/{-list|&|types}",
"incoming_relationships" : "http://localhost:7474/db/data/node/3/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/3/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/3/paged/traverse/{returnType}?pageSize,leaseTime",
"all_relationships" : "http://localhost:7474/db/data/node/3/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/3/relationships/in/{-list|&|types}"
}, {
"outgoing_relationships" : "http://localhost:7474/db/data/node/2/relationships/out",
"data" : {
  "name" : "him"
},
"traverse" : "http://localhost:7474/db/data/node/2/traverse/{returnType}",
"all_typed_relationships" : "http://localhost:7474/db/data/node/2/relationships/all/{-list|&|types}",
"property" : "http://localhost:7474/db/data/node/2/properties/{key}",
"self" : "http://localhost:7474/db/data/node/2",
"properties" : "http://localhost:7474/db/data/node/2/properties",
"outgoing_typed_relationships" : "http://localhost:7474/db/data/node/2/relationships/out/{-list|&|types}",
"incoming_relationships" : "http://localhost:7474/db/data/node/2/relationships/in",
"extensions" : {
},
"create_relationship" : "http://localhost:7474/db/data/node/2/relationships",
"paged_traverse" : "http://localhost:7474/db/data/node/2/paged/traverse/{returnType}?pageSize,leaseTime",
"all_relationships" : "http://localhost:7474/db/data/node/2/relationships/all",
"incoming_typed_relationships" : "http://localhost:7474/db/data/node/2/relationships/in/{-list|&|types}"
}, "him" ] ]

```

22.18.9. Send an arbitrary Groovy script - Lucene sorting

This example demonstrates that you via the Groovy runtime embedded with the server have full access to all of the servers Java APIs. The below example creates Nodes in the database both via the Blueprints and the Neo4j API indexes the nodes via the native Neo4j Indexing API constructs a custom Lucene sorting and searching returns a Neo4j IndexHits result iterator.

Raw script source

```
'***** Additional imports *****'
import org.neo4j.graphdb.index.*
import org.neo4j.graphdb.*
import org.neo4j.index.lucene.*
import org.apache.lucene.search.*

'**** Blueprints API methods on the injected Neo4jGraph at variable g ****'
meVertex = g.addVertex([name:'me'])
meNode = meVertex.getRawVertex()


'*** get the Neo4j raw instance ***'
neo4j = g.getRawGraph()

'***** Neo4j API methods: *****'
tx = neo4j.beginTx()
youNode = neo4j.createNode()
youNode.setProperty('name', 'you')
youNode.createRelationshipTo(meNode, DynamicRelationshipType.withName('knows'))

'*** index using Neo4j APIs ***'
idxManager = neo4j.index()
personIndex = idxManager.forNodes('persons')
personIndex.add(meNode, 'name', meNode.getProperty('name'))
personIndex.add(youNode, 'name', youNode.getProperty('name'))
tx.success()
tx.finish()

'*** Prepare a custom Lucene query context with Neo4j API ***'
query = new QueryContext( 'name:*' ).sort( new Sort(new SortField( 'name',SortField.STRING, true ) ) )
results = personIndex.query( query )
```

Figure 22.64. Final Graph



Example request

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
```

```
"script" : "'***** Additional imports *****';import org.neo4j.graphdb.index.*;import org.neo4j.graphdb.*;import org.neo4j.i
"params" : {
}
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/20/relationships/out",
  "data" : {
    "name" : "you"
  },
  "traverse" : "http://localhost:7474/db/data/node/20/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/20/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/20/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/20",
  "properties" : "http://localhost:7474/db/data/node/20/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/20/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/20/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/20/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/20/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/20/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/20/relationships/in/{-list|&|types}"
}, {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/19/relationships/out",
  "data" : {
    "name" : "me"
  },
  "traverse" : "http://localhost:7474/db/data/node/19/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/19/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/19/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/19",
  "properties" : "http://localhost:7474/db/data/node/19/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/19/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/19/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/19/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/19/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/19/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/19/relationships/in/{-list|&|types}"
} ]
```


22.18.10. Emit a sample graph

Exporting a graph can be done by simple emitting the appropriate String.

Raw script source

```
writer = new GraphMLWriter(g)
out = new java.io.ByteArrayOutputStream()
writer.outputGraph(out)
result = out.toString()
```

Figure 22.65. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```

{
  "script" : "writer = new GraphMLWriter(g);out = new java.io.ByteArrayOutputStream();writer.outputGraph(out);result = out.toString();"
  "params" : {
  }
}

```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```

"<?xml version='1.0' ?><graphml xmlns='http://graphml.graphdrawing.org/xmlns'><key id='name' for='node' attr.name='name' a

```

22.18.11. HyperEdges - find user roles in groups

Imagine a user being part of different groups. A group can have different roles, and a user can be part of different groups. He also can have different roles in different groups apart from the membership. The association of a User, a Group and a Role can be referred to as a *HyperEdge*. However, it can be easily modeled in a property graph as a node that captures this n-ary relationship, as depicted below in the U1G2R1 node.

To find out in what roles a user is for a particular groups (here *Group2*), the following script can traverse this HyperEdge node and provide answers.


Raw script source

```

g.v(39).out('hasRoleInGroup').as('hyperedge').out('hasGroup').filter{it.name=='Group2'}.back('hyperedge').out('hasRole').name

```

Figure 22.66. Final Graph



Example request

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "g.v(39).out('hasRoleInGroup').as('hyperedge').out('hasGroup').filter{it.name=='Group2'}.back('hyperedge').out('hasRole')",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: application/json; charset=UTF-8


```
[ "Role1" ]
```

22.18.12. Group count

This example is showing a group count in Gremlin, for instance the counting of the different relationship types connected to some the start node. The result is collected into a variable that then is returned.

Raw script source

```
m = [:]
g.v(60).bothE().label.groupCount(m).iterate()
m
```

Figure 22.67. Final Graph*Example request*

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "script" : "m = [:];g.v(60).bothE().label.groupCount(m).iterate();m",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
{
  "knows" : 2,
  "likes" : 1
}
```


22.18.13. Collect multiple traversal results

Multiple traversals can be combined into a single result, using splitting and merging pipes in a lazy fashion.

Raw script source

```
g.idx('node_auto_index')[[name:'Peter']].copySplit(_).out('knows'), _().in('likes')).fairMerge.name
```

Figure 22.68. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "script" : "g.idx('node_auto_index')[[name:'Peter']].copySplit(_).out('knows'), _().in('likes')).fairMerge.name",
  "params" : {
  }
}
```

Example response

- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ "Ian", "Marie" ]
```

22.18.14. Collaborative filtering


This example demonstrates basic collaborative filtering - ordering a traversal after occurrence counts and substracting objects that are not interesting in the final result.

Here, we are finding Friends-of-Friends that are not Joes friends already. The same can be applied to graphs of users that LIKE things and others.

Raw script source

```
x=[]
fof=[]
g.v(12).out('knows').aggregate(x).out('knows').except(x).groupCount(fof).iterate()
fof.sort{a,b -> b.value <=> a.value}
```

Figure 22.69. Final Graph

*Example request*

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "x=[];fof=[];g.v(12).out('knows').aggregate(x).out('knows').except(x).groupCount(fof).iterate();fof.sort{a,b -> b.value}
  "params" : {
  }
}
```

Example response


- 200: OK
- Content-Type: application/json; charset=UTF-8

```
{
  "v[10]" : 2,
  "v[9]" : 1,
  "v[11]" : 1
}
```

22.18.15. Chunking and offsetting in Gremlin*Raw script source*

```
g.v(30).out('knows').filter{it.name == 'Sara'}[0..100]
```

Figure 22.70. Final Graph

*Example request*

- POST `http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script`
- Accept: `application/json; charset=UTF-8`
- Content-Type: `application/json`

```
{
  "script" : "g.v(30).out('knows').filter{it.name == 'Sara'}[0..100]",
  "params" : {
  }
}
```

Example response


- 200: OK
- Content-Type: `application/json; charset=UTF-8`

```
[ {
  "outgoing_relationships" : "http://localhost:7474/db/data/node/28/relationships/out",
  "data" : {
    "name" : "Sara"
  },
  "traverse" : "http://localhost:7474/db/data/node/28/traverse/{returnType}",
  "all_typed_relationships" : "http://localhost:7474/db/data/node/28/relationships/all/{-list|&|types}",
  "property" : "http://localhost:7474/db/data/node/28/properties/{key}",
  "self" : "http://localhost:7474/db/data/node/28",
  "properties" : "http://localhost:7474/db/data/node/28/properties",
  "outgoing_typed_relationships" : "http://localhost:7474/db/data/node/28/relationships/out/{-list|&|types}",
  "incoming_relationships" : "http://localhost:7474/db/data/node/28/relationships/in",
  "extensions" : {
  },
  "create_relationship" : "http://localhost:7474/db/data/node/28/relationships",
  "paged_traverse" : "http://localhost:7474/db/data/node/28/paged/traverse/{returnType}{?pageSize,leaseTime}",
  "all_relationships" : "http://localhost:7474/db/data/node/28/relationships/all",
  "incoming_typed_relationships" : "http://localhost:7474/db/data/node/28/relationships/in/{-list|&|types}"
} ]
```

22.18.16. This example shows how to modify the graph while traversing it


In this case, the Peter node is disconnected from all other nodes.

Figure 22.71. Starting Graph

*Raw script source*

```
g.v(77).bothE.each{g.removeEdge(it)}
```

Figure 22.72. Final Graph

*Example request*

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "g.v(77).bothE.each{g.removeEdge(it)}",
  "params" : {
  }
}
```

Example response


- 200: OK
- Content-Type: application/json; charset=UTF-8

```
[ ]
```

22.18.17. Flow algorithms with Gremlin

This is a basic stub example for implementing flow algorithms in for instance [Flow Networks](http://en.wikipedia.org/wiki/Flow_network) <http://en.wikipedia.org/wiki/Flow_network> with a pipes-based approach and scripting, here between source and sink using the capacity property on relationships as the base for the flow function and modifying the graph during calculation.

Figure 22.73. Starting Graph


*Raw script source*

```

source=g.v(79)
sink=g.v(80)
maxFlow = 0
source.outE.inV.loop(2){!it.object.equals(sink)}.paths.each{flow = it.capacity.min()
  maxFlow += flow
it.findAll{it.capacity}.each{it.capacity -= flow}}
maxFlow

```

Figure 22.74. Final Graph

*Example request*

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "source=g.v(79);sink=g.v(80);maxFlow = 0;source.outE.inV.loop(2){!it.object.equals(sink)}.paths.each{flow = it.capacity
  "params" : {
  }
}
```

Example response


- 200: OK
- Content-Type: application/json; charset=UTF-8

4

22.18.18. Script execution errors

Script errors will result in an HTTP error response code.

Figure 22.75. Final Graph



Example request

- POST http://localhost:7474/db/data/ext/GremlinPlugin/graphdb/execute_script
- Accept: application/json; charset=UTF-8
- Content-Type: application/json

```
{
  "script" : "g.addVertex([name: {}])"
}
```

Example response

- 400: Bad Request
- Content-Type: application/json; charset=UTF-8

```
{
  "message": "javax.script.ScriptException: java.lang.IllegalArgumentException: Unknown property type on: Script9$_run_closure1@5e21",
  "exception": "BadInputException",
  "fullname": "org.neo4j.server.rest.repr.BadInputException",
  "stacktrace": [
    "org.neo4j.server.plugin.gremlin.GremlinPlugin.executeScript(GremlinPlugin.java:84)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.plugins.PluginMethod.invoke(PluginMethod.java:57)",
    "org.neo4j.server.plugins.PluginManager.invoke(PluginManager.java:168)",
    "org.neo4j.server.rest.web.ExtensionService.invokeGraphDatabaseExtension(ExtensionService.java:300)",
    "org.neo4j.server.rest.web.ExtensionService.invokeGraphDatabaseExtension(ExtensionService.java:122)",
    "java.lang.reflect.Method.invoke(Method.java:597)",
    "org.neo4j.server.rest.security.SecurityFilter.doFilter(SecurityFilter.java:112)"
  ]
}
```

Chapter 23. Deprecations

This section outlines deprecations in Neo4j 1.9 in order to help you find a smoother transition path to future releases.

Read operations outside of transactions

In future releases, read operations are only allowed within transactions. This only affects embedded usage, where you can already start to wrap read operations in transactions to be prepared for the future.

The ! and ? property operators in Cypher

In future releases, expressions like `node.property = "value"` will not fail when a node without the property is encountered, and will instead return `NULL`. This is the same behavior as `node.property! = "value"`, and thus the `!` property operator is deprecated and will be removed. Expressions using the `?` property operator, such as `node.property? = "value"`, are also deprecated. Please use `not(has(node.property)) OR node.property = "value"` instead.

Cypher CREATE syntax

The `CREATE a={foo:'bar'}` syntax has been deprecated. Instead, use `CREATE (a {foo:'bar'})`.

Using DELETE to remove properties in Cypher

The `DELETE a.prop` syntax has been deprecated. Instead, use `REMOVE a.prop`.

Using parameters for index keys in Cypher

In future releases, parameters can not be used as the key in `START` clauses using indexes (ie. `START n=node:index({key}='value')`). Use the literal key names instead.

Using parameters to identify nodes in patterns

In future releases, parameters can not be used to identify nodes in a pattern (ie. `MATCH ({node})-->(other))`). Note that this form was only possible when mixing the embedded Java API and Cypher, and thus does not affect users of Neo4j Server.

Iteration syntax in FOREACH, EXTRACT, etc

In future releases, the iterating functions will use a `|` instead of a `:` to separate the components of the statement. For example, `EXTRACT(n in ns : n.prop)` is replaced with `EXTRACT(n in ns | n.prop)`. The iterating functions include `FOREACH`, `EXTRACT`, `REDUCE`, `ANY`, `ALL`, `SINGLE` and `NONE`. Note the alternative form using `|` is also possible in Neo4j 1.9.2 and later.

Alternative WITH syntax

In future releases, the alternative `WITH` syntax, `=== <identifiers> ===`, can not be used. Use the `WITH` keyword instead.

Plugins

Plugins will no longer be distributed with Neo4j. Please see individual maintainers about availability. For instance, the source for the Gremlin plugin will be available at: <https://github.com/neo4j-contrib/gremlin-plugin>

Graph Matching

The graph-matching component will be removed in future releases.

The Reference Node

With the introduction of Labels in Neo4j 2.0 the Reference Node becomes obsolete and will be removed. Instead, labeled nodes become the well-known starting points in your graph. See the Neo4j Manual 2.0 for information on how to handle this.

Part V. Operations

This part describes how to install and maintain a Neo4j installation. This includes topics such as backing up the database and monitoring the health of the database as well as diagnosing issues.

Chapter 24. Installation & Deployment

24.1. Deployment Scenarios

Neo4j can be embedded into your application, run as a standalone server or deployed on several machines to provide high availability.

Neo4j deployment options

	Single Instance	Multiple Instances
Embedded	EmbeddedGraphDatabase	HighlyAvailableGraphDatabase
Standalone	Neo4j Server	Neo4j Server high availability mode

24.1.1. Server

Neo4j is normally accessed as a standalone server, either directly through a REST interface or through a language-specific driver. More information about Neo4j server is found in [Chapter 21, Neo4j Server](#). For running the server and embedded installations in high availability mode, see [Chapter 26, High Availability](#).

24.1.2. Embedded

Neo4j can be embedded directly in a server application by including the appropriate Java libraries. When programming, you can refer to the [GraphDatabaseService](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/GraphDatabaseService.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/graphdb/GraphDatabaseService.html>> API. To switch from a single instance to multiple highly available instances, simply switch from the concrete [EmbeddedGraphDatabase](http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/EmbeddedGraphDatabase.html) <<http://components.neo4j.org/neo4j/1.9.7/apidocs/org/neo4j/kernel/EmbeddedGraphDatabase.html>> to the [HighlyAvailableGraphDatabase](http://components.neo4j.org/neo4j-enterprise/1.9.7/apidocs/org/neo4j/kernel/HighlyAvailableGraphDatabase.html) <<http://components.neo4j.org/neo4j-enterprise/1.9.7/apidocs/org/neo4j/kernel/HighlyAvailableGraphDatabase.html>>.

24.2. System Requirements

Memory constrains graph size, disk I/O constrains read/write performance, as always.

24.2.1. CPU

Performance is generally memory or I/O bound for large graphs, and compute bound for graphs which fit in memory.

Minimum

Intel Core i3

Recommended

Intel Core i7

24.2.2. Memory

More memory allows even larger graphs, but runs the risk of inducing larger Garbage Collection operations.

Minimum

2GB

Recommended

16—32GB or more

24.2.3. Disk

Aside from capacity, the performance characteristics of the disk are the most important when selecting storage.

Minimum

10GB SATA

Recommended

SSD w/ SATA

24.2.4. Filesystem

For proper ACID behavior, the filesystem must support flush (fsync, fdatasync).

Minimum

ext4 (or similar)

Recommended

ext4, ZFS

24.2.5. Software

Neo4j is Java-based.

Java

[Oracle Java 7](http://www.oracle.com/technetwork/java/javase/downloads/index.html) <<http://www.oracle.com/technetwork/java/javase/downloads/index.html>>

Operating Systems

Linux, HP UX, Windows 2008 for production; additionally Windows XP, Mac OS X for development.

24.3. Installation

Neo4j can be installed as a server, running either as a headless application or system service. For Java developers, it is also possible to use Neo4j as a library, embedded in your application.

For information on installing Neo4j as a server, see [Section 21.1, “Server Installation”](#).

The following table outlines the available editions and their names for use with dependency management tools.



Tip

Follow the links in the table for details on dependency configuration with Apache Maven, Apache Buildr, Apache Ivy and Groovy Grape!

Neo4j editions

Edition	Dependency	Description	License
Community	org.neo4j:neo4j < http://search.maven.org/#search%3A%22org.neo4j%22%20AND%20a%3A%22neo4j%22 >	a high performance, fully ACID transactional graph database	GPLv3
Advanced	org.neo4j:neo4j-advanced < http://search.maven.org/#search%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-advanced%22 >	adding advanced monitoring	AGPLv3
Enterprise	org.neo4j:neo4j-enterprise < http://search.maven.org/#search%3A%22org.neo4j%22%20AND%20a%3A%22neo4j-enterprise%22 >	adding online backup and High Availability clustering	AGPLv3



Note

The listed dependencies do not contain the implementation, but pulls it in transitively.

For more information regarding licensing, see the [Licensing Guide](http://www.neo4j.org/learn/licensing) <<http://www.neo4j.org/learn/licensing>>.

24.3.1. Embedded Installation

The latest release is always available from <http://neo4j.org/download>, included as part of the Neo4j download packages. After selecting the appropriate version for your platform, embed Neo4j in your Java application by including the Neo4j library jars in your build. Either take the jar files from the *lib/*

directory of the download, or directly use the artifacts available from Maven Central Repository ¹. Stable and milestone releases are available there.

Maven dependency.

```
<project>
...
<dependencies>
  <dependency>
    <groupId>org.neo4j</groupId>
    <artifactId>neo4j</artifactId>
    <version>1.9.7</version>
  </dependency>
  ...
</dependencies>
...
</project>
```

Where *the artifactId* is one of `neo4j`, `neo4j-advanced`, `neo4j-enterprise`.

24.3.2. Server Installation

please refer to [Chapter 21, *Neo4j Server*](#) and [Section 21.1, “Server Installation”](#).

¹<http://repo1.maven.org/maven2/org/neo4j/>

24.4. Upgrading

A database can be upgraded from a minor version to the next, e.g. 1.1 → 1.2, and 1.2 → 1.3, but you can not jump directly from 1.1 → 1.3. For version 1.8 in particular, it is possible to upgrade directly from version 1.5.3 and later, as an explicit upgrade. The upgrade process is a one way step; databases cannot be downgraded.

For most upgrades, only small changes are required to the database store, and these changes proceed automatically when you start up the database using the newer version of Neo4j.

However, some upgrades require more significant changes to the database store. In these cases, Neo4j will refuse to start without explicit configuration to allow the upgrade.

The table below lists recent Neo4j versions, and the type of upgrade required.

Upgrade process for Neo4j version

From Version	To Version	Upgrade Type
1.3	1.4	Automatic
1.4	1.5	Explicit
1.5	1.6	Explicit
1.6	1.7	Automatic
1.7	1.8	Automatic
1.8	1.9	Automatic



Note

Downgrade is supported only between versions which do not have incompatible store layouts. That means that if you did an upgrade where you didn't have to explicitly set the `allow_store_upgrade` flag to false then you can downgrade without any problems to the previous version used. Otherwise downgrading is not supported. In any case, downgrading currently cannot be done in a rolling fashion, even in HA deployments. Instead, the whole cluster must be shutdown and each machine downgraded individually and then service can be resumed.

24.4.1. Automatic Upgrade

To perform a normal upgrade (for minor changes to the database store):

1. download the newer version of Neo4j
2. cleanly shutdown the database to upgrade, if it is running
3. startup the database with the newer version of Neo4j

24.4.2. Explicit Upgrade

To perform a special upgrade (for significant changes to the database store):

1. make sure the database you are upgrading has been cleanly shut down
2. set the Neo4j configuration parameter "`allow_store_upgrade=true`" in your `neo4j.properties` or embedded configuration
3. start the database
4. the upgrade will happen during startup and the process is done when the database has been successfully started

5. "allow_store_upgrade=true" configuration parameter should be removed, set to "false" or commented out

24.4.3. Upgrade 1.8 → 1.9

There are no store format changes between these versions so upgrading standalone instances simply consists of starting the database with the newer version. In the case of High Availability (HA) installations, the communication protocol and the master election algorithm have changed and a new "rolling upgrade" feature has been added, removing the need to shut down the entire cluster. For more information, refer to the "Upgrading a Neo4j HA Cluster" chapter of the HA section of the Neo4j manual.

24.4.4. Upgrade 1.7 → 1.8

There are no store format changes between these versions so upgrading standalone instances simply consists of starting the database with the newer version. In the case of High Availability (HA) installations, the communication protocol and the master election algorithm have changed and a new "rolling upgrade" feature has been added, removing the need to shut down the entire cluster. For more information, refer to the "Upgrading a Neo4j HA Cluster" chapter of the HA section of the Neo4j manual.

24.4.5. Upgrade 1.6 → 1.7

There are no store format changes between these versions, which means there is no particular procedure you need to upgrade a single instance.

In an HA environment these steps need to be performed:

1. shut down all the databases in the cluster
2. shut down the ZooKeeper cluster and clear the *version-2* directories on all the ZooKeeper instances
3. start the ZooKeeper cluster again
4. remove the databases except the master and start the master database with 1.7
5. start up the other databases so that they get a copy from the master

24.4.6. Upgrade 1.5 → 1.6

This upgrade changes lucene version from 3.1 to 3.5. The upgrade itself is done by Lucene by loading an index.

In an HA environment these steps need to be performed:

1. shut down all the databases in the cluster
2. shut down the ZooKeeper cluster and clear the *version-2* directories on all the ZooKeeper instances
3. start the ZooKeeper cluster again
4. start up the other databases so that they get a copy from the master

24.4.7. Upgrade 1.4 → 1.5

This upgrade includes a significant change to the layout of property store files, which reduces their size on disk, and improves IO performance. To achieve this layout change, the upgrade process takes some time to process the whole of the existing database. You should budget for several minutes per gigabyte of data as part of your upgrade planning.



Warning

The upgrade process for this upgrade temporarily requires additional disk space, for the period while the upgrade is in progress. Before starting the upgrade to Neo4j 1.5, you should ensure that the machine performing the upgrade has free space equal to the current size of the database on disk. You can find the current space occupied by the database by inspecting the store file directory (*data/graph.db* is the default location in Neo4j server). Once the upgrade is complete, this additional space is no longer required.

24.5. Usage Data Collector

The Neo4j Usage Data Collector is a sub-system that gathers usage data, reporting it to the UDC-server at udc.neo4j.org. It is easy to disable, and does not collect any data that is confidential. For more information about what is being sent, see below.

The Neo4j team uses this information as a form of automatic, effortless feedback from the Neo4j community. We want to verify that we are doing the right thing by matching download statistics with usage statistics. After each release, we can see if there is a larger retention span of the server software.

The data collected is clearly stated here. If any future versions of this system collect additional data, we will clearly announce those changes.

The Neo4j team is very concerned about your privacy. We do not disclose any personally identifiable information.

24.5.1. Technical Information

To gather good statistics about Neo4j usage, UDC collects this information:

- Kernel version: The build number, and if there are any modifications to the kernel.
- Store id: A randomized globally unique id created at the same time a database is created.
- Ping count: UDC holds an internal counter which is incremented for every ping, and reset for every restart of the kernel.
- Source: This is either "neo4j" or "maven". If you downloaded Neo4j from the Neo4j website, it's "neo4j", if you are using Maven to get Neo4j, it will be "maven".
- Java version: The referrer string shows which version of Java is being used.
- MAC address to uniquely identify instances behind firewalls.
- Registration id: For registered server instances.
- Tags about the execution context (e.g. test, language, web-container, app-container, spring, ejb).
- Neo4j Edition (community, advanced, enterprise).
- A hash of the current cluster name (if any).
- Distribution information for Linux (rpm, dpkg, unknown).
- User-Agent header for tracking usage of REST client drivers

After startup, UDC waits for ten minutes before sending the first ping. It does this for two reasons; first, we don't want the startup to be slower because of UDC, and secondly, we want to keep pings from automatic tests to a minimum. The ping to the UDC servers is done with a HTTP GET.

24.5.2. How to disable UDC

We've tried to make it extremely easy to disable UDC. In fact, the code for UDC is not even included in the kernel jar but as a completely separate component.

There are three ways you can disable UDC:

1. The easiest way is to just remove the `neo4j-udc-*.jar` file. By doing this, the kernel will not load UDC, and no pings will be sent.
2. If you are using Maven, and want to make sure that UDC is never installed in your system, a dependency element like this will do that:

```
<dependency>  
<groupId>org.neo4j</groupId>
```

```
<artifactId>neo4j</artifactId>
<version>${neo4j-version}</version>
<exclusions>
  <exclusion>
    <groupId>org.neo4j</groupId>
    <artifactId>neo4j-udc</artifactId>
  </exclusion>
</exclusions>
</dependency>
```

Where *`${neo4j-version}`* is the Neo4j version in use.

3. Lastly, if you are using a packaged version of Neo4j, and do not want to make any change to the jars, a system property setting like this will also make sure that UDC is never activated: -Dneo4j.ext.udc.disable=true.

Chapter 25. Configuration & Performance

In order to get optimum performance out of Neo4j for your application there are a few parameters that can be tweaked. The two main components that can be configured are the Neo4j caches and the JVM that Neo4j runs in. The following sections describe how to tune these.

25.1. Introduction

To gain good performance, these are the things to look into first:

- Make sure the JVM is not spending too much time performing garbage collection. Monitoring heap usage on an application that uses Neo4j can be a bit confusing since Neo4j will increase the size of caches if there is available memory and decrease if the heap is getting full. The goal is to have a large enough heap to make sure that heavy/peak load will not result in so called GC trashing (performance can drop as much as two orders of magnitude when GC trashing happens).
- Start the JVM with the `-server` flag and a good sized heap (see [Section 25.6, “JVM Settings”](#)). Having too large heap may also hurt performance so you may have to try some different heap sizes.
- Use the parallel/concurrent garbage collector (we found that `-XX:+UseConcMarkSweepGC` works well in most use-cases)

25.1.1. How to add configuration settings

When creating the embedded Neo4j instance it is possible to pass in parameters contained in a map where keys and values are strings, see [the section called “Starting an embedded database with configuration settings”](#) for an example.

If no configuration is provided, the Database Kernel will try to determine suitable settings from the information available via the JVM settings and the underlying operating system.

The JVM is configured by passing command line flags when starting the JVM. The most important configuration parameters for Neo4j are the ones that control the memory and garbage collector, but some of the parameters for configuring the Just In Time compiler are also of interest.

This is an example of starting up your applications main class using 64-bit server VM mode and a heap space of 1GB:

```
java -d64 -server -Xmx1024m -cp /path/to/neo4j-kernel.jar:/path/to/jta.jar:/path/to/your-application.jar com.example.yourapp.MainClass
```

Looking at the example above you will also notice one of the most basic command line parameters: the one for specifying the classpath. The classpath is the path in which the JVM searches for your classes. It is usually a list of jar-files. Specifying the classpath is done by specifying the flag `-cp` (or `-classpath`) and then the value of the classpath. For Neo4j applications this should at least include the path to the Neo4j `neo4j-kernel.jar` and the Java Transaction API (`jta.jar`) as well as the path where the classes for your application are located.



Tip

On Linux, Unix and Mac OS X each element in the path list are separated by a colon symbol (:), on Windows the path elements are separated by a semicolon (;).

When using the Neo4j REST server, see [Section 21.2, “Server Configuration”](#) for how to add configuration settings for the database to the server.

25.2. Performance Guide

This is the Neo4j performance guide. It will attempt to guide you in how to use Neo4j to achieve maximum performance.

25.2.1. Try this first

The first thing is to make sure the JVM is running well and not spending too much time in garbage collection. Monitoring heap usage on an application that uses Neo4j can be a bit confusing since Neo4j will increase the size of caches if there is available memory and decrease if the heap is getting full. The goal is to have a large enough heap so heavy/peak load will not result in so called GC trashing (performance can drop as much as two orders of magnitude when this happens).

Start the JVM with `-server` flag and `-Xmx`<good sized heap> (f.ex. `-Xmx512M` for 512Mb memory or `-Xmx3G` for 3Gb memory). Having too large heap may also hurt performance so you may have to try some different heap sizes. Make sure parallel/concurrent garbage collector is running (`-XX:+UseConcMarkSweepGC` works well in most use-cases).

Finally make sure the OS has some memory to manage proper file system caches meaning if your server has 8GB of RAM don't use all of that RAM for heap (unless you turned off memory mapped buffers) but leave a good size of it to the OS. For more information on this see [Chapter 25, Configuration & Performance](#).

For Linux specific tweaks, see [Section 25.10, "Linux Performance Guide"](#).

25.2.2. Neo4j primitives' lifecycle

Neo4j manages its primitives (nodes, relationships and properties) differently depending on how you use Neo4j. For example if you never get a property from a certain node or relationship that node or relationship will not have its properties loaded into memory. The first time, after loading a node or relationship, any property is accessed all the properties are loaded for that entity. If any of those properties contain an array larger than a few elements or a long string such values are loaded on demand when requesting them individually. Similarly, relationships of a node will only be loaded the first time they are requested for that node.

Nodes and relationships are cached using LRU caches. If you (for some strange reason) only work with nodes the relationship cache will become smaller and smaller while the node cache is allowed to grow (if needed). Working with many relationships and few nodes results in bigger relationship cache and smaller node cache.

The Neo4j API specification does not say anything about order regarding relationships so invoking

```
Node.getRelationships()
```

may return the relationships in a different order than the previous invocation. This allows us to make even heavier optimizations returning the relationships that are most commonly traversed.

All in all Neo4j has been designed to be very adaptive depending on how it is used. The (unachievable) overall goal is to be able to handle any incoming operation without having to go down and work with the file/disk I/O layer.

25.2.3. Configuring Neo4j

In [Chapter 25, Configuration & Performance](#) page there's information on how to configure Neo4j and the JVM. These settings have a lot of impact on performance.

Disks, RAM and other tips

As always, as with any persistence solution, performance is very much depending on the persistence media used. Better disks equals better performance.

If you have multiple disks or persistence media available it may be a good idea to split the store files and transaction logs across those disks. Having the store files running on disks with low seek time can do wonders for non cached read operations. Today a typical mechanical drive has an average seek time of about 5ms, this can cause a query or traversal to be very slow when available RAM is too low or caches and memory mapped settings are badly configured. A new good SATA enabled SSD has an average seek time of <100 microseconds meaning those scenarios will execute at least 50 times faster.

To avoid hitting disk you need more RAM. On a standard mechanical drive you can handle graphs with a few tens of millions of primitives with 1-2GB of RAM. 4-8GB of RAM can handle graphs with hundreds of millions of primitives while you need a good server with 16-32GB to handle billions of primitives. However, if you invest in a good SSD you will be able to handle much larger graphs on less RAM.

Use tools like `vmstat` or equivalent to gather info when your application is running. If you have high I/O waits and not that many blocks going out/in to disks when running write/read transactions its a sign that you need to tweak your Java heap, Neo4j cache and memory mapped settings (maybe even get more RAM or better disks).

Write performance

If you are experiencing poor write performance after writing some data (initially fast, then massive slowdown) it may be the operating system writing out dirty pages from the memory mapped regions of the store files. These regions do not need to be written out to maintain consistency so to achieve highest possible write speed that type of behavior should be avoided.

Another source of writes slow down can be the transaction size. Many small transactions result in a lot of I/O writes to disc and should be avoided. Too big transactions can result in `OutOfMemory` errors, since the uncommitted transaction data is held on the Java Heap in memory. On details about transaction management in Neo4j, please read the [Chapter 17, *Transaction Management*](#) guidelines.

The Neo4j kernel makes use of several store files and a logical log file to store the graph on disk. The store files contain the actual graph and the log contains modifying operations. All writes to the logical log are append-only and when a transaction is committed changes to the logical log will be forced (`fdatasync`) down to disk. The store files are however not flushed to disk and writes to them are not append-only either. They will be written to in a more or less random pattern (depending on graph layout) and writes will not be forced to disk until the log is rotated or the Neo4j kernel is shut down.

Since random writes to memory mapped regions for the store files may happen it is very important that the data does not get written out to disk unless needed. Some operating systems have very aggressive settings regarding when to write out these dirty pages to disk. If the OS decides to start writing out dirty pages of these memory mapped regions, write access to disk will stop being sequential and become random. That hurts performance a lot, so to get maximum write performance when using Neo4j make sure the OS is configured not to write out any of the dirty pages caused by writes to the memory mapped regions of the store files. As an example, if the machine has 8GB of RAM and the total size of the store files is 4GB (fully memory mapped) the OS has to be configured to accept at least 50% dirty pages in virtual memory to make sure we do not get random disk writes.

Note: make sure to read the [Section 25.10, “Linux Performance Guide”](#) as well for more specific information.

Second level caching

While normally building applications and "always assume the graph is in memory", sometimes it is necessary to optimize certain performance critical sections. Neo4j adds a small overhead even if the node, relationship or property in question is cached when you compare to in memory data structures. If this becomes an issue use a profiler to find these hot spots and then add your own second-level caching. We believe second-level caching should be avoided to greatest extend possible since it will force you to take care of invalidation which sometimes can be hard. But when everything else fails you have to use it so here is an example of how it can be done.

We have some POJO that wraps a node holding its state. In this particular POJO we've overridden the equals implementation.

```
public boolean equals( Object obj )
{
    return underlyingNode.getProperty( "some_property" ).equals( obj );
}

public int hashCode()
{
    return underlyingNode.getProperty( "some_property" ).hashCode();
}
```

This works fine in most scenarios but in this particular scenario many instances of that POJO is being worked with in nested loops adding/removing/getting/finding to collection classes. Profiling the applications will show that the equals implementation is being called many times and can be viewed as a hot spot. Adding second-level caching for the equals override will in this particular scenario increase performance.

```
private Object cachedProperty = null;

public boolean equals( Object obj )
{
    if ( cachedProperty == null )
    {
        cachedProperty = underlyingNode.getProperty( "some_property" );
    }
    return cachedProperty.equals( obj );
}

public int hashCode()
{
    if ( cachedProperty == null )
    {
        cachedProperty = underlyingNode.getProperty( "some_property" );
    }
    return cachedProperty.hashCode();
}
```

The problem now is that we need to invalidate the cached property whenever the `some_property` is changed (may not be a problem in this scenario since the state picked for equals and hash code computation often won't change).



Tip

To sum up, avoid second-level caching if possible and only add it when you really need it.

25.3. Kernel configuration

These are the configuration options you can pass to the neo4j kernel. They can either be passed as a map when using the embedded database, or in the neo4j.properties file when using the Neo4j Server.

All stores total mapped memory size

all_stores_total_mapped_memory_size

The size to allocate for a memory mapping pool to be shared between all stores.

Default value: 524288000

Allow store upgrade

allow_store_upgrade

Whether to allow a store upgrade in case the current version of the database starts against an older store version. Setting this to true does not guarantee successful upgrade, just that it allows an attempt at it.

Default value: false

Array block size

array_block_size

Specifies the block size for storing arrays. This parameter is only honored when the store is created, otherwise it is ignored. The default block size is 120 bytes, and the overhead of each block is the same as for string blocks, i.e., 8 bytes.

Default value: 120

Backup slave

backup_slave

Mark this database as a backup slave.

Default value: false

Cache type

cache_type

The type of cache to use for nodes and relationships. Note that the Neo4j Enterprise Edition has the additional 'gcr' cache type. See the chapter on caches in the manual for more information.

Default value: soft

Cypher parser version

cypher_parser_version

Enable this to specify a parser other than the default one.

Dump configuration

dump_configuration

Print out the effective Neo4j configuration after startup.

Default value: false

Forced kernel id

forced_kernel_id

An identifier that uniquely identifies this graph database instance within this JVM. Defaults to an auto-generated number depending on how many instance are started in this JVM.

Gc monitor threshold

gc_monitor_threshold

The amount of time in ms the monitor thread has to be blocked before logging a message it was blocked.

Default value: 200

Gc monitor wait time

gc_monitor_wait_time

Amount of time in ms the GC monitor thread will wait before taking another measurement.

Default value: 100

Gcr cache min log interval

gcr_cache_min_log_interval

The minimal time that must pass in between logging statistics from the cache (when using the 'gcr' cache).

Default value: 60s

Grab file lock

grab_file_lock

Whether to grab locks on files or not.

Default value: true

Intercept committing transactions

intercept_committing_transactions

Determines whether any TransactionInterceptors loaded will intercept prepared transactions before they reach the logical log.

Default value: false

Intercept deserialized transactions

intercept_deserialized_transactions

Determines whether any TransactionInterceptors loaded will intercept externally received transactions (e.g. in HA) before they reach the logical log and are applied to the store.

Default value: false

Keep logical logs

keep_logical_logs

Default value: 7 days

Make Neo4j keep the logical transaction logs for being able to backup the database. Can be used for specifying the threshold to prune logical logs after. For example "10 days" will prune logical logs that only contains transactions older than 10 days from the current time, or "100k txs" will keep the 100k latest transactions and prune any older transactions.

Default value: 7 days

Load kernel extensions

load_kernel_extensions

Enable loading kernel extensions

Default value: true

Log mapped memory stats

log_mapped_memory_stats

Tell Neo4j to regularly log memory mapping statistics.

Default value: false

Log mapped memory stats filename

log_mapped_memory_stats_filename

The file where Neo4j will record memory mapping statistics.

Default value: mapped_memory_stats.log

Log mapped memory stats interval

log_mapped_memory_stats_interval

The number of records to be loaded between regular logging of memory mapping statistics.

Default value: 1000000

Logging.threshold for rotation

logging.threshold_for_rotation

Threshold in bytes for when database logs (text logs, for debugging, that is) are rotated.

Default value: 104857600

Logical log

logical_log

The base name for the logical log files, either an absolute path or relative to the store_dir setting. This should generally not be changed.

Default value: nioneo_logical.log

Logical log rotation threshold

logical_log_rotation_threshold

Specifies at which file size the logical log will auto-rotate. 0 means that no rotation will automatically occur based on file size. Default is 25M

Default value: 26214400

Lucene searcher cache size

lucene_searcher_cache_size

Integer value that sets the maximum number of open lucene index searchers.

Default value: 2147483647

Lucene writer cache size

lucene_writer_cache_size

NOTE: This no longer has any effect. Integer value that sets the maximum number of open lucene index writers.

Default value: 2147483647

Mapped memory page size

mapped_memory_page_size

Target size for pages of mapped memory.

Default value: 1048576

Neo store

neo_store

The base name for the Neo4j Store files, either an absolute path or relative to the store_dir setting. This should generally not be changed.

Default value: neostore

Neostore.nodestore.db.mapped memory

neostore.nodestore.db.mapped_memory

The size to allocate for memory mapping the node store.

Default value: 20971520

Neostore.propertystore.db.arrays.mapped memory

neostore.propertystore.db.arrays.mapped_memory

The size to allocate for memory mapping the array property store.

Default value: 136314880

Neostore.propertystore.db.index.keys.mapped memory

neostore.propertystore.db.index.keys.mapped_memory

The size to allocate for memory mapping the store for property key strings.

Default value: 1048576

Neostore.propertystore.db.index.mapped memory

neostore.propertystore.db.index.mapped_memory

The size to allocate for memory mapping the store for property key indexes.

Default value: 1048576

Neostore.propertystore.db.mapped memory

neostore.propertystore.db.mapped_memory

The size to allocate for memory mapping the property value store.

Default value: 94371840

Neostore.propertystore.db.strings.mapped memory

neostore.propertystore.db.strings.mapped_memory

The size to allocate for memory mapping the string property store.

Default value: 136314880

Neostore.relationshipstore.db.mapped memory

neostore.relationshipstore.db.mapped_memory

The size to allocate for memory mapping the relationship store.

Default value: 104857600

Node auto indexing

node_auto_indexing

Controls the auto indexing feature for nodes. Setting to false shuts it down, while true enables it by default for properties listed in the node_keys_indexable setting.

Default value: false

Node cache array fraction

node_cache_array_fraction

The fraction of the heap (1%-10%) to use for the base array in the node cache (when using the 'gcr' cache).

Default value: 1.0

Node cache size

node_cache_size

The amount of memory to use for the node cache (when using the 'gcr' cache).

Node keys indexable

node_keys_indexable

A list of property names (comma separated) that will be indexed by default. This applies to Nodes only.

Query cache size

query_cache_size

Used to set the number of Cypher query execution plans that are cached.

Default value: 100

Read only database

read_only

Only allow read operations from this Neo4j instance. This mode still requires write access to the directory for lock purposes

Default value: false

Rebuild idgenerators fast

rebuild_idgenerators_fast

Use a quick approach for rebuilding the ID generators. This give quicker recovery time, but will limit the ability to reuse the space of deleted entities.

Default value: true

Relationship auto indexing

relationship_auto_indexing

Controls the auto indexing feature for relationships. Setting to false shuts it down, while true enables it by default for properties listed in the relationship_keys_indexable setting.

Default value: false

Relationship cache array fraction

relationship_cache_array_fraction

The fraction of the heap (1%-10%) to use for the base array in the relationship cache (when using the 'gcr' cache).

Default value: 1.0

Relationship cache size

relationship_cache_size

The amount of memory to use for the relationship cache (when using the 'gcr' cache).

Relationship grab size

relationship_grab_size

How many relationships to read at a time during iteration

Default value: 100

Relationship keys indexable

relationship_keys_indexable

A list of property names (comma separated) that will be indexed by default. This applies to Relationships only.

Remote logging enabled

remote_logging_enabled

Default value: false

Whether to enable logging to a remote server or not.

Default value: false

Remote logging host

remote_logging_host

Host for remote logging using LogBack SocketAppender.

Default value: 127.0.0.1

Remote logging port

remote_logging_port

Port for remote logging using LogBack SocketAppender.

Default value: 4560

Store dir

store_dir

The directory where the database files are located.

String block size

string_block_size

Specifies the block size for storing strings. This parameter is only honored when the store is created, otherwise it is ignored. Note that each character in a string occupies two bytes, meaning that a block size of 120 (the default size) will hold a 60 character long string before overflowing into a second block. Also note that each block carries an overhead of 8 bytes. This means that if the block size is 120, the size of the stored records will be 128 bytes.

Default value: 120

Tx manager impl

tx_manager_impl

The name of the Transaction Manager service to use as defined in the TM service provider constructor.

Default value: native

Use memory mapped buffers

use_memory_mapped_buffers

Tell Neo4j to use memory mapped buffers for accessing the native storage layer.

25.4. Caches in Neo4j

For how to provide custom configuration to Neo4j, see [Section 25.1, “Introduction”](#).

Neo4j utilizes two different types of caches: A file buffer cache and an object cache. The file buffer cache caches the storage file data in the same format as it is stored on the durable storage media.

The object cache caches the nodes, relationships and properties in a format that is optimized for high traversal speeds and transactional writes.

25.4.1. File buffer cache

Quick info

- The file buffer cache is sometimes called *low level cache* or *file system cache*.
- It caches the Neo4j data as stored on the durable media.
- It uses the operating system memory mapping features when possible.
- Neo4j will configure the cache automatically as long as the heap size of the JVM is configured properly.

The file buffer cache caches the Neo4j data in the same format as it is represented on the durable storage media. The purpose of this cache layer is to improve both read and write performance. The file buffer cache improves write performance by writing to the cache and deferring durable write until the logical log is rotated. This behavior is safe since all transactions are always durably written to the logical log, which can be used to recover the store files in the event of a crash.

Since the operation of the cache is tightly related to the data it stores, a short description of the Neo4j durable representation format is necessary background. Neo4j stores data in multiple files and relies on the underlying file system to handle this efficiently. Each Neo4j storage file contains uniform fixed size records of a particular type:

Store file	Record size	Contents
nodestore	9 B	Nodes
relstore	33 B	Relationships
propstore	41 B	Properties for nodes and relationships
stringstore	128 B	Values of string properties
arraystore	128 B	Values of array properties

For strings and arrays, where data can be of variable length, data is stored in one or more 120B chunks, with 8B record overhead. The sizes of these blocks can actually be configured when the store is created using the `string_block_size` and `array_block_size` parameters. The size of each record type can also be used to calculate the storage requirements of a Neo4j graph or the appropriate cache size for each file buffer cache. Note that some strings and arrays can be stored without using the string store or the array store respectively, see [Section 25.7, “Compressed storage of short strings”](#) and [Section 25.8, “Compressed storage of short arrays”](#).

Neo4j uses multiple file buffer caches, one for each different storage file. Each file buffer cache divides its storage file into a number of equally sized windows. Each cache window contains an even

number of storage records. The cache holds the most active cache windows in memory and tracks hit vs. miss ratio for the windows. When the hit ratio of an uncached window gets higher than the miss ratio of a cached window, the cached window gets evicted and the previously uncached window is cached instead.



Important

Note that the block sizes can only be configured at store creation time.

Configuration

Parameter	Possible values	Effect
<code>use_memory_mapped_buffers</code>	true or false	If set to true Neo4j will use the operating systems memory mapping functionality for the file buffer cache windows. If set to false Neo4j will use its own buffer implementation. In this case the buffers will reside in the JVM heap which needs to be increased accordingly. The default value for this parameter is true, except on Windows.
<code>neostore.nodestore.db.mapped_memory</code>	The maximum amount of memory to use for memory mapped buffers for this file buffer cache. The default unit is MiB, for other units use any of the following suffixes: B, k, M or G.	The maximum amount of memory to use for the file buffer cache of the node storage file.
<code>neostore.relationshipstore.db.mapped_memory</code>		The maximum amount of memory to use for the file buffer cache of the relationship store file.
<code>neostore.propertystore.db.index.keys.mapped_memory</code>		The maximum amount of memory to use for the file buffer cache of the something-something file.
<code>neostore.propertystore.db.index.mapped_memory</code>		The maximum amount of memory to use for the file buffer cache of the something-something file.
<code>neostore.propertystore.db.mapped_memory</code>		The maximum amount of memory to use for the file buffer cache of the property storage file.
<code>neostore.propertystore.db.strings.mapped_memory</code>		The maximum amount of memory to use for the file buffer cache of the string property storage file.
<code>neostore.propertystore.db.arrays.mapped_memory</code>		The maximum amount of memory to use for the file buffer cache of the array property storage file.
<code>string_block_size</code>	The number of bytes per block.	Specifies the block size for storing strings. This parameter is only honored when the store is created,

Parameter	Possible values	Effect
		otherwise it is ignored. Note that each character in a string occupies two bytes, meaning that a block size of 120 (the default size) will hold a 60 character long string before overflowing into a second block. Also note that each block carries an overhead of 8 bytes. This means that if the block size is 120, the size of the stored records will be 128 bytes.
array_block_size		Specifies the block size for storing arrays. This parameter is only honored when the store is created, otherwise it is ignored. The default block size is 120 bytes, and the overhead of each block is the same as for string blocks, i.e., 8 bytes.
dump_configuration	true Or false	If set to true the current configuration settings will be written to the default system output, mostly the console or the logfiles.

When memory mapped buffers are used (`use_memory_mapped_buffers = true`) the heap size of the JVM must be smaller than the total available memory of the computer, minus the total amount of memory used for the buffers. When heap buffers are used (`use_memory_mapped_buffers = false`) the heap size of the JVM must be large enough to contain all the buffers, plus the runtime heap memory requirements of the application and the object cache.

When reading the configuration parameters on startup Neo4j will automatically configure the parameters that are not specified. The cache sizes will be configured based on the available memory on the computer, how much is used by the JVM heap, and how large the storage files are.

25.4.2. Object cache

Quick info

- The object cache is sometimes called *high level cache*.
- It caches the Neo4j data in a form optimized for fast traversal.

The object cache caches individual nodes and relationships and their properties in a form that is optimized for fast traversal of the graph. There are two different categories of object caches in Neo4j.

There is the reference caches. Here Neo4j will utilize as much as it can out of the allocated heap memory for the JVM for object caching and relies on garbage collection for eviction from the cache in an LRU manner. Note however that Neo4j is "competing" for the heap space with other objects in the same JVM, such as a your application, if deployed in embedded mode, and Neo4j will let the application "win" by using less memory if the application needs more.

**Note**

The GC resistant cache described below is only available in the Neo4j Enterprise Edition.

The other is the *GC resistant cache* which gets assigned a certain amount of space in the JVM heap and will purge objects whenever it grows bigger than that. It is assigned a maximum amount of memory which the sum of all cached objects in it will not exceed. Objects will be evicted from cache when the maximum size is about to be reached, instead of relying on garbage collection (GC) to make that decision. Here the competition with other objects in the heap as well as GC-pauses can be better controlled since the cache gets assigned a maximum heap space usage. The overhead of the GC resistant cache is also much smaller as well as insert/lookup times faster than for reference caches.

**Tip**

The use of heap memory is subject to the java garbage collector — depending on the cache type some tuning might be needed to play well with the GC at large heap sizes. Therefore, assigning a large heap for Neo4j's sake isn't always the best strategy as it may lead to long GC-pauses. Instead leave some space for Neo4j's filesystem caches. These are outside of the heap and under the kernel's direct control, thus more efficiently managed.

The content of this cache are objects with a representation geared towards supporting the Neo4j object API and graph traversals. Reading from this cache is 5 to 10 times faster than reading from the file buffer cache. This cache is contained in the heap of the JVM and the size is adapted to the current amount of available heap memory.

Nodes and relationships are added to the object cache as soon as they are accessed. The cached objects are however populated lazily. The properties for a node or relationship are not loaded until properties are accessed for that node or relationship. String (and array) properties are not loaded until that particular property is accessed. The relationships for a particular node is also not loaded until the relationships are accessed for that node.

Configuration

The main configuration parameter for the object cache is the `cache_type` parameter. This specifies which cache implementation to use for the object cache. Note that there will exist two cache instances, one for nodes and one for relationships. The available cache types are:

cache_type	Description
none	Do not use a high level cache. No objects will be cached.
soft	Provides optimal utilization of the available memory. Suitable for high performance traversal. May run into GC issues under high load if the frequently accessed parts of the graph does not fit in the cache. This is the default cache implementation.
weak	Provides short life span for cached objects. Suitable for high throughput applications where a larger portion of the graph than what can fit into memory is frequently accessed.
strong	This cache will hold on to all data that gets loaded to never release it again. Provides good performance if your graph is small enough to fit in memory.
gcr	Provides means of assigning a specific amount of memory to dedicate to caching loaded nodes and relationships. Small footprint and fast insert/lookup. Should be the

cache_type	Description
	best option for most scenarios. See below on how to configure it. Note that this option is only available in the Neo4j Enterprise Edition.

GC resistant cache configuration

Since the GC resistant cache operates with a maximum size in the JVM it may be configured per use case for optimal performance. There are two aspects of the cache size.

One is the size of the array referencing the objects that are put in the cache. It is specified as a fraction of the heap, for example specifying 5 will let that array itself take up 5% out of the entire heap. Increasing this figure (up to a maximum of 10) will reduce the chance of hash collisions at the expense of more heap used for it. More collisions means more redundant loading of objects from the low level cache.

configuration option	Description (what it controls)	Example value
node_cache_array_fraction	Fraction of the heap to dedicate to the array holding the nodes in the cache (max 10).	7
relationship_cache_array_fraction	Fraction of the heap to dedicate to the array holding the relationships in the cache (max 10).	5

The other aspect is the maximum size of all the objects in the cache. It is specified as size in bytes, for example 500M for 500 megabytes or 2G for two gigabytes. Right before the maximum size is reached a purge is performed where (currently) random objects are evicted from the cache until the cache size gets below 90% of the maximum size. Optimal settings for the maximum size depends on the size of your graph. The configured maximum size should leave enough room for other objects to coexist in the same JVM, but at the same time large enough to keep loading from the low level cache at a minimum. Predicted load on the JVM as well as layout of domain level objects should also be take into consideration.

configuration option	Description (what it controls)	Example value
node_cache_size	Maximum size of the heap memory to dedicate to the cached nodes.	2G
relationship_cache_size	Maximum size of the heap memory to dedicate to the cached relationships.	800M

You can read about references and relevant JVM settings for Sun HotSpot here:

- [Understanding soft/weak references](http://weblogs.java.net/blog/enicholas/archive/2006/05/understanding_w.html) <http://weblogs.java.net/blog/enicholas/archive/2006/05/understanding_w.html>
- [How Hotspot Decides to Clear SoftReferences](http://jeremymanson.blogspot.com/2009/07/how-hotspot-decides-to-clear_07.html) <http://jeremymanson.blogspot.com/2009/07/how-hotspot-decides-to-clear_07.html>
- [HotSpot FAQ](http://www.oracle.com/technetwork/java/hotspotfaq-138619.html#gc_softrefs) <http://www.oracle.com/technetwork/java/hotspotfaq-138619.html#gc_softrefs>

Heap memory usage

This table can be used to calculate how much memory the data in the object cache will occupy on a 64bit JVM:

Object	Size	Comment
Node	344 B	<i>Size for each node (not counting its relationships or properties).</i>
	48 B	<i>Object overhead.</i>
	136 B	<i>Property storage (ArrayMap 48B, HashMap 88B).</i>
	136 B	<i>Relationship storage (ArrayMap 48B, HashMap 88B).</i>
	24 B	<i>Location of first / next set of relationships.</i>
Relationship	208 B	<i>Size for each relationship (not counting its properties).</i>
	48 B	<i>Object overhead.</i>
	136 B	<i>Property storage (ArrayMap 48B, HashMap 88B).</i>
Property	116 B	<i>Size for each property of a node or relationship.</i>
	32 B	<i>Data element — allows for transactional modification and keeps track of on disk location.</i>
	48 B	<i>Entry in the hash table where it is stored.</i>
	12 B	<i>Space used in hash table, accounts for normal fill ratio.</i>
	24 B	<i>Property key index.</i>
Relationships	108 B	<i>Size for each relationship type for a node that has a relationship of that type.</i>
	48 B	<i>Collection of the relationships of this type.</i>
	48 B	<i>Entry in the hash table where it is stored.</i>
	12 B	<i>Space used in hash table, accounts for normal fill ratio.</i>
Relationships	8 B	<i>Space used by each relationship related to a particular node (both incoming and outgoing).</i>
Primitive	24 B	<i>Size of a primitive property value.</i>
String	64+B	<i>Size of a string property value. 64 + 2*len(string) B (64 bytes, plus two bytes for each character in the string).</i>

25.5. Logical logs

Logical logs in Neo4j are the journal of which operations happens and are the source of truth in scenarios where the database needs to be recovered after a crash or similar. Logs are rotated every now and then (defaults to when they surpass 25 Mb in size) and the amount of legacy logs to keep can be configured. Purpose of keeping a history of logical logs include being able to serve incremental backups as well as keeping an HA cluster running.

For any given configuration at least the latest non-empty logical log will be kept, but configuration can be supplied to control how much more to keep. There are several different means of controlling it and the format in which configuration is supplied is:

```
keep_logical_logs=<true/false>
keep_logical_logs=<amount> <type>
```

For example:

```
# Will keep logical logs indefinitely
keep_logical_logs=true

# Will keep only the most recent non-empty log
keep_logical_logs=false

# Will keep logical logs which contains any transaction committed within 30 days
keep_logical_logs=30 days

# Will keep logical logs which contains any of the most recent 500 000 transactions
keep_logical_logs=500k txs
```

Full list:

Type	Description	Example
files	Number of most recent logical log files to keep	"10 files"
size	Max disk size to allow log files to occupy	"300M size" or "1G size"
txs	Number of latest transactions to keep Keep	"250k txs" or "5M txs"
hours	Keep logs which contains any transaction committed within N hours from current time	"10 hours"
days	Keep logs which contains any transaction committed within N days from current time	"50 days"

25.6. JVM Settings

There are two main memory parameters for the JVM, one controls the heap space and the other controls the stack space. The heap space parameter is the most important one for Neo4j, since this governs how many objects you can allocate. The stack space parameter governs the how deep the call stack of your application is allowed to get.

When it comes to heap space the general rule is: the larger heap space you have the better, but make sure the heap fits in the RAM memory of the computer. If the heap is paged out to disk performance will degrade rapidly. Having a heap that is much larger than what your application needs is not good either, since this means that the JVM will accumulate a lot of dead objects before the garbage collector is executed, this leads to long garbage collection pauses and undesired performance behavior.

Having a larger heap space will mean that Neo4j can handle larger transactions and more concurrent transactions. A large heap space will also make Neo4j run faster since it means Neo4j can fit a larger portion of the graph in its caches, meaning that the nodes and relationships your application uses frequently are always available quickly. The default heap size for a 32bit JVM is 64MB (and 30% larger for 64bit), which is too small for most real applications.

Neo4j works fine with the default stack space configuration, but if your application implements some recursive behavior it is a good idea to increment the stack size. Note that the stack size is shared for all threads, so if you application is running a lot of concurrent threads it is a good idea to increase the stack size.

- The heap size is set by specifying the `-Xmx???m` parameter to hotspot, where `???` is the heap size in megabytes. Default heap size is 64MB for 32bit JVMs, 30% larger (appr. 83MB) for 64bit JVMs.
- The stack size is set by specifying the `-Xss???m` parameter to hotspot, where `???` is the stack size in megabytes. Default stack size is 512kB for 32bit JVMs on Solaris, 320kB for 32bit JVMs on Linux (and Windows), and 1024kB for 64bit JVMs.

Most modern CPUs implement a [Non-Uniform Memory Access \(NUMA\) architecture](http://en.wikipedia.org/wiki/Non-Uniform_Memory_Access) <http://en.wikipedia.org/wiki/Non-Uniform_Memory_Access>, where different parts of the memory have different access speeds. Suns Hotspot JVM is able to allocate objects with awareness of the NUMA structure as of version 1.6.0 update 18. When enabled this can give up to 40% performance improvements. To enable the NUMA awareness, specify the `-XX:+UseNUMA` parameter (works only when using the Parallel Scavenger garbage collector (default or `-XX:+UseParallelGC` not the concurrent mark and sweep one).

Properly configuring memory utilization of the JVM is crucial for optimal performance. As an example, a poorly configured JVM could spend all CPU time performing garbage collection (blocking all threads from performing any work). Requirements such as latency, total throughput and available hardware have to be considered to find the right setup. In production, Neo4j should run on a multi core/CPU platform with the JVM in server mode.

25.6.1. Configuring heap size and GC

A large heap allows for larger node and relationship caches — which is a good thing — but large heaps can also lead to latency problems caused by full garbage collection. The different high level cache implementations available in Neo4j together with a suitable JVM configuration of heap size and garbage collection (GC) should be able to handle most workloads.

The default cache (soft reference based LRU cache) works best with a heap that never gets full: a graph where the most used nodes and relationships can be cached. If the heap gets too full there is a

risk that a full GC will be triggered; the larger the heap, the longer it can take to determine what soft references should be cleared.

Using the strong reference cache means that *all* the nodes and relationships being used must fit in the available heap. Otherwise there is a risk of getting out-of-memory exceptions. The soft reference and strong reference caches are well suited for applications where the overall throughput is important.

The weak reference cache basically needs enough heap to handle the peak load of the application — peak load multiplied by the average memory required per request. It is well suited for low latency requirements where GC interruptions are not acceptable.



Important

When running Neo4j on Windows, keep in mind that the memory mapped buffers are allocated on heap by default, so they need to be taken into account when determining heap size.

Guidelines for heap size

Number of primitives	RAM size	Heap configuration	Reserved RAM for the OS
10M	2GB	512MB	the rest
100M	8GB+	1-4GB	1-2GB
1B+	16GB-32GB+	4GB+	1-2GB



Tip

The recommended garbage collector to use when running Neo4j in production is the Concurrent Mark and Sweep Compactor turned on by supplying `-XX:+UseConcMarkSweepGC` as a JVM parameter.

When having made sure that the heap size is well configured the second thing to tune in order to tune the garbage collector for your application is to specify the sizes of the different generations of the heap. The default settings are well tuned for "normal" applications, and work quite well for most applications, but if you have an application with either really high allocation rate, or a lot of long lived objects you might want to consider tuning the sizes of the heap generation. The ratio between the young and tenured generation of the heap is specified by using the `-XX:NewRatio=#` command line option (where # is replaced by a number). The default ratio is 1:12 for client mode JVM, and 1:8 for server mode JVM. You can also specify the size of the young generation explicitly using the `-Xmn` command line option, which works just like the `-Xmx` option that specifies the total heap space.

GC shortname	Generation	Command line parameter	Comment
Copy	Young	<code>-XX:+UseSerialGC</code>	The Copying collector
MarkSweepCompact	Tenured	<code>-XX:+UseSerialGC</code>	The Mark and Sweep Compactor
ConcurrentMarkSweep	Tenured	<code>-XX:+UseConcMarkSweepGC</code>	The Concurrent Mark and Sweep Compactor
ParNew	Young	<code>-XX:+UseParNewGC</code>	The parallel Young Generation Collector — can only

GC shortname	Generation	Command line parameter	Comment
			be used with the Concurrent mark and sweep compactor.
PS Scavenge	Young	-XX:+UseParallelGC	The parallel object scavenger
PS MarkSweep	Tenured	-XX:+UseParallelGC	The parallel mark and sweep collector

These are the default configurations on some platforms according to our non-exhaustive research:

JVM	-d32 -client	-d32 -server	-d64 -client	-d64 -server
Mac OS X Snow Leopard, 64-bit, Hotspot 1.6.0_17	ParNew and ConcurrentMarkSweep	PS Scavenge and PS MarkSweep	ParNew and ConcurrentMarkSweep	PS Scavenge and PS MarkSweep
Ubuntu, 32-bit, Hotspot 1.6.0_16	Copy and MarkSweepCompact	Copy and MarkSweepCompact	N/A	N/A

25.7. Compressed storage of short strings

Neo4j will try to classify your strings in a short string class and if it manages that it will treat it accordingly. In that case, it will be stored without indirection in the property store, inlining it instead in the property record, meaning that the dynamic string store will not be involved in storing that value, leading to reduced disk footprint. Additionally, when no string record is needed to store the property, it can be read and written in a single lookup, leading to performance improvements and less disk space required.

The various classes for short strings are:

- Numerical, consisting of digits 0..9 and the punctuation space, period, dash, plus, comma and apostrophe.
- Date, consisting of digits 0..9 and the punctuation space dash, colon, slash, plus and comma.
- Hex (lower case), consisting of digits 0..9 and lower case letters a..f
- Hex (upper case), consisting of digits 0..9 and upper case letters a..f
- Upper case, consisting of upper case letters A..Z, and the punctuation space, underscore, period, dash, colon and slash.
- Lower case, like upper but with lower case letters a..z instead of upper case
- E-mail, consisting of lower case letters a..z and the punctuation comma, underscore, period, dash, plus and the at sign (@).
- URI, consisting of lower case letters a..z, digits 0..9 and most punctuation available.
- Alpha-numerical, consisting of both upper and lower case letters a..zA..Z, digits 0..9 and punctuation space and underscore.
- Alpha-symbolical, consisting of both upper and lower case letters a..zA..Z and the punctuation space, underscore, period, dash, colon, slash, plus, comma, apostrophe, at sign, pipe and semicolon.
- European, consisting of most accented european characters and digits plus punctuation space, dash, underscore and period — like latin1 but with less punctuation.
- Latin 1.
- UTF-8.

In addition to the string's contents, the number of characters also determines if the string can be inlined or not. Each class has its own character count limits, which are

Character count limits

String class	Character count limit
Numerical, Date and Hex	54
Uppercase, Lowercase and E-mail	43
URI, Alphanumeric and Alphasympobical	36
European	31
Latin1	27
UTF-8	14

That means that the largest inline-able string is 54 characters long and must be of the Numerical class and also that all Strings of size 14 or less will always be inlined.

Also note that the above limits are for the default 41 byte PropertyRecord layout — if that parameter is changed via editing the source and recompiling, the above have to be recalculated.

25.8. Compressed storage of short arrays

Neo4j will try to store your primitive arrays in a compressed way, so as to save disk space and possibly an I/O operation. To do that, it employs a "bit-shaving" algorithm that tries to reduce the number of bits required for storing the members of the array. In particular:

1. For each member of the array, it determines the position of leftmost set bit.
2. Determines the largest such position among all members of the array
3. It reduces all members to that number of bits
4. Stores those values, prefixed by a small header.

That means that when even a single negative value is included in the array then the natural size of the primitives will be used.

There is a possibility that the result can be inlined in the property record if:

- It is less than 24 bytes after compression
- It has less than 64 members

For example, an array `long[] {0L, 1L, 2L, 4L}` will be inlined, as the largest entry (4) will require 3 bits to store so the whole array will be stored in $4*3=12$ bits. The array `long[] {-1L, 1L, 2L, 4L}` however will require the whole 64 bits for the -1 entry so it needs $64*4 = 32$ bytes and it will end up in the dynamic store.

25.9. Memory mapped IO settings

Each file in the Neo4j store can use memory mapped I/O for reading/writing. Best performance is achieved if the full file can be memory mapped but if there isn't enough memory for that Neo4j will try and make the best use of the memory it gets (regions of the file that get accessed often will more likely be memory mapped).



Important

Neo4j makes heavy use of the `java.nio` package. Native I/O will result in memory being allocated outside the normal Java heap so that memory usage needs to be taken into consideration. Other processes running on the OS will impact the availability of such memory. Neo4j will require all of the heap memory of the JVM plus the memory to be used for memory mapping to be available as physical memory. Other processes may thus not use more than what is available after the configured memory allocation is made for Neo4j.

A well configured OS with large disk caches will help a lot once we get cache misses in the node and relationship caches. Therefore it is not a good idea to use all available memory as Java heap.

If you look into the directory of your Neo4j database, you will find its store files, all prefixed by `neostore`:

- `nodestore` stores information about nodes
- `relationshipstore` holds all the relationships
- `propertystore` stores information of properties and all simple properties such as primitive types (both for relationships and nodes)
- `propertystore strings` stores all string properties
- `propertystore arrays` stores all array properties

There are other files there as well, but they are normally not interesting in this context.

This is how the default memory mapping configuration looks:

```
neostore.nodestore.db.mapped_memory=25M
neostore.relationshipstore.db.mapped_memory=50M
neostore.propertystore.db.mapped_memory=90M
neostore.propertystore.db.strings.mapped_memory=130M
neostore.propertystore.db.arrays.mapped_memory=130M
```

25.9.1. Optimizing for traversal speed example

To tune the memory mapping settings start by investigating the size of the different store files found in the directory of your Neo4j database. Here is an example of some of the files and sizes in a Neo4j database:

```
14M neostore.nodestore.db
510M neostore.propertystore.db
1.2G neostore.propertystore.db.strings
304M neostore.relationshipstore.db
```

In this example the application is running on a machine with 4GB of RAM. We've reserved about 2GB for the OS and other programs. The Java heap is set to 1.5GB, that leaves about 500MB of RAM that can be used for memory mapping.



Tip

If traversal speed is the highest priority it is good to memory map as much as possible of the node- and relationship stores.

An example configuration on the example machine focusing on traversal speed would then look something like:

```
neostore.nodestore.db.mapped_memory=15M
neostore.relationshipstore.db.mapped_memory=285M
neostore.propertystore.db.mapped_memory=100M
neostore.propertystore.db.strings.mapped_memory=100M
neostore.propertystore.db.arrays.mapped_memory=0M
```

25.9.2. Batch insert example

Read general information on batch insertion in [Section 18.1, “Batch Insertion”](#).

The configuration should suit the data set you are about to inject using BatchInsert. Lets say we have a random-like graph with 10M nodes and 100M relationships. Each node (and maybe some relationships) have different properties of string and Java primitive types (but no arrays). The important thing with a random graph will be to give lots of memory to the relationship and node store:

```
neostore.nodestore.db.mapped_memory=90M
neostore.relationshipstore.db.mapped_memory=3G
neostore.propertystore.db.mapped_memory=50M
neostore.propertystore.db.strings.mapped_memory=100M
neostore.propertystore.db.arrays.mapped_memory=0M
```

The configuration above will fit the entire graph (with exception to properties) in memory.

A rough formula to calculate the memory needed for the nodes:

```
number_of_nodes * 9 bytes
```

and for relationships:

```
number_of_relationships * 33 bytes
```

Properties will typically only be injected once and never read so a few megabytes for the property store and string store is usually enough. If you have very large strings or arrays you may want to increase the amount of memory assigned to the string and array store files.

An important thing to remember is that the above configuration will need a Java heap of 3.3G+ since in batch inserter mode normal Java buffers that gets allocated on the heap will be used instead of memory mapped ones.

25.10. Linux Performance Guide

The key to achieve good performance on reads and writes is to have lots of RAM since disks are so slow. This guide will focus on achieving good write performance on a Linux kernel based operating system.

If you have not already read the information available in [Chapter 25, Configuration & Performance](#) do that now to get some basic knowledge on memory mapping and store files with Neo4j.

This section will guide you through how to set up a file system benchmark and use it to configure your system in a better way.

25.10.1. Setup

Create a large file with random data. The file should fit in RAM so if your machine has 4GB of RAM a 1-2GB file with random data will be enough. After the file has been created we will read the file sequentially a few times to make sure it is cached.

```
$ dd if=/dev/urandom of=store bs=1M count=1000
1000+0 records in
1000+0 records out
1048576000 bytes (1.0 GB) copied, 263.53 s, 4.0 MB/s
$
$ dd if=store of=/dev/null bs=100M
10+0 records in
10+0 records out
1048576000 bytes (1.0 GB) copied, 38.6809 s, 27.1 MB/s
$
$ dd if=store of=/dev/null bs=100M
10+0 records in
10+0 records out
1048576000 bytes (1.0 GB) copied, 1.52365 s, 688 MB/s
$ dd if=store of=/dev/null bs=100M
10+0 records in
10+0 records out
1048576000 bytes (1.0 GB) copied, 0.776044 s, 1.4 GB/s
```

If you have a standard hard drive in the machine you may know that it is not capable of transfer speeds as high as 1.4GB/s. What is measured is how fast we can read a file that is cached for us by the operating system.

Next we will use a small utility that simulates the Neo4j kernel behavior to benchmark write speed of the system.

```
$ git clone git@github.com:neo4j-contrib/tooling.git
...
$ cd tooling/write-test/
$ mvn compile
[INFO] Scanning for projects...
...
$ ./run
Usage: <large file> <log file> <[record size] [min tx size] [max tx size] [tx count] <[--nosync | --nowritelog | --nowritestore | --
```

The utility will be given a store file (large file we just created) and a name of a log file. Then a record size in bytes, min tx size, max tx size and transaction count must be set. When started the utility will map the large store file entirely in memory and read (transaction size) records from it randomly and then write them sequentially to the log file. The log file will then force changes to disk and finally the records will be written back to the store file.

25.10.2. Running the benchmark

Lets try to benchmark 100 transactions of size 100-500 with a record size of 33 bytes (same record size used by the relationship store).

```
$ ./run store logfile 33 100 500 100
tx_count[100] records[30759] fdatasyncs[100] read[0.96802425 MB] wrote[1.9360485 MB]
Time was: 4.973
20.108585 tx/s, 6185.2 records/s, 20.108585 fdatasyncs/s, 199.32773 kB/s on reads, 398.65546 kB/s on writes
```

We see that we get about 6185 record updates/s and 20 transactions/s with the current transaction size. We can change the transaction size to be bigger, for example writing 10 transactions of size 1000-5000 records:

```
$ ./run store logfile 33 1000 5000 10
tx_count[10] records[24511] fdatasyncs[10] read[0.77139187 MB] wrote[1.5427837 MB]
Time was: 0.792
12.626263 tx/s, 30948.232 records/s, 12.626263 fdatasyncs/s, 997.35516 kB/s on reads, 1994.7103 kB/s on writes
```

With larger transaction we will do fewer of them per second but record throughput will increase. Lets see if it scales, 10 transactions in under 1s then 100 of them should execute in about 10s:

```
$ ./run store logfile 33 1000 5000 100
tx_count[100] records[308814] fdatasyncs[100] read[9.718763 MB] wrote[19.437527 MB]
Time was: 65.115
1.5357445 tx/s, 4742.594 records/s, 1.5357445 fdatasyncs/s, 152.83751 kB/s on reads, 305.67502 kB/s on writes
```

This is not very linear scaling. We modified a bit more than 10x records in total but the time jumped up almost 100x. Running the benchmark watching vmstat output will reveal that something is not as it should be:

```
$ vmstat 3
procs -----memory----- ---swap-- -----io----- -system-- -----cpu-----
 r b swpd free buff cache si so bi bo in cs us sy id wa
 0 1 47660 298884 136036 2650324 0 0 0 10239 1167 2268 5 7 46 42
 0 1 47660 302728 136044 2646060 0 0 0 7389 1267 2627 6 7 47 40
 0 1 47660 302408 136044 2646024 0 0 0 11707 1861 2016 8 5 48 39
 0 2 47660 302472 136060 2646432 0 0 0 10011 1704 1878 4 7 49 40
 0 1 47660 303420 136068 2645788 0 0 0 13807 1406 1601 4 5 44 47
```

There are a lot of blocks going out to IO, way more than expected for the write speed we are seeing in the benchmark. Another observation that can be made is that the Linux kernel has spawned a process called "flush-x:x" (run top) that seems to be consuming a lot of resources.

The problem here is that the Linux kernel is trying to be smart and write out dirty pages from the virtual memory. As the benchmark will memory map a 1GB file and do random writes it is likely that this will result in 1/4 of the memory pages available on the system to be marked as dirty. The Neo4j kernel is not sending any system calls to the Linux kernel to write out these pages to disk however the Linux kernel decided to start doing so and it is a very bad decision. The result is that instead of doing sequential like writes down to disk (the logical log file) we are now doing random writes writing regions of the memory mapped file to disk.

It is possible to observe this behavior in more detail by looking at /proc/vmstat "nr_dirty" and "nr_writeback" values. By default the Linux kernel will start writing out pages at a very low ratio of dirty pages (10%).

```
$ sync
$ watch grep -A 1 dirty /proc/vmstat
...
nr_dirty 22
nr_writeback 0
```

The "sync" command will write out all data (that needs writing) from memory to disk. The second command will watch the "nr_dirty" and "nr_writeback" count from vmstat. Now start the benchmark again and observe the numbers:

```
nr_dirty 124947
nr_writeback 232
```

The "nr_dirty" pages will quickly start to rise and after a while the "nr_writeback" will also increase meaning the Linux kernel is scheduling a lot of pages to write out to disk.

25.10.3. Fixing the problem

As we have 4GB RAM on the machine and memory map a 1GB file that does not need its content written to disk (until we tell it to do so because of logical log rotation or Neo4j kernel shutdown) it should be possible to do endless random writes to that memory with high throughput. All we have to do is to tell the Linux kernel to stop trying to be smart. Edit the /etc/sysctl.conf (need root access) and add the following lines:

```
vm.dirty_background_ratio = 50
vm.dirty_ratio = 80
```

Then (as root) execute:

```
# sysctl -p
```

The "vm.dirty_background_ratio" tells at what ratio should the linux kernel start the background task of writing out dirty pages. We increased this from the default 10% to 50% and that should cover the 1GB memory mapped file. The "vm.dirty_ratio" tells at what ratio all IO writes become synchronous, meaning that we can not do IO calls without waiting for the underlying device to complete them (which is something you never want to happen).

Rerun the benchmark:

```
$ ./run store logfile 33 1000 5000 100
tx_count[100] records[265624] fdatasyncs[100] read[8.35952 MB] wrote[16.71904 MB]
Time was: 6.781
14.7470875 tx/s, 39171.805 records/s, 14.7470875 fdatasyncs/s, 1262.3726 kB/s on reads, 2524.745 kB/s on writes
```

Results are now more in line with what can be expected, 10x more records modified results in 10x longer execution time. The vmstat utility will not report any absurd amount of IO blocks going out (it reports the ones caused by the fdatasync to the logical log) and Linux kernel will not spawn a "flush-x:x" background process writing out dirty pages caused by writes to the memory mapped store file.

25.11. Linux specific notes

25.11.1. File system tuning for high IO

In order to support the high IO load of small transactions from a database, the underlying file system should be tuned. Symptoms for this are low CPU load with high iowait. In this case, there are a couple of tweaks possible on Linux systems:

- Disable access-time updates: `noatime,nodiratime` flags for disk mount command or in the `/etc/fstab` for the database disk volume mount.
- Tune the IO scheduler for high disk IO on the database disk.

25.11.2. Setting the number of open files

Linux platforms impose an upper limit on the number of concurrent files a user may have open. This number is reported for the current user and session with the command

```
user@localhost:~$ ulimit -n
1024
```

The usual default of 1024 is often not enough, especially when many indexes are used or a server installation sees too many connections (network sockets count against that limit as well). Users are therefore encouraged to increase that limit to a healthy value of 40000 or more, depending on usage patterns. Setting this value via the `ulimit` command is possible only for the root user and that for that session only. To set the value system wide you have to follow the instructions for your platform.

What follows is the procedure to set the open file descriptor limit to 40k for user `neo4j` under Ubuntu 10.04 and later. If you opted to run the `neo4j` service as a different user, change the first field in step 2 accordingly.

1. Become root since all operations that follow require editing protected system files.

```
user@localhost:~$ sudo su -
Password:
root@localhost:~$
```

2. Edit `/etc/security/limits.conf` and add these two lines:

```
neo4j soft nofile 40000
neo4j hard nofile 40000
```

3. Edit `/etc/pam.d/su` and uncomment or add the following line:

```
session required pam_limits.so
```

4. A restart is required for the settings to take effect.

After the above procedure, the `neo4j` user will have a limit of 40000 simultaneous open files. If you continue experiencing exceptions on `Too many open files` or `Could not stat() directory` then you may have to raise that limit further.

Chapter 26. High Availability



Note

The High Availability features are only available in the Neo4j Enterprise Edition.

Neo4j High Availability or “Neo4j HA” provides the following two main features:


1. It enables a *fault-tolerant database architecture*, where several Neo4j slave databases can be configured to be exact replicas of a single Neo4j master database. This allows the end-user system to be fully functional and both read and write to the database in the event of hardware failure.
2. It enables a *horizontally scaling read-mostly architecture* that enables the system to handle more read load than a single Neo4j database instance can handle.

26.1. Architecture

Neo4j HA has been designed to make the transition from single machine to multi machine operation simple, by not having to change the already existing application.

Consider an existing application with Neo4j embedded and running on a single machine. To deploy such an application in a multi machine setup the only required change is to switch the creation of the `GraphDatabaseService` from `GraphDatabaseFactory` to `HighlyAvailableGraphDatabaseFactory`. Since both implement the same interface, no additional changes are required.

Figure 26.1. Multiple Neo4j instances in HA mode



When running Neo4j in HA mode there is always a single master and zero or more slaves. Compared to other master-slave replication setups Neo4j HA can handle write requests on all machines so there is no need to redirect those to the master specifically.

A slave will handle writes by synchronizing with the master to preserve consistency. Writes to master can be configured to be optimistically pushed to 0 or more slaves. By optimistically we mean the master will try to push to slaves before the transaction completes but if it fails the transaction will still be successful (different from normal replication factor). All updates will however propagate from the master to other slaves eventually so a write from one slave may not be immediately visible on all other slaves. This is the only difference between multiple machines running in HA mode compared to single machine operation. All other ACID characteristics are the same.

26.2. Setup and configuration

Neo4j HA can be set up to accommodate differing requirements for load, fault tolerance and available hardware.

Within a cluster, Neo4j HA uses its own Paxos implementation for all cluster membership related tasks, from instances joining/leaving a cluster to master election and availability information propagation. Read operations through the GraphDatabaseService API will always work, whereas write operations requires an available master.

For reaching quorum among the members in the cluster the number of members should be odd.

26.2.1. Installation Notes

For installation instructions of a High Availability cluster see [Section 26.6, “High Availability setup tutorial”](#).

Note that while the HA version of Neo4j supports the same API as the single instance embedded version, it does have additional configuration parameters. Although there are many parameters, most of them have defaults that should work in most cases.

HA configuration parameters

Parameter Name	Description	Example value	Required?
ha.server_id	Id for a cluster instance. Must be unique within the cluster.	1	yes
ha.initial_hosts	A comma-separated list of other members of the cluster to join.	server1:5001, server2:5001	yes
ha.cluster_server	Host & port to bind the cluster management communication.	server1:5001	no
ha.allow_init_cluster	Whether to allow this instance to create a cluster if unable to join.	true	no
ha.default_timeout	Default timeout used for clustering timeouts. Override specific timeout settings with proper values if necessary. This value is the default value for settings ha.heartbeat_interval, ha.paxos_timeout and ha.learn_timeout.	5s	no
ha.heartbeat_interval	How often heartbeat messages should	5s	no

Parameter Name	Description	Example value	Required?
	be sent. Defaults to ha.default_timeout.		
ha.heartbeat_timeout	Timeout for heartbeats between cluster members. Should be at least twice that of ha.heartbeat_interval.	11s	no
ha.broadcast_timeout	Timeout for broadcasting values in cluster. Must consider end-to-end duration of Paxos algorithm. This value is the default value for settings ha.join_timeout and ha.leave_timeout.	30s	no
ha.join_timeout	Timeout for joining a cluster. Defaults to ha.broadcast_timeout.	30s	no
ha.configuration_timeout	Timeout for waiting for configuration from an existing cluster member during cluster join.	1s	no
ha.leave_timeout	Timeout for waiting for cluster leave to finish. Defaults to ha.broadcast_timeout.	30s	no
ha.paxos_timeout	Default timeout for all Paxos timeouts. Defaults to ha.default_timeout. This value is the default value for settings ha.phase1_timeout, ha.phase2_timeout and ha.election_timeout.	5s	no
ha.phase1_timeout	Timeout for Paxos phase 1. Defaults to ha.paxos_timeout.	5s	no
ha.phase2_timeout	Timeout for Paxos phase 2. Defaults to ha.paxos_timeout.	5s	no

Parameter Name	Description	Example value	Required?
ha.learn_timeout	Timeout for learning values. Defaults to ha.default_timeout.	5s	no
ha.election_timeout	Timeout for waiting for other members to finish a role election. Defaults to ha.paxos_timeout.	5s	no
ha.read_timeout	How long a slave will wait for response from master before giving up.	20s	no
ha.state_switch_timeout	Timeout for waiting for instance to become master or slave.	20s	no
ha.lock_read_timeout	Timeout for taking remote (write) locks on slaves. Defaults to ha.read_timeout.	20s	no
ha.max_concurrent_channels_per_slave	Maximum number of connections a slave can have to the master.	20	no
ha.server	Hostname and port to bind the HA server.	my-domain.com:6001	no
ha.slave_only	Whether this instance should only participate as slave in cluster. If set to true, it will never be elected as master.	false	no
ha.branched_data_policy	Policy for how to handle branched data.	keep_last or keep_all or keep_none	no
ha.com_chunk_size	Max size of the data chunks that flows between master and slaves in HA. Bigger size may increase throughput, but may be more sensitive to variations in bandwidth, whereas lower size increases tolerance for bandwidth variations.	2M	no

Parameter Name	Description	Example value	Required?
ha.pull_interval	Interval of pulling updates from master.	10s	no
ha.tx_push_factor	The amount of slaves the master will ask to replicate a committed transaction.	1	no
ha.tx_push_strategy	Push strategy of a transaction to a slave during commit.	fixed or round_robin	no



Note

Note that the `org.neo4j.server.database.mode` setting in the `neo4j-server.properties` file has to be set to HA to run Neo4j in High Availability mode.

26.2.2. Different methods for participating in a cluster

There are currently multiple ways of telling your database instance which cluster to create/join, different ways for different occasions.

Knowing at least one other member

A database instance can join a cluster by supplying `ha.initial_hosts` a comma-separated list of URLs to at least one other cluster member of the cluster to join. It's called initial hosts since it's only for joining the cluster. After an instance joins it gets aware of all the members of the cluster. It will take turn contacting each one in the list until it gets a response, where it will await a decision by the existing cluster members to have it join. If it cannot contact any of the members in the URL list it will create a new cluster with itself as master.

26.3. How Neo4j HA operates

A Neo4j HA cluster operates cooperatively — each database instance contains the logic needed in order to coordinate with the other members of the cluster. On startup a Neo4j HA database instance will try to connect to an existing cluster specified by configuration. If the cluster exists, the instance will join it as a slave. Otherwise the cluster will be created and the instance will become its master.

When performing a write transaction on a slave each write operation will be synchronized with the master (locks will be acquired on both master and slave). When the transaction commits it will first be committed on the master and then, if successful, on the slave. To ensure consistency, a slave has to be up to date with the master before performing a write operation. This is built into the communication protocol between the slave and master, so that updates will be applied to a slave communicating with its master automatically.

Write transactions performed directly through the master will execute in the same way as running in normal non-HA mode. On success the transaction will be pushed out to a configurable number of slaves (default one slave). This is done optimistically meaning if the push fails the transaction will still be successful. It's also possible to configure push factor to 0 for higher write performance when writing directly through the master, although increasing the risk of losing any transaction not yet pulled by another slave if the master goes down.

Slaves can also be configured to pull updates asynchronously by setting the `ha.pull_interval` option.

Whenever a Neo4j database becomes unavailable, by means of for example hardware failure or network outages, the other database instances in the cluster will detect that and mark it as temporarily failed. A database instance that becomes available after being unavailable will automatically catch up with the cluster. If the master goes down another (best suited) member will be elected and have its role switched from slave to master after a quorum has been reached within the cluster. When the new master has performed its role switch it will broadcast its availability to all the other members of the cluster. Normally a new master is elected and started within just a few seconds and during this time no writes can take place (the writes will block or in rare cases throw an exception). The only time this is not true is when an old master had changes that did not get replicated to any other member before becoming unavailable. If the new master is elected and performs changes before the old master recovers, there will be two "branches" of the database after the point where the old master became unavailable. The old master will move away its database (its "branch") and download a full copy from the new master, to become available as a slave in the cluster.

All this can be summarized as:

- Write transactions can be performed on any database instance in a cluster.
- Neo4j HA is fault tolerant and can continue to operate from any number of machines down to a single machine.
- Slaves will be automatically synchronized with the master on write operations.
- If the master fails a new master will be elected automatically.
- The cluster automatically handles instances becoming unavailable (for example due to network issues), and also makes sure to accept them as members in the cluster when they are available again.
- Transactions are atomic, consistent and durable but eventually propagated out to other slaves.
- Updates to slaves are eventual consistent by nature but can be configured to be pushed optimistically from master during commit.
- If the master goes down any running write transaction will be rolled back and new transactions will block or fail until a new master has become available.

- Reads are highly available and the ability to handle read load scales with more database instances in the cluster.

26.4. Arbiter Instances

A typical deployment of Neo4j will use a cluster of 3 machines to provide fault-tolerance and read scalability. This setup is described in [Section 26.6, “High Availability setup tutorial”](#).

While having at least 3 instances is necessary for failover to happen in case the master becomes unavailable, it is not required for all instances to run the full Neo4j stack, which includes the database engine. Instead, what is called arbiter instances can be deployed. They can be regarded as cluster participants in that their role is to take part in master elections with the single purpose of breaking ties in the election process. That makes possible a scenario where you have a cluster of 2 Neo4j database instances and an additional arbiter instance and still enjoy tolerance of a single failure of either of the 3 instances.

Arbiter instances are configured in the same way as Neo4j HA members are — through the *neo4j.properties* file in the installation’s *conf/* directory. Settings that are not cluster specific are of course ignored, so you can easily start up an arbiter instance in place of a properly configured Neo4j instance.

To start an arbiter instance, call

```
neo4j_home$ ./bin/neo4j-arbiter start
```

You can also stop, install and remove it as a service and ask for it’s status in exactly the same way as for Neo4j instances. See also [Section 21.1, “Server Installation”](#).

26.5. Upgrading a Neo4j HA Cluster

This document describes the steps required to upgrade a Neo4j cluster from a previous version to 1.9 without disrupting its operation, a process referred to as a rolling upgrade. The starting assumptions are that there exists a cluster running Neo4j version 1.8 or newer with the corresponding ZooKeeper instances and that the machine which is currently the master is known. It is also assumed that on each machine the Neo4j service and the neo4j coordinator service is installed under a directory which from here on is assumed to be `/opt/old-neo4j`

26.5.1. Overview

The process consists of upgrading each machine in turn by removing it from the cluster, moving over the database and starting it back up again. Configuration settings also have to be transferred. It is important to note that the last machine to be upgraded must be the master. In general, the "cluster version" is defined by the version of the master, providing the master is of the older version the cluster as a whole can operate (the 1.9 instances running in compatibility mode). When a 1.9 instance is elected master however, the older instances are not capable of communicating with it, so we have to make sure that the last machine upgraded is the old master. The upgrade process is detected automatically from the joining 1.9 instances and they will not participate in a master election while even a single old instance is part of the cluster.

26.5.2. Step 1: On each slave perform the upgrade

Download and unpack the new version. Copy over any configuration settings you run your instances with, taking care for deprecated settings and API changes that can occur between versions. Also, ensure that newly introduced settings have proper values (see [Section 26.2, "Setup and configuration"](#)). The most important thing about the settings setup is the `ha.coordinators` setting in `neo4j.properties` which must be set to the value the existing 1.8 instances are using. You also have to make sure that all but one instance have the `ha.allow_init_cluster` setting to `false` - the machine that has it set to `true` should be the one that is to become the new master. In addition, it is necessary that the last machine to be upgraded (the 1.8 master) does not have the `ha.coordinators` setting present in its configuration file. Finally, don't forget to copy over any server plugins you may have. First, shutdown the neo4j instance with

```
service neo4j-service stop
```

Next, uninstall it

```
service neo4j-service remove
```

Now you can copy over the database. Assuming the old instance is at `/opt/old-neo4j` and the newly unpacked under `/opt/neo4j-enterprise-1.9` the proper command would be

```
cp -R /opt/old-neo4j/data/graph.db /opt/neo4j-enterprise-1.9/data/
```

Next install the neo4j service, which also starts it

```
/opt/neo4j-enterprise-1.9/bin/neo4j install
```

Done. Now check that the services are running and that webadmin reports the version 1.9. Transactions should also be applied from the master as usual.

26.5.3. Step 2: Upgrade the master, complete the procedure



Warning

Make sure that the installation that will replace the current master instance does not have `ha.coordinators` setting present in the `neo4j.properties` file.

Go to the current master and execute step 1. The moment it will be stopped another instance will take over (the one with the `allow_init_cluster` setting set to true), transitioning the cluster to 1.9. Finish Step 1 on this machine as well and you will have completed the process.

26.5.4. Step 3: Cleanup, removing the coordinator services

Each 1.8 installation still has a coordinator service installed and running. To have those removed you need to execute at every upgraded instance

```
service neo4j-coordinator stop
service neo4j-coordinator remove
```

After that, the 1.8 instances are no longer active or needed and can be removed or archived.

26.5.5. Upgrading within the 1.9.x series



Warning

Due to a bug in the HA code, it may not be possible to do a rolling (i.e. uninterrupted) upgrade to Neo4j 1.9.2 on 1.9 or 1.9.1 clusters. Attempting to do so may lead to an unstable cluster and data loss may occur. It is suggested that upgrades to 1.9.2 happen offline, where all instances are shutdown, upgraded and restarted. Upgrading from 1.9.2 to any other version works as described elsewhere in this guide.

Upgrading between 1.9.x versions follows the same general pattern as described in the first part of this guide, but is much simpler because of the compatibility of the configuration options between 1.9.x releases. We will describe a step by step procedure, aimed at reducing the master switches to a single change.

26.5.6. Step 1: On each slave perform the upgrade

Download and unpack the new version. Copy over any configuration settings, ensuring that newly introduced settings have proper values (see [Section 26.2, “Setup and configuration”](#)). Don't forget to copy over any server plugins you may have.

First, shutdown the neo4j instance with

```
service neo4j-service stop
```

Next, uninstall it

```
service neo4j-service remove
```

Now you can copy over the database. Assuming the old instance is at `/opt/old-neo4j` and the newly unpacked under `/opt/neo4j-enterprise-1.9.x` the proper command would be

```
cp -R /opt/old-neo4j/data/graph.db /opt/neo4j-enterprise-1.9.x/data/
```

Next install the neo4j service, which also starts it

```
/opt/neo4j-enterprise-1.9.x/bin/neo4j install
```

Now check that the services are running and that webadmin reports the version 1.9.x. Transactions should also be applied from the master as usual.

26.5.7. Step 2: Upgrade the master, complete the procedure

Go to the current master and execute step 1. The moment it will be stopped another instance will take over, transitioning the cluster to the new 1.9.x version. Finish Step 1 on this machine as well and you will have completed the process.

26.6. High Availability setup tutorial

This guide will help you understand how to configure and deploy a Neo4j High Availability cluster. Two scenarios will be considered:

- Configuring 3 instances to be deployed on 3 separate machines, in a setting similar to what might be encountered in a production environment.
- Modifying the former to make it possible to run a cluster of 3 instances on the same physical machine, which is particularly useful during development.

26.6.1. Background

Each instance in a Neo4j HA cluster must be assigned an integer ID, which serves as its unique identifier. At startup, a Neo4j instance contacts the other instances specified in the `ha.initial_hosts` configuration option.

When an instance establishes a connection to any other, it determines the current state of the cluster and ensures that it is eligible to join. To be eligible the Neo4j instance must host the same database store as other members of the cluster (although it is allowed to be in an older state), or be a new deployment without a database store.



Explicitly configure IP Addresses/Hostnames for a cluster

Neo4j will attempt to configure IP addresses for itself in the absence of explicit configuration. However in typical operational environments where machines have multiple network cards and support IPv4 and IPv6 it is **strongly** recommended that the operator explicitly sets the IP address/hostname configuration for each machine in the cluster.

Let's examine the available settings and the values they accept.

ha.server_id

`ha.server_id` is the cluster identifier for each instance. It must be a positive integer and must be unique among all Neo4j instances in the cluster.

For example, `ha.server_id=1`.

ha.cluster_server

`ha.cluster_server` is an address/port setting that specifies where the Neo4j instance will listen for cluster communications (like heartbeat messages). The default port is 5001. In the absence of a specified IP address, Neo4j will attempt to find a valid interface for binding. While this behavior typically results in a well-behaved server, it is **strongly** recommended that users explicitly choose an IP address bound to the network interface of their choosing to ensure a coherent cluster deployment.

For example, `ha.cluster_server=192.168.33.22:5001` will listen for cluster communications on the network interface bound to the 192.168.33.0 subnet on port 5001.

ha.initial_hosts

`ha.initial_hosts` is a comma separated list of address/port pairs, which specify how to reach other Neo4j instances in the cluster (as configured via their `ha.cluster_server` option). These hostname/ports will be used when the Neo4j instances starts, to allow it up to find and join the cluster. Specifying an instance's own address is permitted.



Warning

Do **not** use any whitespace in this configuration option.

For example, `ha.initial_hosts=192.168.33.22:5001,192.168.33.21:5001` will attempt to reach Neo4j instances listening on 192.168.33.22 on port 5001 and 192.168.33.21 on port 5001 on the 192.168.33.0 subnet.

ha.server

`ha.server` is an address/port setting that specifies where the Neo4j instance will listen for transactions (changes to the graph data) from the cluster master. The default port is 6001. In the absence of a specified IP address, Neo4j will attempt to find a valid interface for binding. While this behavior typically results in a well-behaved server, it is **strongly** recommended that users explicitly choose an IP address bound to the network interface of their choosing to ensure a coherent cluster topology.

`ha.server` must use a different port to `ha.cluster_server`.

For example, `ha.server=192.168.33.22:6001` will listen for cluster communications on the network interface bound to the 192.168.33.0 subnet on port 6001.



Address/port format

The `ha.cluster_server` and `ha.server` configuration options are specified as `<IP address>:<port>`.

For `ha.server` the IP address **MUST** be the address assigned to one of the host's network interfaces.

For `ha.cluster_server` the IP address **MUST** be the address assigned to one of the host's network interfaces, or the value `0.0.0.0`, which will cause Neo4j to listen on every network interface.

Either the address or the port can be omitted, in which case the default for that part will be used. If the address is omitted, then the port must be preceded with a colon (eg. `:5001`).

The syntax for setting the port range is: `<hostname>:<first port>[-<second port>]`. In this case, Neo4j will test each port in sequence, and select the first that is unused. Note that this usage is not permitted when the hostname is specified as `0.0.0.0` (the "all interfaces" address).

26.6.2. Getting started: Setting up a production cluster

Download and unpack Neo4j Enterprise

Download Neo4j Enterprise from [the Neo4j download site](http://neo4j.org/download) `<http://neo4j.org/download>`, and unpack on 3 separate machines.

Configure HA related settings for each installation

The following settings should be configured for each Neo4j installation. Note that all 3 installations have the same configuration, except for the `ha.server_id` property.

Neo4j instance #1 — neo4j-01.local

`conf/neo4j.properties`:

```
# Unique server id for this Neo4j instance
# can not be negative id and must be unique
ha.server_id = 1

# List of other known instances in this cluster
ha.initial_hosts = neo4j-01.local:5001,neo4j-02.local:5001,neo4j-03.local:5001
# Alternatively, use IP addresses:
#ha.initial_hosts = 192.168.0.20:5001,192.168.0.21:5001,192.168.0.22:5001
```

conf/neo4j-server.properties

```
# Let the webservice only listen on the specified IP.
org.neo4j.server.webservice.address=0.0.0.0

# HA - High Availability
# SINGLE - Single mode, default.
org.neo4j.server.database.mode=HA
```

Neo4j instance #2 — neo4j-02.local

conf/neo4j.properties:

```
# Unique server id for this Neo4j instance
# can not be negative id and must be unique
ha.server_id = 2

# List of other known instances in this cluster
ha.initial_hosts = neo4j-01.local:5001,neo4j-02.local:5001,neo4j-03.local:5001
# Alternatively, use IP addresses:
#ha.initial_hosts = 192.168.0.20:5001,192.168.0.21:5001,192.168.0.22:5001
```

conf/neo4j-server.properties

```
# Let the webservice only listen on the specified IP.
org.neo4j.server.webservice.address=0.0.0.0

# HA - High Availability
# SINGLE - Single mode, default.
org.neo4j.server.database.mode=HA
```

Neo4j instance #3 — neo4j-03.local

conf/neo4j.properties:

```
# Unique server id for this Neo4j instance
# can not be negative id and must be unique
ha.server_id = 3

# List of other known instances in this cluster
ha.initial_hosts = neo4j-01.local:5001,neo4j-02.local:5001,neo4j-03.local:5001
# Alternatively, use IP addresses:
#ha.initial_hosts = 192.168.0.20:5001,192.168.0.21:5001,192.168.0.22:5001
```

conf/neo4j-server.properties

```
# Let the webservice only listen on the specified IP.
org.neo4j.server.webservice.address=0.0.0.0

# HA - High Availability
# SINGLE - Single mode, default.
org.neo4j.server.database.mode=HA
```

Start the Neo4j Servers

Start the Neo4j servers as normal. Note the startup order does not matter.

```
neo4j-01$ ./bin/neo4j start
```

```
neo4j-02$ ./bin/neo4j start
```

```
neo4j-03$ ./bin/neo4j start
```



Startup Time

When running in HA mode, the startup script returns immediately instead of waiting for the server to become available. This is because the instance does not accept any requests

until a cluster has been formed. In the example above this happens when you startup the second instance. To keep track of the startup state you can follow the messages in console.log - the path to that is printed before the startup script returns.

Now, you should be able to access the 3 servers and check their HA status:

<http://neo4j-01.local:7474/webadmin/#/info/org.neo4j/High%20Availability/>

<http://neo4j-02.local:7474/webadmin/#/info/org.neo4j/High%20Availability/>

<http://neo4j-03.local:7474/webadmin/#/info/org.neo4j/High%20Availability/>



Tip

You can replace database #3 with an *arbiter* instance, see [Arbiter Instances](#).

That's it! You now have a Neo4j HA cluster of 3 instances running. You can start by making a change on any instance and those changes will be propagated between them. For more HA related configuration options take a look at [HA Configuration](#).

26.6.3. Alternative setup: Creating a local cluster for testing

If you want to start a cluster similar to the one described above, but for development and testing purposes, it is convenient to run all Neo4j instances on the same machine. This is easy to achieve, although it requires some additional configuration as the defaults will conflict with each other.

Download and unpack Neo4j Enterprise

Download Neo4j Enterprise from [the Neo4j download site](http://neo4j.org/download) <<http://neo4j.org/download>>, and unpack into 3 separate directories on your test machine.

Configure HA related settings for each installation

The following settings should be configured for each Neo4j installation.

Neo4j instance #1 — ~/neo4j-01

conf/neo4j.properties:

```
# Unique server id for this Neo4j instance
# can not be negative id and must be unique
ha.server_id = 1

# IP and port for this instance to bind to for communicating data with the
# other neo4j instances in the cluster.
ha.server = 127.0.0.1:6363
online_backup_server = 127.0.0.1:6366

# IP and port for this instance to bind to for communicating cluster information
# with the other neo4j instances in the cluster.
ha.cluster_server = 127.0.0.1:5001

# List of other known instances in this cluster
ha.initial_hosts = 127.0.0.1:5001,127.0.0.1:5002,127.0.0.1:5003
```

conf/neo4j-server.properties

```
# http port (for all data, administrative, and UI access)
org.neo4j.server.webserver.port=7474

# https port (for all data, administrative, and UI access)
org.neo4j.server.webserver.https.port=7484
```

```
# HA - High Availability
# SINGLE - Single mode, default.
org.neo4j.server.database.mode=HA
```

Neo4j instance #2 — ~/neo4j-02

conf/neo4j.properties:

```
# Unique server id for this Neo4j instance
# can not be negative id and must be unique
ha.server_id = 2

# IP and port for this instance to bind to for communicating data with the
# other neo4j instances in the cluster.
ha.server = 127.0.0.1:6364
online_backup_server = 127.0.0.1:6367

# IP and port for this instance to bind to for communicating cluster information
# with the other neo4j instances in the cluster.
ha.cluster_server = 127.0.0.1:5002

# List of other known instances in this cluster
ha.initial_hosts = 127.0.0.1:5001,127.0.0.1:5002,127.0.0.1:5003
```

conf/neo4j-server.properties

```
# http port (for all data, administrative, and UI access)
org.neo4j.server.webserver.port=7475

# https port (for all data, administrative, and UI access)
org.neo4j.server.webserver.https.port=7485

# HA - High Availability
# SINGLE - Single mode, default.
org.neo4j.server.database.mode=HA
```

Neo4j instance #3 — ~/neo4j-03

conf/neo4j.properties:

```
# Unique server id for this Neo4j instance
# can not be negative id and must be unique
ha.server_id = 3

# IP and port for this instance to bind to for communicating data with the
# other neo4j instances in the cluster.
ha.server = 127.0.0.1:6365
online_backup_server = 127.0.0.1:6368

# IP and port for this instance to bind to for communicating cluster information
# with the other neo4j instances in the cluster.
ha.cluster_server = 127.0.0.1:5003

# List of other known instances in this cluster
ha.initial_hosts = 127.0.0.1:5001,127.0.0.1:5002,127.0.0.1:5003
```

conf/neo4j-server.properties

```
# http port (for all data, administrative, and UI access)
org.neo4j.server.webserver.port=7476

# https port (for all data, administrative, and UI access)
org.neo4j.server.webserver.https.port=7486

# HA - High Availability
```

```
# SINGLE - Single mode, default.  
org.neo4j.server.database.mode=HA
```

Start the Neo4j Servers

Start the Neo4j servers as normal. Note the startup order does not matter.

```
localhost:~/neo4j-01$ ./bin/neo4j start
```

```
localhost:~/neo4j-02$ ./bin/neo4j start
```

```
localhost:~/neo4j-03$ ./bin/neo4j start
```

Now, you should be able to access the 3 servers and check their HA status:

<http://127.0.0.1:7474/webadmin/#/info/org.neo4j/High%20Availability/>

<http://127.0.0.1:7475/webadmin/#/info/org.neo4j/High%20Availability/>

<http://127.0.0.1:7476/webadmin/#/info/org.neo4j/High%20Availability/>

26.7. REST endpoint for HA status information

26.7.1. Introduction

A common use case for Neo4j HA clusters is to direct all write requests to the master while using slaves for read operations, distributing the read load across the cluster and gain failover capabilities for your deployment. The most common way to achieve this is to place a load balancer in front of the HA cluster, an example being shown with [HA Proxy](#). As you can see in that guide, it makes use of a REST endpoint to discover which instance is the master and direct write load to it. In this section, we'll deal with this REST endpoint and explain its semantics.

26.7.2. The endpoints

Each HA instance comes with 3 endpoints regarding its HA status. They are complimentary but each may be used depending on your load balancing needs and your production setup. Those are:

- `/db/manage/server/ha/master`
- `/db/manage/server/ha/slave`
- `/db/manage/server/ha/available`

The `/master` and `/slave` endpoints can be used to direct write and non-write traffic respectively to specific instances. This is the optimal way to take advantage of Neo4j's scaling characteristics. The `/available` endpoint exists for the general case of directing arbitrary request types to instances that are available for transaction processing.

To use the endpoints, perform an HTTP GET operation on either and the following will be returned:

HA REST endpoint responses

Endpoint	Instance State	Returned Code	Body text
<code>/db/manage/server/ha/master</code>	Master	200 OK	true
	Slave	404 Not Found	false
	Unknown	404 Not Found	UNKNOWN
<code>/db/manage/server/ha/slave</code>	Master	404 Not Found	false
	Slave	200 OK	true
	Unknown	404 Not Found	UNKNOWN
<code>/db/manage/server/ha/available</code>	Master	200 OK	master
	Slave	200 OK	slave
	Unknown	404 Not Found	UNKNOWN

26.7.3. Examples

From the command line, a common way to ask those endpoints is to use `curl`. With no arguments, `curl` will do an HTTP GET on the URI provided and will output the body text, if any. If you also want to get the response code, just add the `-v` flag for verbose output. Here are some examples:

- Requesting master endpoint on a running master with verbose output

```
#> curl -v localhost:7474/db/manage/server/ha/master
```

```
* About to connect() to localhost port 7474 (#0)
* Trying ::1...
* connected
* Connected to localhost (::1) port 7474 (#0)
> GET /db/manage/server/ha/master HTTP/1.1
> User-Agent: curl/7.24.0 (x86_64-apple-darwin12.0) libcurl/7.24.0 OpenSSL/0.9.8r zlib/1.2.5
> Host: localhost:7474
> Accept: */*
>
< HTTP/1.1 200 OK
< Content-Type: text/plain
< Access-Control-Allow-Origin: *
< Transfer-Encoding: chunked
< Server: Jetty(6.1.25)
<
* Connection #0 to host localhost left intact
true* Closing connection #0
```

- Requesting slave endpoint on a running master without verbose output:

```
#> curl localhost:7474/db/manage/server/ha/slave
false
```

- Finally, requesting the master endpoint on a slave with verbose output

```
#> curl -v localhost:7475/db/manage/server/ha/master
* About to connect() to localhost port 7475 (#0)
* Trying ::1...
* connected
* Connected to localhost (::1) port 7475 (#0)
> GET /db/manage/server/ha/master HTTP/1.1
> User-Agent: curl/7.24.0 (x86_64-apple-darwin12.0) libcurl/7.24.0 OpenSSL/0.9.8r zlib/1.2.5
> Host: localhost:7475
> Accept: */*
>
< HTTP/1.1 404 Not Found
< Content-Type: text/plain
< Access-Control-Allow-Origin: *
< Transfer-Encoding: chunked
< Server: Jetty(6.1.25)
<
* Connection #0 to host localhost left intact
false* Closing connection #0
```



Unknown status

The UNKNOWN status exists to describe when a Neo4j instance is neither master nor slave. For example, the instance could be transitioning between states (master to slave in a recovery scenario or slave being promoted to master in the event of failure). If the UNKNOWN status is returned, the client should not treat the instance as a master or a slave and should instead pick another instance in the cluster to use, wait for the instance to transit from the UNKNOWN state, or undertake restorative action via systems admin.

26.8. Setting up HAProxy as a load balancer

In the Neo4j HA architecture, the cluster is typically fronted by a load balancer. In this section we will explore how to set up HAProxy to perform load balancing across the HA cluster.

26.8.1. Installing HAProxy

For this tutorial we will assume a Linux environment. We will also be installing HAProxy from source, and we'll be using version 1.4.18. You need to ensure that your Linux server has a development environment set up. On Ubuntu/apt systems, simply do:

```
aptitude install build-essential
```

And on CentOS/yum systems do:

```
yum -y groupinstall 'Development Tools'
```

Then download the tarball from the [HAProxy website](http://haproxy.1wt.eu/) <http://haproxy.1wt.eu/>. Once you've downloaded it, simply build and install HAProxy:

```
tar -zxvf haproxy-1.4.18.tar.gz
cd haproxy-1.4.18
make
cp haproxy /usr/sbin/haproxy
```

Or specify a target for make (TARGET=linux26 for linux kernel 2.6 or above or linux24 for 2.4 kernel)

```
tar -zxvf haproxy-1.4.18.tar.gz
cd haproxy-1.4.18
make TARGET=linux26
cp haproxy /usr/sbin/haproxy
```

26.8.2. Configuring HAProxy

HAProxy can be configured in many ways. The full documentation is available at their website.

For this example, we will configure HAProxy to load balance requests to three HA servers. Simply write the follow configuration to /etc/haproxy.cfg:

```
global
    daemon
    maxconn 256

defaults
    mode http
    timeout connect 5000ms
    timeout client 50000ms
    timeout server 50000ms

frontend http-in
    bind *:80
    default_backend neo4j

backend neo4j
    option httpchk GET /db/manage/server/ha/available
    server s1 10.0.1.10:7474 maxconn 32
    server s2 10.0.1.11:7474 maxconn 32
    server s3 10.0.1.12:7474 maxconn 32

listen admin
    bind *:8080
```

```
stats enable
```

HAProxy can now be started by running:

```
/usr/sbin/haproxy -f /etc/haproxy.cfg
```

You can connect to <http://<ha-proxy-ip>:8080/haproxy?stats> to view the status dashboard. This dashboard can be moved to run on port 80, and authentication can also be added. See the HAProxy documentation for details on this.

26.8.3. Optimizing for reads and writes

Neo4j provides a catalogue of *health check URLs* (see [Section 26.7, “REST endpoint for HA status information”](#)) that HAProxy (or any load balancer for that matter) can use to distinguish machines using HTTP response codes. In the example above we used the `/available` endpoint, which directs requests to machines that are generally available for transaction processing (they are alive!).

However, it is possible to have requests directed to slaves only, or to the master only. If you are able to distinguish in your application between requests that write, and requests that only read, then you can take advantage of two (logical) load balancers: one that sends all your writes to the master, and one that sends all your read-only requests to a slave. In HAProxy you build logical load balancers by adding multiple backends.

The trade-off here is that while Neo4j allows slaves to proxy writes for you, this indirection unnecessarily ties up resources on the slave and adds latency to your write requests. Conversely, you don't particularly want read traffic to tie up resources on the master; Neo4j allows you to scale out for reads, but writes are still constrained to a single instance. If possible, that instance should exclusively do writes to ensure maximum write performance.

The following example excludes the master from the set of machines using the `/slave` endpoint.

```
global
  daemon
  maxconn 256

defaults
  mode http
  timeout connect 5000ms
  timeout client 50000ms
  timeout server 50000ms

frontend http-in
  bind *:80
  default_backend neo4j-slaves

backend neo4j-slaves
  option httpchk GET /db/manage/server/ha/slave
  server s1 10.0.1.10:7474 maxconn 32 check
  server s2 10.0.1.11:7474 maxconn 32 check
  server s3 10.0.1.12:7474 maxconn 32 check

listen admin
  bind *:8080
  stats enable
```



Note

In practice, writing to a slave is uncommon. While writing to slaves has the benefit of ensuring that data is persisted in two places (the slave and the master), it comes at a cost. The cost is that the slave must immediately become consistent with the master by applying any missing transactions and then synchronously apply the new transaction with

the master. This is a more expensive operation than writing to the master and having the master push changes to one or more slaves.

26.8.4. Cache-based sharding with HAProxy

Neo4j HA enables what is called cache-based sharding. If the dataset is too big to fit into the cache of any single machine, then by applying a consistent routing algorithm to requests, the caches on each machine will actually cache different parts of the graph. A typical routing key could be user ID.

In this example, the user ID is a query parameter in the URL being requested. This will route the same user to the same machine for each request.

```
global
  daemon
  maxconn 256

defaults
  mode http
  timeout connect 5000ms
  timeout client 50000ms
  timeout server 50000ms

frontend http-in
  bind *:80
  default_backend neo4j-slaves

backend neo4j-slaves
  balance url_param user_id
  server s1 10.0.1.10:7474 maxconn 32
  server s2 10.0.1.11:7474 maxconn 32
  server s3 10.0.1.12:7474 maxconn 32

listen admin
  bind *:8080
  stats enable
```

Naturally the health check and query parameter-based routing can be combined to only route requests to slaves by user ID. Other load balancing algorithms are also available, such as routing by source IP (source), the URI (uri) or HTTP headers(hdr()).

Chapter 27. Backup



Note

The Backup features are only available in the Neo4j Enterprise Edition.

Backups are performed over the network live from a running graph database onto a local copy. There are two types of backup: full and incremental.

A *full backup* copies the database files without acquiring any locks, allowing for continued operations on the target instance. This of course means that while copying, transactions will continue and the store will change. For this reason, the transaction that was running when the backup operation started is noted and, when the copy operation completes, all transactions from the latter down to the one happening at the end of the copy are replayed on the backup files. This ensures that the backed up data represent a consistent and up-to-date snapshot of the database storage.

In contrast, an *incremental backup* does not copy store files — instead it copies the logs of the transactions that have taken place since the last full or incremental backup which are then replayed over an existing backup store. This makes incremental backups far more efficient than doing full backups every time but they also require that a *full backup* has taken place before they are executed.

The backup tool will detect whether you are trying to run a full backup or an incremental one by inspecting the target directory. Regardless of the mode a backup is created with, the resulting files represent a consistent database snapshot and they can be used to boot up a Neo4j instance.

The database to be backed up is specified using a URI with syntax

```
<running mode>://<host>[:port]{,<host>[:port]*}
```

Running mode must be defined and is either *single* for non-HA or *ha* for HA clusters. The `<host>[:port]` part points to a host running the database, on port *port* if not the default. The additional *host:port* arguments are useful for passing multiple cluster members.



Important

As of version 1.9, backups are enabled by default. That means that the configuration parameter `online_backup_enabled` defaults to true and that makes the backup service available on the default port (6362). To enable the backup service on a different port use `online_backup_server=:9999`.

27.1. Embedded and Server

To perform a backup from a running embedded or server database run:

```
# Performing a full backup: create a blank directory and run the backup tool
mkdir /mnt/backup/neo4j-backup
./neo4j-backup -from single://192.168.1.34 -to /mnt/backup/neo4j-backup

# Performing an incremental backup: just specify the location of your previous backup
./neo4j-backup -from single://192.168.1.34 -to /mnt/backup/neo4j-backup

# Performing an incremental backup where the service is registered on a custom port
./neo4j-backup -from single://192.168.1.34:9999 -to /mnt/backup/neo4j-backup
```

27.2. Online Backup from Java

In order to programmatically backup your data full or subsequently incremental from a JVM based program, you need to write Java code like

```
OnlineBackup backup = OnlineBackup.from( "127.0.0.1" );  
backup.full( backupPath.getPath() );  
backup.incremental( backupPath.getPath() );
```

For more information, please see [the Javadocs for OnlineBackup](http://components.neo4j.org/neo4j-enterprise/1.9.7/apidocs/org/neo4j/backup/OnlineBackup.html) <<http://components.neo4j.org/neo4j-enterprise/1.9.7/apidocs/org/neo4j/backup/OnlineBackup.html>>

27.3. High Availability

To perform a backup on an HA cluster you specify one or more members of the target HA cluster. Note that the addresses you must provide are the cluster server addresses and not the HA server addresses. That is, use the value of the `ha.cluster_server` setting in the configuration.

```
# Performing a backup from HA cluster, specifying only one cluster member
./neo4j-backup -from ha://192.168.1.15:5001 -to /mnt/backup/neo4j-backup

# Performing a backup from HA cluster, specifying two possible cluster members
./neo4j-backup -from ha://192.168.1.15:5001,192.168.1.16:5002 -to /mnt/backup/neo4j-backup
```

27.4. Restoring Your Data

The Neo4j backups are fully functional databases. To use a backup, all you need to do replace your database folder with the backup. Just make sure the database isn't running while replacing the folder.

Chapter 28. Security

Neo4j in itself does not enforce security on the data level. However, there are different aspects that should be considered when using Neo4j in different scenarios.

28.1. Securing access to the Neo4j Server

28.1.1. Secure the port and remote client connection accepts

By default, the Neo4j Server is bundled with a Web server that binds to host localhost on port 7474, answering only requests from the local machine.

This is configured in the `conf/neo4j-server.properties` file:

```
# http port (for all data, administrative, and UI access)
org.neo4j.server.webserver.port=7474

#let the webservice only listen on the specified IP. Default
#is localhost (only accept local connections). Uncomment to allow
#any connection.
#org.neo4j.server.webserver.address=0.0.0.0
```

If you need to enable access from external hosts, configure the Web server in the `conf/neo4j-server.properties` by setting the property `org.neo4j.server.webserver.address=0.0.0.0` to enable access from any host.

28.1.2. Arbitrary code execution

By default, the Neo4j Server comes with some places where arbitrary code execution can happen. These are the [Section 22.14, “Traversals”](#) REST endpoints. To secure these, either disable them completely by removing offending plugins from the server classpath, or secure access to these URLs through proxies or Authorization Rules. Also, the [Java Security Manager](http://docs.oracle.com/javase/1.4.2/docs/api/java/lang/SecurityManager.html) can be used to secure parts of the codebase.

28.1.3. HTTPS support

The Neo4j server includes built in support for SSL encrypted communication over HTTPS. The first time the server starts, it automatically generates a self-signed SSL certificate and a private key. Because the certificate is self signed, it is not safe to rely on for production use, instead, you should provide your own key and certificate for the server to use.

To provide your own key and certificate, replace the generated key and certificate, or change the `neo4j-server.properties` file to set the location of your certificate and key:

```
# Certificate location (auto generated if the file does not exist)
org.neo4j.server.webserver.https.cert.location=ssl/snakeoil.cert

# Private key location (auto generated if the file does not exist)
org.neo4j.server.webserver.https.key.location=ssl/snakeoil.key
```

Note that the key should be unencrypted. Make sure you set correct permissions on the private key, so that only the Neo4j server user can read/write it.

Neo4j also supports chained SSL certificates. This requires to have all certificates in PEM format combined in one file and the private key needs to be in DER format.

You can set what port the HTTPS connector should bind to in the same configuration file, as well as turn HTTPS off:

```
# Turn https-support on/off
org.neo4j.server.webserver.https.enabled=true

# https port (for all data, administrative, and UI access)
```

```
org.neo4j.server.webserver.https.port=443
```

28.1.4. Server Authorization Rules

Administrators may require more fine-grained security policies in addition to IP-level restrictions on the Web server. Neo4j server supports administrators in allowing or disallowing access the specific aspects of the database based on credentials that users or applications provide.

To facilitate domain-specific authorization policies in Neo4j Server, [SecurityRules](http://components.neo4j.org/neo4j-server/1.9.7/apidocs/org/neo4j/server/rest/security/SecurityRule.html) <http://components.neo4j.org/neo4j-server/1.9.7/apidocs/org/neo4j/server/rest/security/SecurityRule.html> can be implemented and registered with the server. This makes scenarios like user and role based security and authentication against external lookup services possible.

Enforcing Server Authorization Rules

In this example, a (dummy) failing security rule is registered to deny access to all URIs to the server by listing the rules class in *neo4j-server.properties*:

```
org.neo4j.server.rest.security_rules=my.rules.PermanentlyFailingSecurityRule
```

with the rule source code of:

```
public class PermanentlyFailingSecurityRule implements SecurityRule
{
    public static final String REALM = "WallyWorld"; // as per RFC2617 :-)

    @Override
    public boolean isAuthorized( HttpServletRequest request )
    {
        return false; // always fails - a production implementation performs
                       // deployment-specific authorization logic here
    }

    @Override
    public String forUriPath()
    {
        return "/*";
    }

    @Override
    public String wwwAuthenticateHeader()
    {
        return SecurityFilter.basicAuthenticationResponse(REALM);
    }
}
```

With this rule registered, any access to the server will be denied. In a production-quality implementation the rule will likely lookup credentials/claims in a 3rd-party directory service (e.g. LDAP) or in a local database of authorized users.

Example request

- POST http://localhost:7474/db/data/node
- Accept: application/json; charset=UTF-8

Example response

- 401: Unauthorized
- WWW-Authenticate: Basic realm="WallyWorld"

Using Wildcards to Target Security Rules

In this example, a security rule is registered to deny access to all URIs to the server by listing the rule(s) class(es) in *neo4j-server.properties*. In this case, the rule is registered using a wildcard URI path (where * characters can be used to signify any part of the path). For example `/users*` means the rule will be bound to any resources under the `/users` root path. Similarly `/users*type*` will bind the rule to resources matching URIs like `/users/fred/type/premium`.

```
org.neo4j.server.rest.security_rules=my.rules.PermanentlyFailingSecurityRuleWithWildcardPath
```

with the rule source code of:

```
public String forUriPath()
{
    return "/protected/*";
}
```

With this rule registered, any access to URIs under `/protected/` will be denied by the server. Using wildcards allows flexible targeting of security rules to arbitrary parts of the server's API, including any unmanaged extensions or managed plugins that have been registered.

Example request

- GET `http://localhost:7474/protected/tree/starts/here/dummy/more/stuff`
- Accept: `application/json`

Example response

- 401: Unauthorized
- WWW-Authenticate: `Basic realm="WallyWorld"`

Using Complex Wildcards to Target Security Rules

In this example, a security rule is registered to deny access to all URIs matching a complex pattern. The config looks like this:

```
org.neo4j.server.rest.security_rules=my.rules.PermanentlyFailingSecurityRuleWithComplexWildcardPath
```

with the rule source code of:

```
public class PermanentlyFailingSecurityRuleWithComplexWildcardPath implements SecurityRule
{
    public static final String REALM = "WallyWorld"; // as per RFC2617 :-)

    @Override
    public boolean isAuthorized( HttpServletRequest request )
    {
        return false;
    }

    @Override
    public String forUriPath()
    {
        return "/protected/*/something/else/*/final/bit";
    }

    @Override
    public String wwwAuthenticateHeader()
    {
        return SecurityFilter.basicAuthenticationResponse(REALM);
    }
}
```

```
}
}
```

Example request

- GET http://localhost:7474/protected/wildcard_replacement/x/y/z/something/else/more_wildcard_replacement/a/b/c/final/bit/more/stuff
- Accept: application/json

Example response

- 401: Unauthorized
- WWW-Authenticate: Basic realm="WallyWorld"

28.1.5. Hosted Scripting



Important

The neo4j server exposes remote scripting functionality by default that allow full access to the underlying system. Exposing your server without implementing a security layer poses a substantial security vulnerability.

28.1.6. Security in Depth

Although the Neo4j server has a number of security features built-in (see the above chapters), for sensitive deployments it is often sensible to front against the outside world it with a proxy like Apache `mod_proxy`¹.

This provides a number of advantages:

- Control access to the Neo4j server to specific IP addresses, URL patterns and IP ranges. This can be used to make for instance only the `/db/data` namespace accessible to non-local clients, while the `/db/admin` URLs only respond to a specific IP address.

```
<Proxy *>
  Order Deny,Allow
  Deny from all
  Allow from 192.168.0
</Proxy>
```

While equivalent functionality can be implemented with Neo4j's `SecurityRule` plugins (see above), for operations professionals configuring servers like Apache is often preferable to developing plugins. However it should be noted that where both approaches are used, they will work harmoniously providing the behavior is consistent across proxy server and `SecurityRule` plugins.

- Run Neo4j Server as a non-root user on a Linux/Unix system on a port < 1000 (e.g. port 80) using

```
ProxyPass /neo4jdb/data http://localhost:7474/db/data
ProxyPassReverse /neo4jdb/data http://localhost:7474/db/data
```

- Simple load balancing in a clustered environment to load-balance read load using the Apache `mod_proxy_balancer`² plugin

```
<Proxy balancer://mycluster>
```

¹http://httpd.apache.org/docs/2.2/mod/mod_proxy.html

²http://httpd.apache.org/docs/2.2/mod/mod_proxy_balancer.html

```
BalancerMember http://192.168.1.50:80  
BalancerMember http://192.168.1.51:80  
</Proxy>  
ProxyPass /test balancer://mycluster
```

Chapter 29. Monitoring



Note

Most of the monitoring features are only available in the Advanced and Enterprise editions of Neo4j.

In order to be able to continuously get an overview of the health of a Neo4j database, there are different levels of monitoring facilities available. Most of these are exposed through [JMX](http://www.oracle.com/technetwork/java/javase/tech/javamanagement-140525.html) <<http://www.oracle.com/technetwork/java/javase/tech/javamanagement-140525.html>>.

29.1. Adjusting remote JMX access to the Neo4j Server

Per default, the Neo4j Advanced Server and Neo4j Enterprise Server editions do not allow remote JMX connections, since the relevant options in the *conf/neo4j-wrapper.conf* configuration file are commented out. To enable this feature, you have to remove the # characters from the various `com.sun.management.jmxremote` options there.

When commented in, the default values are set up to allow remote JMX connections with certain roles, refer to the *conf/jmx.password*, *conf/jmx.access* and *conf/wrapper.conf* files for details.

Make sure that *conf/jmx.password* has the correct file permissions. The owner of the file has to be the user that will run the service, and the permissions should be read only for that user. On Unix systems, this is `0600`.

On Windows, follow the tutorial at <http://docs.oracle.com/javase/6/docs/technotes/guides/management/security-windows.html> to set the correct permissions. If you are running the service under the Local System Account, the user that owns the file and has access to it should be SYSTEM.

With this setup, you should be able to connect to JMX monitoring of the Neo4j server using `<IP-OF-SERVER>:3637`, with the username `monitor` and the password `Neo4j`.

Note that it is possible that you have to update the permissions and/or ownership of the *conf/jmx.password* and *conf/jmx.access* files — refer to the relevant section in *conf/wrapper.conf* for details.



Warning

For maximum security, please adjust at least the password settings in *conf/jmx.password* for a production installation.

For more details, see: <http://download.oracle.com/javase/6/docs/technotes/guides/management/agent.html>

29.2. How to connect to a Neo4j instance using JMX and JConsole

First, start your embedded database or the Neo4j Server, for instance using


```
$NEO4j_HOME/bin/neo4j start
```

Now, start JConsole with

```
$JAVA_HOME/bin/jconsole
```

Connect to the process running your Neo4j database instance:


Figure 29.1. Connecting JConsole to the Neo4j Java process



Now, beside the MBeans exposed by the JVM, you will see an `org.neo4j` section in the MBeans tab. Under that, you will have access to all the monitoring information exposed by Neo4j.

For opening JMX to remote monitoring access, please see [Section 29.1, “Adjusting remote JMX access to the Neo4j Server”](#) and [the JMX documentation](http://docs.oracle.com/javase/6/docs/technotes/guides/management/agent.html#gdenl) <<http://docs.oracle.com/javase/6/docs/technotes/guides/management/agent.html#gdenl>>. When using Neo4j in embedded mode, make sure to pass the `com.sun.management.jmxremote.port=portNum` or other configuration as JVM parameters to your running Java process.

Figure 29.2. Neo4j MBeans View



29.3. How to connect to the JMX monitoring programmatically

In order to programmatically connect to the Neo4j JMX server, there are some convenience methods in the Neo4j Management component to help you find out the most commonly used monitoring attributes of Neo4j. See [Section 4.8, “Reading a management attribute”](#) for an example.

Once you have access to this information, you can use it to for instance expose the values to [SNMP](#) <http://en.wikipedia.org/wiki/Simple_Network_Management_Protocol> or other monitoring systems.

29.4. Reference of supported JMX MBeans

MBeans exposed by Neo4j

- **Branched Store**: Information about the branched stores present in this HA cluster member
- **Cache/NodeCache**: Information about the caching in Neo4j
- **Cache/RelationshipCache**: Information about the caching in Neo4j
- **Configuration**: The configuration parameters used to configure Neo4j
- **Diagnostics**: Diagnostics provided by Neo4j
- **High Availability**: Information about an instance participating in a HA cluster
- **Kernel**: Information about the Neo4j kernel
- **Locking**: Information about the Neo4j lock status
- **Memory Mapping**: The status of Neo4j memory mapping
- **Primitive count**: Estimates of the numbers of different kinds of Neo4j primitives
- **Store file sizes**: Information about the sizes of the different parts of the Neo4j graph store
- **Transactions**: Information about the Neo4j transaction manager
- **XA Resources**: Information about the XA transaction manager



Note

For additional information on the primitive datatypes (int, long etc.) used in the JMX attributes, please see [Property value types](#).

MBean Branched Store (org.neo4j.management.BranchedStore) Attributes

Name	Description	Type	Read	Write
<i>Information about the branched stores present in this HA cluster member</i>				
BranchedStores	A list of the branched stores	org.neo4j.management.BranchedStoreInfo[] as CompositeData[]	yes	no

MBean Cache/NodeCache (org.neo4j.management.Cache) Attributes

Name	Description	Type	Read	Write
<i>Information about the caching in Neo4j</i>				
CacheSize	The size of this cache (nr of entities or total size in bytes)	long	yes	no
CacheType	The type of cache used by Neo4j	String	yes	no
HitCount	The number of times a cache query returned a result	long	yes	no
MissCount	The number of times a cache query did not return a result	long	yes	no

MBean Cache/NodeCache (org.neo4j.management.Cache) Operations

Name	Description	ReturnType	Signature
clear	Clears the Neo4j caches	void	(no parameters)

MBean Cache/RelationshipCache (org.neo4j.management.Cache) Attributes

Name	Description	Type	Read	Write
<i>Information about the caching in Neo4j</i>				
CacheSize	The size of this cache (nr of entities or total size in bytes)	long	yes	no
CacheType	The type of cache used by Neo4j	String	yes	no
HitCount	The number of times a cache query returned a result	long	yes	no
MissCount	The number of times a cache query did not return a result	long	yes	no

MBean Cache/RelationshipCache (org.neo4j.management.Cache) Operations

Name	Description	ReturnType	Signature
clear	Clears the Neo4j caches	void	(no parameters)

MBean Configuration (org.neo4j.jmx.impl.ConfigurationBean) Attributes

Name	Description	Type	Read	Write
<i>The configuration parameters used to configure Neo4j</i>				
all_stores_total_mapped_memory_size	The size to allocate for a memory mapping pool to be shared between all stores.	String	yes	no
allow_store_upgrade	Whether to allow a store upgrade in case the current version of the database starts against an older store version. Setting this to true does not guarantee successful upgrade, just that it allows an attempt at it.	String	yes	no
array_block_size	Specifies the block size for storing arrays. This parameter is only honored when the store is created, otherwise it is ignored. The default block size is 120 bytes, and the overhead of each block is the same as for string blocks, i.e., 8 bytes.	String	yes	no
backup_slave	Mark this database as a backup slave.	String	yes	no
cache_type	The type of cache to use for nodes and relationships. Note that the Neo4j Enterprise Edition has the additional <i>gcr</i> cache type. See the chapter on caches in the manual for more information.	String	yes	no
cypher_parser_version	Enable this to specify a parser other than the default one.	String	yes	no
dump_configuration	Print out the effective Neo4j configuration after startup.	String	yes	no

Name	Description	Type	Read	Write
ephemeral	Configuration attribute	String	yes	no
forced_kernel_id	An identifier that uniquely identifies this graph database instance within this JVM. Defaults to an auto-generated number depending on how many instance are started in this JVM.	String	yes	no
gc_monitor_threshold	The amount of time in ms the monitor thread has to be blocked before logging a message it was blocked.	String	yes	no
gc_monitor_wait_time	Amount of time in ms the GC monitor thread will wait before taking another measurement.	String	yes	no
gcr_cache_min_log_interval	The minimal time that must pass in between logging statistics from the cache (when using the <i>gcr</i> cache).	String	yes	no
grab_file_lock	Whether to grab locks on files or not.	String	yes	no
ha.allow_init_cluster	Whether to allow this instance to create a cluster if unable to join.	String	yes	no
ha.branched_data_policy	Policy for how to handle branched data.	String	yes	no
ha.broadcast_timeout	Timeout for broadcasting values in cluster. Must consider end-to-end duration of Paxos algorithm. This value is the default value for settings <i>ha.join_timeout</i> and <i>ha.leave_timeout</i> .	String	yes	no
ha.cluster_name	The name of a cluster.	String	yes	no
ha.cluster_server	Host & port to bind the cluster management communication.	String	yes	no
ha.com_chunk_size	Max size of the data chunks that flows between master and slaves in HA. Bigger size may increase throughput, but may be more sensitive to variations in bandwidth, whereas lower size increases tolerance for bandwidth variations.	String	yes	no
ha.configuration_timeout	Timeout for waiting for configuration from an existing cluster member during cluster join.	String	yes	no
ha.default_timeout	Default timeout used for clustering timeouts. Override specific timeout settings with proper values if necessary. This value is the default value	String	yes	no

Name	Description	Type	Read	Write
	for settings <code>ha.heartbeat_interval</code> , <code>ha.paxos_timeout</code> and <code>ha.learn_timeout</code> .			
<code>ha.election_timeout</code>	Timeout for waiting for other members to finish a role election. Defaults to <code>ha.paxos_timeout</code> .	String	yes	no
<code>ha.heartbeat_interval</code>	How often heartbeat messages should be sent. Defaults to <code>ha.default_timeout</code> .	String	yes	no
<code>ha.heartbeat_timeout</code>	Timeout for heartbeats between cluster members. Should be at least twice that of <code>ha.heartbeat_interval</code> .	String	yes	no
<code>ha.initial_hosts</code>	A comma-separated list of other members of the cluster to join.	String	yes	no
<code>ha.join_timeout</code>	Timeout for joining a cluster. Defaults to <code>ha.broadcast_timeout</code> .	String	yes	no
<code>ha.learn_timeout</code>	Timeout for learning values. Defaults to <code>ha.default_timeout</code> .	String	yes	no
<code>ha.leave_timeout</code>	Timeout for waiting for cluster leave to finish. Defaults to <code>ha.broadcast_timeout</code> .	String	yes	no
<code>ha.lock_read_timeout</code>	Timeout for taking remote (write) locks on slaves. Defaults to <code>ha.read_timeout</code> .	String	yes	no
<code>ha.max_concurrent_channels_per_slave</code>	Maximum number of connections a slave can have to the master.	String	yes	no
<code>ha.paxos_timeout</code>	Default timeout for all Paxos timeouts. Defaults to <code>ha.default_timeout</code> . This value is the default value for settings <code>ha.phase1_timeout</code> , <code>ha.phase2_timeout</code> and <code>ha.election_timeout</code> .	String	yes	no
<code>ha.phase1_timeout</code>	Timeout for Paxos phase 1. Defaults to <code>ha.paxos_timeout</code> .	String	yes	no
<code>ha.phase2_timeout</code>	Timeout for Paxos phase 2. Defaults to <code>ha.paxos_timeout</code> .	String	yes	no
<code>ha.pull_interval</code>	Interval of pulling updates from master.	String	yes	no
<code>ha.read_timeout</code>	How long a slave will wait for response from master before giving up.	String	yes	no
<code>ha.server_id</code>	Id for a cluster instance. Must be unique within the cluster.	String	yes	no
<code>ha.server</code>	Hostname and port to bind the HA server.	String	yes	no
<code>ha.slave_only</code>	Whether this instance should only participate as slave in cluster. If set to true, it will never be elected as master.	String	yes	no

Name	Description	Type	Read	Write
ha.state_switch_timeout	Timeout for waiting for instance to become master or slave.	String	yes	no
ha.tx_push_factor	The amount of slaves the master will ask to replicate a committed transaction.	String	yes	no
ha.tx_push_strategy	Push strategy of a transaction to a slave during commit.	String	yes	no
ha.zk_session_timeout	ZooKeeper session timeout. Only needed for rolling upgrade from 1.8 to 1.9.	String	yes	no
intercept_committing_transactions	Determines whether any TransactionInterceptors loaded will intercept prepared transactions before they reach the logical log.	String	yes	no
intercept_deserialized_transactions	Determines whether any TransactionInterceptors loaded will intercept externally received transactions (e.g. in HA) before they reach the logical log and are applied to the store.	String	yes	no
jmx.port	Configuration attribute	String	yes	no
keep_logical_logs	Make Neo4j keep the logical transaction logs for being able to backup the database. Can be used for specifying the threshold to prune logical logs after. For example "10 days" will prune logical logs that only contains transactions older than 10 days from the current time, or "100k txs" will keep the 100k latest transactions and prune any older transactions.	String	yes	no
load_kernel_extensions	Enable loading kernel extensions	String	yes	no
log_mapped_memory_stats_filename	The file where Neo4j will record memory mapping statistics.	String	yes	no
log_mapped_memory_stats_interval	The number of records to be loaded between regular logging of memory mapping statistics.	String	yes	no
log_mapped_memory_stats	Tell Neo4j to regularly log memory mapping statistics.	String	yes	no
logging.threshold_for_rotation	Threshold in bytes for when database logs (text logs, for debugging, that is) are rotated.	String	yes	no

Name	Description	Type	Read	Write
logical_log_rotation_threshold	Specifies at which file size the logical log will auto-rotate. 0 means that no rotation will automatically occur based on file size. Default is 25M	String	yes	no
logical_log	The base name for the logical log files, either an absolute path or relative to the store_dir setting. This should generally not be changed.	String	yes	no
mapped_memory_page_size	Target size for pages of mapped memory.	String	yes	no
neo_store	The base name for the Neo4j Store files, either an absolute path or relative to the store_dir setting. This should generally not be changed.	String	yes	no
neostore.nodestore.db.mapped_memory	The size to allocate for memory mapping the node store.	String	yes	no
neostore.propertystore.db.arrays.mapped_memory	The size to allocate for memory mapping the array property store.	String	yes	no
neostore.propertystore.db.index.keys.mapped_memory	The size to allocate for memory mapping the store for property key strings.	String	yes	no
neostore.propertystore.db.index.mapped_memory	The size to allocate for memory mapping the store for property key indexes.	String	yes	no
neostore.propertystore.db.mapped_memory	The size to allocate for memory mapping the property value store.	String	yes	no
neostore.propertystore.db.strings.mapped_memory	The size to allocate for memory mapping the string property store.	String	yes	no
neostore.relationshipstore.db.mapped_memory	The size to allocate for memory mapping the relationship store.	String	yes	no
node_auto_indexing	Controls the auto indexing feature for nodes. Setting to false shuts it down, while true enables it by default for properties listed in the node_keys_indexable setting.	String	yes	no

Name	Description	Type	Read	Write
node_cache_array_fraction	The fraction of the heap (1%-10%) to use for the base array in the node cache (when using the <i>gcr</i> cache).	String	yes	no
node_cache_size	The amount of memory to use for the node cache (when using the <i>gcr</i> cache).	String	yes	no
node_keys_indexable	A list of property names (comma separated) that will be indexed by default. This applies to Nodes only.	String	yes	no
online_backup_enabled	Enable support for running online backups	String	yes	no
online_backup_port	Listening port for online backups. Replaced by <code>online_backup_server</code>	String	yes	no
online_backup_server	Listening server for online backups	String	yes	no
read_only	Only allow read operations from this Neo4j instance. This mode still requires write access to the directory for lock purposes	String	yes	no
rebuild_idgenerators_fast	Use a quick approach for rebuilding the ID generators. This give quicker recovery time, but will limit the ability to reuse the space of deleted entities.	String	yes	no
relationship_auto_indexing	Controls the auto indexing feature for relationships. Setting to false shuts it down, while true enables it by default for properties listed in the <code>relationship_keys_indexable</code> setting.	String	yes	no
relationship_cache_array_fraction	The fraction of the heap (1%-10%) to use for the base array in the relationship cache (when using the <i>gcr</i> cache).	String	yes	no
relationship_cache_size	The amount of memory to use for the relationship cache (when using the <i>gcr</i> cache).	String	yes	no
relationship_grab_size	How many relationships to read at a time during iteration	String	yes	no
relationship_keys_indexable	A list of property names (comma separated) that will be indexed by default. This applies to Relationships only.	String	yes	no
remote_logging_enabled	Whether to enable logging to a remote server or not.	String	yes	no
remote_logging_host	Host for remote logging using LogBack SocketAppender.	String	yes	no

Name	Description	Type	Read	Write
remote_logging_port	Port for remote logging using LogBack SocketAppender.	String	yes	no
remote_shell_enabled	Enable a remote shell server which shell clients can log in to	String	yes	no
store_dir	The directory where the database files are located.	String	yes	no
string_block_size	Specifies the block size for storing strings. This parameter is only honored when the store is created, otherwise it is ignored. Note that each character in a string occupies two bytes, meaning that a block size of 120 (the default size) will hold a 60 character long string before overflowing into a second block. Also note that each block carries an overhead of 8 bytes. This means that if the block size is 120, the size of the stored records will be 128 bytes.	String	yes	no
tx_manager_impl	The name of the Transaction Manager service to use as defined in the TM service provider constructor.	String	yes	no
use_memory_mapped_buffers	Tell Neo4j to use memory mapped buffers for accessing the native storage layer.	String	yes	no

MBean Configuration (org.neo4j.jmx.impl.ConfigurationBean) Operations

Name	Description	ReturnType	Signature
apply	Apply settings	void	(no parameters)

MBean Diagnostics (org.neo4j.management.Diagnostics) Attributes

Name	Description	Type	Read	Write
<i>Diagnostics provided by Neo4j</i>				
DiagnosticsProviders	A list of the ids for the registered diagnostics providers.	List (java.util.List)	yes	no

MBean Diagnostics (org.neo4j.management.Diagnostics) Operations

Name	Description	ReturnType	Signature
dumpAll	Dump diagnostics information to JMX	String	(no parameters)
dumpToLog	Dump diagnostics information to the log.	void	(no parameters)
dumpToLog	Dump diagnostics information to the log.	void	java.lang.String
extract	Operation exposed for management	String	java.lang.String

MBean High Availability (org.neo4j.management.HighAvailability) Attributes

Name	Description	Type	Read	Write
<i>Information about an instance participating in a HA cluster</i>				
Alive	Whether this instance is alive or not	boolean	yes	no
Available	Whether this instance is available or not	boolean	yes	no
InstanceId	The identifier used to identify this server in the HA cluster	String	yes	no
InstancesInCluster	Information about all instances in this cluster	org.neo4j.management.ClusterMemberInfo[] as CompositeData[]	yes	no
LastCommittedTxId	The latest transaction id present in this instance's store	long	yes	no
LastUpdateTime	The time when the data on this instance was last updated from the master	String	yes	no
Role	The role this instance has in the cluster	String	yes	no

MBean High Availability (org.neo4j.management.HighAvailability) Operations

Name	Description	ReturnType	Signature
update	(If this is a slave) Update the database on this instance with the latest transactions from the master	String	(no parameters)

MBean Kernel (org.neo4j.jmx.Kernel) Attributes

Name	Description	Type	Read	Write
<i>Information about the Neo4j kernel</i>				
KernelStartTime	The time from which this Neo4j instance was in operational mode.	Date (java.util.Date)	yes	no
KernelVersion	The version of Neo4j	String	yes	no
MBeanQuery	An ObjectName that can be used as a query for getting all management beans for this Neo4j instance.	javax.management.ObjectName	yes	no
ReadOnly	Whether this is a read only instance	boolean	yes	no
StoreCreationDate	The time when this Neo4j graph store was created.	Date (java.util.Date)	yes	no
StoreDirectory	The location where the Neo4j store is located	String	yes	no
StoreId	An identifier that, together with store creation time, uniquely identifies this Neo4j graph store.	String	yes	no
StoreLogVersion	The current version of the Neo4j store logical log.	long	yes	no

MBean Locking (org.neo4j.management.LockManager) Attributes

Name	Description	Type	Read	Write
<i>Information about the Neo4j lock status</i>				
Locks	Information about all locks held by Neo4j	java.util.List<org.neo4j.kernel.info.LockInfo> as CompositeData[]	yes	no
NumberOfAvertedDeadlocks	The number of lock sequences that would have lead to a deadlock situation that Neo4j has detected and averted (by throwing DeadlockDetectedException).	long	yes	no

MBean Locking (org.neo4j.management.LockManager) Operations

Name	Description	ReturnType	Signature
getContendedLocks	getContendedLocks	java.util.List<org.neo4j.kernel.info.LockInfo> as CompositeData[]	long

MBean Memory Mapping (org.neo4j.management.MemoryMapping) Attributes

Name	Description	Type	Read	Write
<i>The status of Neo4j memory mapping</i>				
MemoryPools	Get information about each pool of memory mapped regions from store files with memory mapping enabled	org.neo4j.management.WindowPoolInfo[] as CompositeData[]	yes	no

MBean Primitive count (org.neo4j.jmx.Primitives) Attributes

Name	Description	Type	Read	Write
<i>Estimates of the numbers of different kinds of Neo4j primitives</i>				
NumberOfNodeIds InUse	An estimation of the number of nodes used in this Neo4j instance	long	yes	no
NumberOfPropertyIds InUse	An estimation of the number of properties used in this Neo4j instance	long	yes	no
NumberOfRelationshipIds InUse	An estimation of the number of relationships used in this Neo4j instance	long	yes	no
NumberOfRelationshipTypeIds InUse	The number of relationship types used in this Neo4j instance	long	yes	no

MBean Store file sizes (org.neo4j.management.StoreFile) Attributes

Name	Description	Type	Read	Write
<i>Information about the sizes of the different parts of the Neo4j graph store</i>				

Name	Description	Type	Read	Write
ArrayStoreSize	The amount of disk space used to store array properties, in bytes.	long	yes	no
LogicalLogSize	The amount of disk space used by the current Neo4j logical log, in bytes.	long	yes	no
NodeStoreSize	The amount of disk space used to store nodes, in bytes.	long	yes	no
PropertyStoreSize	The amount of disk space used to store properties (excluding string values and array values), in bytes.	long	yes	no
RelationshipStoreSize	The amount of disk space used to store relationships, in bytes.	long	yes	no
StringStoreSize	The amount of disk space used to store string properties, in bytes.	long	yes	no
TotalStoreSize	The total disk space used by this Neo4j instance, in bytes.	long	yes	no

MBean Transactions (org.neo4j.management.TransactionManager) Attributes

Name	Description	Type	Read	Write
<i>Information about the Neo4j transaction manager</i>				
LastCommittedTxId	The id of the latest committed transaction	long	yes	no
NumberOfCommittedTransactions	The total number of committed transactions on this instance since the database started	long	yes	no
NumberOfOpenedTransactions	The total number started transactions on this instance since the database started	int	yes	no
NumberOfOpenTransactions	The number of currently open transactions	int	yes	no
NumberOfRolledBackTransactions	The total number of rolled back transactions on this instance since the database started	long	yes	no
PeakNumberOfConcurrentTransactions	The highest number of transactions ever opened concurrently on this instance since the database started	int	yes	no

MBean XA Resources (org.neo4j.management.XaManager) Attributes

Name	Description	Type	Read	Write
<i>Information about the XA transaction manager</i>				
XaResources	Information about all XA resources managed by the transaction manager	org.neo4j.management.XaResourceInfo[] as CompositeData[]	yes	no

Part VI. Tools

The Tools part describes available Neo4j tools and how to use them.

Chapter 30. Web Interface

The Neo4j Web Interface is the primary user interface for Neo4j. With it, you can:

- monitor the Neo4j Server
- manipulate and browse data
- interact with the database via various consoles
- view raw data management objects (JMX MBeans)

The tool is available at <http://127.0.0.1:7474/> after you have installed the [Neo4j Server](#).

See the tool itself for more information!

Chapter 31. Neo4j Shell

Neo4j shell is a command-line shell for browsing the graph, much like how the Unix shell along with commands like `cd`, `ls` and `pwd` can be used to browse your local file system. It consists of two parts:

- a lightweight client that sends commands via RMI and
- a server that processes those commands and sends the result back to the client.

It's a nice tool for development and debugging. This guide will show you how to get it going!

31.1. Starting the shell

When used together with Neo4j started as a server, simply issue the following at the command line:

```
./bin/neo4j-shell
```

For the full list of options, see the reference in the [Shell manual page](#).

To connect to a running Neo4j database, use [Section 31.1.4, “Read-only mode”](#) for local databases and see [Section 31.1.1, “Enabling the shell server”](#) for remote databases.

You need to make sure that the shell jar file is on the classpath when you start up your Neo4j instance.

31.1.1. Enabling the shell server

Shell is enabled from the configuration of the Neo4j kernel, see [Section 21.2, “Server Configuration”](#). Here’s some sample configurations:

```
# Using default values
enable_remote_shell = true
# ...or specify custom port, use default values for the others
enable_remote_shell = port=1234
```

When using the Neo4j server, see [Section 21.2, “Server Configuration”](#) for how to add configuration settings in that case.

There are two ways to start the shell, either by connecting to a remote shell server or by pointing it to a Neo4j store path.

31.1.2. Connecting to a shell server

To start the shell and connect to a running server, run:

```
neo4j-shell
```

Alternatively supply `-port` and `-name` options depending on how the remote shell server was enabled. Then you’ll get the shell prompt like this:

```
neo4j-sh (0)$
```

31.1.3. Pointing the shell to a path

To start the shell by just pointing it to a Neo4j store path you run the shell jar file. Given that the right `neo4j-kernel-<version>.jar` and `jta.jar` files are in the same path as your `neo4j-shell-<version>.jar` file you run it with:

```
$ neo4j-shell -path path/to/neo4j-db
```

31.1.4. Read-only mode

By issuing the `-readonly` switch when starting the shell with a store path, changes cannot be made to the database during the session.

```
$ neo4j-shell -readonly -path path/to/neo4j-db
```

31.1.5. Run a command and then exit

It is possible to tell the shell to just start, execute a command and then exit. This opens up for uses of background jobs and also handling of huge output of f.ex. an `ls` command where you then could pipe the output to `less` or another reader of your choice, or even to a file. So some examples of usage:

```
$ neo4j-shell -c "cd -a 24 && set name Mattias"
```

```
$ neo4j-shell -c "trav -r KNOWS" | less
```

31.2. Passing options and arguments

Passing options and arguments to your commands is very similar to many CLI commands in an *nix environment. Options are prefixed with a - and can contain one or more options. Some options expect a value to be associated with it. Arguments are string values which aren't prefixed with -. Let's look at `ls` as an example:

`ls -r -f KNOWS:out -v 12345` will make a verbose listing of node 12345's outgoing relationships of type `KNOWS`. The node id, 12345, is an argument to `ls` which tells it to do the listing on that node instead of the current node (see `pwd` command). However a shorter version of this can be written:

`ls -rfv KNOWS:out 12345`. Here all three options are written together after a single - prefix. Even though `f` is in the middle it gets associated with the `KNOWS:out` value. The reason for this is that the `ls` command doesn't expect any values associated with the `r` or `v` options. So, it can infer the right values for the rights options.

31.3. Enum options

Some options expects a value which is one of the values in an enum, f.ex. direction part of relationship type filtering where there's INCOMING, OUTGOING and BOTH. All such values can be supplied in an easier way. It's enough that you write the start of the value and the interpreter will find what you really meant. F.ex. out, in, i or even INCOMING.

31.4. Filters

Some commands makes use of filters for varying purposes. F.ex. `-f` in `ls` and in `trav`. A filter is supplied as a [json](http://www.json.org/) `<http://www.json.org/>` object (w/ or w/o the surrounding `{}` brackets). Both keys and values can contain regular expressions for a more flexible matching. An example of a filter could be `.*url.*:http.*neo4j.*,name:Neo4j`. The filter option is also accompanied by the options `-i` and `-l` which stands for `ignore case` (ignore casing of the characters) and `loose matching` (it's considered a match even if the filter value just matches a part of the compared value, not necessarily the entire value). So for a case-insensitive, loose filter you can supply a filter with `-f -i -l` or `-fil` for short.

31.5. Node titles

To make it easier to navigate your graph the shell can display a title for each node, f.ex. in `ls -r`. It will display the relationships as well as the nodes on the other side of the relationships. The title is displayed together with each node and its best suited property value from a list of property keys.

If you're standing on a node which has two `KNOWS` relationships to other nodes it'd be difficult to know which friend is which. The title feature addresses this by reading a list of property keys and grabbing the first existing property value of those keys and displays it as a title for the node. So you may specify a list (with or without regular expressions), f.ex: `name,title.*,caption` and the title for each node will be the property value of the first existing key in that list. The list is defined by the client (you) using the `TITLE_KEYS` [environment variable](#) and the default being `.*name.*,.*title.*`

31.6. How to use (individual commands)

The shell is modeled after Unix shells like bash that you use to walk around your local file system. It has some of the same commands, like `cd` and `ls`. When you first start the shell (see instructions above), you will get a list of all the available commands. Use `man <command>` to get more info about a particular command. Some notes:

31.6.1. Comments

Single line comments, which will be ignored, can be made by using the prefix `///`. Example:

```
// This is a comment
```

31.6.2. Current node/relationship and path

You have a current node/relationship and a "current path" (like a current working directory in bash) that you've traversed so far. You start at the [reference node](http://api.neo4j.org/current/org/neo4j/graphdb/GraphDatabaseService.html#getReferenceNode()>) `<http://api.neo4j.org/current/org/neo4j/graphdb/GraphDatabaseService.html#getReferenceNode()>` and can then `cd` your way through the graph (check your current path at any time with the `pwd` command). `cd` can be used in different ways:

- `cd <node-id>` will traverse one relationship to the supplied node id. The node must have a direct relationship to the current node.
- `cd -a <node-id>` will do an absolute path change, which means the supplied node doesn't have to have a direct relationship to the current node.
- `cd -r <relationship-id>` will traverse to a relationship instead of a node. The relationship must have the current node as either start or end point. To see the relationship ids use the `ls -vr` command on nodes.
- `cd -ar <relationship-id>` will do an absolute path change which means the relationship can be any relationship in the graph.
- `cd` will take you back to the reference node, where you started in the first place.
- `cd ..` will traverse back one step to the previous location, removing the last path item from your current path (`pwd`).
- `cd start` (*only if your current location is a relationship*). Traverses to the start node of the relationship.
- `cd end` (*only if your current location is a relationship*). Traverses to the end node of the relationship.

31.6.3. Listing the contents of a node/relationship

List contents of the current node/relationship (or any other node) with the `ls` command. Please note that it will give an empty output if the current node/relationship has no properties or relationships (for example in the case of a brand new graph). `ls` can take a node id as argument as well as filters, see [Section 31.4, "Filters"](#) and for information about how to specify direction see [Section 31.3, "Enum options"](#). Use `man ls` for more info.

31.6.4. Creating nodes and relationships

You create new nodes by connecting them with relationships to the current node. For example, `mkrel -t A_RELATIONSHIP_TYPE -d OUTGOING -c` will create a new node (`-c`) and draw to it an `OUTGOING` relationship of type `A_RELATIONSHIP_TYPE` from the current node. If you already have two nodes which you'd like to draw a relationship between (without creating a new node) you can do for example, `mkrel -t A_RELATIONSHIP_TYPE -d OUTGOING -n <other-node-id>` and it will just create a new relationship between the current node and that other node.

31.6.5. Setting, renaming and removing properties

Property operations are done with the `set`, `mv` and `rm` commands. These commands operates on the current node/relationship. * `set <key> <value>` with optionally the `-t` option (for value type) sets a property. Supports every type of value that Neo4j supports. Examples of a property of type `int`:

```
$ set -t int age 29
```

And an example of setting a `double[]` property:

```
$ set -t double[] my_values [1.4,12.2,13]
```

Example of setting a `String` property containing a JSON string:

```
mkrel -c -d i -t DOMAIN_OF --np '{"app': 'foobar'}"
```

- `rm <key>` removes a property.
- `mv <key> <new-key>` renames a property from one key to another.

31.6.6. Deleting nodes and relationships

Deletion of nodes and relationships is done with the `rmnode` and `rmrel` commands. `rmnode` can delete nodes, if the node to be deleted still has relationships they can also be deleted by supplying `-f` option. `rmrel` can delete relationships, it tries to ensure connectedness in the graph, but relationships can be deleted regardless with the `-f` option. `rmrel` can also delete the node on the other side of the deleted relationship if it's left with no more relationships, see `-d` option.

31.6.7. Environment variables

The shell uses environment variables a-la bash to keep session information, such as the current path and more. The commands for this mimics the bash commands `export` and `env`. For example you can at anytime issue a `export STACKTRACES=true` command to set the `STACKTRACES` environment variable to `true`. This will then result in `stacktraces` being printed if an exception or error should occur. List environment variables using `env`

31.6.8. Executing groovy/python scripts

The shell has support for executing scripts, such as [Groovy](http://groovy.codehaus.org) <http://groovy.codehaus.org> and [Python](http://www.python.org) <http://www.python.org> (via [Jython](http://www.jython.org) <http://www.jython.org>). As of now the scripts (*.groovy, *.py) must exist on the server side and gets called from a client with for example, `gsh --renamePerson 1234 "Mathias" "Mattias" --doSomethingElse` where the scripts `renamePerson.groovy` and `doSomethingElse.groovy` must exist on the server side in any of the paths given by the `GSH_PATH` environment variable (defaults to `.:src:src/script`). This variable is like the java classpath, separated by a `..`. The python/jython scripts can be executed with the `jsh` in a similar fashion, however the scripts have the `.py` extension and the environment variable for the paths is `JSH_PATH`.

When writing the scripts assume that there's made available an `args` variable (a `String[]`) which contains the supplied arguments. In the case of the `renamePerson` example above the array would contain `["1234", "Mathias", "Mattias"]`. Also please write your outputs to the `out` variable, such as `out.println("My tracing text")` so that it will be printed at the shell client instead of the server.

31.6.9. Traverse

You can traverse the graph with the `trav` command which allows for simple traversing from the current node. You can supply which relationship types (w/ regex matching) and optionally direction as well as property filters for matching nodes. In addition to that you can supply a command line to execute for each match. An example: `trav -o depth -r KNOWS:both,HAS_.*:incoming -c "ls $n"`.

Which means traverse depth first for relationships with type `KNOWS` disregarding direction and incoming relationships with type matching `HAS_.*` and do a `ls <matching node>` for each match. The node filtering is supplied with the `-f` option, see [Section 31.4, “Filters”](#). See [Section 31.3, “Enum options”](#) for the traversal order option. Even relationship types/directions are supplied using the same format as filters.

31.6.10. Query with Cypher

You can use Cypher to query the graph. For that, use the `start` command.



Tip

Cypher queries need to be terminated by a semicolon `;`.

- `start n = node(0) return n;` will give you a listing of the node with ID 0
- `cypher 1.8 start n = node(0) return n;` will execute the query with the Cypher parser version 1.8

31.6.11. Indexing

It's possible to query and manipulate indexes via the `index` command. Example: `index -i persons name` (will index the name for the current node or relationship in the "persons" index).

- `-g` will do exact lookup in the index and display hits. You can supply `-c` with a command to be executed for each hit.
- `-q` will ask the index a query and display hits. You can supply `-c` with a command to be executed for each hit.
- `--cd` will change current location to the hit from the query. It's just a convenience for using the `-c` option.
- `--ls` will do a listing of the contents for each hit. It's just a convenience for using the `-c` option.
- `-i` will index a key-value pair in an index for the current node/relationship. If no value is given the property value for that key for the current node is used as value.
- `-r` will remove a key-value pair (if it exists) from an index for the current node/relationship. Key and value is optional.
- `-t` will set the index type to work with, for example `index -t Relationship --delete friends` will delete the `friends` relationship index.

31.6.12. Transactions

It is useful to be able to test changes, and then being able to commit or rollback said changes.

Transactions can be nested. With a nested transaction, a commit does not write any changes to disk, except for the top level transaction. A rollback, however works regardless of the level of the transaction. It will roll back all open transactions.

- `begin` transaction Starts a transaction.
- `commit` Commits a transaction.
- `rollback` Rollbacks all open transactions.

31.7. Extending the shell: Adding your own commands

Of course the shell is extendable and has a generic core which has nothing to do with Neo4j... only some of the [commands](http://components.neo4j.org/neo4j-shell/1.9.7/apidocs/org/neo4j/shell/App.html) <<http://components.neo4j.org/neo4j-shell/1.9.7/apidocs/org/neo4j/shell/App.html>> do.

So you say you'd like to start a Neo4j [graph database](http://api.neo4j.org/current/org/neo4j/graphdb/GraphDatabaseService.html) <<http://api.neo4j.org/current/org/neo4j/graphdb/GraphDatabaseService.html>>, [enable the remote shell](#) and add your own apps to it so that your apps and the standard Neo4j apps co-exist side by side? Well, here's an example of how an app could look like:

```
import org.neo4j.helpers.Service;
import org.neo4j.shell.kernel.apps.GraphDatabaseApp;

@Service.Implementation( App.class )
public class LsRelTypes extends GraphDatabaseApp
{
    @Override
    protected String exec( AppCommandParser parser, Session session, Output out )
        throws ShellException, RemoteException
    {
        GraphDatabaseService graphDb = getServer().getDb();
        out.println( "Types:" );
        for ( RelationshipType type : graphDb.getRelationshipTypes() )
        {
            out.println( type.name() );
        }
        return null;
    }
}
```

And you could now use it in the shell by typing `lsreltypes` (its name is based on the class name) if `getName` method isn't overridden.

If you'd like it to display some nice help information when using the `help` (or `man`) app, override the `getDescription` method for a general description and use `addValueType` method to add descriptions about (and logic to) the options you can supply when using your app.

Know that the apps reside server-side so if you have a running server and starts a remote client to it from another JVM you can't add your apps on the client.

31.8. An example shell session

```

# where are we?
neo4j-sh (0)$ pwd
Current is (0)
(0)

# On the current node, set the key "name" to value "Jon"
neo4j-sh (0)$ set name "Jon"

# send a cypher query
neo4j-sh (Jon,0)$ start n=node(0) return n;
+-----+
| n      |
+-----+
| Node[0]{name:"Jon"} |
+-----+
1 row

24 ms

# make an incoming relationship of type LIKES, create the end node with the node properties specified.
neo4j-sh (Jon,0)$ mkrel -c -d i -t LIKES --np '{"app':'foobar'}"

# where are we?
neo4j-sh (Jon,0)$ ls
*name =[Jon]
(me)<-[:LIKES]-(1)

# change to the newly created node
neo4j-sh (Jon,0)$ cd 1

# list relationships, including relationship id
neo4j-sh (1)$ ls -avr
(me)-[:LIKES,0]->(Jon,0)

# create one more KNOWS relationship and the end node
neo4j-sh (1)$ mkrel -c -d i -t KNOWS --np '{"name':'Bob'}"

# print current history stack
neo4j-sh (1)$ pwd
Current is (1)
(Jon,0)-->(1)


# verbose list relationships
neo4j-sh (1)$ ls -avr
(me)-[:LIKES,0]->(Jon,0)
(me)<-[:KNOWS,1]-(Bob,2)

```

31.9. A Matrix example

This example is creating a graph of the characters in the Matrix via the shell and then executing Cypher queries against it:

Figure 31.1. Shell Matrix Example



Neo4j is configured for autoindexing, in this case with the following in the Neo4j configuration file:

```

node_auto_indexing=true
node_keys_indexable=name,age

relationship_auto_indexing=true
relationship_keys_indexable=ROOT,KNOWS,CODED_BY

```

The following is a sample shell session creating the Matrix graph and querying it.

```

# create the Thomas Andersson node
neo4j-sh (0)$ mkrel -t ROOT -c -v
Node (1) created

```

```
Relationship [:ROOT,0] created

# go to the new node
neo4j-sh (0)$ cd 1

# set the name property
neo4j-sh (1)$ set name "Thomas Andersson"

# create Thomas direct friends
neo4j-sh (Thomas Andersson,1)$ mkrel -t KNOWS -cv
Node (2) created
Relationship [:KNOWS,1] created

# go to the new node
neo4j-sh (Thomas Andersson,1)$ cd 2

# set the name property
neo4j-sh (2)$ set name "Trinity"

# go back in the history stack
neo4j-sh (Trinity,2)$ cd ..

# create Thomas direct friends
neo4j-sh (Thomas Andersson,1)$ mkrel -t KNOWS -cv
Node (3) created
Relationship [:KNOWS,2] created

# go to the new node
neo4j-sh (Thomas Andersson,1)$ cd 3

# set the name property
neo4j-sh (3)$ set name "Morpheus"

# create relationship to Trinity
neo4j-sh (Morpheus,3)$ mkrel -t KNOWS 2

# list the relationships of node 3
neo4j-sh (Morpheus,3)$ ls -rv
(me)-[:KNOWS,3]->(Trinity,2)
(me)<-[:KNOWS,2]-(Thomas Andersson,1)

# change the current position to relationship #2
neo4j-sh (Morpheus,3)$ cd -r 2

# set the age property on the relationship
neo4j-sh [:KNOWS,2]$ set -t int age 3

# back to Morpheus
neo4j-sh [:KNOWS,2]$ cd ..

# next relationship
neo4j-sh (Morpheus,3)$ cd -r 3

# set the age property on the relationship
neo4j-sh [:KNOWS,3]$ set -t int age 90

# position to the start node of the current relationship
neo4j-sh [:KNOWS,3]$ cd start

# new node
neo4j-sh (Morpheus,3)$ mkrel -t KNOWS -c
```

```

# list relationships on the current node
neo4j-sh (Morpheus,3)$ ls -r
(me)-[:KNOWS]->(Trinity,2)
(me)-[:KNOWS]->(4)
(me)-[:KNOWS]-(Thomas Andersson,1)

# go to Cypher
neo4j-sh (Morpheus,3)$ cd 4

# set the name
neo4j-sh (4)$ set name Cypher

# create new node from Cypher
neo4j-sh (Cypher,4)$ mkrel -ct KNOWS

# list relationships
neo4j-sh (Cypher,4)$ ls -r
(me)-[:KNOWS]->(5)
(me)-[:KNOWS]-(Morpheus,3)

# go to the Agent Smith node
neo4j-sh (Cypher,4)$ cd 5

# set the name
neo4j-sh (5)$ set name "Agent Smith"

# outgoing relationship and new node
neo4j-sh (Agent Smith,5)$ mkrel -cvt CODED_BY
Node (6) created
Relationship [[:CODED_BY,6] created

# go there
neo4j-sh (Agent Smith,5)$ cd 6

# set the name
neo4j-sh (6)$ set name "The Architect"

# go to the first node in the history stack
neo4j-sh (The Architect,6)$ cd

# Morpheus' friends, looking up Morpheus by name in the Neo4j autoindex
neo4j-sh (0)$ start morpheus = node:node_auto_index(name='Morpheus') match morpheus-[:KNOWS]-zionist return zionist.name;
+-----+
| zionist.name      |
+-----+
| "Trinity"        |
| "Cypher"         |
| "Thomas Andersson" |
+-----+
3 rows

121 ms

# Morpheus' friends, looking up Morpheus by name in the Neo4j autoindex
neo4j-sh (0)$ cypher 1.9 start morpheus = node:node_auto_index(name='Morpheus') match morpheus-[:KNOWS]-zionist return zionist.name;
+-----+
| zionist.name      |
+-----+
| "Trinity"        |
| "Cypher"         |

```

```
| "Thomas Andersson" |  
+-----+  
3 rows  
39 ms
```

Part VII. Community

The Neo4j project has a strong community around it. Read about how to get help from the community and how to contribute to it.

Chapter 32. Community Support

You can learn a lot about Neo4j on different *events*. To get information on upcoming Neo4j events, have a look here:

- <http://www.neo4j.org/>
- <http://neo4j.meetup.com/>

Get help from the Neo4j open source community; here are some starting points.

- The neo4j tag at stackoverflow: <http://stackoverflow.com/questions/tagged/neo4j>
- Neo4j Community Discussions: <https://groups.google.com/forum/#!forum/neo4j>
- Twitter: <https://twitter.com/neo4j>
- IRC channel: <irc://irc.freenode.net/neo4j> web chat <<http://webchat.freenode.net/?randomnick=1&channels=neo4j>>.

Report a *bug* or add a *feature request*:

- <https://github.com/neo4j/neo4j/issues>

Questions regarding the *documentation*: The Neo4j Manual is published online with a comment function, please use that to post any questions or comments regarding the documentation. See <http://docs.neo4j.org/>.

Chapter 33. Contributing to Neo4j

The Neo4j project is an Open Source effort to bring fast complex data storage and processing to life. Every form of help is highly appreciated by the community - and you are not alone, see [Section 33.6, “Contributors”](#)!

One crucial aspect of contributing to the Neo4j project is the [Section 33.1, “Contributor License Agreement”](#).

In short: make sure to sign the CLA and send in the email, or the Neo4j project won't be able to accept your contribution.

Note that you can contribute to Neo4j also by contributing documentation or giving feedback on the current documentation. Basically, at all the places where you can get help, there's also room for contributions.

If you want to contribute, there are some good areas to start with, especially for getting in contact with the community, [Chapter 32, *Community Support*](#).

To document your efforts, we highly recommend to read [Section 33.3, “Writing Neo4j Documentation”](#).

33.1. Contributor License Agreement

33.1.1. Summary

We require all source code that is hosted on the Neo4j infrastructure to be contributed through the [Neo4j Contributor License Agreement](http://dist.neo4j.org/neo4j-cla.pdf) (CLA). The purpose of the Neo4j Contributor License Agreement is to protect the integrity of the code base, which in turn protects the community around that code base: the founding entity Neo Technology, the Neo4j developer community and the Neo4j users. This kind of contributor agreement is common amongst free software and open source projects (it is in fact very similar to the widely signed [Oracle Contributor Agreement](http://www.oracle.com/technetwork/community/oca-486395.html)).

Please see the below or send a mail to [admins \[at\] neofourjay.org](mailto:admins@neofourjay.org) if you have any other questions about the intent of the CLA. If you have a legal question, please ask a lawyer.

33.1.2. Common questions

Am I losing the rights to my own code?

No, the [Neo4j CLA](http://dist.neo4j.org/neo4j-cla.pdf) only asks you to *share* your rights, not relinquish them. Unlike some contribution agreements that require you to transfer copyrights to another organization, the CLA does not take away your rights to your contributed intellectual property. When you agree to the CLA, you grant us joint ownership in copyright, and a patent license for your contributions. You retain all rights, title, and interest in your contributions and may use them for any purpose you wish. Other than revoking our rights, you can still do whatever you want with your code.

What can you do with my contribution?

We may exercise all rights that a copyright holder has, as well as the rights you grant in the [Neo4j CLA](http://dist.neo4j.org/neo4j-cla.pdf) to use any patents you have in your contributions. As the CLA provides for joint copyright ownership, you may exercise the same rights as we in your contributions.

What are the community benefits of this?

Well, it allows us to sponsor the Neo4j projects and provide an infrastructure for the community, while making sure that we can include this in software that we ship to our customers without any nasty surprises. Without this ability, we as a small company would be hard pressed to release all our code as free software.

Moreover, the CLA lets us protect community members (both developers and users) from hostile intellectual property litigation should the need arise. This is in line with how other free software stewards like the [Free Software Foundation - FSF](http://www.fsf.org) defend projects (except with the FSF, there's no shared copyright but instead you completely sign it over to the FSF). The contributor agreement also includes a "free software covenant," or a promise that a contribution will remain available as free software.

At the end of the day, you still retain all rights to your contribution and we can stand confident that we can protect the Neo4j community and the Neo Technology customers.

Can we discuss some items in the CLA?

Absolutely! Please give us feedback! But let's keep the legalese off the mailing lists. Please mail your feedback directly to [cla \(@\) neotechnology dot com](mailto:cla@neotechnology.com) and we'll get back to you.

I still don't like this CLA.

That's fine. You can still host it anywhere else, of course. Please do! We're only talking here about the rules for the infrastructure that we provide.

33.1.3. How to sign

When you've read through the CLA, please send a mail to `cla (@t) neotechnology dot com`. Include the following information:

- Your full name.
- Your e-mail address.
- An attached copy of the [Neo4j CLA](http://dist.neo4j.org/neo4j-cla.pdf) <<http://dist.neo4j.org/neo4j-cla.pdf>>.
- That you agree to its terms.

For example:

```
Hi. My name is John Doe (john@doe.com).  
I agree to the terms in the attached Neo4j Contributor License Agreement.
```

33.2. Areas for contribution

Neo4j is a project with a vast ecosystem and a lot of space for contributions. Where you can and want to pitch in depends of course on your time, skill set and interests. Below are some of the areas that might interest you:

33.2.1. Neo4j Distribution

- [The Neo4j open issues](https://github.com/neo4j/neo4j/issues) <https://github.com/neo4j/neo4j/issues> for some starting points for contribution
- See the [GitHub Neo4j area](https://github.com/neo4j/) <https://github.com/neo4j/> for a list of projects

33.2.2. Maintaining Neo4j Documentation

Some parts of the documentation need extra care from the community to stay up to date. They typically refer to different kinds of community contributions. Below is a list of such places, feel free to look into them and check for outdated or missing content! The easiest way to contribute fixes is to comment at the [online HTML version](http://docs.neo4j.org/chunked/snapshot/) <http://docs.neo4j.org/chunked/snapshot/>.

- [Chapter 5, Neo4j Remote Client Libraries](#)

33.2.3. Drivers and bindings to Neo4j

- REST: see [Chapter 5, Neo4j Remote Client Libraries](#) for a list of active projects

33.3. Writing Neo4j Documentation



Note

Other than writing documentation, you can help out by providing comments - head over to the [online HTML version](http://docs.neo4j.org/chunked/snapshot/) <<http://docs.neo4j.org/chunked/snapshot/>> to do that!

For how to build the manual see: [readme](https://github.com/neo4j/neo4j/blob/master/manual/README.asciidoc) <<https://github.com/neo4j/neo4j/blob/master/manual/README.asciidoc>>

The documents use the asciidoc format, see:

- [Aciidoc Reference](http://www.methods.co.nz/asciidoc/) <<http://www.methods.co.nz/asciidoc/>>
- [AsciiDoc cheatsheet](http://powerman.name/doc/asciidoc/) <<http://powerman.name/doc/asciidoc/>>

The cheatsheet is really useful!

33.3.1. Overall Flow

Each (sub)project has its own documentation, which will produce a *docs.jar* file. By default this file is assembled from the contents in *src/docs/*. AsciiDoc documents have the *.asciidoc* file extension.

The documents can use code snippets which will extract code from the project. The corresponding code must be deployed to the *sources.jar* or *test-sources.jar* file.

By setting up a unit test accordingly, documentation can be written directly in the JavaDoc comment.

The above files are all consumed by the build of the manual (by adding them as dependencies). To get content included in the manual, it has to be explicitly included by a document in the manual as well.

Note that different ways to add documentation works best for different cases:

- For detail level documentation, it works well to write the documentation as part of unit tests (in the JavaDoc comment). In this case, you typically do not want to link to the source code in the documentation.
- For tutorial level documentation, the result will be best by writing a *.asciidoc* file containing the text. Source snippets and output examples can then be included from there. In this case you typically want to link to the source code, and users should be able to run it without any special setup.

33.3.2. File Structure in *docs.jar*

Directory	Contents
<i>dev/</i>	content aimed at developers
<i>dev/images/</i>	images used by the dev docs
<i>ops/</i>	content aimed at operations
<i>ops/images/</i>	images used by the ops docs
<i>man/</i>	manpages

Additional subdirectories are used as needed to structure the documents, like *dev/tutorial/*, *ops/tutorial/* etc.

33.3.3. Headings and document structure

Each document starts over with headings from level zero (the document title). Each document should have an id. In some cases sections in the document need to have id's as well, this depends on where they fit in the overall structure. To be able to link to content, it has to have an id. Missing id's in mandatory places will fail the build.

This is how a document should start:

```
[[unique-id-verbose-is-ok]]
= The Document Title =
```

To push the headings down to the right level in the output, the `leveloffset` attribute is used when including the document inside of another document.

Subsequent headings in a document should use the following syntax:

```
== Subheading ==
... content here ...
=== Subsubheading ===
content here ...
```

AsciiDoc comes with one more syntax for headings, but in this project it's not used.

33.3.4. Writing

Put one sentence on each line. This makes it easy to move content around, and also easy to spot (too) long sentences.

33.3.5. Gotchas

- A chapter can't be empty. (the build will fail on the docbook xml validity check)
- The document title should be "underlined" by the same number of = as there are characters in the title.
- Always leave a blank line at the end of documents (or the title of the next document might end up in the last paragraph of the document)
- As {} are used for AsciiDoc attributes, everything inside will be treated as an attribute. What you have to do is to escape the opening brace: \{. If you don't, the braces and the text inside them will be removed without any warning being issued!

33.3.6. Links

To link to other parts of the manual the id of the target is used. This is how such a reference looks:

```
<<community-docs-overall-flow>>
```

Which will render like: [Section 33.3.1, "Overall Flow"](#)



Note

Just write "see <<target-id>>" and similar, that should suffice in most cases.

If you need to link to another document with your own link text, this is what to do:

```
<<target-id, link text that fits in the context>>
```

**Note**

Having lots of linked text may work well in a web context but is a pain in print, and we aim for both!

External links are added like this:

```
http://neo4j.org/[Link text here]
```

Which renders like: [Link text here](http://neo4j.org/) <http://neo4j.org/>

For short links it may be better not to add a link text, just do:

```
http://neo4j.org/
```

Which renders like: <http://neo4j.org/>

**Note**

It's ok to have a dot right after the URL, it won't be part of the link.

33.3.7. Text Formatting

- `_Italics_` is rendered as *Italics* and used for emphasis.
- `*Bold*` is rendered as **Bold** and used sparingly, for strong emphasis only.
- `+methodName()+` is rendered as `methodName()` and is used for literals as well (note: the content between the + signs *will* be parsed).
- ``command`` is rendered as `command` (typically used for command-line) (note: the content between the ` signs *will not* be parsed).
- `'my/path/` is rendered as `my/path/` (used for file names and paths).
- ```Double quoted"` (that is two grave accents to the left and two acute accents to the right) renders as "Double quoted".
- ``Single quoted'` (that is a single grave accent to the left and a single acute accent to the right) renders as 'Single quoted'.

33.3.8. Admonitions

These are very useful and should be used where appropriate. Choose from the following (write all caps and no, we can't easily add new ones):

**Note**

Note.

**Tip**

Tip.

**Important**

Important

**Caution**

Caution



Warning

Warning

Here's how it's done:

```
NOTE: Note.
```

A multiline variation:

```
[TIP]
Tiptext.
Line 2.
```

Which is rendered as:



Tip

Tiptext. Line 2.

33.3.9. Images



Important

All images in the entire manual share the same namespace. You know how to handle that.

Images Files

To include an image file, make sure it resides in the *images/* directory relative to the document you're including it from. Then go:

```
image::neo4j-logo.png[]
```

Which is rendered as:




Static Graphviz/DOT

We use the Graphviz/DOT language to describe graphs. For documentation see <http://graphviz.org/>.

This is how to include a simple example graph:

```
["dot", "community-docs-graphdb-rels.svg"]
----
"Start node" -> "End node" [label="relationship"]
----
```


Which is rendered as:



Here's an example using some predefined variables available in the build:

```
["dot", "community-docs-graphdb-rels-overview.svg", "meta"]
----
"A Relationship" [fillcolor="NODEHIGHLIGHT"]
"Start node" [fillcolor="NODE2HIGHLIGHT"]
"A Relationship" -> "Start node" [label="has a"]
"A Relationship" -> "End node" [label="has a"]
"A Relationship" -> "Relationship type" [label="has a"]
Name [TEXTNODE]
"Relationship type" -> Name [label="uniquely identified by" color="EDGEHIGHLIGHT" fontcolor="EDGEHIGHLIGHT"]
----
```

Which is rendered as:



The optional second argument given to the dot filter defines the style to use:

- when not defined: Default styling for nodespace examples.
- `neoviz`: Nodespace view generated by Neoviz.
- `meta`: Graphs that don't resemble db contents, but rather concepts.



Caution

Keywords of the DOT language have to be surrounded by double quotes when used for other purposes. The keywords include *node*, *edge*, *graph*, *digraph*, *subgraph*, and *strict*.

33.3.10. Attributes

Common attributes you can use in documents:

- `{neo4j-version}` - rendered as "1.9.7"
- `{neo4j-git-tag}` - rendered as "1.9.7"
- `{lucene-version}` - rendered as "3_6_2"

These can substitute part of URLs that point to for example API docs or source code. Note that `neo4j-git-tag` also handles the case of snapshot/master.

Sample AsciiDoc attributes which can be used:

- `{docdir}` - root directory of the documents

- `{nbsp}` - non-breaking space

33.3.11. Comments

There's a separate build including comments. The comments show up with a yellow background. This build doesn't run by default, but after a normal build, you can use `make annotated` to build it. You can also use the resulting page to search for content, as the full manual is on a single page.

Here's how to write a comment:

```
// this is a comment
```

The comments are not visible in the normal build. Comment blocks won't be included in the output of any build at all. Here's a comment block:

```
////
Note that includes in here will still be processed, but not make it into the output.
That is, missing includes here will still break the build!
////
```

33.3.12. Code Snippets

Explicitly defined in the document



Warning

Use this kind of code snippets as little as possible. They are well known to get out of sync with reality after a while.

This is how to do it:

```
[source,cypher]
----
start n=(2, 1) where (n.age < 30 and n.name = "Tobias") or not(n.name = "Tobias") return n
----
```

Which is rendered as:

```
start n=(2, 1) where (n.age < 30 and n.name = "Tobias") or not(n.name = "Tobias") return n
```

If there's no suitable syntax highlighter, just omit the language: `[source]`.

Currently the following syntax highlighters are enabled:

- Bash
- Cypher
- Groovy
- Java
- JavaScript
- Python
- XML

For other highlighters we could add see <http://alexgorbatchev.com/SyntaxHighlighter/manual/brushes/>.

Fetches from source code

Code can be automatically fetched from source files. You need to define:

- **component:** the artifactId of the Maven coordinates,
- **source:** path to the file inside the jar it's deployed to,
- **classifier:** sources or test-sources or any other classifier pointing to the artifact,
- **tag:** tag name to search the file for,
- the language of the code, if a corresponding syntax highlighter is available.

Note that the artifact has to be included as a Maven dependency of the Manual project so that the files can be found.

Be aware of that the tag "abc" will match "abcd" as well. It's a simple on/off switch, meaning that multiple occurrences will be assembled into a single code snippet in the output. This behavior can be user to hide away assertions from code examples sourced from tests.

This is how to define a code snippet inclusion:

```
[snippet,java]
----
component=neo4j-examples
source=org/neo4j/examples/JmxDocTest.java
classifier=test-sources
tag=getStartTime
----
```

This is how it renders:

```
private static Date getStartTimeFromManagementBean(
    GraphDatabaseService graphDbService )
{
    ObjectName objectName = JmxUtils.getObjectNames( graphDbService, "Kernel" );
    Date date = JmxUtils.getAttribute( objectName, "KernelStartTime" );
    return date;
}
```

Query Results

There's a special filter for Cypher query results. This is how to tag a query result:

```
.Result
[queryresult]
----
+-----+
| friend_of_friend.name | count(*) |
+-----+
| Ian                    | 2        |
| Derrick                | 1        |
| Jill                   | 1        |
+-----+
3 rows, 12 ms
----
```

This is how it renders:

Result

friend_of_friend.name	count(*)
Ian	2
Derrick	1
3 rows, 12 ms	

friend_of_friend.name	count(*)
Jill	1
3 rows, 12 ms	

33.3.13. A sample Java based documentation test

For Java, there are a couple of premade utilities that keep code and documentation together in Javadocs and code snippets that generate AsciiDoc for the rest of the toolchain.

To illustrate this, look at the following documentation that generates the AsciiDoc file `hello-world-title.asciidoc` with a content of:

```
[[examples-hello-world-sample-chapter]]
```

```
Hello world Sample Chapter
```

```
=====
```

```
This is a sample documentation test, demonstrating different ways of
bringing code and other artifacts into AsciiDoc form. The title of the
generated document is determined from the method name, replacing "+_" with
" ".
```

```
Below you see a number of different ways to generate text from source,
inserting it into the JavaDoc documentation (really being AsciiDoc markup)
via the +@@" snippet markers and programmatic adding with runtime data
in the Java code.
```

```
- The annotated graph as http://www.graphviz.org/[GraphViz]-generated visualization:
```

```
.Hello World Graph
```

```
["dot", "Hello-World-Graph-hello-world-Sample-Chapter.svg", "neoviz", ""]
```

```
----
```

```
N1 [
  label = "{Node\[1\]|name = \'you\'\l}"
]
N2 [
  label = "{Node\[2\]|name = \'I\'\l}"
]
N2 -> N1 [
  color = "#2e3436"
  fontcolor = "#2e3436"
  label = "know\n"
]
----
```

```
- A sample Cypher query:
```

```
[source,cypher]
```

```
----
```

```
START n = node(2)
RETURN n
----
```

```
- A sample text output snippet:
```

```
[source]
```

```
----
```

```
Hello graphy world!
----
```

```
- a generated source link to the original GitHub source for this test:
```

```
https://github.com/neo4j/neo4j/blob/{neo4j-git-tag}/community/embedded-examples/src/test/java/org/neo4j/examples/DocumentationDocTest.java
```

- The full source for this example as a source snippet, highlighted as Java code:

```
[snippet,java]
----
component=neo4j-examples
source=org/neo4j/examples/DocumentationDocTest.java
classifier=test-sources
tag=sampleDocumentation
----
```

This is the end of this chapter.

this file is included in this documentation via

```
:leveloffset: 3
include::{importdir}/neo4j-examples-docs-jar/dev/examples/hello-world-sample-chapter.asciidoc[]
```

which renders the following chapter:


33.3.14. Hello world Sample Chapter

This is a sample documentation test, demonstrating different ways of bringing code and other artifacts into AsciiDoc form. The title of the generated document is determined from the method name, replacing "_" with " ".

Below you see a number of different ways to generate text from source, inserting it into the JavaDoc documentation (really being AsciiDoc markup) via the @@ snippet markers and programmatic adding with runtime data in the Java code.

- The annotated graph as [GraphViz](http://www.graphviz.org/) <http://www.graphviz.org/>-generated visualization:

Figure 33.1. Hello World Graph



- A sample Cypher query:

```
START n = node(2)
RETURN n
```

- A sample text output snippet:

```
Hello graphy world!
```

- a generated source link to the original GitHub source for this test:

[DocumentationDocTest.java](https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/DocumentationDocTest.java) <https://github.com/neo4j/neo4j/blob/1.9.7/community/embedded-examples/src/test/java/org/neo4j/examples/DocumentationDocTest.java>

- The full source for this example as a source snippet, highlighted as Java code:

```
// START SNIPPET: _sampleDocumentation
package org.neo4j.examples;

import org.junit.Test;
import org.neo4j.kernel.impl.annotations.Documented;
import org.neo4j.test.GraphDescription.Graph;

import static org.neo4j.visualization.asciidoc.AsciidocHelper.*;

public class DocumentationDocTest extends AbstractJavaDocTestbase
{
    /**
     * This is a sample documentation test, demonstrating different ways of
     * bringing code and other artifacts into AsciiDoc form. The title of the
     * generated document is determined from the method name, replacing "+_" with
     * " ".
     *
     * Below you see a number of different ways to generate text from source,
     * inserting it into the JavaDoc documentation (really being AsciiDoc markup)
     * via the +@@+ snippet markers and programmatic adding with runtime data
     * in the Java code.
     *
     * - The annotated graph as http://www.graphviz.org/\[GraphViz\]-generated visualization:
     *
     * @@graph
     *
     * - A sample Cypher query:
     *
     * @@cypher
     *
     * - A sample text output snippet:
     *
     * @@output
     *
     * - a generated source link to the original GitHub source for this test:
     *
     * @@github
     *
     * - The full source for this example as a source snippet, highlighted as Java code:
     *
     * @@sampleDocumentation
     *
     * This is the end of this chapter.
     */
    @Test
    // signaling this to be a documentation test
    @Documented
    // the graph data setup as simple statements
    @Graph( "I know you" )
    // title is determined from the method name
    public void hello_world_Sample_Chapter()
    {
        // initialize the graph with the annotation data
        data.get();
        gen.get().addTestSourceSnippets( this.getClass(), "sampleDocumentation" );
        gen.get()
            .addGithubTestSourceLink( "github", this.getClass(),
                "community/embedded-examples" );

        gen.get().addSnippet( "output",
            createOutputSnippet( "Hello graphy world!" ) );
    }
}
```

```

gen.get().addSnippet(
    "graph",
    createGraphVizWithNodeId( "Hello World Graph", graphdb(),
        gen.get().getTitle() ) );
// A cypher snippet referring to the generated graph in the start clause
gen.get().addSnippet(
    "cypher",
    createCypherSnippet( "start n = node(" + data.get().get( "I" ).getId()
        + ") return n" ) );
}
}
// END SNIPPET: _sampleDocumentation

```

This is the end of this chapter.

33.3.15. Integrated Live Console

An interactive console can be added and will show up in the online HTML version. An optional title can be added, which will be used for the text of the button.

This is how to do it, using Geoff to define the data, with an empty line to separate it from the query:

```

.Interactive Matrix Example
[console]
----
(A) {"name" : "Neo"};
(B) {"name" : "Trinity"};
(A)-[:LOVES]->(B)

start n = node(*)
return n
----

```

And this is the result:

33.3.16. Toolchain

Useful links when configuring the docbook toolchain:

- <http://www.docbook.org/tdg/en/html/docbook.html>
- <http://www.sagehill.net/docbookxsl/index.html>
- <http://docbook.sourceforge.net/release/xsl/1.76.1/doc/html/index.html>
- <http://docbook.sourceforge.net/release/xsl/1.76.1/doc/fo/index.html>

33.4. Translating the Neo4j Manual

To translate the Neo4j Manual, there's a special project setup to use. See the French translation project for an example: <https://github.com/neo4j/manual-french>

The project contains:

- *conf/* — configuration for the project.
- *docs/* — translated files for content provided by Neo4j modules.
- *po/* — translation files and po4a configuration files.
- *src/* — translated files for content provided by the original manual.
- *Makefile* — a makefile with project-specific configuration.
- *pom.xml* — Maven build configuration.

33.4.1. Prerequisites

- Apache Maven
- GNU Make
- Python
- Perl
- Perl module: `Unicode::GCString`

To check if you have the `Unicode::GCString` module installed, you can issue the following command:

```
perl -MUnicode::GCString -e ''
```

If there's no error, the module has been successfully installed on your system.

To install the module, you can use `cpanminus` <<http://search.cpan.org/dist/App-cpanminus/lib/App-cpanminus.pm>>. For a convenient way to install it, see <http://cpanmin.us>. With `cpanminus` installed, execute this command:

```
cpanm Unicode::GCString
```

You will probably want to use a *.po* file editor as well, see [Section 33.4.5, “Translation tools”](#).

33.4.2. Build flow and file layout

The build is essentially a two-step process. The first step generates or copies translated documents, while the second step is an ordinary AsciiDoc build using the output from the first step as sources.

Other than the *src/* and *docs/* directories of the project, the build generates files with the same layout in two more places:

1. *target/original/(src|docs)/* — the contents of the original manual. Note that's it easier to look for content here than to dig into the original manual itself.
2. *target/(src|docs)/* — the translated source to use for the AsciiDoc build.

The translated documents in *target/(src|docs)/* are generated in three steps:

1. It starts out as a copy of the original manual.
2. Next, any static translated files from the *src/* and *docs/* directories of the project are copied.
3. Finally, the translation files in the *po/* directory are used to generate translated documents.

Files produced by later steps will overwrite existing files from earlier steps.

33.4.3. Adding a chapter to a translation file

The translation is split over multiple translation files, one per “part” of the manual. It’s all about making the translation easier to manage and the tools to perform well. The basic rule of thumb is that if some content is moved, it should likely still end up in the same translation file. In that case, the tools will even detect this and the translation will be moved automatically.

To add a document to a translation file, do like this:

```
make add DOCUMENT="src/introduction/the-neo4j-graphdb.asciidoc" PART="introduction"
```

If the translation file does not already exist, it will be created. The document will be added to the translation build configuration file as well. (The configuration is in the corresponding *.conf* file in the *po/* directory.)

If there exists a translated copy of the document at the location the `DOCUMENT` parameter points to, the script will attempt to populate the translation file with translated paragraphs from that document. Note that the structure of the document has to be a *perfect match*, or it will fail. However, the error messages are helpful, so just fix and try again until it works! Translation file and configuration are only changed when the first part succeeds.



Note

Only documents that need to be translated should be added. For example Cypher queries and query results should not be translated. In general, documents residing in a directory named *includes* should not be translated.

Also note that AsciiDoc `include::` lines are normally not part of the translation at all, but handled automatically. In case they need to be handled differently in a document, this has to be configured in the corresponding *.conf* file. For example a normal document entry in such a file can look like this:

```
[type: asciidoc] target/original/src/operations/index.asciidoc fr:target/src/operations/index.asciidoc
```

To configure a single document not to handle `include::` lines automatically, add the following at the end of the line:

```
opt: "-o definitions=target/tools/main/resources/conf/translate-includes"
```

33.4.4. Workflow

First, use Maven to set up the environment and download the original manual and documentation tools:

```
mvn clean package
```

To refresh the original manual and the tools, use the maven command again. For the sake of keeping in sync with the original manual, a daily run of this command is recommended.

Once things are set up, use `make` during work.

- `make` — same as `make preview`.
- `make add` — add a document to a translation file.
- `make preview` — refresh and build preview of the manual.
- `make refresh` — refresh translation files from the original and generated translated documents.

The preview of the translated manual is found in the *target/html/* directory.

The actual work on translation is done by editing translation files. Suggested tools for that are found below.

33.4.5. Translation tools

There are different editors for *.po* files containing the translations Below is a list of editors.

- [Gtranslator](http://projects.gnome.org/gtranslator/) <http://projects.gnome.org/gtranslator/>
- [Lokalize](http://userbase.kde.org/Lokalize) <http://userbase.kde.org/Lokalize>
- [Virtaal](http://translate.sourceforge.net/wiki/virtaal/index) <http://translate.sourceforge.net/wiki/virtaal/index>
- [Poedit](http://www.poedit.net/) <http://www.poedit.net/>

33.5. Contributing Code to Neo4j

33.5.1. Intro

The Neo4j community is a free software and open source community centered around software and components for the Neo4j Graph Database. It is sponsored by [Neo Technology](http://neotechnology.com) <http://neotechnology.com>, which provides infrastructure (different kinds of hosting, documentation, etc) as well as people to manage it. The Neo4j community is an open community, in so far as it welcomes any member that accepts the basic criterias of contribution and adheres to the community's Code of Conduct.

Contribution can be in many forms (documentation, discussions, bug reports). This document outlines the rules of governance for a contributor of code.

33.5.2. Governance fundamentals

In a nutshell, you need to be aware of the following fundamentals if you wish to contribute code:

- All software published by the Neo4j project must have been contributed under the [Neo4j Code Contributor License Agreement](#).
- Neo4j is a free software and open source community. As a contributor, you are free to place your work under any license that has been approved by either the [Free Software Foundation](http://fsf.org) <http://fsf.org> or the [Open Source Initiative](http://opensource.org) <http://opensource.org>. You still retain copyright, so in addition to that license you can of course release your work under any other license (for example a fully proprietary license), just not on the Neo4j infrastructure.
- The Neo4j software is split into components. A Git repository holds either a single or multiple components.
- The source code should follow the Neo4j [Code Style](#) and “fit in” with the Neo4j infrastructure as much as is reasonable for the specific component.

33.5.3. Contributor roles

Every individual that contributes code does so in the context of a role (a single individual can have multiple roles). The role defines their responsibilities and privileges:

- A *patch submitter* is a person who wishes to contribute a patch to an existing component. See [Workflow](#) below.
- A *committer* can contribute code directly to one or more components.
- A *component maintainer* is in charge of a specific component. They can:
 - commit code in their component's repository,
 - manage tickets for the repository,
 - grant push rights to the repository.
- A *Neo4j admin* manages the Neo4j infrastructure. They:
 - define new components and assign component maintainership,
 - drive, mentor and coach Neo4j component development.

33.5.4. Contribution workflow

Code contributions to Neo4j are normally done via Github Pull Requests, following the workflow shown below. Please check the [pull request checklist](#) before sending off a pull request as well.

1. Fork the appropriate Github repository.

2. Create a new branch for your specific feature or fix.
3. [Write unit tests for your code.](#)
4. [Write code.](#)
5. Write appropriate Javadocs and [Manual entries.](#)
6. [Commit changes.](#)
7. [Send pull request.](#)

33.5.5. Pull request checklist

1. [Sign the CLA.](#)
2. [Ensure that you have not added any merge commits.](#)
3. [Squash all your changes into a single commit.](#)
4. [Rebase against the latest master.](#)
5. [Run all relevant tests.](#)
6. Send the request!

33.5.6. Unit Tests

You have a much higher chance of getting your changes accepted if you supply us with small, readable unit tests covering the code you've written. Also, make sure your code doesn't break any existing tests. *Note that there may be downstream components that need to be tested as well, depending on what you change.*

To run tests, use Maven rather than your IDE to ensure others can replicate your test run. The command for running Neo4j tests in any given component is `mvn clean validate`.

33.5.7. Code Style

The Neo4j Code style is maintained on GitHub in [styles for the different IDEs](https://github.com/neo4j/neo4j.github.com/tree/master/code-style) <<https://github.com/neo4j/neo4j.github.com/tree/master/code-style>>.

33.5.8. Commit messages

Please take some care in providing good commit messages. Use your common sense. In particular:

- Use *english*. This includes proper punctuation and correct spelling. Commit messages are supposed to convey some information at a glance — they're not a chat room.
- Remember that a commit is a *changeset*, which describes a cohesive set of changes across potentially many files. Try to group every commit as a logical change. Explain what it changes. If you have to redo work, you might want to clean up your commit log before doing a pull request.
- If you fix a bug or an issue that's related to a ticket, then refer to the ticket in the message. For example, *'Added this and then changed that. This fixes #14.'* Just mentioning #xxx in the commit will connect it to the GitHub issue with that number, see [GitHub issues](https://github.com/blog/831-issues-2-0-the-next-generation) <<https://github.com/blog/831-issues-2-0-the-next-generation>>. Any of these synonyms will also work:
 - fixes #xxx
 - fixed #xxx
 - fix #xxx
 - closes #xxx
 - close #xxx
 - closed #xxx.

- Remember to convey *intent*. Don't be too brief but don't provide too much detail, either. That's what `git diff` is for.

33.5.9. Signing the CLA

One crucial aspect of contributing is the [Contributor License Agreement](#). In short: make sure to sign the CLA, or the Neo4j project won't be able to accept your contribution.

33.5.10. Don't merge, use rebase instead!

Because we would like each contribution to be contained in a single commit, merge commits are not allowed inside a pull request. Merges are messy, and should only be done when necessary, eg. when merging a branch into master to remember where the code came from.

If you want to update your development branch to incorporate the latest changes from master, use `git rebase`. For details on how to use rebase, see Git manual on rebase: [the Git Manual <http://git-scm.com/book/en/Git-Branching-Rebasing>](http://git-scm.com/book/en/Git-Branching-Rebasing).

33.5.11. Single commit

If you have multiple commits, you should squash them into a single one for the pull request, unless there is some extraordinary reason not to. Keeping your changes in a single commit makes the commit history easier to read, it also makes it easy to revert and move features around.

One way to do this is to, while standing on your local branch with your changes, create a new branch and then interactively rebase your commits into a single one.

Interactive rebasing with Git.

```
# On branch mychanges
git checkout -b mychanges-clean

# Assuming you have 4 commits, rebase the last four commits interactively:
git rebase -i HEAD~4

# In the dialog git gives you, keep your first commit, and squash all others into it.
# Then reword the commit description to accurately depict what your commit does.
# If applicable, include any issue numbers like so: #760
```

For more details, see the git manual: <http://git-scm.com/book/en/Git-Tools-Rewriting-History#Changing-Multiple-Commit-Messages>

If you are asked to modify parts of your code, work in your original branch (the one with multiple commits), and follow the above process to create a fixed single commit.

33.6. Contributors

As an Open Source Project, the Neo4j User community extends its warmest thanks to all the contributors who have signed the [Section 33.1, “Contributor License Agreement”](#) to date and are contributing to this collective effort.

name	GITHub ID
Johan Svensson	johan-neo < https://github.com/johan-neo >
Emil Eifrem	emileifrem < https://github.com/emileifrem >
Peter Neubauer	peterneubauer < https://github.com/peterneubauer >
Mattias Persson	tinwelint < https://github.com/tinwelint >
Tobias Lindaaker	thobe < https://github.com/thobe >
Anders Nawroth	nawroth < https://github.com/nawroth >
Andrés Taylor	systay < https://github.com/systay >
Jacob Hansson	jakewins < https://github.com/jakewins >
Jim Webber	jimwebber < https://github.com/jimwebber >
Josh Adell	jadell < https://github.com/jadell >
Andreas Kollegger	akollegger < https://github.com/akollegger >
Chris Gioran	digitalstain < https://github.com/digitalstain >
Thomas Baum	tbaum < https://github.com/tbaum >
Alistair Jones	apcj < https://github.com/apcj >
Michael Hunger	jexp < https://github.com/jexp >
Jesper Nilsson	jespernilsson < https://github.com/jespernilsson >
Tom Sulston	tomsulston < https://github.com/tomsulston >
David Montag	dmontag < https://github.com/dmontag >
Marlon Richert	marlonrichert < https://github.com/marlonrichert >
Hugo Josefson	hugojosefson < https://github.com/hugojosefson >
Vivek Prahlaad	vivekprahlaad < https://github.com/vivekprahlaad >
Adriano Almeida	adrianoalmeida7 < https://github.com/adrianoalmeida7 >
Benjamin Gehrels	BGehrels < https://github.com/BGehrels >
Christopher Schmidt	FaKod < https://github.com/FaKod >
Pascal Rehfeldt	prehfeldt < https://github.com/prehfeldt >
Björn Söderqvist	cybear < https://github.com/cybear >
Abdul Azeez Shaik	abdulazezsk < https://github.com/abdulazezsk >
James Thornton	espeed < https://github.com/espeed >
Radhakrishna Kalyan	nrkkalyan < https://github.com/nrkkalyan >
Michel van den Berg	promontis < https://github.com/promontis >

name	GIThub ID
Brandon McCauslin	bm3780 <https://github.com/bm3780>
Hendy Irawan	ceefour <https://github.com/ceefour>
Luanne Misquitta	luanne <https://github.com/luanne>
Jim Radford	radford <https://github.com/radford>
Axel Morgner	amorgner <https://github.com/amorgner>
Taylor Buley	editor <https://github.com/editor>
Alex Smirnov	alexsmirnov <https://github.com/alexsmirnov>
Johannes Mockenhaupt	jotomo <https://github.com/jotomo>
Pablo Pareja Tobes	pablopereja <https://github.com/pablopereja>
Björn Granvik	bjorngranvik <https://github.com/bjorngranvik>
Julian Simpson	simpsonjulian <https://github.com/simpsonjulian>
Pablo Pareja Tobes	pablopereja <https://github.com/pablopereja>
Rickard Öberg	rickardoberg <https://github.com/rickardoberg>
Stefan Armbruster	sarmbruster <https://github.com/sarmbruster>
Stephan Hagemann	shageman <https://github.com/shageman>
Linan Wang	wangii <https://github.com/wangii>
Ian Robinson	iansrobinson <https://github.com/iansrobinson>
Marko Rodriguez	okram <https://github.com/okram>
Saikat Kanjilal	skanjila <https://github.com/skanjila>
Craig Taverner	craigtaverner <https://github.com/craigtaverner>
David Winslow	dwins <https://github.com/dwins>
Patrick Fitzgerald	paddydub <https://github.com/paddydub>
Stefan Berder	hrbonz <https://github.com/hrbonz>
Michael Kanner	SepiaGroup <https://github.com/SepiaGroup>
Lin Zhemin	miaoski <https://github.com/miaoski>
Christophe Willemsen	kwattro <https://github.com/kwattro>
Tony Liu	kooyeed <https://github.com/kooyeed>
Michael Klishin	michaelklishin <https://github.com/michaelklishin>
Wes Freeman	wfreeman <https://github.com/wfreeman>
Chris Leishman	chrisleishman <https://github.com/chrisleishman>
Brian Levine	blevine <https://github.com/blevine>
Ben Day	benday280412 <https://github.com/benday280412>
Davide Savazzi	svzdvd <https://github.com/svzdvd>

name	GitHub ID
Nigel Small	nigelsmall <https://github.com/nigelsmall>
Lasse Westh-Nielsen	lassewesth <https://github.com/lassewesth>
Wujek Srujek	wujek-srujek <https://github.com/wujek-srujek>
Alexander Yastrebov	AlexanderYastrebov <https://github.com/AlexanderYastrebov>
Mike Bryant	mikesname <https://github.com/mikesname>
Klaus Grossmann	iKlaus <https://github.com/iKlaus>
Pablo Lalloni	plalloni <https://github.com/plalloni>
Stefan Plantikow	boggle <https://github.com/boggle>
Trenton Strong	trentonstrong <https://github.com/trentonstrong>
Maciej Mazur	mamciek <https://github.com/mamciek>
German Borbolla	germanborbolla <https://github.com/germanborbolla>
Laurent Raufaste	lra <https://github.com/lra>
Thomas Häfele	Perfect-Pixel <https://github.com/Perfect-Pixel>
Sevki Hasirci	Sevki <https://github.com/Sevki>
Max De Marzi	maxdemarzi <https://github.com/maxdemarzi>
Pieter-Jan Van Aeken	PieterJanVanAeken <https://github.com/PieterJanVanAeken>
Shane Gibbs	sgibbs-kellermed <https://github.com/sgibbs-kellermed>
Volker Lanting	VolkerL <https://github.com/VolkerL>
Mark Needham	mneedham <https://github.com/mneedham>
Chris Vest	chrisvest <https://github.com/chrisvest>
Ben Butler-Cole	benbc <https://github.com/benbc>

Appendix A. Manpages

The Neo4j Unix manual pages are included on the following pages.

- [neo4j](#)
- [neo4j-shell](#)
- [neo4j-backup](#)
- [neo4j-arbiter](#)

Name

neo4j — Neo4j Server control and management

Synopsis

neo4j <command>

DESCRIPTION

Neo4j is a graph database, perfect for working with highly connected data.

COMMANDS**console**

Start the server as an application, running as a foreground proces. Stop the server using CTRL-C.

start

Start server as daemon, running as a background process.

stop

Stops a running daemonized server.

restart

Restarts the server.

status

Current running state of the server.

install

Installs the server as a platform-appropriate system service.

remove

Uninstalls the system service.

info

Displays configuration information, such as the current NEO4J_HOME and CLASSPATH.

Usage - Windows**Neo4j.bat**

Double-clicking on the Neo4j.bat script will start the server in a console. To quit, just press control-C in the console window.

Neo4j.bat install/remove

Neo4j can be installed and run as a Windows Service, running without a console window. You'll need to run the scripts with Administrator priveleges. Just use the Neo4j.bat script with the proper argument:

- Neo4j.bat install - install as a Windows service
 - will install the service
- Neo4j.bat remove - remove the Neo4j service
 - will stop and remove the Neo4j service
- Neo4j.bat start - will start the Neo4j service
 - will start the Neo4j service if installed or a console session otherwise.
- Neo4j.bat stop - stop the Neo4j service if running

- Neo4j.bat restart - restart the Neo4j service if installed
- Neo4j.bat status - report on the status of the Neo4j service
 - returns RUNNING, STOPPED or NOT INSTALLED

FILES

conf/neo4j-server.properties

Server configuration.

conf/neo4j-wrapper.conf

Configuration for service wrapper.

conf/neo4j.properties

Tuning configuration for the database.

Name

neo4j-shell — a command-line tool for exploring and manipulating a graph database

Synopsis

neo4j-shell [*REMOTE OPTIONS*]

neo4j-shell [*LOCAL OPTIONS*]

DESCRIPTION

Neo4j shell is a command-line shell for browsing the graph, much like how the Unix shell along with commands like `cd`, `ls` and `pwd` can be used to browse your local file system. The shell can connect directly to a graph database on the file system. To access local a local database used by other processes, use the readonly mode.

REMOTE OPTIONS

-port *PORT*

Port of host to connect to (default: 1337).

-host *HOST*

Domain name or IP of host to connect to (default: localhost).

-name *NAME*

RMI name, i.e. `rmi://<host>:<port>/<name>` (default: shell).

-readonly

Access the database in read-only mode.

LOCAL OPTIONS

-path *PATH*

The path to the database directory. If there is no database at the location, a new one will e created.

-pid *PID*

Process ID to connect to.

-readonly

Access the database in read-only mode.

-c *COMMAND*

Command line to execute. After executing it the shell exits.

-config *CONFIG*

The path to the Neo4j configuration file to be used.

EXAMPLES

Examples for remote:

```
neo4j-shell
neo4j-shell -port 1337
neo4j-shell -host 192.168.1.234 -port 1337 -name shell
neo4j-shell -host localhost -readonly
```

Examples for local:

```
neo4j-shell -path /path/to/db
neo4j-shell -path /path/to/db -config /path/to/neo4j.properties
neo4j-shell -path /path/to/db -readonly
```

Name

neo4j-backup — Neo4j Backup Tool

Synopsis

neo4j-backup -from SourceURI -to Directory

DESCRIPTION

A tool to perform live backups over the network from a running Neo4j graph database onto a local filesystem. Backups can be either full or incremental. The first backup must be a full backup, after that incremental backups can be performed.

The source(s) are given as URIs in a special format, the target is a filesystem location.

BACKUP TYPE

-full

copies the entire database to a directory.

-incremental

copies the changes that have taken place since the last full or incremental backup to an existing backup store.

The backup tool will automatically detect whether it needs to do a full or an incremental backup.

SOURCE URI

Backup sources are given in the following format:

<running mode>://<host>[:<port>][,<host>[:<port>]]...

Note that multiple hosts can be defined.

running mode

'single' or 'ha'. 'ha' is for instances in High Availability mode, 'single' is for standalone databases.

host

In single mode, the host of a source database; in ha mode, the cluster address of a cluster member. Note that multiple hosts can be given when using High Availability mode.

port

In single mode, the port of a source database backup service; in ha mode, the port of a cluster instance. If not given, the default value 6362 will be used for single mode, 5001 for HA

IMPORTANT

Backups can only be performed on databases which have the configuration parameter `enable_online_backup=true` set. That will make the backup service available on the default port (6362). To enable the backup service on a different port use for example `enable_online_backup=port=9999` instead.

Usage - Windows

The `Neo4jBackup.bat` script is used in the same way.

EXAMPLES

```
# Performing a backup the first time: create a blank directory and run the backup tool
```

```
mkdir /mnt/backup/neo4j-backup
neo4j-backup -from single://192.168.1.34 -to /mnt/backup/neo4j-backup

# Subsequent backups using the same _target_directory will be incremental and therefore quick
neo4j-backup -from single://freja -to /mnt/backup/neo4j-backup

# Performing a backup where the service is registered on a custom port
neo4j-backup -from single://freja:9999 -to /mnt/backup/neo4j-backup

# Performing a backup from HA cluster, specifying only one cluster member
./neo4j-backup -from ha://oden:5002 -to /mnt/backup/neo4j-backup

# Performing a backup from HA cluster, specifying two cluster members
./neo4j-backup -from ha://oden:5001,loke:5002 -to /mnt/backup/neo4j-backup
```

RESTORE FROM BACKUP

The Neo4j backups are fully functional databases. To use a backup, replace your database directory with the backup.

Name

neo4j-arbiter — Neo4j Arbiter for High-Availability clusters

Synopsis

neo4j-arbiter <command>

DESCRIPTION

Neo4j Arbiter is a service that can help break ties in Neo4j clusters that have an even number of cluster members.

COMMANDS

console

Start the server as an application, running as a foreground process. Stop the server using CTRL-C.

start

Start server as daemon, running as a background process.

stop

Stops a running daemonized server.

restart

Restarts a running server.

status

Current running state of the server

install

Installs the server as a platform-appropriate system service.

remove

Uninstalls the system service

FILES

conf/arbiter.cfg

Arbiter server configuration.

conf/arbiter-wrapper.cfg

Configuration for service wrapper.