

# Lab 11 – Cassandra

CC5212-1

May 25, 2016

So today we get to play around with a NoSQL store! In particular, we will play around with Apache Cassandra, which is a column family/tabular store. We will have a look at adding and removing keyspaces, adding and removing tables, managing schema and types, how data are inserted and queried, etc. We are going to use the Cassandra Query Language (CQL) interface, which is a bit like a lightweight version of SQL. You may find the documentation here helpful: <https://cassandra.apache.org/doc/cq13/CQL.html>.

You can use Linux or Windows, as you prefer (you just need SSH today). To start with, log into the cluster (username `uhadoop`, password on the board).

- You can use the following commands to get some info on the Cassandra nodes:

```
/data/hadoop/cassandra/bin/nodetool -h 10.10.10.1 info
/data/hadoop/cassandra/bin/nodetool -h 10.10.10.11 info
/data/hadoop/cassandra/bin/nodetool -h 10.10.10.12 info
/data/hadoop/cassandra/bin/nodetool -h 10.10.10.13 info
/data/hadoop/cassandra/bin/nodetool -h 10.10.10.14 info
/data/hadoop/cassandra/bin/nodetool -h 10.10.10.15 info
```

- Next we are going to launch the CQL (Cassandra Query Language) client. Run the following command (the first argument is the IP of the server, the second is the port):

```
/data/hadoop/cassandra/bin/cqlsh 10.10.10.1 9160
```

- **Note that you will be submitting the (working) commands you enter for the following parts so please copy and paste them into a text file.**
- Now we first create a keyspace and some tables.

- First we will create a keyspace with a simple replication strategy (hashing on key) and 3 copies:

```
CREATE keyspace YOURUSERNAME with replication =
{'class' : 'SimpleStrategy', 'replication_factor' : 3 };
```

We can see the list of keyspaces:

```
DESCRIBE keyspaces;
```

- Next we want to move to that keyspace: `USE YOURUSERNAME;`
- Now we can *try* create a table in that keyspace with two columns, one called `username` of type `text`, one called `age` of type `int`:

```
CREATE TABLE UNIQUETABLENAME ( username text, age int);
```

But wait, that didn't work! Cassandra tables are really maps and every map needs keys. We need to specify a primary key, which gives a unique identifier for each row:

```
CREATE TABLE UNIQUETABLENAME ( username text PRIMARY KEY, age int);
```

That's better. Note that if you are not in a given namespace, you can still reference a table using `namespace.table`. If you are inside the right namespace, you can just use the table name.

- Now we can list the tables in the keyspace we are in:

```
DESCRIBE tables;
```

On second thought, let's add an alive column to the table:

```
ALTER TABLE UNIQUETABLENAME ADD alive boolean;
```

- Okay, now you have a taste for how to create keyspaces and tables. After all your hard work above, you're going to delete everything you just made. Please only drop your stuff! We will have more interesting data if we work from now on in one common keyspace/table.

```
DROP TABLE UNIQUETABLENAME;
```

```
DROP KEYSPACE YOURUSERNAME;
```

- Let's move to a common keyspace I made earlier where we can see how to update and query data.

- First load the keyspace: `USE cc5212_2016`;
- Next, see what tables there are: `DESCRIBE tables`;
- What's inside the table? `SELECT * FROM students`;

- Now you will add your own data. In the following, I use my own data as an example. *Please replace my details with yours.* Please don't overwrite my row. If you do, I will be sad.

- First we will insert a name with initials:

```
INSERT INTO students( username, name ) VALUES ( 'ahogan', 'A. Hogan');
```

Query the table again to make sure it's in there: `SELECT * FROM students`;

- What happens if we try to insert a bad type:

```
INSERT INTO students( username, name ) VALUES ( 'ahogan', true);
```

(Doesn't work! Type of `name` is text, not boolean.)

- Let's just delete the name

```
DELETE name FROM students WHERE username = 'ahogan';
```

- Actually let's delete the entire row

```
DELETE FROM students WHERE username = 'ahogan';
```

- Now let's try give our full name instead:

```
INSERT INTO students( username, name ) VALUES ( 'ahogan', 'Aidan Hogan');
```

If you query the table again, you will see that the old value is overwritten. Unlike BigTable/Dynamo, etc., values are not versioned in Cassandra.

- Let's try add more information, like our age:

```
INSERT INTO students( age ) VALUES ( 30 );
```

Doesn't work? Why not? (Whose age should we update? We need to give the key!)

```
INSERT INTO students( username, age ) VALUES ( 'ahogan', 30 );
```

Try query the table again. We see that age is added alongside name.

- Fill in the rest of the details in the table for yourself following the above idea.

- Okay, so let's try some queries over the table now; hopefully it's filling up.

- Let's get the number of folks in the table:

```
SELECT COUNT(*) AS student_count FROM students;
```

- The name of a given user: `SELECT name FROM students WHERE username = 'ahogan';`
- People who are a given age: `SELECT name FROM students WHERE age = 30;`  
Oops, didn't work.
- Actually we are very limited in what we can do by default in Cassandra. Like a map, we can only really do lookups by keys: in this case, the primary key is `username`. To do more, we need to create indexes.
  - As a first step, you should make a personal copy of the table by saving it to a file and reloading it (all one command here):  
`COPY students (username, age, fav_colour, female, message, name)  
TO 'MYNAME.csv';`
  - Now create a new table with your username that fits the data, and load the data from the file:  
`COPY MYUSERNAME (username, age, fav_colour, female, message, name)  
FROM 'MYNAME.csv';`  
This is your own copy of the table to play with.
  - Let's say we're building an application that needs to lookup users with a given age. In Cassandra, we need to build an index for that.  
`CREATE INDEX ageIndex ON MYUSERNAME (age);`
  - Now we can try find users by age:  
`SELECT name FROM MYUSERNAME WHERE age = 30;`  
Okay great, now it works!
  - Now try to figure out how to get the following query to work (check out the documentation linked at the top):  
`SELECT name FROM MYUSERNAME WHERE age > 25 AND female = false;`
- Submit all the working CQL commands you typed in the lab into a text file and submit it to u-cursos by next Monday.

As we see, the default query power of Cassandra is very limited. This can be improved by adding default indexes, up to a point. The scalability of Cassandra trades off against its flexibility.