

Exercise 2 – CQL

In this exercise, you will:

- Create a keyspace for KillrVideo
- Create a table to store video metadata
- Load the data for the video table from a CSV file

Steps

Welcome to the KillrVideo company! KillrVideo hired you to build the latest and greatest video sharing application on the Internet. Your task is to ramp up on the domain and become acquainted with Apache Cassandra™. To start, you decided to look into creating a table schema and to load some video data.

The video metadata is made up of:

Column Name	Data Type
video_id	timeuuid
added_date	timestamp
title	text

1) Back in your Terminal window, make sure DataStax Enterprise is still running with `./dsetool status`. If not, restart DataStax Enterprise.

2) In the terminal window, start `cqlsh`:

```
/home/ubuntu/node/resources/cassandra/bin/cqlsh
```

3) In `cqlsh`, create a keyspace called `killrvideo`. Use `SimpleStrategy` for the replication class with a replication factor of one.

NOTE: You can press the tab key within the `CREATE KEYSPACE` command to have `cqlsh`` autocomplete the replication parameters.

```
CREATE KEYSPACE killrvideo
WITH replication = {
  'class': 'SimpleStrategy',
  'replication_factor': 1
};
```

4) In cqlsh switch to the newly created keyspace with the USE command.

```
USE killrvideo;
```

5) Create a single table called videos with the same structure as shown above. video_id is the primary key.

```
CREATE TABLE videos (  
  video_id TIMEUUID,  
  added_date TIMESTAMP,  
  title TEXT,  
  PRIMARY KEY (video_id)  
);
```

6) Manually insert a single record into the table using INSERT command. Use the first row from the table below:

video_id	added_date	title
1645ea59-14bd-11e5-a993-8138354b7e31	2014-01-29	Cassandra History
245e8024-14bd-11e5-9743-8238356b7e32	2012-04-03	Cassandra & SSDs
3452f7de-14bd-11e5-855e-8738355b7e3a	2013-03-17	Cassandra Intro
4845ed97-14bd-11e5-8a40-8338255b7e33	2013-10-16	DataStax DevCenter
5645f8bd-14bd-11e5-af1a-8638355b8e3a	2013-04-16	What is DataStax Enterprise?

```
INSERT INTO videos (video_id, added_date, title)  
VALUES (1645ea59-14bd-11e5-a993-8138354b7e31, '2014-01-29', 'Cassandra History');
```

7) Write a select statement to verify your record was inserted.

```
SELECT *  
FROM videos;
```

8) Insert the second record as well and run a select statement to verify it's there.

```
INSERT INTO videos (video_id, added_date, title)  
VALUES (245e8024-14bd-11e5-9743-8238356b7e32, '2012-04-03', 'Cassandra & SSDs');
```

NOTE: You should now see two records in your videos table.

9) Let's remove the data you inserted using the TRUNCATE command.

```
TRUNCATE videos;
```

10) Execute the following command to import data into your videos table.

```
COPY videos(video_id, added_date, title)
FROM '/home/ubuntu/labwork/data-files/videos.csv'
WITH HEADER=TRUE;
```

11) Use SELECT to verify the data loaded correctly.

```
SELECT *
FROM videos;
```

12) Use SELECT to COUNT(*) the number of imported rows. It should match the number of rows COPY reported as imported.

```
SELECT COUNT(*)
FROM videos;
```

13) To leave CQLSH execute this command:

```
QUIT
```