cqlsh Cassandra query shell

CONCEPTS

Statements end with a semicolon; shell commands such as DESCRIBE or CONSISTENCY do not need one.

Cluster nodes, partitioner
- **Keyspace** replication settings
  - **Table** partition key + clustering
    - **Row** columns

CONNECT

cqlsh *<host> <port>*
Connect to a node (default port 9042)

cqlsh -u *<user>* **-p** *<password> <host>*
Authenticate on connect

cqlsh -k *<keyspace> <host>*
Start in a keyspace

cqlsh --ssl *<host>*
Connect over TLS

cqlsh -e "*<statement>***"** *<host>*
Run one statement and quit

cqlsh -f *<file.cql> <host>*
Run a script file and quit

cqlsh --request-timeout=*<sec> <host>*
Raise the request timeout (default 10)

cqlsh --version
Print the cqlsh version

EXPLORE SCHEMA

DESCRIBE CLUSTER
Cluster name and partitioner

DESCRIBE KEYSPACES
List all keyspaces

DESCRIBE KEYSPACE *<keyspace>*
DDL of one keyspace

DESCRIBE TABLES
Tables in current keyspace

DESCRIBE TABLE *<table>*
DDL of a table

DESCRIBE SCHEMA
DDL for the whole schema

DESC *<object>*
Short form of DESCRIBE

SESSION CONTROL

USE *<keyspace>*;
Make a keyspace current

CONSISTENCY *<level>*
Set consistency level

CONSISTENCY
Show current consistency level

TRACING ON
Print a trace after each query

TRACING OFF
Stop tracing

EXPAND ON
One column per line

EXPAND OFF
Back to table output

PAGING ON
Page results with a prompt

PAGING *<rows>*
Set the page size

PAGING OFF
Print all rows at once

EXIT
Quit the shell

SERVER AND FILES

SHOW VERSION
cqlsh, Cassandra and protocol versions

SHOW HOST
Connected node and cluster name

SOURCE '*<file.cql>***'**
Run a CQL file from the prompt

CAPTURE '*<file>***'**
Append query output to a file

CAPTURE OFF
Stop capturing output

HELP *<topic>*
Built-in help on a command

BASIC CQL

CREATE KEYSPACE *<ks>* **WITH replication = {'class': 'SimpleStrategy', 'replication_factor': 1};**
Create a keyspace

CREATE TABLE *<table>* **(***<id>* **uuid PRIMARY KEY,** *<col>* **text);**
Create a table

BASIC CQL (cont.)

INSERT INTO *<table>* **(***<id>*, *<col>***) VALUES (***<id>*, **'***<val>***');**
Insert a row

SELECT * FROM *<table>* **WHERE** *<id>* **=** *<value>*;
Read by partition key

SELECT * FROM *<table>* **LIMIT** *<n>*;
Limit rows returned

UPDATE *<table>* **SET** *<col>* **= '***<val>***' WHERE** *<id>* **=** *<value>*;
Change a column

DELETE FROM *<table>* **WHERE** *<id>* **=** *<value>*;
Delete a row

TRUNCATE *<table>*;
Remove all rows

ALTER TABLE *<table>* **ADD** *<col> <type>*;
Add a column

DROP TABLE *<table>*;
Delete a table

COPY IMPORT AND EXPORT

COPY *<table>* **TO '***<file.csv>***' WITH HEADER = true;**
Export a table to CSV

COPY *<table>* **FROM '***<file.csv>***' WITH HEADER = true;**
Import CSV into a table

COPY *<table>* **(***<col1>*, *<col2>***) TO STDOUT;**
Export chosen columns

CONFIG

~/.cassandra/cqlshrc
Default config file

RECIPES

Debug a slow query

```
cqlsh <host> -k <keyspace>
TRACING ON
CONSISTENCY LOCAL_QUORUM
SELECT * FROM <table> WHERE <id>
= <value>;
# trace prints after the result
```

Script a schema over TLS

```
cqlsh --ssl -u <user> -p
<password> -f schema.cql <host>
```