

NDBI040: Big Data Management and NoSQL Databases

<http://www.ksi.mff.cuni.cz/~svoboda/courses/2016-1-NDBI040/>

Practical Class 9

Cassandra Column-Family Store

Martin Svoboda

svoboda@ksi.mff.cuni.cz

20. 12. 2016

Charles University in Prague, Faculty of Mathematics and Physics

Czech Technical University in Prague, Faculty of Electrical Engineering

Data Model

Database system structure

Instance → **keyspaces** → **tables** → **rows** → **columns**

- Keyspace
- Table (column family)
 - **Collection of (similar) rows**
 - Table schema must be specified, yet can be modified later on
- Row
 - **Collection of columns**
 - Rows in a table do not need to have the same columns
 - Each row is **uniquely identified** by a **primary key**
- Column
 - **Name-value pair** + additional data

Data Model

Column values

- Empty value
 - `null`
- Atomic value
 - **Native data types** such as texts, integers, dates, ...
 - **Tuples**
 - Tuple of anonymous fields, each of any type (even different)
 - **User defined types (UDT)**
 - Set of named fields of any type
- Collections
 - **Lists, sets, and maps**
 - Nested tuples, UDTs, or collections are allowed, but currently only in **frozen mode** (such elements are serialized when stored)

Query Language

CQL = Cassandra Query Language

- **DDL statements**

- CREATE KEYSPACE – creates a new keyspace
- CREATE TABLE – creates a new table
- ...

- **DML statements**

- SELECT – selects and projects rows from a single table
- INSERT – inserts rows into a table
- UPDATE – updates columns of rows in a table
- DELETE – removes rows from a table
- ...

Tutorial: Apache Cassandra



First Steps

Remotely connect to our NoSQL server

- SSH and SFTP access
- PuTTY and WinSCP on Windows
- **nosql.ms.mff.cuni.cz:42222**

Start CQLSH shell

- `cqlsh`

Basic useful commands

- **CLEAR**
 - Clear the shell terminal window
- **EXIT**
QUIT
 - Terminates the current database connection

Keyspace

Create your personal keyspace

```
CREATE KEYSPACE login
WITH
    replication = {'class': 'SimpleStrategy', 'replication_factor': 3}
```

- Use your login name as a name of your keyspace

List all existing keyspaces

- DESCRIBE KEYSPACES

Switch to your keyspace

- USE login

Tables

Create a new table for users

- Columns: integer identifier, first name, last name

View table definition

- `DESCRIBE TABLE users`

Insert new users into the table of users

- 1, Irena, Holubova
- 2, Martin, Svoboda

Browse existing users

- Find all users
- Find a specific user with identifier *1*

Filtering

Try to find a particular user according to their last name

- `lname = 'Holubova'`

Create a secondary index for last names

- `CREATE INDEX ON ...`

Try to find a particular user once again

- Enable filtering

Types

Create a user-defined type for names of people

- CREATE TYPE ...
- Fields: first, last

Create a new table for contacts

- Columns
 - id: integer identifier
 - name: first and last name
 - address: triple of street, city and ZIP code
 - emails: set of e-mail addresses
 - apps: list of preferred messenger applications
 - phones: map of phone numbers (work, home, ...)

Insertion

Insert new records into the table of contacts

- 1
Irena Holubova
Malostranske namesti, Praha, 11800
holubova@ksi.mff.cuni.cz
WhatsApp, Messenger
work +420951554316
- 2
Martin Svoboda
svoboda@ksi.mff.cuni.cz, martin.svoboda@mff.cuni.cz
Viber, WhatsApp
work +420951554250, fax +420951554323

Update

Modify existing contact records

- Replace certain columns of a person with id 1
 - Replace address: Malostranske namesti 25, Praha, 11800
 - Replace applications: Hangouts
- Modify certain columns of a person with id 1
 - Add new e-mail address: holubova@ksi.mff.cuni.cz
 - Add new applications: Messenger and WhatsApp
 - Add new phone number: home +420123456789
- Modify certain columns of a person with id 1
 - Remove e-mail address: irena.holubova@mff.cuni.cz
 - Remove applications: Hangouts and Messenger
 - Remove phone number: home

Deletion

Modify columns of existing contact records

- Remove / update certain columns of a person with id *1*
 - Remove address column
 - Remove the first application
 - Remove phone number to work

References

CQL – Cassandra Query Language

- <http://cassandra.apache.org/doc/latest/cql/>