

## Assignment 2 - HBase

1. Start HBase daemon

```
start-hbase.sh
```

2. Start HBase shell

```
hbase shell
```

3. Create a table called *Book* whose schema will be able to house book's title, description, author's first and last names. Book's title and description should be grouped as they will be saved/retrieved together. Author's first and last name should also be grouped (**hint:** since title and description need to be grouped together and so do author's first and last name, it would be wise to place them into 2 families such as *info* and *author*. Then *title* and *description* will become columns of *info* family and *first* and *last* columns of *author* family).

```
create 'Book', {NAME=>'info'}, {NAME=>'author'}
```

4. Add the following information to *Book* table:

ID	Title	Description	First name	Last name
1	Faster than the speed love	Long book about love	Brian	Dog
2	Long day	Story about Monday	Emily	Blue
3	Flying Car	Novel about airplanes	Phil	High

```

put 'Book', '1', 'info:title', 'Faster than the speed love'
put 'Book', '1', 'info:description', 'Long book about love'
put 'Book', '1', 'author:first', 'Brian'
put 'Book', '1', 'author:last', 'Dog'

put 'Book', '2', 'info:title', 'Long day'
put 'Book', '2', 'info:description', 'Story about Monday'
put 'Book', '2', 'author:first', 'Emily'
put 'Book', '2', 'author:last', 'Blue'

put 'Book', '3', 'info:title', 'Flying Car'
put 'Book', '3', 'info:description', 'Novel about airplanes'
put 'Book', '3', 'author:first', 'Phil'
put 'Book', '3', 'author:last', 'High'

```

- Count the number of rows. Make sure that every row is printed to the screen as it being counted.

```
count 'Book', INTERVAL => 1
```

- Retrieve an entire record with ID 1

```
get 'Book', '1'
```

- Only retrieve title and description for record with ID 3.

```
get 'Book', '3', {COLUMNS=>['info:title', 'info:description']}
```

- Change the last name of an author for the record with title *Long Day* to *Happy*.

- Display the record on the screen to verify the change.
- Display both new and old value. You should be able to see both *Blue* and *Happy*. Why is that?

```

put 'Book', '2', 'author:last', 'Happy'

# to verify select the record
get 'Book', '2', {COLUMNS=>'author:last'}

# to display both versions
get 'Book', '2', {COLUMNS=>'author:last', VERSIONS=>3}get 'Book', '3',
  {COLUMNS=>['info:title', 'info:description']}

```

9. Display all the records to the screen.

```
scan 'Book'
```

10. Display title and author's last name for all the records.

```
scan 'Book', {COLUMNS=>['info:title', 'author:last']}
```

11. Display title and description for the first 2 records.

```
scan 'Book', {COLUMNS=>['info:title', 'info:description'], LIMIT=>'2'}  
or  
scan 'Book', {COLUMNS=>['info:title', 'info:description'], STOPROW=>'3'}
```

12. Explore HBase Web-based management console, try and learn as much as you can about your new table.

Book table is hosted via 1 Region Server and there is only 1 Region. There are no start or end keys for that region because there is only 1 region. It has 2 families *info* and *author*. There is no compression set for both families, and replication is set to 3.

13. Check the *detailed* status of your cluster via HBase shell.

```
status 'detailed'
```

14. Delete a record whose title is *Flying Car*, and validate the record was deleted by scanning all the records or by attempting to select the record.

```
delete 'Book', '3', 'info:title'  
delete 'Book', '3', 'info:description'  
delete 'Book', '3', 'author:first'  
delete 'Book', '3', 'author:last'
```

15. Drop the table *Book*.

```
disable 'Book'  
drop 'Book'
```