Hadoop HDFS

Sang Shin
JPassion.com
"Learn with Passion!"

Topics

- What is and Why HDFS?
- HDFS Architecture
- HDFS Features
- HDFS Commands
- HDFS Web UI
- Hue web UI

What is and Why HDFS?

What is HDFS?

- HDFS is a virtual FS (File System) built on top of local FS
  - When you start writing data into HDFS, it eventually gets written onto the local FS (of distributed machines)
- You can't browse HDFS like you do with the local FS
  - You need to use the HDFS commands (similar to local FS commands, however) or
  - Or you can use the HDFS Web UI
  - Or the available APIs
- HDFS stores data as blocks in a replicated fashion
  - Management and replication of blocks are handled by HDFS
- HDFS is the primary distributed storage used by Hadoop applications
  - Scalability, Reliability, Automatic distribution of data

HDFS Architecture

HDFS Architecture

- For each Hadoop cluster, there is a single Name node, a single Secondary name node, and multiple data nodes
- Each block is replicated among data nodes

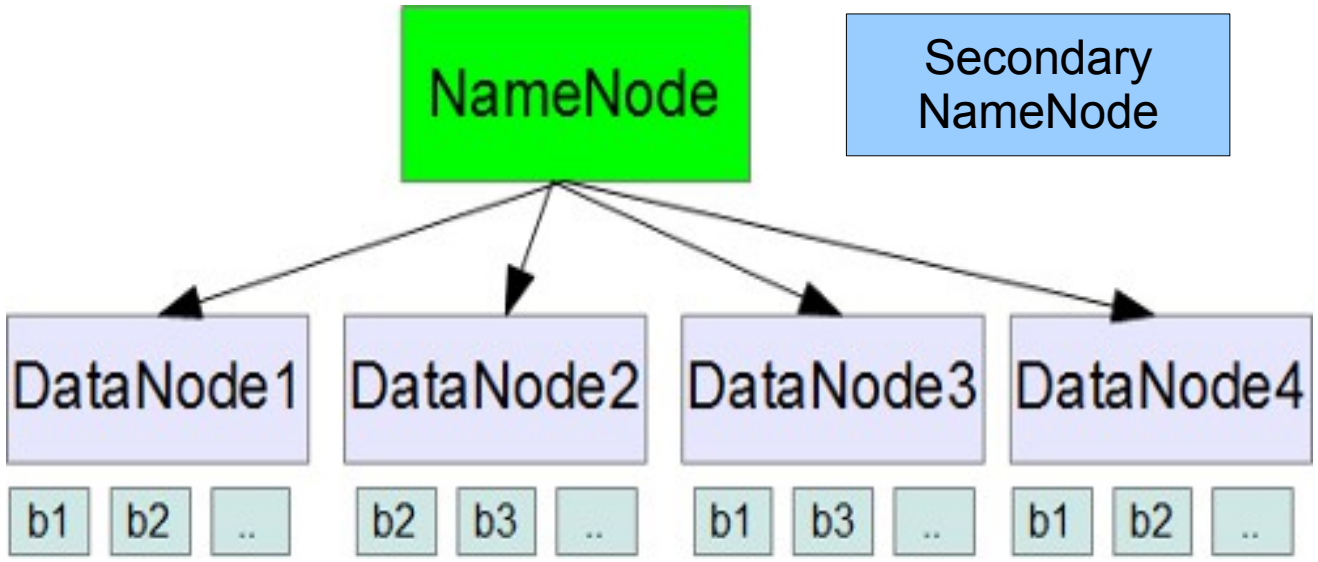

In the example above, block b1 is replicated in DataNode1, DataNode3, DataNode4

Nodes in HDFS Architecture

- Name node
  - Contains file system medadata
  - Monitors health of the data nodes
  - Clients communicate with Name node
- Data node
  - Handles client requests
  - Sends heartbeats to namenode (default is 3 seconds)
  - Rack aware – A set of data nodes can be located in a rack, blocks will be replicated across the racks
- Secondary name node
  - Not for high availability of Name node (as its name may sound like it)
  - Performs periodic checkpoints - Reads periodically the filesystem changes log of Name node and apply them into the Name node's metadata file

How Client Reads Data in HDFS

- When access to a particular block of data is required, the following occurs
  - > Metadata stored in the Name node is searched
  - > The location of that block on a particular datanode is returned to the client
  - > The client then reads data directly from there

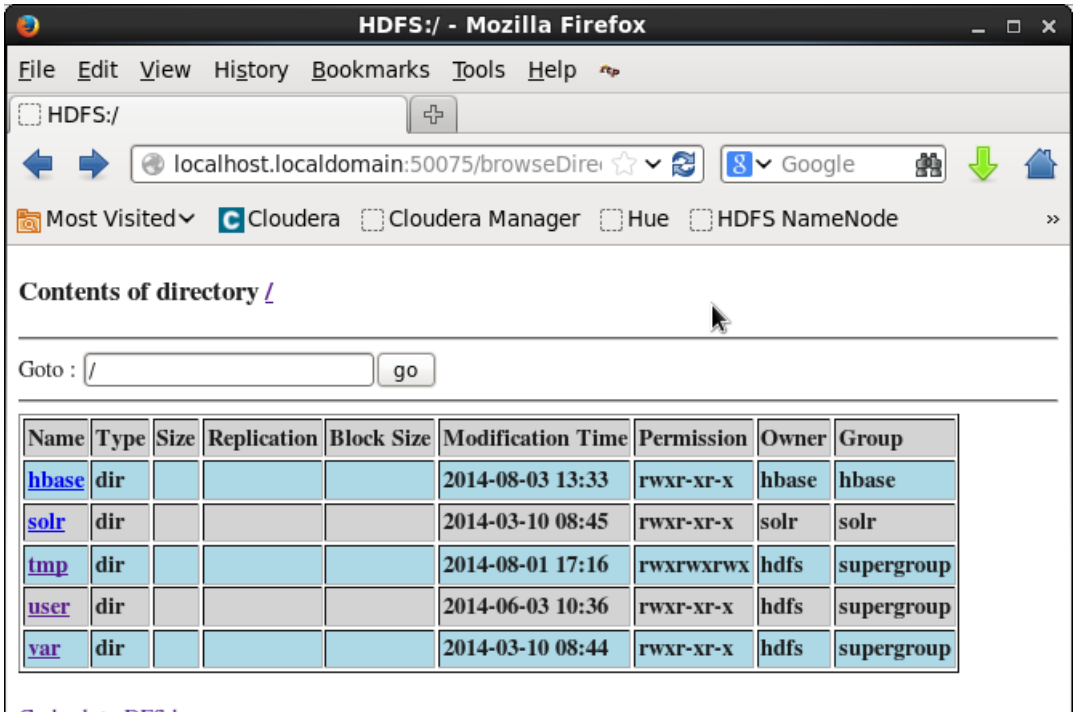

HDFS Architecture

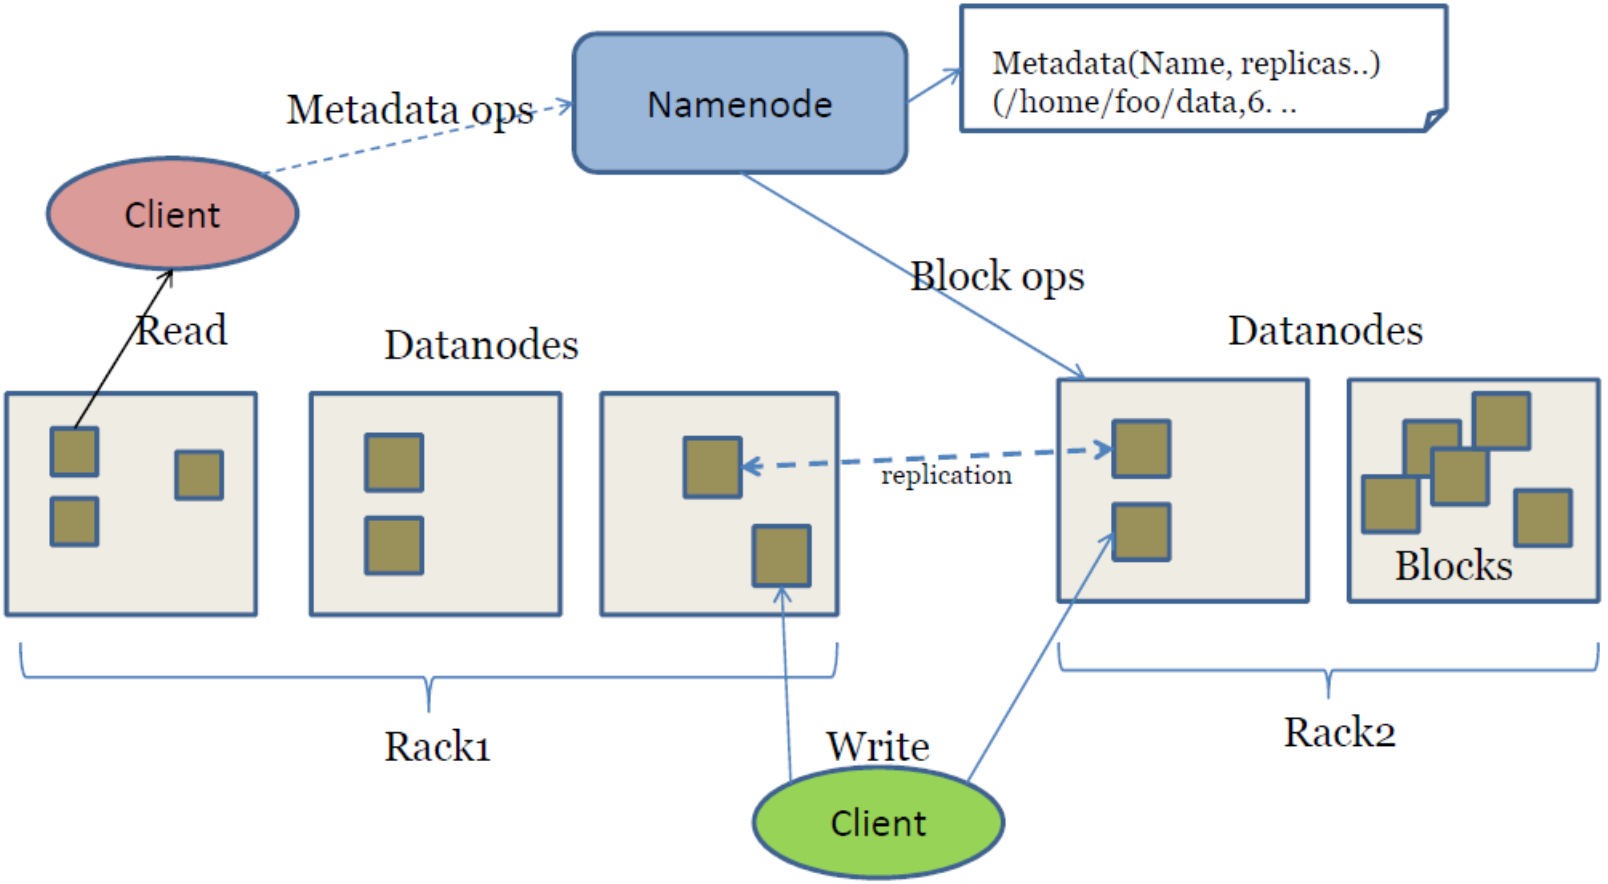

HDFS Internal Directories

- Directory where namenode stores its metadata
  - Set with default value ${hadoop.tmp.dir}/dfs/name
  - Can be specified explicitly by *dfs.name.dir*
- Directory where HDFS data blocks are stored
  - Set with default value ${hadoop.tmp.dir}/dfs/data
  - Can be specified explicitly by *dfs.data.dir*
- Directory where secondary namenode store its checkpoints
  - Set with default value is ${hadoop.tmp.dir}/dfs/namesecondary
  - Can be specified explicitly by *fs.checkpoint.dir*

HDFS Features

Blocks in Data nodes

- Blocks are 64M bytes (default)
  - In Cloudera QuickStart VM, it is set to 128M bytes
- HDFS handles
  - Block placement
  - Rebalancing
  - Replication management

Block Replication

- Blocks are replicated for reliability

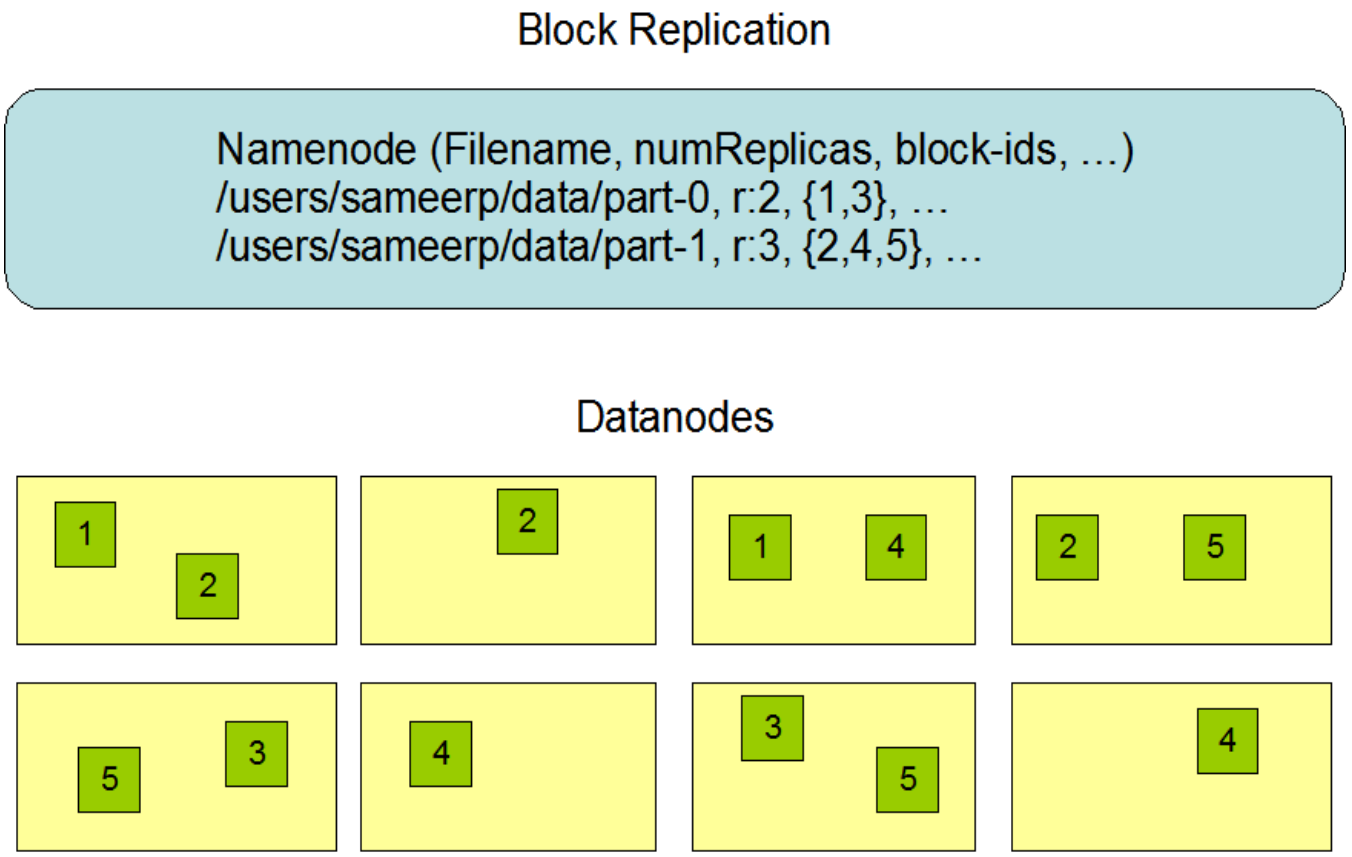

HDFS Data Distribution

- Number of data nodes a block is placed on is controlled by replication factor - If your replication is set to 3, it will be put on 3 separate data nodes

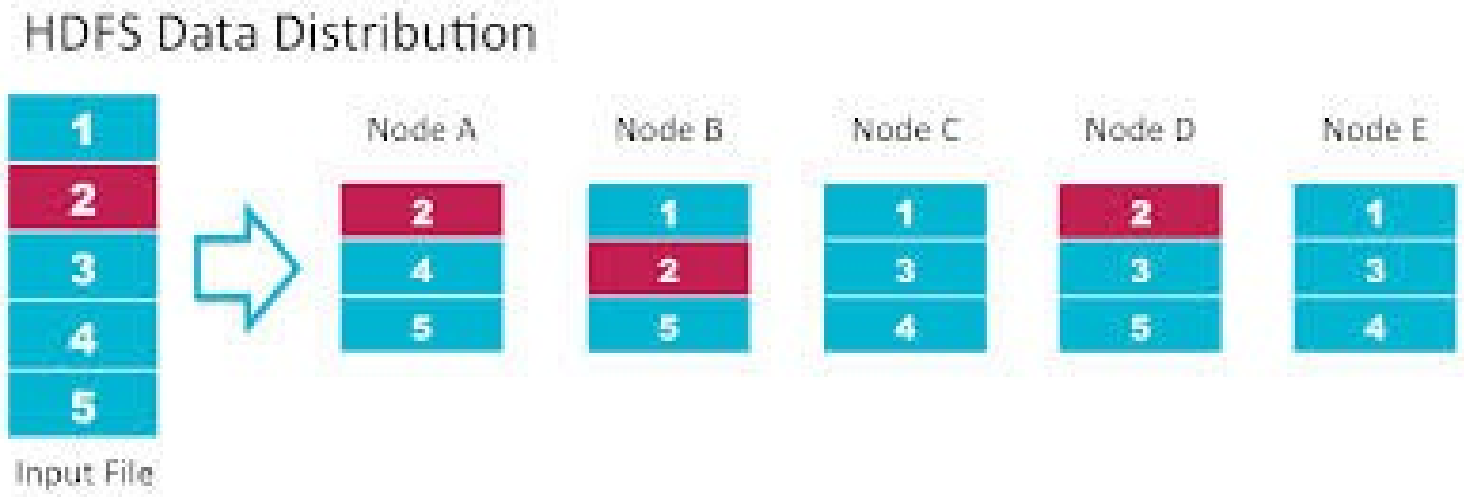

- Block is also placed across different racks
  - > If you set replication of 3 and have 2 racks, two blocks will be in one rack and the third block will be placed in the other rack

Lab:

Exercise 1: Study the HDFS Architecture
5906_hadoop_hdfs.zip

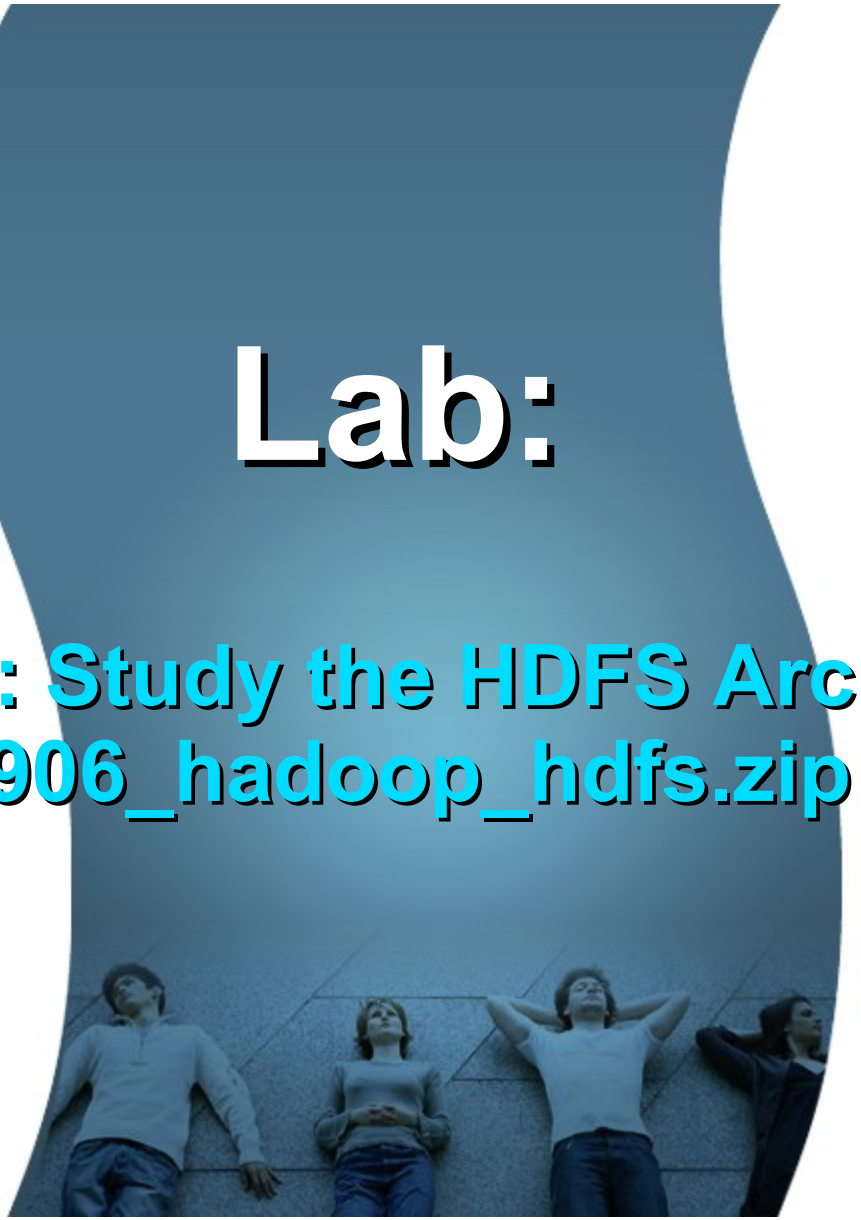

HDFS Commands

HDFS Commands

- http://hadoop.apache.org/docs/r2.4.1/hadoop-project-dist/hadoop-common/FileSystemShell.html

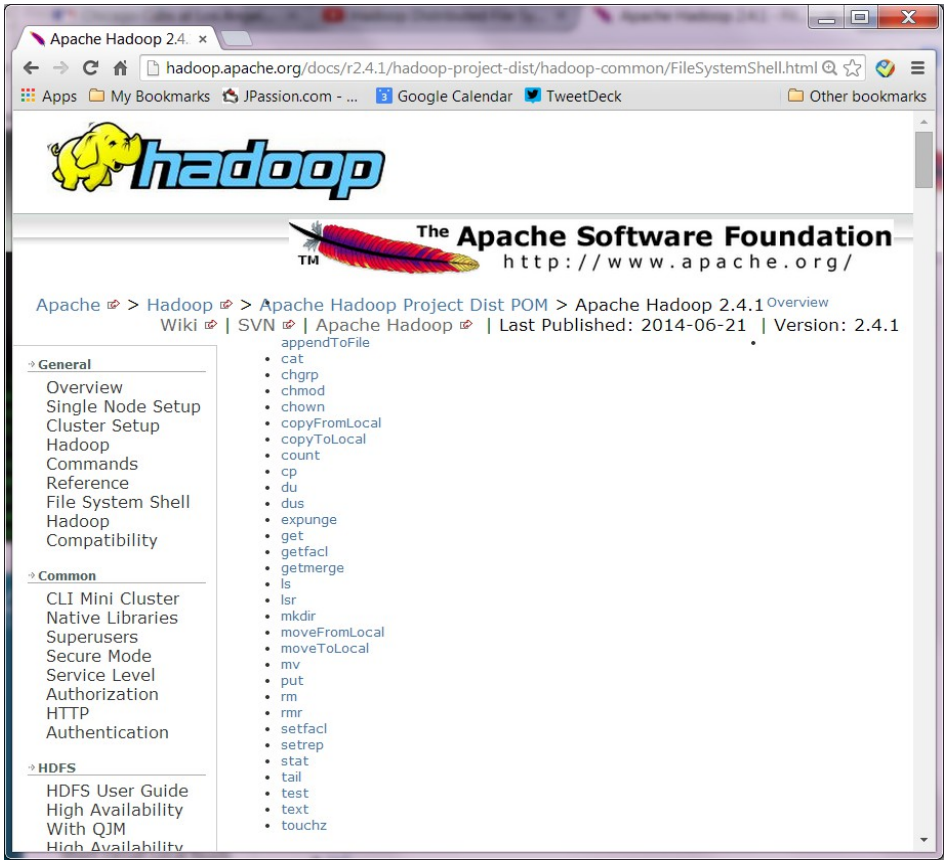

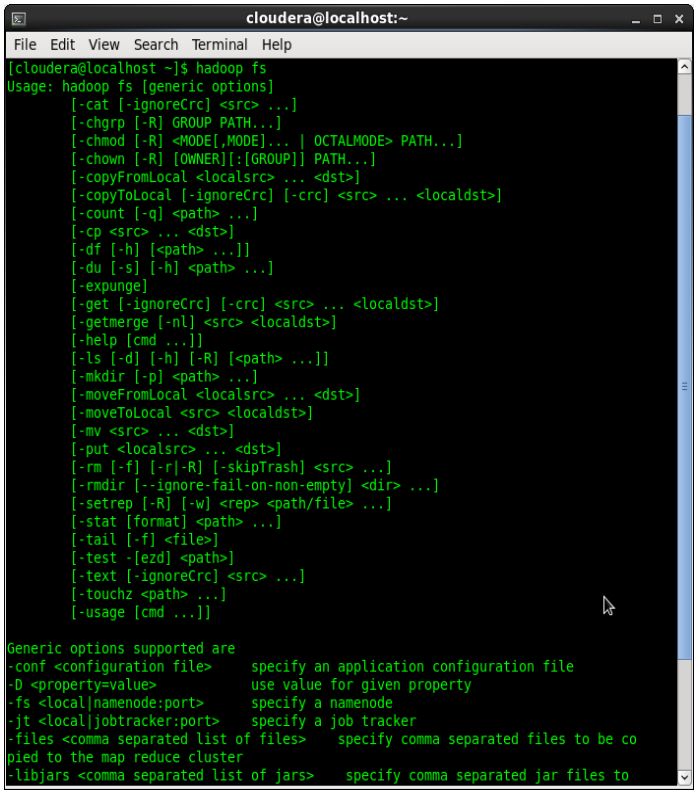

HDFS Commands

- Start with “hadoop fs -<command>
  - *hadoop fs -ls*
  - *hadoop fs -cat*
  - *hadoop fs -mkdir test_dir1*
  - *hadoop fs -rm -r test_dir1*
- Copying local directory/files to HDFS
  - *hadoop fs -copyFromLocal <local-directory/files> <HDFS-directory/files>*
- Copying HDFS directory/files to local file system
  - *hadoop fs -copyToLocal <HDFS-directory/files> <local-directory/files>*

Lab:

Exercise 2: Command-line HDFS Operations
5906_hadoop_hdfs.zip

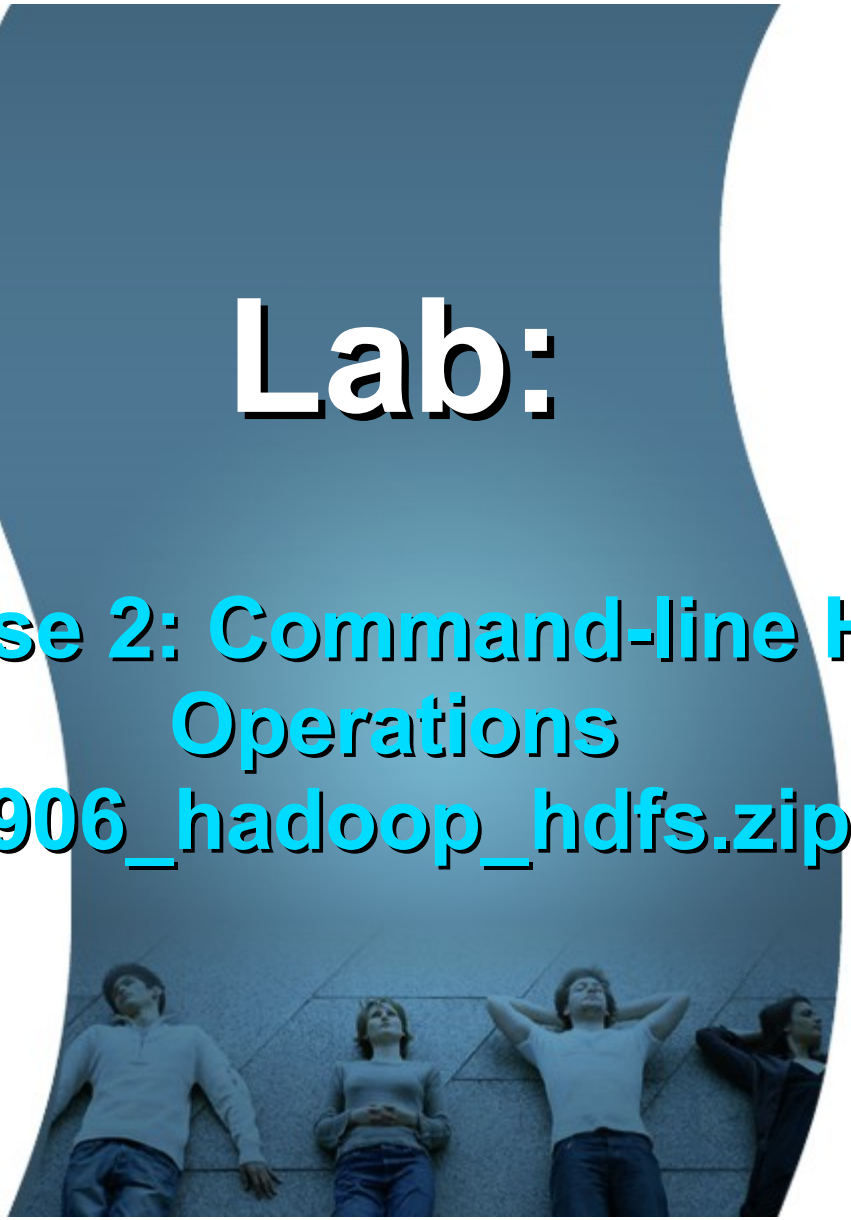

HDFS Admin Commands

HDFS Admin Commands

- http://hadoop.apache.org/docs/r1.0.4/commands_manual.html#Administration+Commands

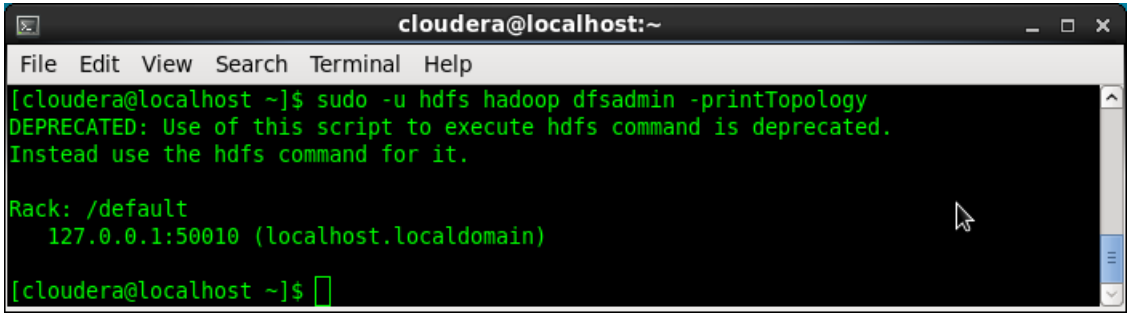

Lab:

Exercise 3: Command-line HDFS Admin Operations
5906_hadoop_hdfs.zip

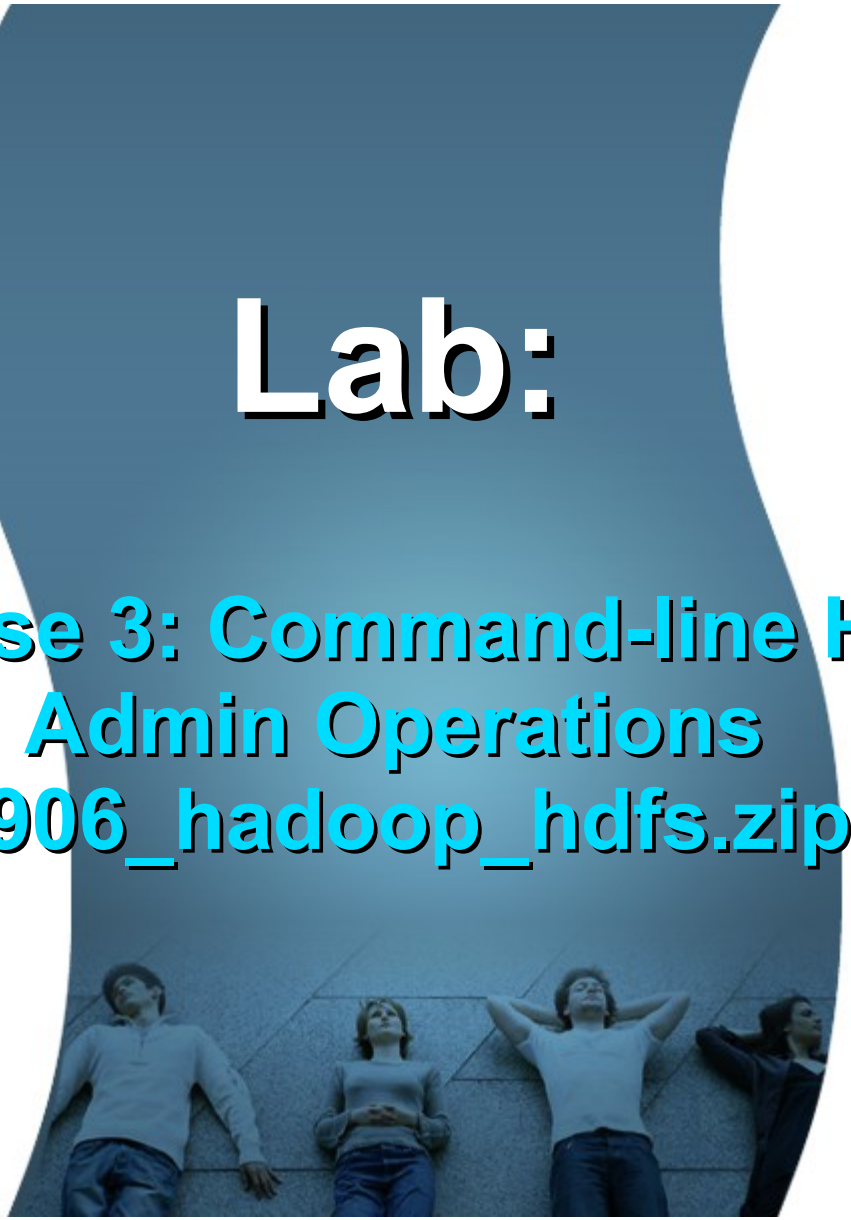

HDFS Web UI

HDFS Web UI

- NameNode and DataNode each run an internal web server in order to display basic information about the current status of the cluster.
  - With the default configuration, the NameNode front page is at http://namenode-name:50070/
  - It lists the DataNodes in the cluster and basic statistics of the cluster
- The web interface can also be used to browse the file system (using "Browse the file system" link on the NameNode front page).

HDFS Web UI

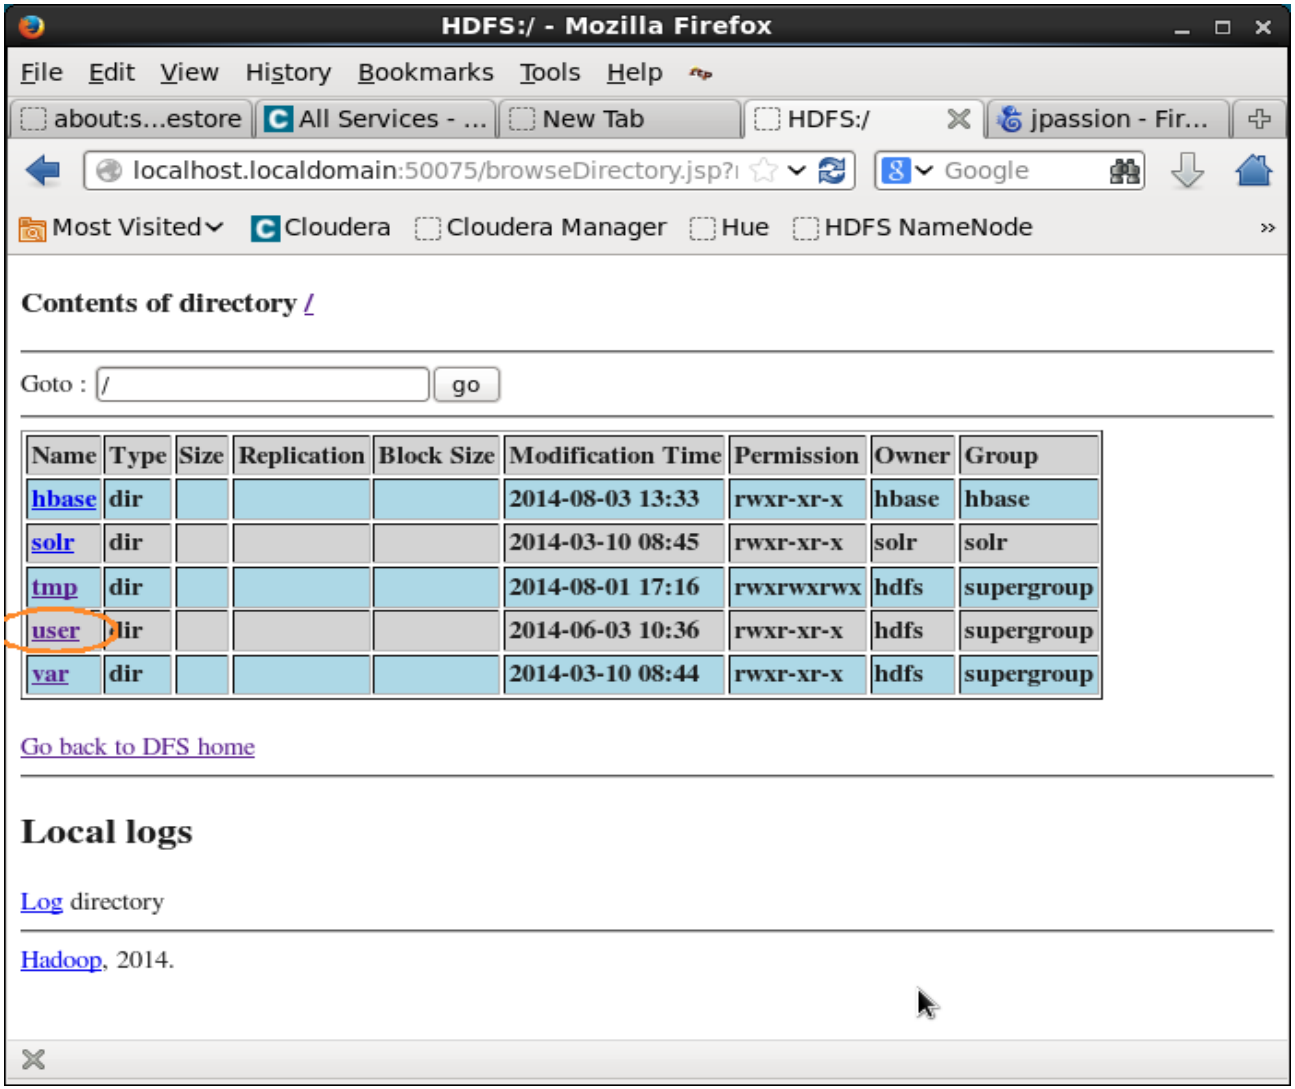
HDFS:/ - Mozilla Firefox
File Edit View History Bookmarks Tools Help
about:s...estore
All Services - ...
New Tab
HDFS:/
jpassion - Fir...
localhost.localdomain:50075/browseDirectory.jsp?
Google
Most Visited
Cloudera
Cloudera Manager
Hue
HDFS NameNode
Contents of directory /
Goto : /
go
Name Type Size Replication Block Size Modification Time Permission Owner Group
hbase dir 2014-08-03 13:33 rwxr-xr-x hbase hbase
solr dir 2014-03-10 08:45 rwxr-xr-x solr solr
tmp dir 2014-08-01 17:16 rwxrwxrwx hdfs supergroup
user dir 2014-06-03 10:36 rwxr-xr-x hdfs supergroup
var dir 2014-03-10 08:44 rwxr-xr-x hdfs supergroup
Go back to DFS home
Local logs
Log directory
Hadoop, 2014.

Hue Web UI

What is Hue?

- Hue is a Web interface for all Hadoop related operations

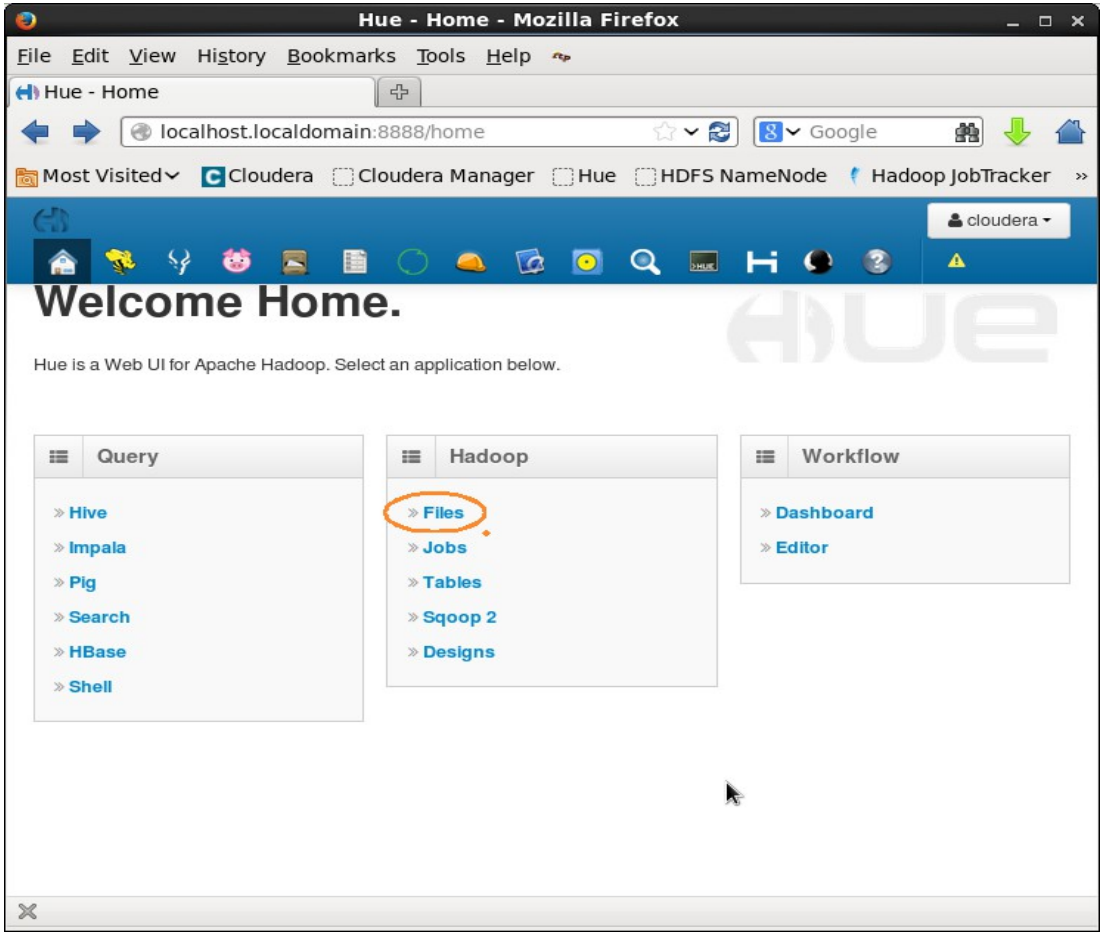

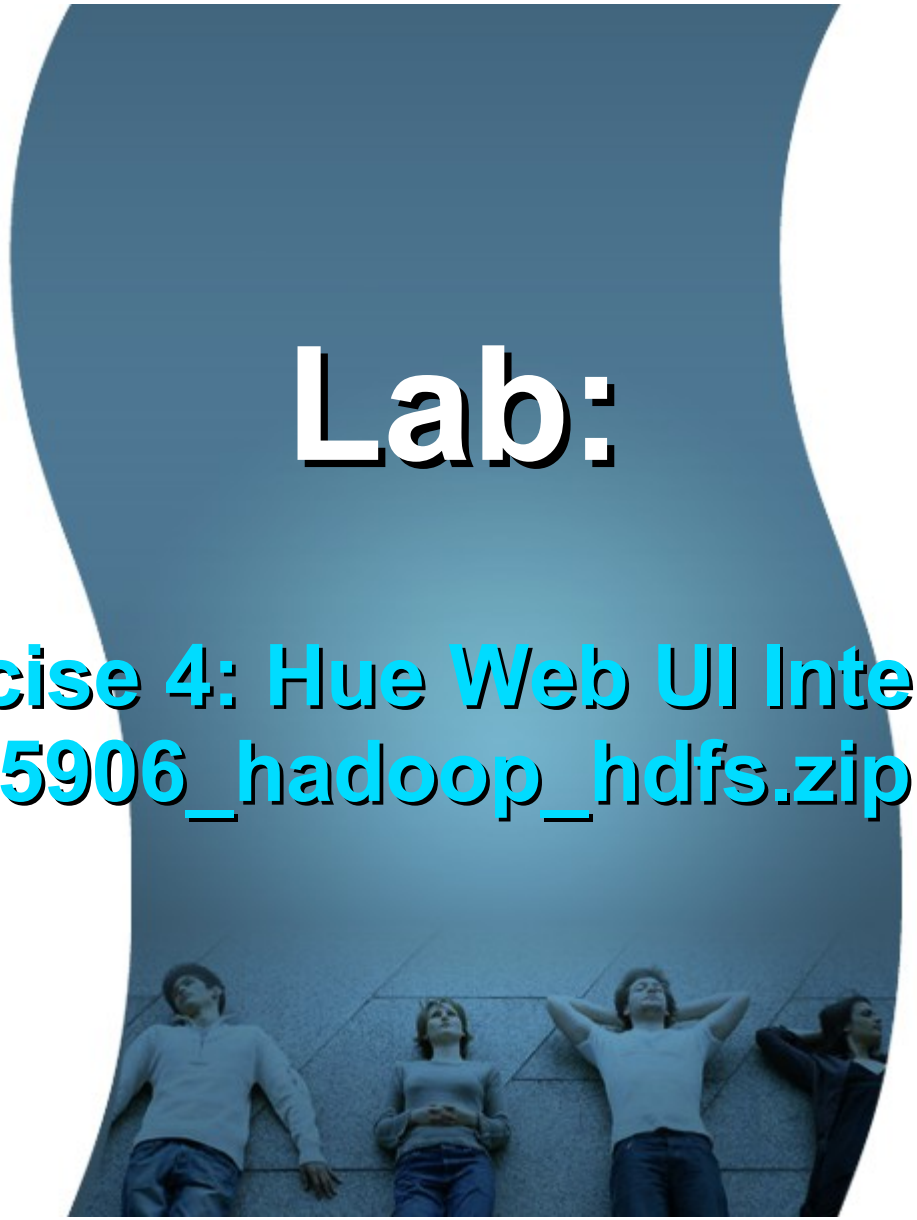

Lab:

Exercise 4: Hue Web UI Interface
5906_hadoop_hdfs.zip

Learn with Passion!
JPassion.com