



101

Ahbaid Gaffoor

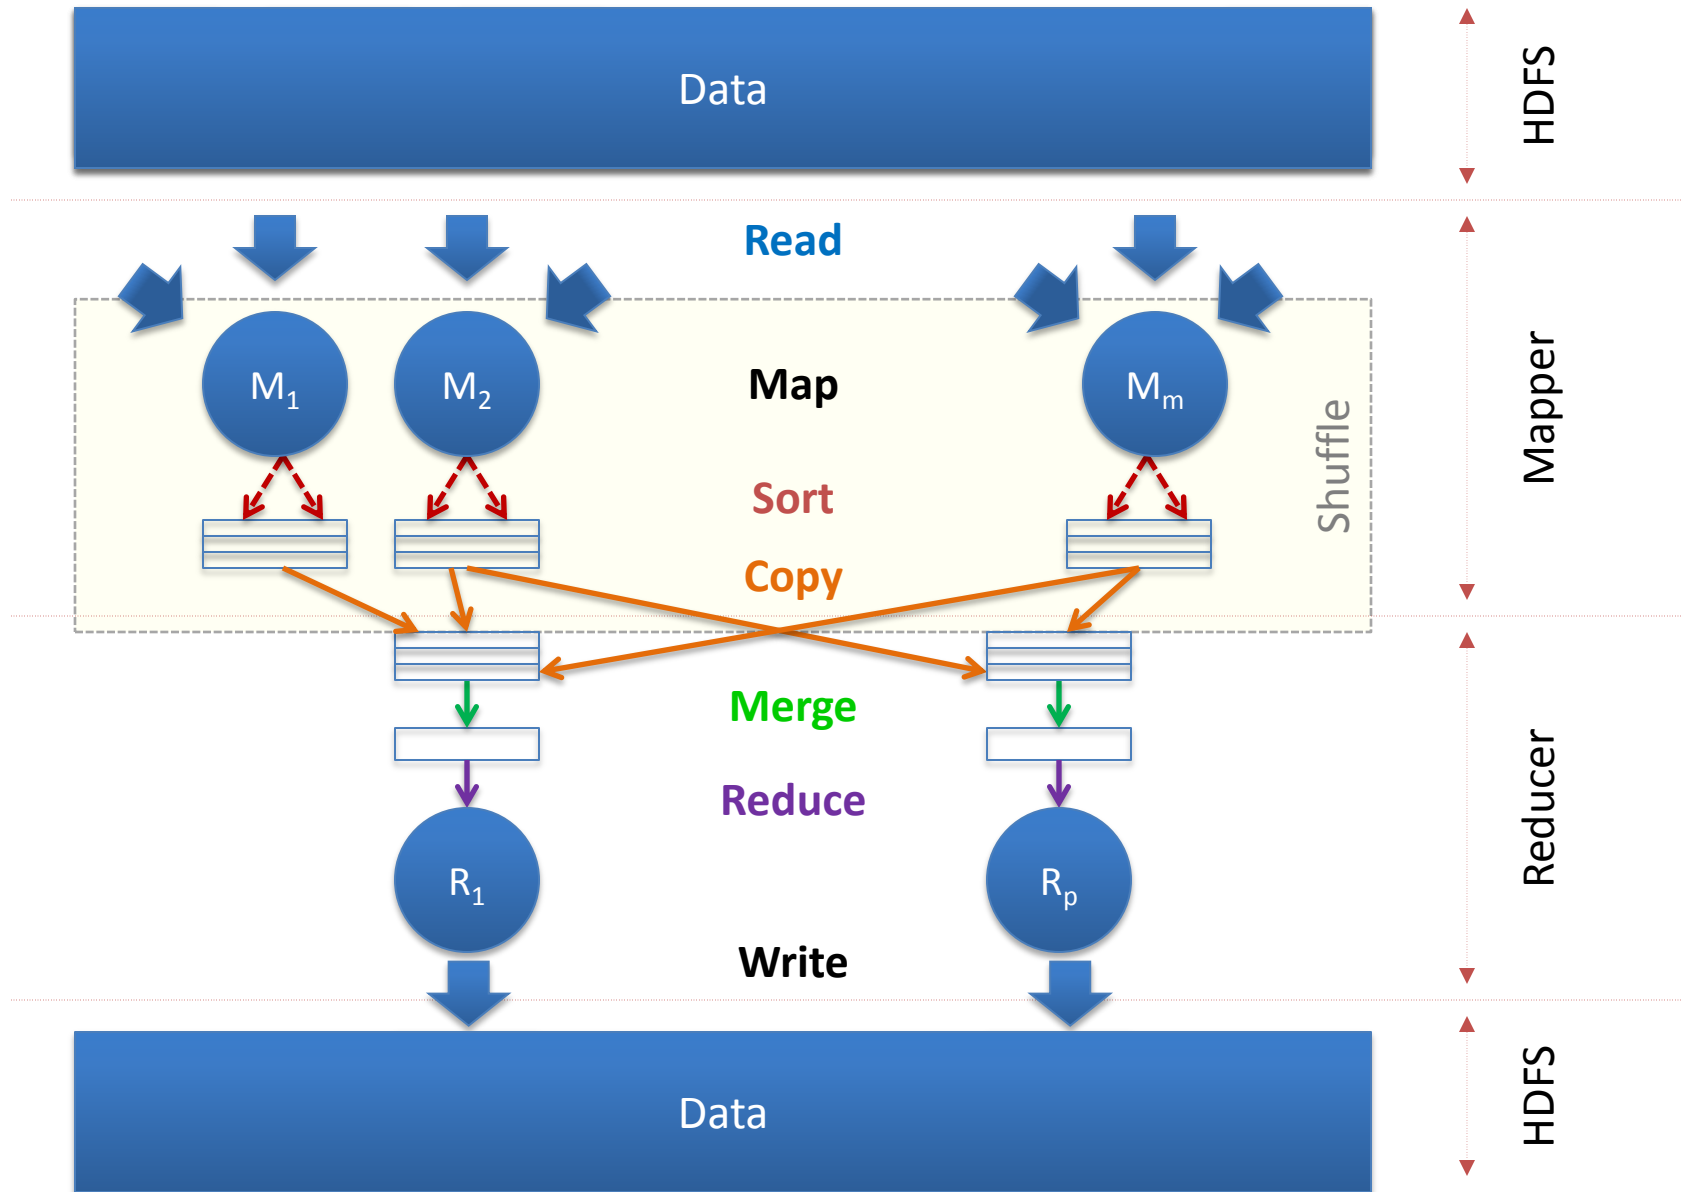
Updated Book List added!
05/23/2013 - Slide 34.

Overview

1. Map Reduce
2. HDFS
3. Hadoop Versions
4. Hadoop Installation
5. Word Count
6. Hadoop Web Interfaces




Map Reduce – An Overview

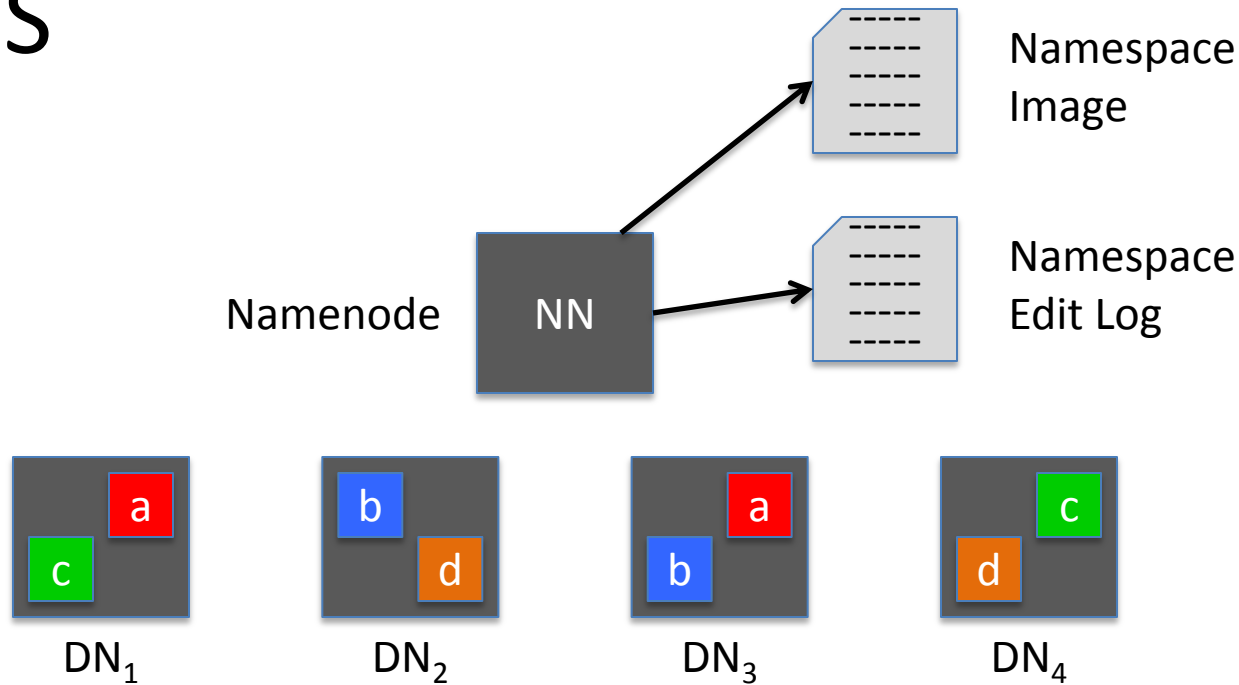


Map Reduce – Job / Task Tracker

- Job Tracker
 - MapReduce jobs submitted by clients to Job Tracker
 - Job Tracker divides and allocates work
 - Data locality used in Task node selection
 - Re-tries and re-assigns failed jobs a configurable no. of times
 - Can re-assign jobs to different task trackers
- Task Tracker
 - One per Task Tracker node / Datanode
 - Executes Job Tracker requests



HDFS



- Namenode
 - Knows datanodes on which file blocks exist
 - Does not know location of blocks on datanodes
 - Loss renders datanodes useless
- Datanodes
 - Track physical block location
 - Storage & Retrieval requests from Namenode and clients

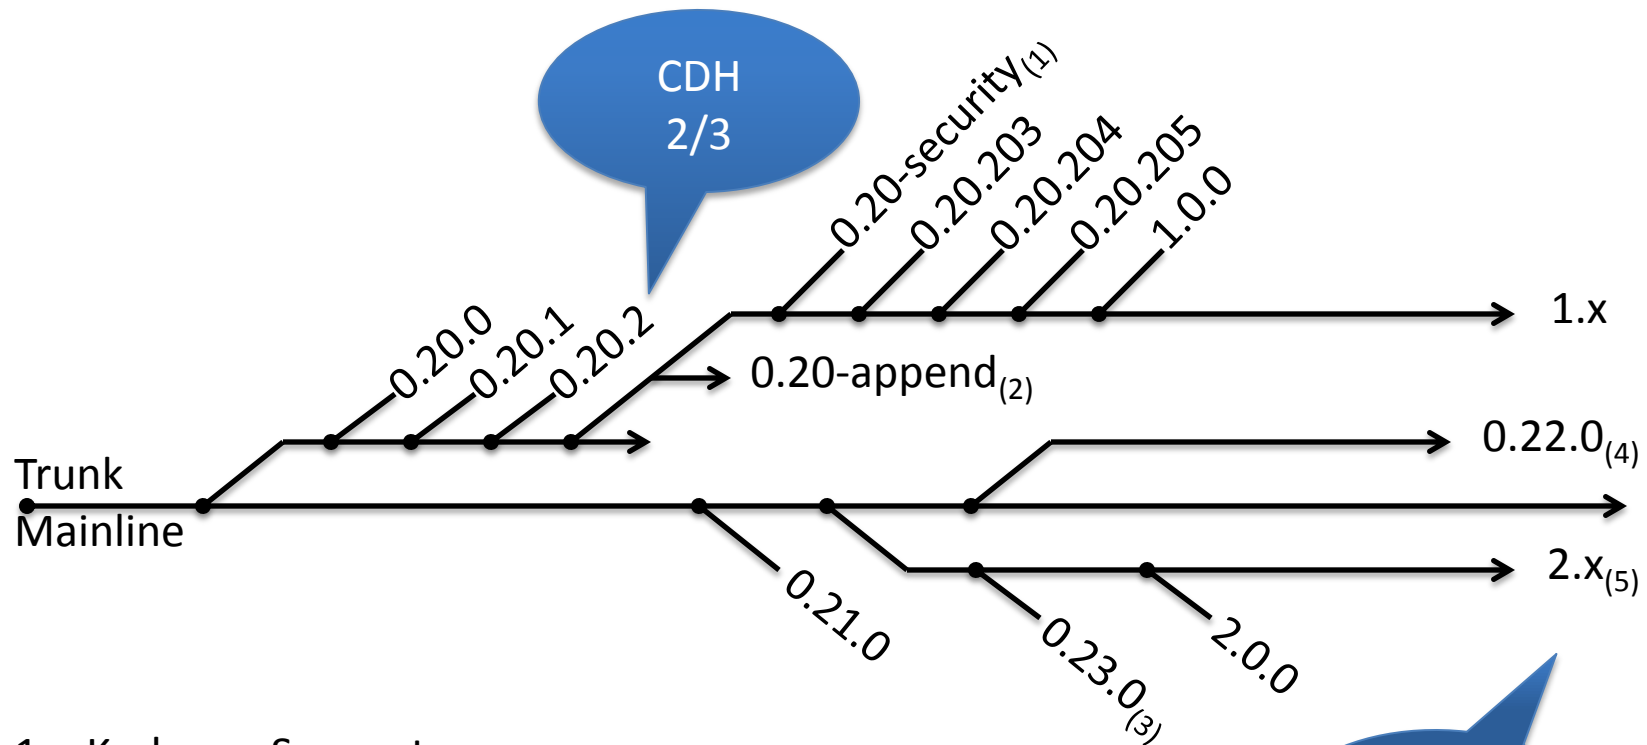
HDFS

- Hadoop Distributed File System
 - Write once, read many optimal
 - Minimal seek cost
- Blocks
 - Default Block Size of 64M
 - Seek time to transfer time of 1%
 - Allow for replication (3 by default)
 - Replication Factor
 - Rack awareness
 - Block locality
- JBOD vs. RAID
 - JBOD, replication done by HDFS
 - RAID as fast as slowest disk
 - JBOD average speed is faster than slowest disk

HDFS

- Master Namenode
 - Namespace image
 - Namespace edit log
- Namenode Protection / Resilience
 - Configure Namenode to write persistent state to multiple filesystems
 - Include an NFS and local FS
 - Writes are synchronous and atomic
- Secondary Namenodes
 - Periodic merge of namespace image & namespace edit log
 - Usually runs on a separate host
 - A copy of the merged namespace image is kept

Hadoop Versions



1. Kerberos Support
2. HDFS Append
3. Nov 2011 – v0.23
4. Dec 2011 – v0.22
5. MRV2, YARN
6. CDH2/CDH3 0.20 + (0.21 & 0.22)

CDH
4

Hadoop Versions

Feature / Version	0.20	0.21	0.22	0.23	1.0	2.0	CDH3	CDH4	
Production Ready	X				X		X	X	Hadoop 1.0 CDH2/CDH3
HDFS Append		X	X	X	X	X	X	X	
HDFS Symlink		X	X	X		X		X	
Kerberos		X	X	X	X	X	X	X	
MRV1	X	X	X		X		X	X	
MRV2/YARN				X			X	X	Hadoop 2.0 CDH4
Namenode Federation							X	X	
Namenode HA							X	X	

Hadoop Installation –Processes

- HDFS
 - Namenode
 - Secondary Namenode(s) (if configured)
- MapReduce Jobs
 - Job Tracker
 - Task Tracker

Hadoop Installation – Config Files

- `$HADOOP_HOME` (for reference)
- `$HADOOP_HOME/conf/hadoop-env.sh`
 - Environment variables affecting the JDK
- **`$HADOOP_HOME/conf/core-site.xml`**
 - Parameters related to all Hadoop daemons and clients
- **`$HADOOP_HOME/conf/hdfs-site.xml`**
 - Parameters used by HDFS daemons and clients
- **`$HADOOP_HOME/conf/mapred-site.xml`**
 - Parameters used by MapReduce daemons and clients
- `$HADOOP_HOME/conf/log4j.properties`
 - Java property file
 - All logging configuration information
- **`$HADOOP_HOME/conf/masters`**
 - Hosts that run a secondary Namenode (**optional, misleading name**)
 - Used by “start-*.sh” scripts
- **`$HADOOP_HOME/conf/slaves`**
 - Hosts running datanode / tasktracker daemons
 - Used by “start-*.sh” scripts
- There are more files!

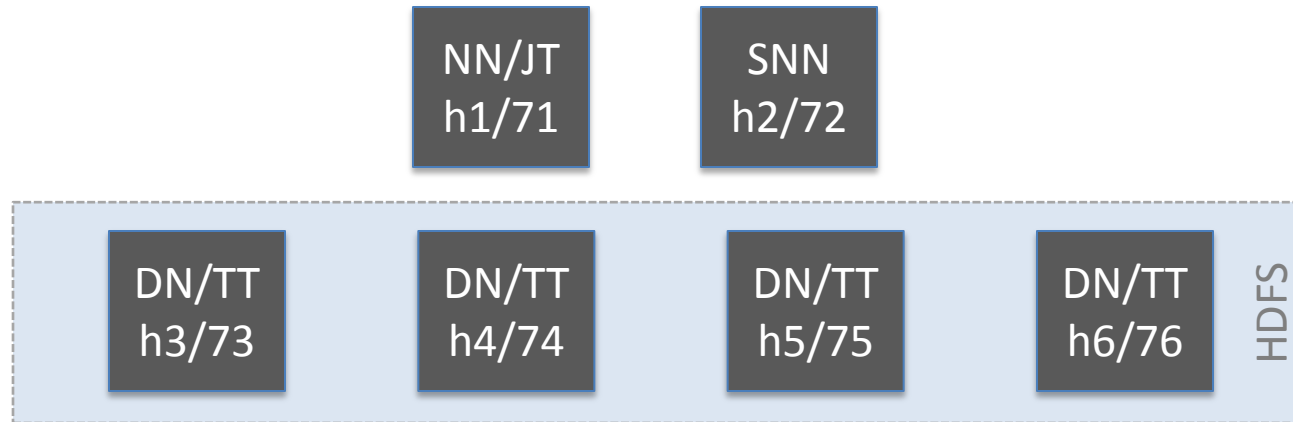
Hadoop Installation

- Software
 - Java 6 (Oracle)
 - Apache Hadoop 1.0.4
 - Centos Linux 6.2 (2.6 Kernel)
 - Download `hadoop-1.0.4-bin.tar.gz` (hadoop.apache.org)
 - `cd <wdir>; tar -zxvf hadoop-1.0.4-bin.tar.gz`
- Configure `$HADOOP_HOME/conf/core-site.xml`
 - `fs.default.name` (Set the Master Namenode host and port)
- Configure `$HADOOP_HOME/conf/mapred-site.xml`
 - `mapred.job.tracker` (Set the Job Tracker host and port)
- Configure `$HADOOP_HOME/conf/hdfs-site.xml`
 - `dfs.name.dir` (Namenode fs location)
 - `dfs.data.dir` (Dir used for HDFS on each node)
- Configure `$HADOOP_HOME/conf/slaves`
- Configure `$HADOOP_HOME/conf/masters` (optional)

Hadoop Installation – Lab Layout

- Hardware
 - Intel i7-3930K (6 cores, 12 virtual cpus) - \$700
 - Motherboard - \$200
 - 64G RAM - \$400
 - 13TB SAS Array (RAID 10 of 3TB SATA drives)
 - Controller Card - \$700
 - Disks - \$400
 - Total ~ \$2400
- Software
 - Windows 7 64-bit OS
 - Oracle Virtual Box
- Virtual Hosts
 - 6 hosts (h1 – h6)
 - Bridged NIC, Static IPs 199.1.1.[71-76]
 - 8G root “/”, 36G “/h” for hadoop
 - “/h/hdfs” for HDFS
 - 2G RAM
 - 2 CPUs

Hadoop Installation – Lab Layout



- NN: Namenode 199.1.1.71
- JT: JobTracker 199.1.1.71
- SNN: Secondary Namenode 199.1.1.72
- DN: Datanode 199.1.1.[73-76] (HDFS)
- TT: TaskTracker 199.1.1.[73-76]

Hadoop Installation – Sample Cfg. Files

- \$HADOOP_HOME/conf/core-site.xml

```
<property>
  <name>fs.default.name</name>
  <value>hdfs://h1:9000/</value>
</property>
```

- \$HADOOP_HOME/conf/mapred-site.xml

```
<property>
  <name>mapred.job.tracker</name>
  <value>h1:9001</value>
</property>
```

Hadoop Installation – Sample Cfg. Files

- \$HADOOP_HOME/conf/hdfs-site.xml

```
<property>
  <name>dfs.name.dir</name>
  <value>/h/hdfs/current/name</value>
</property>
<property>
  <name>dfs.data.dir</name>
  <value>/h/hdfs/current/data</value>
</property>
<property>
  <name>dfs.replication</name>
  <value>2</value>
</property>
```


Hadoop Installation – Control Scripts

- `$HADOOP_HOME/bin/start-dfs.sh`
 - Starts namenode on the local host.
 - Starts secondary namenodes on hosts in “masters” file.
 - Starts a datanode on each host in the “slaves” file.
- `$HADOOP_HOME/bin/start-mapred.sh`
 - Starts a Jobtracker on the local host.
 - Starts a tasktracker on each host in the “slaves” file.
- `$HADOOP_HOME/bin/stop-dfs.sh`
- `$HADOOP_HOME/bin/stop-mapred.sh`
- `$HADOOP_HOME/bin/start-all.sh`
- `$HADOOP_HOME/bin/stop-all.sh`

Hadoop Installation – Format NN

- “\$HADOOP_HOME/bin/hadoop namenode -format” (master NN only)

```
[hadoop@h1 ~]$ hadoop namenode -format
13/05/20 18:35:32 INFO namenode.NameNode: STARTUP_MSG:
/*****
STARTUP_MSG: Starting NameNode
STARTUP_MSG: host = h1.ilmtech.com/199.1.1.71
STARTUP_MSG: args = [-format]
STARTUP_MSG: version = 1.0.4
STARTUP_MSG: build = https://svn.apache.org/repos/asf/hadoop/common/branches/branch-1.0 -r 1393290; compiled by 'hortonfo' on Wed Oct 3 05:13:58
UTC 2012
*****/
13/05/20 18:35:33 INFO util.GSet: VM type    = 64-bit
13/05/20 18:35:33 INFO util.GSet: 2% max memory = 17.77875 MB
13/05/20 18:35:33 INFO util.GSet: capacity    = 2^21 = 2097152 entries
13/05/20 18:35:33 INFO util.GSet: recommended=2097152, actual=2097152
13/05/20 18:35:33 INFO namenode.FSNamesystem: fsOwner=hadoop
13/05/20 18:35:33 INFO namenode.FSNamesystem: supergroup=supergroup
13/05/20 18:35:33 INFO namenode.FSNamesystem: isPermissionEnabled=true
13/05/20 18:35:33 INFO namenode.FSNamesystem: dfs.block.invalidate.limit=100
13/05/20 18:35:33 INFO namenode.FSNamesystem: isAccessTokenEnabled=false accessKeyUpdateInterval=0 min(s), accessTokenLifetime=0 min(s)
13/05/20 18:35:33 INFO namenode.NameNode: Caching file names occuring more than 10 times
13/05/20 18:35:33 INFO common.Storage: Image file of size 112 saved in 0 seconds.
13/05/20 18:35:33 INFO common.Storage: Storage directory /h/hdfs/current/name has been successfully formatted.
13/05/20 18:35:33 INFO namenode.NameNode: SHUTDOWN_MSG:
/*****
SHUTDOWN_MSG: Shutting down NameNode at h1.ilmtech.com/199.1.1.71
*****/
```

Hadoop Installation – hdfs Startup

- “\$HADOOP_HOME/bin/start-dfs.sh” (master NN only)

```
[hadoop@h1 bin]$ ./start-dfs.sh
starting namenode, logging to /h/hadoop/hadoop-1.0.4/libexec/../logs/hadoop-hadoop-namenode-h1.ilmtech.com.out
h3: starting datanode, logging to /h/hadoop/hadoop-1.0.4/libexec/../logs/hadoop-hadoop-datanode-h3.ilmtech.com.out
h6: starting datanode, logging to /h/hadoop/hadoop-1.0.4/libexec/../logs/hadoop-hadoop-datanode-h6.ilmtech.com.out
h5: starting datanode, logging to /h/hadoop/hadoop-1.0.4/libexec/../logs/hadoop-hadoop-datanode-h5.ilmtech.com.out
h4: starting datanode, logging to /h/hadoop/hadoop-1.0.4/libexec/../logs/hadoop-hadoop-datanode-h4.ilmtech.com.out
h2: starting secondarynamenode, logging to /h/hadoop/hadoop-1.0.4/libexec/../logs/hadoop-hadoop-secondarynamenode-h2.ilmtech.com.out
```

```
[hadoop@h1 bin]$ jpsall
Checking h1....
13232 NameNode
Checking h2....
11542 SecondaryNameNode
Checking h3....
11697 DataNode
Checking h4....
11653 DataNode
Checking h5....
11655 DataNode
Checking h6....
11610 DataNode
```

- “\$HADOOP_HOME/bin/jpsall” (Helper Script)

```
[hadoop@h1 bin]$ cat jpsall
for h in `seq 1 6`; do
    echo "Checking h$h...."
    ssh h$h jps | grep -iv jps
done
```

Hadoop Installation – Job/Task Startup

- “\$HADOOP_HOME/bin/start-mapred.sh” (master NN only)

```
[hadoop@h1 bin]$ ./start-mapred.sh
starting jobtracker, logging to /h/hadoop/hadoop-1.0.4/libexec/./logs/hadoop-hadoop-jobtracker-h1.ilmtech.com.out
h3: starting tasktracker, logging to /h/hadoop/hadoop-1.0.4/libexec/./logs/hadoop-hadoop-tasktracker-h3.ilmtech.com.out
h4: starting tasktracker, logging to /h/hadoop/hadoop-1.0.4/libexec/./logs/hadoop-hadoop-tasktracker-h4.ilmtech.com.out
h5: starting tasktracker, logging to /h/hadoop/hadoop-1.0.4/libexec/./logs/hadoop-hadoop-tasktracker-h5.ilmtech.com.out
h6: starting tasktracker, logging to /h/hadoop/hadoop-1.0.4/libexec/./logs/hadoop-hadoop-tasktracker-h6.ilmtech.com.out

[hadoop@h1 bin]$ jpsall
Checking h1....
13463 JobTracker
13232 NameNode
Checking h2....
11542 SecondaryNameNode
Checking h3....
11697 DataNode
11834 TaskTracker
Checking h4....
11791 TaskTracker
11653 DataNode
Checking h5....
11791 TaskTracker
11655 DataNode
Checking h6....
11748 TaskTracker
11610 DataNode
```

Word Count – “hadoop fs”

- The Count of Monte Cristo - <http://www.gutenberg.org/ebooks/1184.txt.utf-8>
- Word <count>
- Get the file

```
[hadoop@h1 testdata]$ pwd
/home/hadoop/testdata
[hadoop@h1 testdata]$ wget http://www.gutenberg.org/ebooks/1184.txt.utf-8
hadoop@h1 testdata]$ mv 1184.txt.utf8 CountOfMonteCristo.txt
[hadoop@h1 testdata]$ ls -alh CountOfMonteCristo.txt
-rw-rw-r--. 1 hadoop hadoop 2.6M May 20 19:25 CountOfMonteCristo.txt
```

- “\$HADOOP_HOME/bin/hadoop fs”

```
$HADOOP_HOME/bin/hadoop fs -help

hadoop fs [-fs <local | file system URI>] [-conf <configuration file>]
  [-D <property=value>] [-ls <path>] [-lsr <path>] [-du <path>]
  [-dus <path>] [-mv <src> <dst>] [-cp <src> <dst>] [-rm [-skipTrash] <src>]
  [-rmr [-skipTrash] <src>] [-put <localsrc> ... <dst>] [-copyFromLocal <localsrc> ... <dst>]
  [-moveFromLocal <localsrc> ... <dst>] [-get [-ignoreCrc] [-crc] <src> <localdst>]
  [-getmerge <src> <localdst> [addnl]] [-cat <src>]
  [-copyToLocal [-ignoreCrc] [-crc] <src> <localdst>] [-moveToLocal <src> <localdst>]
  [-mkdir <path>] [-report] [-setrep [-R] [-w] <rep> <path/file>]
  [-touchz <path>] [-test [-ezd] <path>] [-stat [format] <path>]
  [-tail [-f] <path>] [-text <path>]
  [-chmod [-R] <MODE[,MODE]... | OCTALMODE> PATH...]
  [-chown [-R] [OWNER][:[GROUP]] PATH...]
  [-chgrp [-R] GROUP PATH...]
  [-count[-q] <path>]
  [-help [cmd]]
```

```
[hadoop@h1]$ alias hdfs="hadoop fs"
```

Word Count – “hadoop fs”

- Make a directory in HDFS

```
[hadoop@h1 testdata]$ hdfs -ls /  
Found 1 items  
drwxr-xr-x - hadoop supergroup      0 2013-05-20 18:38 /tmp  
  
[hadoop@h1 testdata]$ hdfs -mkdir /data  
  
[hadoop@h1 testdata]$ hdfs -ls /  
Found 2 items  
drwxr-xr-x - hadoop supergroup      0 2013-05-20 21:02 /data  
drwxr-xr-x - hadoop supergroup      0 2013-05-20 18:38 /tmp
```

- Upload the file from a local fs to HDFS

```
[hadoop@h1 testdata]$ pwd; ls -alh CountOfMonteCristo.txt  
/home/hadoop/testdata  
-rw-rw-r--. 1 hadoop hadoop 2.6M May 20 19:25 CountOfMonteCristo.txt  
  
[hadoop@h1 testdata]$ hdfs -put ./CountOfMonteCristo.txt /data/book.txt  
  
[hadoop@h1 testdata]$ hdfs -ls /data  
Found 1 items  
-rw-r--r--  2 hadoop supergroup 2666130 2013-05-20 21:02 /data/book.txt  
  
[hadoop@h1 testdata]$ ssh h6 hadoop fs -ls /data  
Found 1 items  
-rw-r--r--  2 hadoop supergroup 2666130 2013-05-20 21:02 /data/book.txt
```

Word Count – map.py

- Credit: <http://www.michael-noll.com>
- map.py

```
#!/usr/bin/env python
# http://www.michael-noll.com/tutorials/writing-an-hadoop-mapreduce-program-in-python/
import sys

for line in sys.stdin:
    # Remove EOL
    line=line.strip()

    # Split the line of words into an array of words
    words=line.split()

    for w in words:
        print '%s\t%s' % (w,1)
```

Word Count – map.py

- Sample output (k,v) : (word,1)

```
[hadoop@h1 testdata]$ head -70 CountOfMonteCristo.txt | tail -4

"Fell into the sea?"

"No, sir, he died of brain-fever in dreadful agony." Then turning to the

[hadoop@h1 testdata]$ head -70 CountOfMonteCristo.txt | tail -4 | ./map.py
"Fell 1
into 1
the 1
sea?" 1
"No, 1
sir, 1
he 1
died 1
of 1
brain-fever 1
in 1
dreadful 1
agony." 1
Then 1
turning 1
to 1
the 1
```


Word Count – reduce.py

- Credit <http://www.michael-noll.com>
- reduce.py (1 of 2)

```
#!/usr/bin/env python
# http://www.michael-noll.com/tutorials/writing-an-hadoop-mapreduce-program-in-python/

from operator import itemgetter
import sys

current_word = None
current_count = 0
word = None

# input comes from STDIN
for line in sys.stdin:
    # remove leading and trailing whitespace
    line = line.strip()

    # parse the input we got from mapper.py
    word, count = line.split('\t', 1)

    # convert count (currently a string) to int
    try:
        count = int(count)
    except ValueError:
        # count was not a number, so silently
        # ignore/discard this line
        continue
```

Word Count – reduce.py

- reduce.py (2 of 2)

```
# this IF-switch only works because Hadoop sorts map output
# by key (here: word) before it is passed to the reducer
if current_word == word:
    current_count += count
else:
    if current_word:
        # write result to STDOUT
        print '%s\t%s' % (current_word, current_count)
    current_count = count
    current_word = word

# do not forget to output the last word if needed!
if current_word == word:
    print '%s\t%s' % (current_word, current_count)
```

Word Count – Local Test

- Map, sort and reduce – Map

```
[hadoop@h1 testdata]$ head -70 CountOfMonteCristo.txt | tail -4 | ./map.py
"Fell 1
into 1
the 1
sea?" 1
"No, 1
sir, 1
he 1
died 1
of 1
brain-fever 1
in 1
dreadful 1
agony." 1
Then 1
turning 1
to 1
the 1
```

Word Count – Local Test

- Map, sort and reduce – Map & Sort

```
[hadoop@h1 testdata]$ head -70 CountOfMonteCristo.txt | tail -4 | ./map.py | sort
agony." 1
brain-fever 1
died 1
dreadful 1
"Fell 1
he 1
in 1
into 1
"No, 1
of 1
sea?" 1
sir, 1
the 1
the 1
Then 1
to 1
turning 1
```

Word Count – Local Test

- Map, sort and reduce – Map, Sort & Reduce

```
[hadoop@h1 testdata]$ head -70 CountOfMonteCristo.txt | tail -4 | ./map.py | sort | ./reduce.py
agony." 1
brain-fever 1
died 1
dreadful 1
"Fell 1
he 1
in 1
into 1
"No, 1
of 1
sea?" 1
sir, 1
the 2
Then 1
to 1
turning 1
```

Word Count – Hadoop

- Hadoop Word Count

```
[hadoop@h1 testdata]$ cat wc.sh  
export H=/h/hadoop/current
```

```
$H/bin/hadoop jar $H/contrib/streaming/hadoop-streaming-1.0.4.jar \  
-file /home/hadoop/testdata/map.py -mapper /home/hadoop/testdata/map.py \  
-file /home/hadoop/testdata/reduce.py -reducer /home/hadoop/testdata/reduce.py \  
-input /data/book.txt -output /results
```

```
[hadoop@h1 testdata]$ ./wc.sh  
packageJobJar: [/home/hadoop/testdata/map.py, /home/hadoop/testdata/reduce.py, /tmp/hadoop-hadoop/hadoop-  
unjar8581470847816847548/] [] /tmp/streamjob6203005843654881832.jar tmpDir=null  
13/05/20 22:01:54 INFO util.NativeCodeLoader: Loaded the native-hadoop library  
13/05/20 22:01:54 WARN snappy.LoadSnappy: Snappy native library not loaded  
13/05/20 22:01:54 INFO mapred.FileInputFormat: Total input paths to process : 1  
13/05/20 22:01:55 INFO streaming.StreamJob: getLocalDirs(): [/tmp/hadoop-hadoop/mapred/local]  
13/05/20 22:01:55 INFO streaming.StreamJob: Running job: job_201305201838_0002  
13/05/20 22:01:55 INFO streaming.StreamJob: To kill this job, run:  
13/05/20 22:01:55 INFO streaming.StreamJob: /h/hadoop/hadoop-1.0.4/libexec/./bin/hadoop job -  
Dmapred.job.tracker=h1:9001 -kill job_201305201838_0002  
13/05/20 22:01:55 INFO streaming.StreamJob: Tracking URL:  
http://h1.ilmtech.com:50030/jobdetails.jsp?jobid=job\_201305201838\_0002  
13/05/20 22:01:56 INFO streaming.StreamJob: map 0% reduce 0%  
13/05/20 22:02:08 INFO streaming.StreamJob: map 100% reduce 0%  
13/05/20 22:02:20 INFO streaming.StreamJob: map 100% reduce 100%  
13/05/20 22:02:26 INFO streaming.StreamJob: Job complete: job_201305201838_0002  
13/05/20 22:02:26 INFO streaming.StreamJob: Output: /results
```

Word Count – Hadoop

- Hadoop Word Count - Results

```
[hadoop@h1 testdata]$ hdfs -ls /results
Found 3 items
-rw-r--r--  2 hadoop supergroup      0 2013-05-20 22:02 /results/_SUCCESS
drwxr-xr-x  - hadoop supergroup      0 2013-05-20 22:01 /results/_logs
-rw-r--r--  2 hadoop supergroup 440663 2013-05-20 22:02 /results/part-00000

[hadoop@h1 testdata]$ hdfs -cat /results/part-00000 | head -10000 | tail -5
bolt, 2
bolt. 1
bolted 5
bolted, 1
bolted. 1
```

- Hadoop Word Count – Top 10 words

```
[hadoop@h1 testdata]$ hdfs -cat /results/part-00000 | awk '{print $2"\t"$1}' | sort -n | tail -10
5430  he
5690  you
5753  his
6095  in
6524  I
8888  a
11094 and
12557 to
12601 of
26065 the
```

Hadoop Web Interfaces

Daemon	Default Port	Parameter	File
Namenode	50070	dfs.http.address	hdfs-site.xml
Secondary Namenode	50090	dfs.secondary.http.address	hdfs-site.xml
Datanodes	50075	dfs.datanode.http.address	hdfs-site.xml
Job Tracker	50030	mapred.job.tracker.http.address	mapred-site.xml
Task Trackers	50060	mapred.task.tracker.http.address	mapred-site.xml

Thank You

ahbaid@{a9.com|amazon.com|gmail.com}

ahbaid.gaffoor.org/nocoug

Reading Material

- Hadoop – The Definitive Guide - [Amazon](#)
- Hadoop Operations - [Amazon](#)
- Hadoop in Action - [Amazon](#)
- Hadoop Real World Solutions Cookbook - [Amazon](#)
- Programming Hive - [Amazon](#)
- HBase – The Definitive Guide - [Amazon](#)
- HBase in Action – [Amazon](#)
- Hbase Administration CookBook - [Amazon](#)
- Programming Pig - [Amazon](#)

Credit / References

- <http://www.michael-noll.com>
 - Python Map Reduce Example