

TP1-2: Analyzing *Hadoop* Logs

Shadi Ibrahim
January 26th, 2017

MapReduce has emerged as a leading programming model for data-intensive computing. It was originally proposed by Google to simplify development of web search applications on a large number of machines.

Hadoop is a java open source implementation of MapReduce sponsored by Yahoo! The Hadoop project is a collection of various subprojects for reliable, scalable distributed computing. The two fundamental subprojects are the Hadoop MapReduce framework and the HDFS. HDFS is a distributed file system that provides high throughput access to application data. It is inspired by the GFS. HDFS has master/slave architecture. The master server, called NameNode, splits files into blocks and distributes them across the cluster with replications for fault tolerance. It holds all metadata information about stored files. The HDFS slaves, the actual store of the data blocks called DataNodes, serve read/write requests from clients and propagate replication tasks as directed by the NameNode.

The Hadoop MapReduce is a software framework for distributed processing of large data sets on compute clusters. It runs on the top of the HDFS. Thus data processing is collocated with data storage. It also has master/slave architecture. The master, called Job Tracker (JT), is responsible of : (a) Querying the NameNode for the block locations, (b) considering the information retrieved by the NameNode, JT schedule the tasks on the slaves, called Task Trackers (TT), and (c) monitoring the success and failures of the tasks.

The goal of this TP is to study the operation of Hadoop platform through exploring the logs of the execution of MapReduce applications. We will examine these logs by scanning their content, by looking at Hadoop's web GUI, and using another Log file visualization tool.

Exercise 1: Accessing Hadoop's Logs

Hadoop keeps track of several logs of the execution of your programs. They are located in the logs sub-directory in the Hadoop directory. There are two important files for each job: the configuration file and the log file.

```
job_201410292315_0002_1414658868472_useranme_word+count
job_201410292315_0002_conf.xml
```

`job_201410292315_0002_conf.xml` file includes all configurations of *Hadoop* platform when running the application.

`job_201410292315_0002_1414658868472_useranme_word+count` contains the start and end time of all the tasks that run during the execution of our Hadoop program. It contains several different types of lines:

Lines starting with "Job", listing information about the job (priority, submit time, configuration, number of map tasks, number of reduce tasks, etc.

```
Job JOBID = "job_201410292315_0002" JOBNAME = "word count" USER =
"xxxxxxx" SUBMIT_TIME = "1414658868472" JOBCONF = "hdfs://localhost:8020/
Users/$Hadoop_Home/hdfs/tmp/mapred/staging/xxxxx/staging/job_201410292315_0002/job.xml"
VIEW_JOB = "*" MODIFY_JOB = "*" JOB_QUEUE = "default" WORKFLOW_ID = ""
WORKFLOW_NAME = " " WORKFLOW_NODE_NAME = " " WORKFLOW_ADJACENCIES =
" " WORKFLOW_TAGS = "".
Job JOBID = "job_201410292315_0002" JOB_PRIORITY = "NORMAL". Job JOBID =
"job_201410292315_0002" LAUNCH_TIME = "1414658869065" TOTAL_MAPS = "115"
TOTAL_REDUCES = "1" JOB_STATUS = "PREP"
```

Lines starting with "Task" referring to the creation or completion of Map or Reduce tasks, indicating which host they start on, and which split they work on (i.e., replica). On completion, all the counters associated with the task are listed.

```
Task TASKID = "task_201410292315_0002_m_000000" TASK_TYPE = "MAP"
START_TIME = "1414658870533" SPLITS = "/default-rack/shadis-mbp".
```

Lines starting with "MapAttempt", reporting mostly status update, except if they contain the keywords `SUCCESS` and/or `FINISH_TIME`, indicating that the task has completed. The final time when the task finished is included in this line.

```

MapAttempt TASK_TYPE="MAP" TASKID="task_201410292315_0002_m_000000" TASK_ATTEMPT_ID="attempt_201410292315_0002_m_000000_0"
START_TIME="1414658870536" TRACKER_NAME="tracker_shadis-mbp:localhost/127.0.0.1:59456" HTTP_PORT="50060" LOCALITY="NODE_LOCAL"
AVATAAR="VIRGIN" .
MapAttempt TASK_TYPE="MAP" TASKID="task_201410292315_0002_m_000000" TASK_ATTEMPT_ID="attempt_201410292315_0002_m_000000_0"
TASK_STATUS="SUCCESS" FINISH_TIME="1414658871453" HOSTNAME="/default-rack/shadis-mbp" STATE_STRING="" COUNTERS="{(FileSystemCounters)
(FileSystemCounters){(HDFS_BYTES_READ)(HDFS_BYTES_READ)(528495)}{(FILE_BYTES_WRITTEN)(FILE_BYTES_WRITTEN)(1049138)}{(org.apache.hadoop
.mapreduce.lib.input.FileInputFormat$Counter)(File Input Format Counters ){(BYTES_READ)(Bytes Read)(528384)}{(org.apache.hadoop
.mapred.Task$Counter)(Map-Reduce Framework){(MAP_OUTPUT_MATERIALIZED_BYTES)(Map output materialized bytes)(981142)}
{(COMBINE_OUTPUT_RECORDS)(Combine output records)(0)}{(MAP_INPUT_RECORDS)(Map input records)(6903)}{(SPILLED_RECORDS)(Spilled Records)
(76181)}{(MAP_OUTPUT_BYTES)(Map output bytes)(828774)}{(COMMITTED_HEAP_BYTES)(Total committed heap usage \\\(bytes\\))(201326592)}
{(COMBINE_INPUT_RECORDS)(Combine input records)(0)}{(MAP_OUTPUT_RECORDS)(Map output records)(76181)}{(SPLIT_RAW_BYTES)(SPLIT_RAW_BYTES)
(111)}" .

```

Figure 1

Lines starting with “ReduceAttempt”, similar to the MapAttempt tasks, report on the intermediary status of the tasks, and when the keyword SUCCESS is included, the finish time of the sort and shuffle phases will also be included.

```

ReduceAttempt TASK_TYPE="REDUCE" TASKID="task_201410292315_0001_r_000000" TASK_ATTEMPT_ID="attempt_201410292315_0001_r_000000_0"
START_TIME="1414621199172" TRACKER_NAME="tracker_shadis-mbp:localhost/127.0.0.1:59456" HTTP_PORT="50060" LOCALITY="OFF_SWITCH"
AVATAAR="VIRGIN" .
ReduceAttempt TASK_TYPE="REDUCE" TASKID="task_201410292315_0001_r_000000" TASK_ATTEMPT_ID="attempt_201410292315_0001_r_000000_0"
TASK_STATUS="SUCCESS" SHUFFLE_FINISHED="1414621299424" SORT_FINISHED="1414621300450" FINISH_TIME="1414621301579" HOSTNAME="/default-
rack/shadis-mbp" STATE_STRING="reduce > reduce" COUNTERS="{(org.apache.hadoop.mapreduce.lib.output.FileOutputFormat$Counter)(File
Output Format Counters ){(BYTES_WRITTEN)(Bytes Written)(2402)}{(FileSystemCounters)(FileSystemCounters){(FILE_BYTES_READ)
(FILE_BYTES_READ)(275201)}{(FILE_BYTES_WRITTEN)(FILE_BYTES_WRITTEN)(343365)}{(HDFS_BYTES_WRITTEN)(HDFS_BYTES_WRITTEN)(2402)}
{(org.apache.hadoop.mapred.Task$Counter)(Map-Reduce Framework){(REDUCE_INPUT_GROUPS)(Reduce input groups)(182)}
{(COMBINE_OUTPUT_RECORDS)(Combine output records)(0)}{(REDUCE_SHUFFLE_BYTES)(Reduce shuffle bytes)(275885)}{(REDUCE_OUTPUT_RECORDS)
(Reduce output records)(182)}{(SPILLED_RECORDS)(Spilled Records)(20930)}{(COMMITTED_HEAP_BYTES)(Total committed heap usage \\\(bytes
\\))(201326592)}{(COMBINE_INPUT_RECORDS)(Combine input records)(0)}{(REDUCE_INPUT_RECORDS)(Reduce input records)(20930)}" .
Task TASKID="task_201410292315_0001_r_000000" TASK_TYPE="REDUCE" TASK_STATUS="SUCCESS" FINISH_TIME="1414621301655" COUNTERS="
{(org.apache.hadoop.mapreduce.lib.output.FileOutputFormat$Counter)(File Output Format Counters ){(BYTES_WRITTEN)(Bytes Written)
(2402)}{(FileSystemCounters)(FileSystemCounters){(FILE_BYTES_READ)(FILE_BYTES_READ)(275201)}{(FILE_BYTES_WRITTEN)(FILE_BYTES_WRITTEN)
(343365)}{(HDFS_BYTES_WRITTEN)(HDFS_BYTES_WRITTEN)(2402)}{(org.apache.hadoop.mapred.Task$Counter)(Map-Reduce Framework)
{(REDUCE_INPUT_GROUPS)(Reduce input groups)(182)}{(COMBINE_OUTPUT_RECORDS)(Combine output records)(0)}{(REDUCE_SHUFFLE_BYTES)(Reduce
shuffle bytes)(275885)}{(REDUCE_OUTPUT_RECORDS)(Reduce output records)(182)}{(SPILLED_RECORDS)(Spilled Records)(20930)}
{(COMMITTED_HEAP_BYTES)(Total committed heap usage \\\(bytes\\))(201326592)}{(COMBINE_INPUT_RECORDS)(Combine input records)
(0)}{(REDUCE_INPUT_RECORDS)(Reduce input records)(20930)}" .

```

Figure 2

Question 1.1

Run Wordcount application using the data set 75MB and the default block size (64MB). Check the configuration file and see all the properties which were discussed in TP2. Check the number of map tasks, reduce tasks, split size and location in your log files.

Question 1.2

Use the Data set 75MB and run both benchmarks using different block sizes (i.e., 5, 10, and 20MB). Check the configuration and the log files, what can you observe?

Exercise 2: Accessing Hadoop’s Web GUI

Same logs and conf files are also available from the hadoop Web GUI:
<http://localhost:50030/>

localhost Hadoop Map/Reduce Administration

[Quick Links](#)

State: RUNNING
 Started: Fri Oct 31 04:23:22 CET 2014
 Version: 1.2.1, r1503152
 Compiled: Mon Jul 22 15:23:09 PDT 2013 by mattf
 Identifier: 201410310423
 SafeMode: OFF

Cluster Summary (Heap Size is 239 MB/1.74 GB)

Running Map Tasks	Running Reduce Tasks	Total Submissions	Nodes	Occupied Map Slots	Occupied Reduce Slots	Reserved Map Slots	Reserved Reduce Slots	Map Task Capacity	Reduce Task Capacity	Avg. Tasks/Node	Blacklisted Nodes	Graylisted Nodes	Excluded Nodes
0	0	0	1	0	0	0	0	2	2	4.00	0	0	0

Scheduling Information

Queue Name	State	Scheduling Information
default	running	N/A

Filter (JobId, Priority, User, Name)
 Example: 'user:smith 3200' will filter by 'smith' only in the user field and '3200' in all fields

Running Jobs

Retired Jobs

Local Logs

[Log directory](#), [Job Tracker History](#)
 This is [Apache Hadoop](#) release 1.2.1

Figure 3: Hadoop Main page

shadis-mbp Hadoop Map/Reduce History Viewer

Filter (username:jobname)

Specify [user][:jobname keyword(s)][,MM/DD/YYYY]. Each of the three components is optional. Filter components are conjunctive.

Example: 'smith' will display jobs submitted by user 'smith', 'smith:sort' will display jobs from user 'smith' having a 'sort' keyword in the jobname, ':07/04/2010' restricts to July 4, 2010

Available Jobs in History (*Displaying 10 jobs from 1 to 10 out of 10 jobs*) [\[get more results\]](#) [\[show in one page\]](#) [\[first page\]](#) [\[last page\]](#)

Job submit time	Job Id	Name	User
Thu Oct 30 13:47:36 CET 2014	job_201410301137_0005	word count	shadiibrahim
Thu Oct 30 13:33:30 CET 2014	job_201410301137_0004	word count	shadiibrahim
Thu Oct 30 13:06:24 CET 2014	job_201410301137_0003	WordsCount	shadiibrahim
Thu Oct 30 12:35:20 CET 2014	job_201410301137_0002	WordsCount	shadiibrahim
Thu Oct 30 11:38:00 CET 2014	job_201410301137_0001	LineCount	shadiibrahim
Thu Oct 30 10:14:43 CET 2014	job_201410301001_0004	wordcount	shadiibrahim
Thu Oct 30 10:11:17 CET 2014	job_201410301001_0003	word count	shadiibrahim
Thu Oct 30 10:09:15 CET 2014	job_201410301001_0002	word count	shadiibrahim
Thu Oct 30 10:03:36 CET 2014	job_201410301001_0001	word count	shadiibrahim
Wed Oct 29 23:19:52 CET 2014	job_201410292315_0001	word count	shadiibrahim

Figure 4: All logs History

Hadoop Job job_201410292315_0001 on History Viewer

User: shadibrahim
 JobName: word count
 JobConf: hdfs://localhost:80201/users/shadibrahim/Desktop/hadoop/hadoop/shadi/hdfs/tmp/mapred/staging/shadibrahim/.staging/job_201410292315_0001.job.xml
 Job-ACLs: All users are allowed
 Submitted At: 29-Oct-2014 23:19:52
 Launched At: 29-Oct-2014 23:19:52 (0sec)
 Finished At: 29-Oct-2014 23:21:42 (1mins, 50sec)
 Status: SUCCESS
 Failure Info:
[Analyse This Job](#)

Kind	Total Tasks(successful+failed+killed)	Successful tasks	Failed tasks	Killed tasks	Start Time	Finish Time
Setup	1	1	0	0	29-Oct-2014 23:19:52	29-Oct-2014 23:19:53 (1sec)
Map	115	115	0	0	29-Oct-2014 23:19:54	29-Oct-2014 23:21:36 (1mins, 42sec)
Reduce	1	1	0	0	29-Oct-2014 23:19:59	29-Oct-2014 23:21:41 (1mins, 42sec)
Cleanup	1	1	0	0	29-Oct-2014 23:21:41	29-Oct-2014 23:21:42 (0sec)

Figure 5: The log file of one application : Overview

	Counter	Map	Reduce	Total
Job Counters	Launched reduce tasks	0	0	1
	SLOTS_MILLIS_MAPS	0	0	190,139
	Total time spent by all reduces waiting after reserving slots (ms)	0	0	0
	Total time spent by all maps waiting after reserving slots (ms)	0	0	0
	Launched map tasks	0	0	115
	Data-local map tasks	0	0	115
	SLOTS_MILLIS_REDUCE	0	0	102,407
File Output Format Counters	Bytes Written	0	2,402	2,402
FileSystemCounters	FILE_BYTES_READ	0	275,201	275,201
	HDFS_BYTES_READ	60,539,143	0	60,539,143
	FILE_BYTES_WRITTEN	8,132,920	343,365	8,476,285
	HDFS_BYTES_WRITTEN	0	2,402	2,402
File Input Format Counters	Bytes Read	60,526,378	0	60,526,378
Map-Reduce Framework	Reduce input groups	0	182	182
	Map output materialized bytes	275,885	0	275,885
	Combine output records	20,930	0	20,930
	Map input records	790,642	0	790,642
	Reduce shuffle bytes	0	275,885	275,885
	Reduce output records	0	182	182
	Spilled Records	20,930	20,930	41,860
	Map output bytes	94,935,486	0	94,935,486
	Total committed heap usage (bytes)	23,152,558,080	201,326,592	23,353,884,672
	Combine input records	8,726,334	0	8,726,334
	Map output records	8,726,334	0	8,726,334
	SPLIT_RAW_BYTES	12,765	0	12,765
	Reduce input records	0	20,930	20,930

Figure 6: The log file of one application: Details

Hadoop Job [job_201410292315_0001](#)

User : shadiibrahim
 JobName : word count
 JobConf : hdfs://localhost:8020/Users/shadiibrahim/Desktop/hadoop/hadoop/shadi/hdfs/tmp/mapred/staging/shadiibrahim/staging/job_201410292315_0001/job.xml
 Submitted At : 29/10 23:19:52
 Launched At : 29/10 23:19:52 (0sec)
 Finished At : 29/10 23:21:42 (1mins, 50sec)
 Status : SUCCESS

Time taken by best performing Map task [task_201410292315_0001_m_000098](#) : 0sec

Average time taken by Map tasks: 1sec

Worse performing map tasks

Task Id	Time taken
task_201410292315_0001_m_000091	2sec
task_201410292315_0001_m_000006	2sec
task_201410292315_0001_m_000023	2sec
task_201410292315_0001_m_000036	2sec
task_201410292315_0001_m_000051	2sec
task_201410292315_0001_m_000059	2sec
task_201410292315_0001_m_000079	2sec
task_201410292315_0001_m_000076	2sec
task_201410292315_0001_m_000112	2sec
task_201410292315_0001_m_000048	2sec

The last Map task [task_201410292315_0001_m_000113](#) finished at (relative to the Job launch time): 29/10 23:21:36 (1mins, 44sec)

Figure 7: The log file of one application: Tasks executions 1

Time taken by best performing shufflejobId [task_201410292315_0001_r_000000](#) : 1mins, 40sec

Average time taken by Shuffle: 1mins, 40sec

Worse performing Shuffle(s)

Task Id	Time taken
task_201410292315_0001_r_000000	1mins, 40sec

The last Shuffle [task_201410292315_0001_r_000000](#) finished at (relative to the Job launch time): 29/10 23:21:39 (1mins, 47sec)

Time taken by best performing Reduce task : [task_201410292315_0001_r_000000](#) : 2sec

Average time taken by Reduce tasks: 2sec

Worse performing reduce tasks

Task Id	Time taken
task_201410292315_0001_r_000000	2sec

The last Reduce task [task_201410292315_0001_r_000000](#) finished at (relative to the Job launch time): 29/10 23:21:41 (1mins, 49sec)

Figure 8: The log file of one application: Tasks executions 2

[task_201410292315_0001_m_000091](#) attempts for [job_201410292315_0001](#)

Task Id	Start Time	Finish Time	Host	Error	Task Logs	Counters
attempt_201410292315_0001_m_000091_0	29/10 23:21:14	29/10 23:21:17 (2sec)	/default-rack/shadis-mbp		Last 4KB Last 8KB All	12

Input Split Locations

[/default-rack/shadis-mbp](#)

Figure 9: The log file of one application: Map task execution 1

Question 2.1

Check the configuration and the log files of your previous jobs through Hadoop GUI, what can you observe?

Exercise 3: Accessing VisHadoop GUI

We have developed a new visualization tool to show the progress of different tasks during the execution of MapReduce application. Check this site <http://hadoop-log.irisa.fr>.




Figure 10: An Example of VisHadoop

Question 3.1

Check the pattern of the three jobs (waves), Can you explain it?

Question 3.2

Check the two log files in the (resources), and identify values of the following properties, using any of the three log analysis methods:

- Total Map tasks
- Total Reduce tasks
- Total local map tasks
- Replication factor
- Total killed tasks

How to distinguish speculated task from normal one?