

## **Bench marking and Stress Testing on HDP Hadoop Cluster with Terasort, TestDFSIO:**

In this article I would like to introduce to some of the benchmarking and testing tools that are included in the Hortonworks Hadoop Distribution. Namely, we look at the benchmarks TestDFSIO and Terasort.

Hence knowing how to run these tools will help you to shake out your cluster in terms of architecture, hardware and software, to measure its performance and also to share and compare your results with other people.

Here are the differences in executing the same teragen/terasort on HDP and Cloudera.

### **HDP measures on a 4 node cluster with m3.medium:**

```
Total time spent by all maps in occupied slots (ms)=19716
Total time spent by all reduces in occupied slots (ms)=23438
Total time spent by all map tasks (ms)=19716
Total time spent by all reduce tasks (ms)=11719
Total vcore-seconds taken by all map tasks=19716
Total vcore-seconds taken by all reduce tasks=11719
Total megabyte-seconds taken by all map tasks=6723156
Total megabyte-seconds taken by all reduce tasks=7992358
```

### **CDH measures on a 5 node cluster:**

```
Total time spent by all maps in occupied slots (ms)=12838
Total time spent by all reduces in occupied slots (ms)=36156
Total time spent by all map tasks (ms)=12838
Total time spent by all reduce tasks (ms)=36156
Total vcore-seconds taken by all map tasks=12838
Total vcore-seconds taken by all reduce tasks=36156
Total megabyte-seconds taken by all map tasks=13146112
Total megabyte-seconds taken by all reduce tasks=37023744
```

### **HDP measures on a 4 node cluster:**

```
Map-Reduce Framework
Map input records=10000
Map output records=10000
Map output bytes=1020000
Map output materialized bytes=1040012
Input split bytes=296
Combine input records=0
Combine output records=0
Reduce input groups=10000
Reduce shuffle bytes=1040012
Reduce input records=10000
Reduce output records=10000
Spilled Records=20000
Shuffled Maps =2
Failed Shuffles=0
```

Merged Map outputs=2  
GC time elapsed (ms)=753  
CPU time spent (ms)=5590  
Physical memory (bytes) snapshot=772042752  
Virtual memory (bytes) snapshot=3504418816  
Total committed heap usage (bytes)=572661760  
16/05/13 01:20:18 INFO terasort.TeraSort: done

#### **CDH measures on a 5 node cluster with m3.large flavor:**

Map-Reduce Framework  
Map input records=10000  
Map output records=10000  
Map output bytes=1020000  
Map output materialized bytes=430163  
Input split bytes=294  
Combine input records=0  
Combine output records=0  
Reduce input groups=10000  
Reduce shuffle bytes=430163  
Reduce input records=10000  
Reduce output records=10000  
Spilled Records=20000  
Shuffled Maps =8  
Failed Shuffles=0  
Merged Map outputs=8  
GC time elapsed (ms)=335  
CPU time spent (ms)=10870  
Physical memory (bytes) snapshot=1622138880  
Virtual memory (bytes) snapshot=9454505984  
Total committed heap usage (bytes)=1480589312  
16/05/13 01:33:43 INFO terasort.TeraSort: done

#### **Prerequisites:**

Running HDP (Single-Node Cluster) or  
Running HDP (Multi-Node Cluster)

#### **Overview of Benchmarks and Testing Tools on HDP:**

##### **1. Browse to the ResourceManager:**

[http://\\$resourcemanager.full.hostname:8088/](http://$resourcemanager.full.hostname:8088/)

##### **2.Create a \$CLIENT\_USER in all of the nodes and add it to the users groups.**

```
useradd client  
usermod -a -G users client
```

##### **3. As the HDFS user, create a /user/client**

```
sudo su - hdfs
```

```

hdfs dfs -mkdir /user/client
hdfs dfs -chown client:client /user/client
hdfs dfs -chmod -R 755 /user/client

```



4. Run the smoke test as the \$CLIENT\_USER. Using Terasort, sort 10GB of data.  
 su - \$CLIENT\_USER

**Teragen on HDP:**

```

/usr/hdp/current/hadoop-client/bin/hadoop jar /usr/hdp/current/hadoop-mapreduce-client/hadoop-mapreduce-examples-*.jar teragen 10000 tmp/teragenout

```



**Terasort on HDP:**

```

/usr/hdp/current/hadoop-client/bin/hadoop jar /usr/hdp/current/hadoop-mapreduce-client/hadoop-mapreduce-examples-*.jar terasort tmp/teragenout tmp/terasortout

```



**Teravalidate on HDP:**

```

/usr/hdp/current/hadoop-client/bin/hadoop jar /usr/hdp/current/hadoop-mapreduce-client/hadoop-mapreduce-examples.jar teravalidate tmp/terasortout tmp/teravalidate

```



```
hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-examples.jar teragen 10000 /user/cdhuser/teragenout
```

```
hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-examples.jar terasort /user/cdhuser/teragenout /user/cdhuser/terasortout
```

```
hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-examples.jar teravalidate /user/cdhuser/terasortout /user/cdhuser/teravalidate
```

```
hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-examples.jar teragen 10000 /user/cdhuser/teragenout
16/05/13 14:23:18 DMO Client:MPProxy: Connecting to ResourceManager at suj11ha022.cloud.hortonworks.com/372.24.64.54:8032
16/05/13 14:23:19 DMO Terasort:Terasort: Generating 1000000000 records
16/05/13 14:23:19 DMO mapreduce.JobSubmitter: number of splits:1
16/05/13 14:24:00 DMO mapreduce.JobSubmitter: Submitting tokens for job: job_1463126249715_0001
16/05/13 14:24:00 DMO LocalFileSystem: Submitted application application_1463126249715_0001
16/05/13 14:24:00 DMO mapreduce.Job: The url to track the job: http://suj11ha022.cloud.hortonworks.com:8088/proxy/application_1463126249715_0001/
16/05/13 14:24:00 DMO mapreduce.Job: Running job: job_1463126249715_0001

-----
16/05/13 01:29:45 INFO mapreduce.Job: Job job_1463126249715_0001 completed successfully
16/05/13 01:29:46 INFO mapreduce.Job: Counters: 31
  File System Counters
    FILE: Number of bytes read=0
    FILE: Number of bytes written=234512
    FILE: Number of read operations=0
    FILE: Number of large read operations=0
    FILE: Number of write operations=0
    HDFS: Number of bytes read=164
    HDFS: Number of bytes written=1000000
    HDFS: Number of read operations=8
    HDFS: Number of large read operations=0
    HDFS: Number of write operations=4
  Job Counters
    Launched map tasks=2
    Other local map tasks=2
    Total time spent by all maps in occupied slots (ms)=15287
    Total time spent by all reduces in occupied slots (ms)=0
    Total time spent by all map tasks (ms)=15287
    Total vcore-seconds taken by all map tasks=15287
    Total megabyte-seconds taken by all map tasks=15653888
  Map-Reduce Framework
    Map input records=100000
    Map output records=10000
    Input split bytes=164
    Spilled Records=0
    Failed Shuffles=0
    Merged Map outputs=0
    GC time elapsed (ms)=102
    CPU time spent (ms)=1908
    Physical memory (bytes) snapshot=349241344
    Virtual memory (bytes) snapshot=3120852992
    Total committed heap usage (bytes)=355467264
  org.apache.hadoop.examples.Terasort.TerasortCounters
    CHECKSUM=21555390172850
  File Input Format Counters
    Bytes Read=0
  File Output Format Counters
    Bytes Written=1000000
-----

16/05/13 01:30:10 INFO mapreduce.Job: map 100% reduce 100%
16/05/13 01:30:10 INFO mapreduce.Job: map 100% reduce 100%
16/05/13 01:30:11 INFO mapreduce.Job: map 100% reduce 100%
16/05/13 01:30:12 INFO mapreduce.Job: Job job_1463126249715_0001 completed successfully
16/05/13 01:30:13 INFO mapreduce.Job: Counters: 30
  File System Counters
    FILE: Number of bytes read=42256
    FILE: Number of bytes written=114421
    FILE: Number of read operations=0
    FILE: Number of large read operations=0
    FILE: Number of write operations=0
    HDFS: Number of bytes read=100000
    HDFS: Number of bytes written=100000
    HDFS: Number of read operations=0
    HDFS: Number of large read operations=0
  Job Counters
    Launched map tasks=2
    Other local map tasks=2
    Total time spent by all maps in occupied slots (ms)=12800
    Total time spent by all reduces in occupied slots (ms)=9056
    Total time spent by all map tasks (ms)=12800
    Total time spent by all reduce tasks (ms)=9056
    Total vcore-seconds taken by all map tasks=12800
    Total megabyte-seconds taken by all map tasks=10000000
    Total megabyte-seconds taken by all reduce tasks=9056256
  Map-Reduce Framework
    Map input records=100000
    Map output records=100000
    Map output bytes=100000
    Map output shuffle bytes=40464
    Input split bytes=164
    Canceled input records=0
    Canceled input records=0
    Reduce input groups=12800
    Reduce shuffle bytes=40464
    Reduce input records=10000
    Reduce output records=10000
    Spilled Records=0
    Failed Shuffles=0
    Merged Map outputs=0
    GC time elapsed (ms)=125
    CPU time spent (ms)=2878
    Physical memory (bytes) snapshot=1822130880
    Virtual memory (bytes) snapshot=16143808
    Total committed heap usage (bytes)=16800000
  Shuffle Counters
    map_time
      COMBINED
      I_COUNT=0
      WR_COUNT=0
      WR_BYTES=0
    File Input Format Counters
      Bytes Read=1000000
    File Output Format Counters
      Bytes Written=100000
-----
16/05/13 01:30:13 INFO Teragen:Teragen: done
```

```
hdfs@sankudh01 ~$ hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-examples.jar teravalidate /user/cdhuser/teravalidate
16/05/17 12:47:28 INFO client: Connecting to ResourceManager at sankudh01.cloudfortworks.com/172.24.04.101:8032
16/05/17 12:47:33 INFO InputFileSystem: Total input paths to process : 4
Spent 0ms computing base-split.
Spent 0ms computing TerDeclarator splits.
16/05/17 12:47:33 INFO mapreduce.JobSubmitter: number of splits=4
16/05/17 12:47:33 INFO mapreduce.JobSubmitter: Submitting tokens for job: job_1463126249713_0005
16/05/17 12:47:34 INFO api.YarnClientImpl: Submitted application application_1463126249713_0005
16/05/17 12:47:34 INFO mapreduce.Job: The url to track the job: http://sankudh01.cloudfortworks.com:8080/rows/application_1463126249713_0005
16/05/17 12:47:43 INFO mapreduce.Job: Running job: job_1463126249713_0005 running in uber mode : false
16/05/17 12:47:43 INFO mapreduce.Job: map 0% reduce 0%
16/05/17 12:47:53 INFO mapreduce.Job: map 50% reduce 0%
16/05/17 12:47:53 INFO mapreduce.Job: map 100% reduce 0%
16/05/17 12:48:00 INFO mapreduce.Job: map 100% reduce 100%
16/05/17 12:48:00 INFO mapreduce.Job: Job job_1463126249713_0005 completed successfully
16/05/17 12:48:00 INFO mapreduce.Job: Counter: 40
File System Counters:
FILE: Number of bytes read=267
FILE: Number of bytes written=8009
FILE: Number of read operations=0
FILE: Number of large read operations=0
FILE: Number of write operations=0
HDFS: Number of bytes read=100302
HDFS: Number of bytes written=22
HDFS: Number of read operations=1
```

```
Found 2 items
-rw-r--r-- 3 hdfs supergroup 0 2016-05-17 12:08 /user/cdhuser/teravalidate/_SUCCESS
-rw-r--r-- 3 hdfs supergroup 22 2016-05-17 12:08 /user/cdhuser/teravalidate/part-r-0001
[hdfs@sankudh01 ~]$ hadoop fs -cat /user/cdhuser/teravalidate/part-r-0000
checksum 139abef74b2
[hdfs@sankudh01 ~]$
```

### TestDFSIO on cloudera: write

```
hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-mapreduce-client-jobclient-2.6.0-cdh5.7.0-tests.jar TestDFSIO -write -nrFiles 10 -fileSize 40
```

```
Bytes Written=78
16/05/17 15:09:48 INFO fs.TestDFSIO: ----- TestDFSIO ----- : write
16/05/17 15:09:48 INFO fs.TestDFSIO: Date & time: Tue May 17 15:09:48 PDT 2016
16/05/17 15:09:48 INFO fs.TestDFSIO: Number of files: 10
16/05/17 15:09:48 INFO fs.TestDFSIO: Total MBytes processed: 400.0
16/05/17 15:09:48 INFO fs.TestDFSIO: Throughput mb/sec: 11.576083810846791
16/05/17 15:09:48 INFO fs.TestDFSIO: Average IO rate mb/sec: 35.46977996826172
16/05/17 15:09:48 INFO fs.TestDFSIO: IO rate std deviation: 25.08206855713185
16/05/17 15:09:48 INFO fs.TestDFSIO: Test exec time sec: 79.861
16/05/17 15:09:48 INFO fs.TestDFSIO:
```

### TestDFSIO on cloudera: read

```
hadoop jar /opt/cloudera/parcels/CDH-5.7.0-1.cdh5.7.0.p0.45/jars/hadoop-mapreduce-client-jobclient-2.6.0-cdh5.7.0-tests.jar TestDFSIO -read -nrFiles 10 -fileSize 40
```

```
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
16/05/17 12:30:24 INFO LocalizedFile: 1684 0x8c 1786 86c: 42*001
```

Reference: <http://www.michael-noll.com/blog/2011/04/09/benchmarking-and-stress-testing-an-hadoop-cluster-with-terasort-testdfsio-nnbench-mrbench/>