



# CERN Log Analysis and Optimization Suggestion

Intel SSG/STO/PAC

# Cluster Hardware Platform

- 3 types of nodes
- Node type 1: SandyBridge 8C16T
  - physical thread num : memory = 1 : 4GB
- Node type 2: BroadWell 12C24T
  - physical thread num : memory = 1 : 21GB
- Node type 3: BroadWell 10C20T
  - physical thread num : memory = 1 : 6.4GB

## Node Type 1:

Node Count: 21

CPU: Intel(R) Xeon(R) CPU E5-2650 0 @ 2.00GHz

RAM: 64 GB

Network: 10Gbps x 1

Storage: 66.79 TB

## Node Type 2:

Node Count: 12

CPU: Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz

RAM: 503 GB

Network: 10 Gbps X 2

Storage: 260.95 TB

## Node Type 3:

Node Count: 5

CPU: Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz

RAM: 128 GB

Network: 10 Gbps X 2

Storage: 260.95 TB

# Log Summary

- Single application: 16 jobs, 32 stages
- Application not **finished**: stop at stage 32
- Stage 0 ~ 31 looks like CPU intensive

## Stages for All Jobs

Active Stages: 1

Completed Stages: 32

### Active Stages (1)

Stage Id ▾	Description	Submitted	Duration	Tasks: Succeeded/Total	Input	Output	Shuffle Read	Shuffle Write
32	<a href="#">save at DimuonReductionAODMultiDataset.scala:369</a> <a href="#">+details</a>	Unknown	160.5 h	<div style="background-color: #0070C0; color: white; padding: 2px;">16825/25397 (36 failed)</div>				

### Completed Stages (32)

Stage Id ▾	Description	Submitted	Duration	Tasks: Succeeded/Total	Input	Output	Shuffle Read	Shuffle Write
31	<a href="#">count at DimuonReductionAODMultiDataset.scala:331</a> <a href="#">+details</a>	2018/05/30 21:38:42	8 s	<div style="background-color: #0070C0; color: white; padding: 2px;">1/1</div>			355.5 KB	
30	<a href="#">count at DimuonReductionAODMultiDataset.scala:331</a> <a href="#">+details</a>	2018/05/30 21:27:30	11 min	<div style="background-color: #0070C0; color: white; padding: 2px;">6171/6171</div>				355.5 KB
29	<a href="#">count at DimuonReductionAODMultiDataset.scala:331</a> <a href="#">+details</a>	2018/05/30 21:27:09	10 s	<div style="background-color: #0070C0; color: white; padding: 2px;">1/1</div>			245.0 KB	
28	<a href="#">count at DimuonReductionAODMultiDataset.scala:331</a> <a href="#">+details</a>	2018/05/30 21:24:41	2.5 min	<div style="background-color: #0070C0; color: white; padding: 2px;">4252/4252</div>				245.0 KB
27	<a href="#">count at DimuonReductionAODMultiDataset.scala:331</a> <a href="#">+details</a>	2018/05/30 21:24:30	2 s	<div style="background-color: #0070C0; color: white; padding: 2px;">1/1</div>			98.8 KB	
26	<a href="#">count at DimuonReductionAODMultiDataset.scala:331</a> <a href="#">+details</a>	2018/05/30 21:22:44	1.8 min	<div style="background-color: #0070C0; color: white; padding: 2px;">1715/1715</div>				98.8 KB

# Stage 32 Issues

- Many tasks hang up
  - root cause still not very clear
- Concurrent multiple HDFS file writing
  - Trigger some errors
- data skew happens on output data size among tasks

Log of application 1524481314062\_31603

Index	ID	Attempt	Status	Locality Level	Executor ID	Host	Launch Time	Duration	GC Time	Errors
15090	40439	0	RUNNING	PROCESS_LOCAL	413	hdp1107.cem.ch	959u 2018/05/31 01:30:59	104.0 h		
14990	40439	0	RUNNING	PROCESS_LOCAL	372	hdp1106.cem.ch	959u 2018/05/31 01:58:46	103.6 h		
15015	40464	0	RUNNING	PROCESS_LOCAL	383	hdp1111.cem.ch	959u 2018/05/31 01:59:39	103.6 h		
15021	40470	0	RUNNING	PROCESS_LOCAL	386	p06636710d25480.cem.ch	959u 2018/05/31 01:59:50	103.6 h		
15002	40481	0	RUNNING	PROCESS_LOCAL	373	hdp1108.cem.ch	959u 2018/05/31 02:00:19	103.6 h		
15038	40487	0	RUNNING	PROCESS_LOCAL	430	p06636710d41523.cem.ch	959u 2018/05/31 02:00:35	103.5 h		
15099	40548	0	RUNNING	PROCESS_LOCAL	411	p06636710d62992.cem.ch	959u 2018/05/31 02:03:26	103.5 h		
15101	40550	0	RUNNING	PROCESS_LOCAL	408	hdp1111.cem.ch	959u 2018/05/31 02:03:32	103.5 h		

Log of application 1524481314062\_23239

```
org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.hdfs.server.namenode.DFSClient): Failed to close inode 1308396373
org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.hdfs.server.namenode.DFSClient): Failed to close inode 1308396373
org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.hdfs.server.namenode.DFSClient): Failed to close inode 1308394670
org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.hdfs.server.namenode.DFSClient): Failed to close inode 1308396977
org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.hdfs.server.namenode.DFSClient): Failed to close inode 1308394921
org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.hdfs.server.namenode.DFSClient): Failed to close inode 1308394102
```

```
18/05/23 12:09:00 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 12,196
18/05/23 12:09:00 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 11,884
18/05/23 12:09:37 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 10,944
18/05/23 12:09:58 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 10,576
18/05/23 12:10:07 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 9,628
18/05/23 12:11:09 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 148
18/05/23 12:11:12 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 76
18/05/23 12:12:14 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 64
18/05/23 12:12:29 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 124
18/05/23 12:12:52 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 76
18/05/23 12:13:05 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 100
18/05/23 12:14:14 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 40
18/05/23 12:14:17 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 100
18/05/23 11:53:51 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 19,200
18/05/23 11:54:52 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 20,008
18/05/23 11:54:55 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 19,372
18/05/23 11:55:46 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 19,380
18/05/23 11:56:09 INFO InternalParquetRecordWriter: Flushing mem columnStore to file. allocated memory 19,192
```

# Potential Issues of Other Stages

- Huge task execution time diversity
  - Data skew? (need input size to judge)
  - Slow machine?
  - Executor memory not enough? (4 tasks share 2g memory/executor)
- Tasks failed in stage
  - Executor lost due to heartbeat time out
- Max parallel executor number is 64
  - Cluster resource not fully used

Summary Metrics for 6171 Completed Tasks

Metric	Min	25th percentile	Median	75th percentile	Max
Duration	2 s	3 s	4 s	5 s	9.9 min
GC Time	34 ms	0.2 s	0.3 s	0.4 s	1 s
Shuffle Write Size / Records	58.0 B / 1	59.0 B / 1	59.0 B / 1	59.0 B / 1	59.0 B / 1

Index	ID	Attempt	Status	Locality Level	Executor ID	Host	Launch Time	Duration	GC Time	Write Time	Shuffle Write Size / Records	Errors
189	2609	0	FAILED	PROCESS_LOCAL	96	ithdp1112.cern.ch	2018/05/30 21:01:07	2.4 min			0.0 B / 0	ExecutorLostFailure (executor 96 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 145734 ms
432	2852	0	FAILED	PROCESS_LOCAL	96	ithdp1112.cern.ch	2018/05/30 21:01:16	2.3 min			0.0 B / 0	ExecutorLostFailure (executor 96 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 145734 ms
433	2853	0	FAILED	PROCESS_LOCAL	96	ithdp1112.cern.ch	2018/05/30 21:01:16	2.3 min			0.0 B / 0	ExecutorLostFailure (executor 96 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 145734 ms
434	2854	0	FAILED	PROCESS_LOCAL	96	ithdp1112.cern.ch	2018/05/30 21:01:16	2.3 min			0.0 B / 0	ExecutorLostFailure (executor 96 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 145734 ms
2452	4872	0	FAILED	PROCESS_LOCAL	117	p05151113452631.cern.ch	2018/05/30 21:02:03	6.5 min			0.0 B / 0	ExecutorLostFailure (executor 117 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 153911 ms
2453	4873	0	FAILED	PROCESS_LOCAL	117	p05151113452631.cern.ch	2018/05/30 21:02:03	6.5 min			0.0 B / 0	ExecutorLostFailure (executor 117 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 153911 ms
2454	4874	0	FAILED	PROCESS_LOCAL	117	p05151113452631.cern.ch	2018/05/30 21:02:03	6.5 min			0.0 B / 0	ExecutorLostFailure (executor 117 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 153911 ms
2455	4875	0	FAILED	PROCESS_LOCAL	117	p05151113452631.cern.ch	2018/05/30 21:02:03	6.5 min			0.0 B / 0	ExecutorLostFailure (executor 117 exited caused by one of the running tasks) Reason: Executor heartbeat timed out after 153911 ms

# Optimization Suggestion

- Node of type 1: disable HT
- Node of type 2: enable HT
- Node of type 3: enable HT
- Spark parameters setting
  - Spark.executor.memory 16g
  - Spark.executor.cores 4
  - Spark.executor.memory.overhead 4g
- Application:
  - If data skew exists, considering increase partition number or repartition
  - Stage 32 stores output in local file system or Hbase?

## Node Type 1:

yarn.nodemanager.resource.memory-mb:  
51200  
yarn.nodemanager.resource.cpu-vcores:  
8

## Node Type 2:

yarn.nodemanager.resource.memory-mb:  
131072  
yarn.nodemanager.resource.cpu-vcores:  
26

## Node Type 3:

yarn.nodemanager.resource.memory-mb:  
102400  
yarn.nodemanager.resource.cpu-vcores:  
20

# Next Step Work

- Get spark history log and detail log of the same job
  - Identify the root cause of application failure
  - Provide parameters setting optimization package
- Cluster usage scenarios: single user with single application or multiple users with multiple applications
- Scale up/out analysis after successful execution