Performance Analysis of Large-Scale Online Data Processing Applications like Apache Spark

POLYTECHNIQUE MONTRÉAL

UNIVERSITÉ D'INGÉNIERIE

Reza Rouh

December 7th, 2023

TABLE OF CONTENTS

Introduction/ Tools

Apache Spark

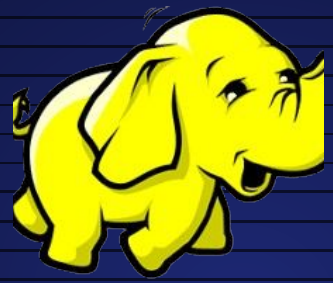

Apache Hadoop

Apache Flink

Apache Storm

Apache Kafka

Apache Beam

Introduction/ Why spark?

Speed

Fast for both batch and interactive queries

Ease of Use

Consistent APIs in Python, Java, Scala, and R

Unified Engine

Combine SQL, streaming, and complex analytics

In-Memory Processing

The RDD, allows for in-memory processing

Fault Tolerance

Can recover the lost data automatically

ML Libraries

Spark's MLlib offers a powerful set of ML

Introduction/ Apache tools logging system

	Spark	Hadoop	Flink	Storm	Kafka	Beam
Support log4j ?	✓	✓	✓	✓	✓	✓

They share similar logging mechanisms!

Introduction/Log4j

Type

Logging framework for Java logging ecosystem

Functionality

ERROR, WARN, INFO, DEBUG, TRACE

Customization

Allows detailed customization

Versions

Log4j 1.x and Log4j 2.x

Introduction/ Spark Architecture

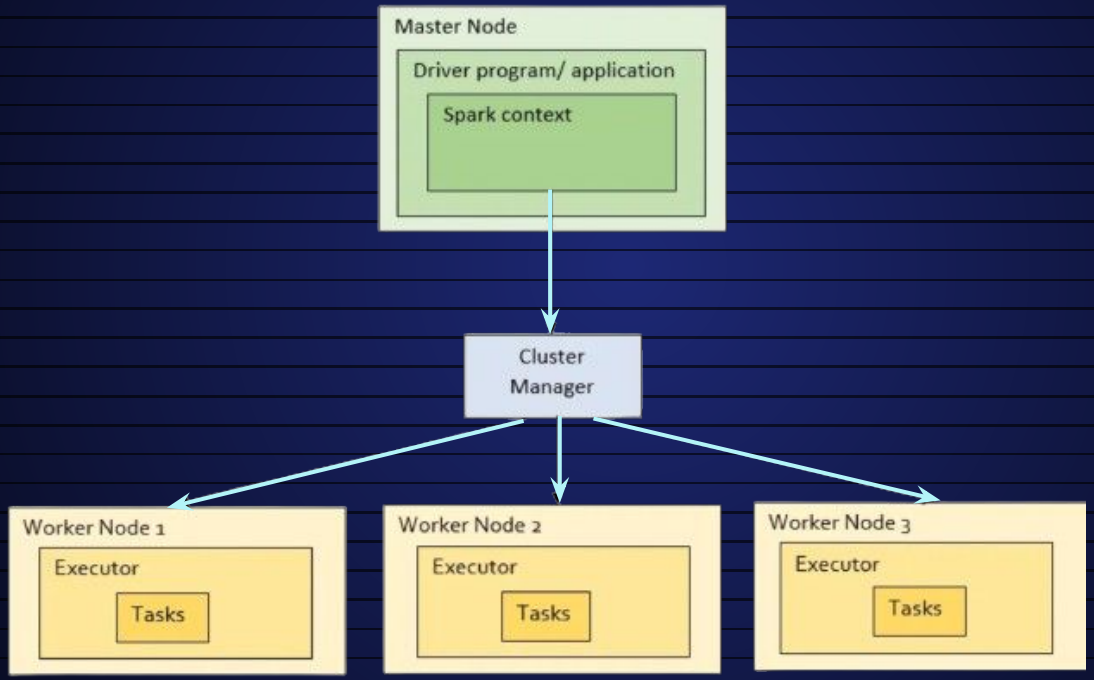

Introduction/ Spark Arcitecture

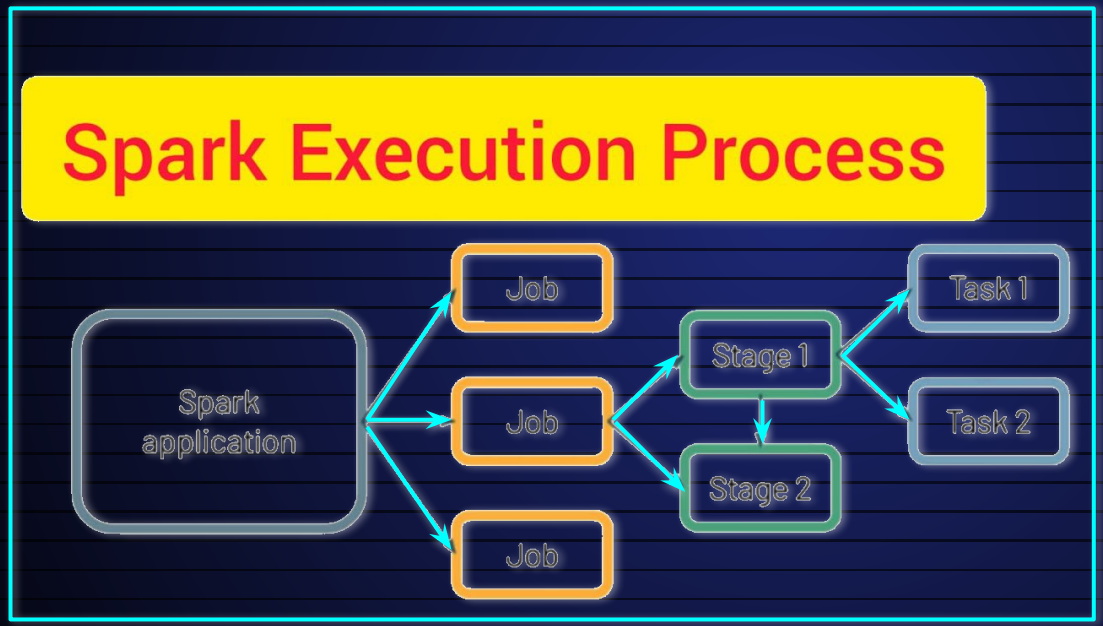

Objectives and Scope

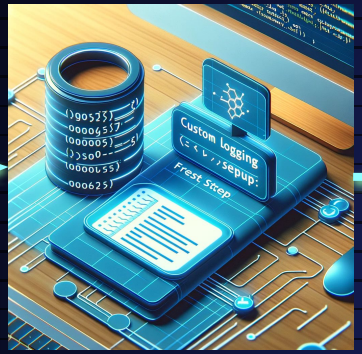

Set up

Logging During Execution

Analyze Traces

Methodology

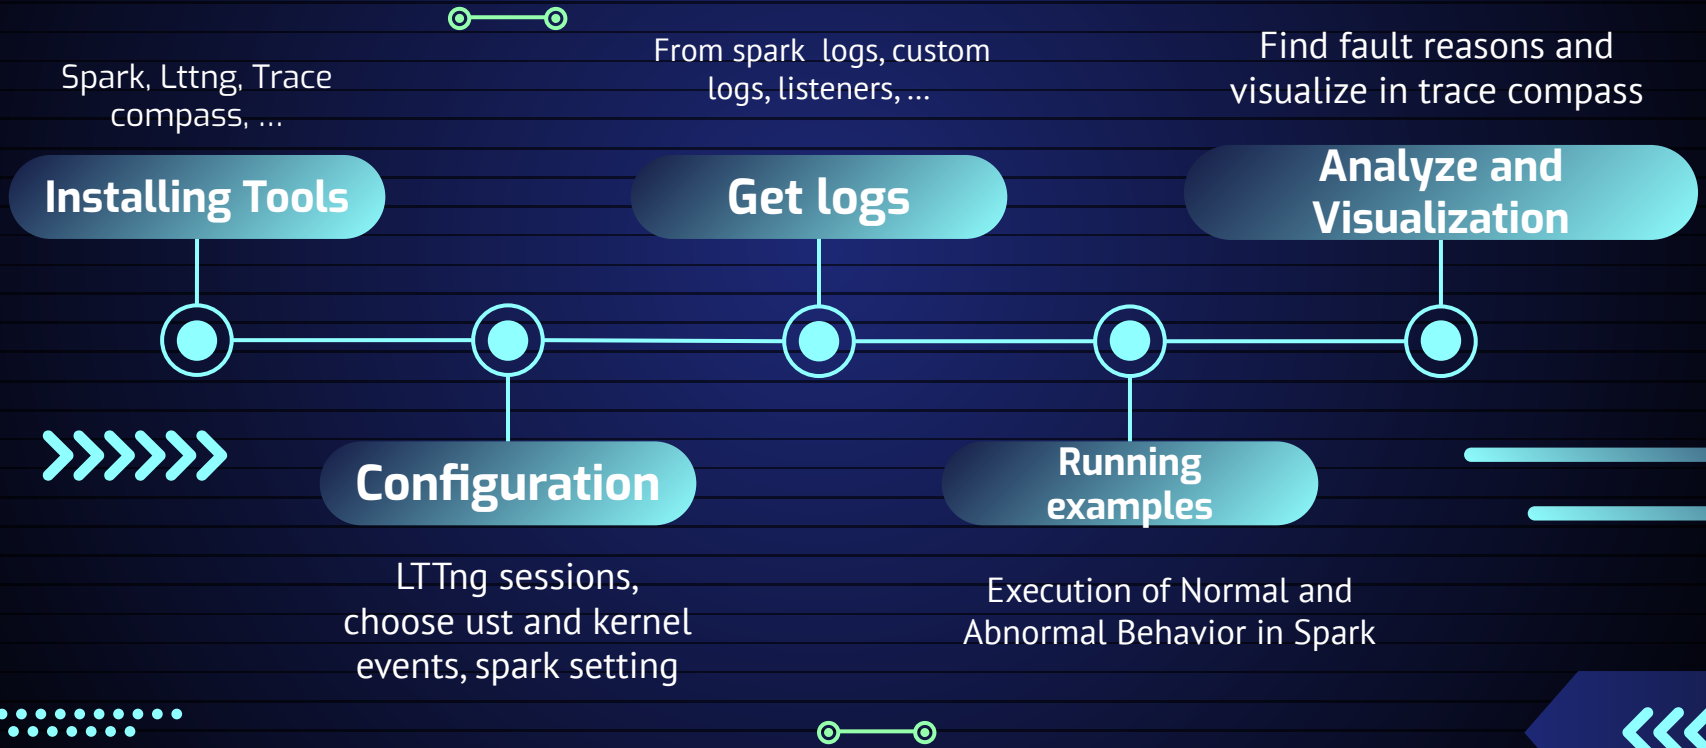

Methodology/ capturing logs

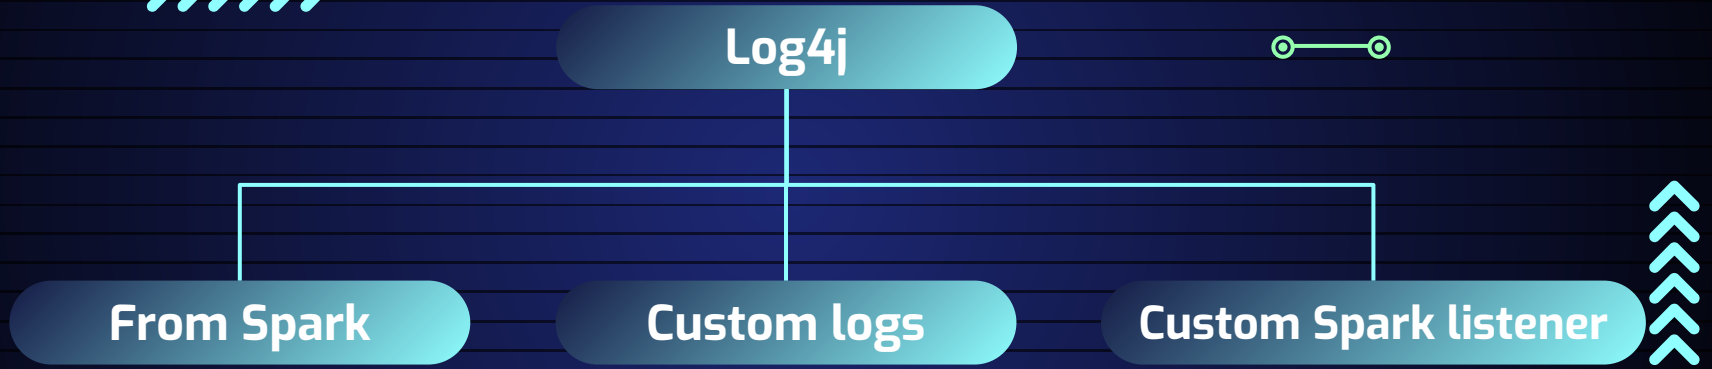

Methodology/ capturing logs

1) Log4j from Spark

1. Adding log4j2.xml
2. Create a pattern for logs in Console
3. There is no pattern layout setting here from lttng
4. Add Thread_ID to LTTNG UST source code
5. Enable log4j logs and run Spark example

```
<?xml version="1.0" encoding="UTF-8"?>
<Configuration status="WARN" strict="true">
 <Appenders>
   <Lttng name="Lttng1" domain="LOG4J">
   </Lttng>
   <Lttng name="Lttng2" domain="LOG4J2">
   </Lttng>
 </Appenders>
 <Loggers>
   <Root level="INFO">
     <AppenderRef ref="Lttng1"/>
     <AppenderRef ref="Lttng2"/>
   </Root>
 </Loggers>
</Configuration>
```

Methodology/ a little change in LTTNG-UST

Logs result in lttng session

```
[19:51:23.497752838] (+0.018001959) Rezghool
lttng_log4j:event: { cpu_id = 7 },
{ msg = "Starts! Parent method: main by Reza",
logger_name = "org.apache.spark.sql.SparkSession",
class_name = "org.apache.spark.sql.SparkSession",
method_name = "read", filename = "SparkSession.scala",
line_number = 725, timestamp = 1700268683497,
int_loglevel = 20000,
thread_name = "main +
1 +
null +
org.apache.spark.sql.SparkSession.read(SparkSession.scala:725) +
5 +
null + org.apache.logging.slf4j.Log4jLogger" }
```

LTTNG UST add new code to LttngLogAppender class

```
        event.getThreadName() +
" + " + event.getMarker()  +
" + " + event.getSource() +
" + " + event.getThreadPriority() +
" + " + event.getThrown() +
" + " + event.getLoggerFqcn()
```

Methodology/ capturing logs

2) Custom logs

```
/**
 * Persist this RDD with the de
 */
def cache(): JavaRDD[T] = {

  wrapRDD(rdd.cache())

}
```

After

```
2
3   /**
4    * Persist this RDD with the default storage level (`MEMORY_ONLY`).
     */
    def cache(): JavaRDD[T] = {

      val callerMethodName = Thread.currentThread.getStackTrace()(2).getMethodName
      logger.info("Starts! Parent method: " + callerMethodName + " by Reza");

11    val res = wrapRDD(rdd.cache())
12
13    logger.info("End! by Reza")
14    res
15  }
16
```

Methodology/ capturing logs

2) Custom logs Result:

```
lttng_log4j:event: { cpu_id = 7 },
{ msg = "Starts! Parent method: main by Reza",
logger_name = "org.apache.spark.sql.SparkSession",
class_name = "org.apache.spark.sql.SparkSession",
method_name = "read", filename = "SparkSession.scala",
line_number = 725, timestamp = 1700268683497,
int_loglevel = 20000,
thread_name = "main +
1 +
null +
org.apache.spark.sql.SparkSession.read(SparkSession.scala:725) +
5 +
null + org.apache.logging.slf4j.Log4jLogger" }
```

Methodology/ capturing logs

3) Logs result in lttng session

```
msg = "Task end failed with error reason: Task
end info
- Stage ID: 35, Task ID: 65,
Executor ID: driver,
Duration: 93,
Task end reason::
ExceptionFailure(java.lang.NegativeArraySizeExcep
tion,
-727379968,[Ljava.lang.StackTraceElement;@5ed6f37
1,java.lang.NegativeArraySizeException:
-727379968
```

3) New Spark Listener

```
public class MyCustomSparkListener extends SparkListener {}
```

```
public void onJobEnd(SparkListenerJobEnd jobEnd) {
     String jobInfo = jobEnd.jobId() + ", Result: " + jobEnd.jobResult();
     if (!jobEnd.jobResult().toString().equals("JobSucceeded")) {
        logger.error("jobs end failed with error with Job ID: " + jobInfo);
     } else {
        logger.info("Job ended successfully with Job ID: " + jobInfo);
     }

  }
```

Key Findings/ Running the Spark example

1. Using kmeans ML example with a 10GB data input
2. Restricting Acphe resources
3. Input some other bottleneck in the code
4. Compare time overhead

spark.executor.memory	512M
total.executor.cores	1
spark.driver.memory	512M
Spark.yarn.am.memory	512m

Key Findings/ Running the Spark example

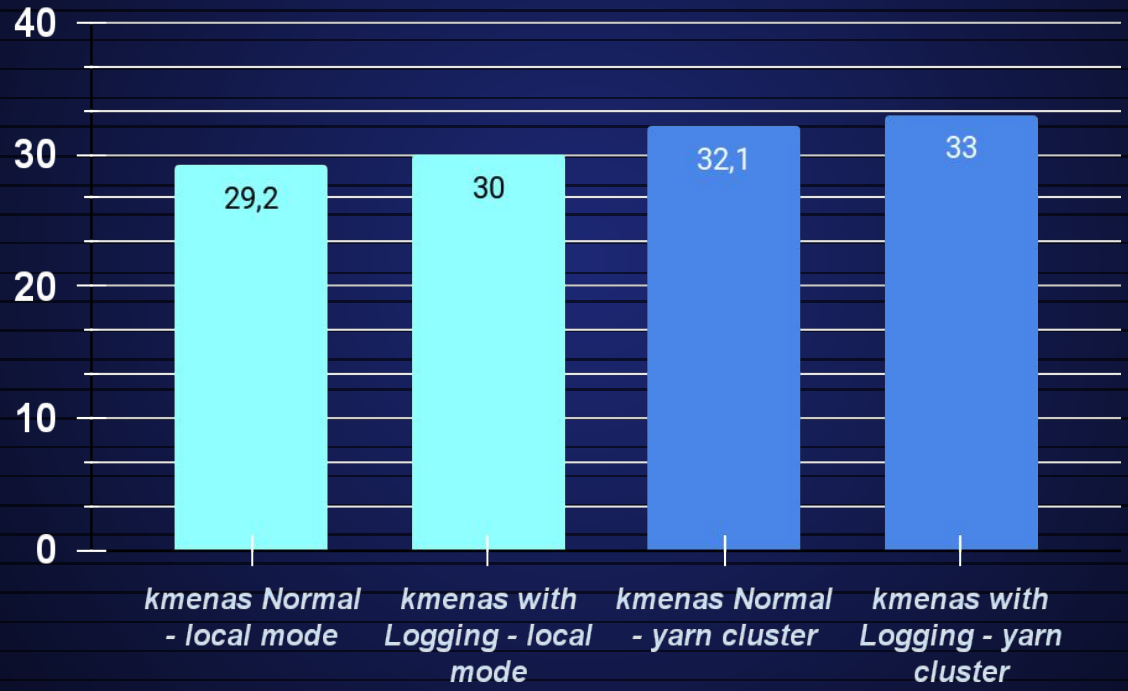

Key Findings/ Showing the fault in Spark UI

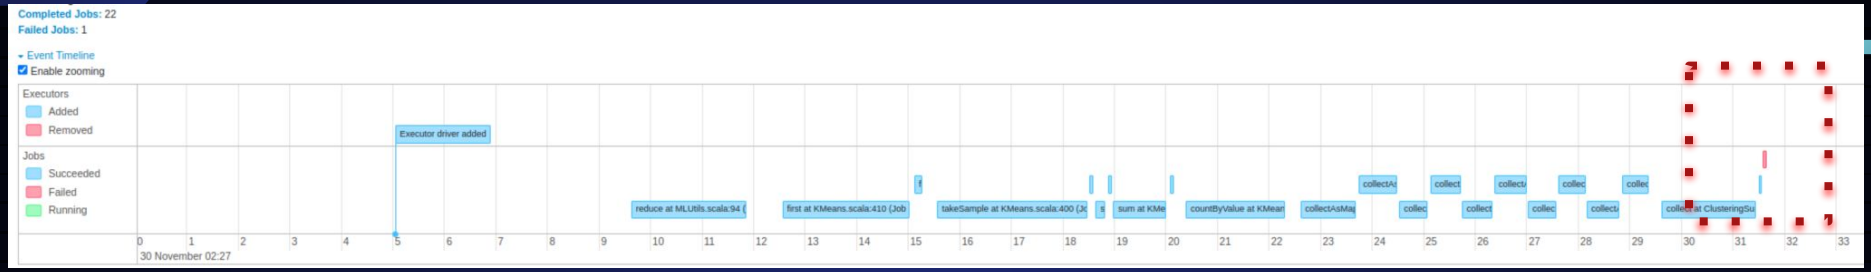

Failed Stages (1)

Page: 1

1 Pages. Jump to 1 . Show 100 items in a page. Go

Stage Id ▾	Description	Submitted	Duration	Tasks: Succeeded/Total	Input	Output	Shuffle Read	Shuffle Write	Failure Reason
35	foreachPartition at JavaKMeansExample.java:102 +details	2023/12/05 03:59:52	51 ms	0/2 (1 failed) (1 killed:	64.0 KiB				Job aborted due to stage failure: Task 0 in stage 35.0 failed 1 times, most recent failure: Lost task 0.0 in stage 35.0 (TID 64) (Rezghool.ht.home executor driver): org.apache.spark.SparkException: Intentional failure in stage +details

Page: 1

1 Pages. Jump to 1 . Show 100 items in a page. Go

Index	Task ID	Attempt	Status	Locality level	Executor ID	Host	Logs	Launch Time	Duration	GC Time	Input Size / Records	Errors
0	64	0	FAILED	PROCESS_LOCAL	driver	Rezghool.ht.home		2023-12-05 03:59:52	27.0 ms			org.apache.spark.SparkException: Intentional failure in stage +details

```
org.apache.spark.SparkException: Intentional failure in stage
        at
org.apache.spark.examples.ml.JavaKMeansExample.lambda$main$dc6d3fb3$1(JavaKMeansExample.java:104)
        at
org.apache.spark.sql.Dataset.$anonfun$foreachPartition$2(Dataset.scala:3370)
        at
org.apache.spark.sql.Dataset.$anonfun$foreachPartition$2$adapted(Dataset.scala:3370)
        at
```

Key Findings/ Showing the fault in Trace Compass

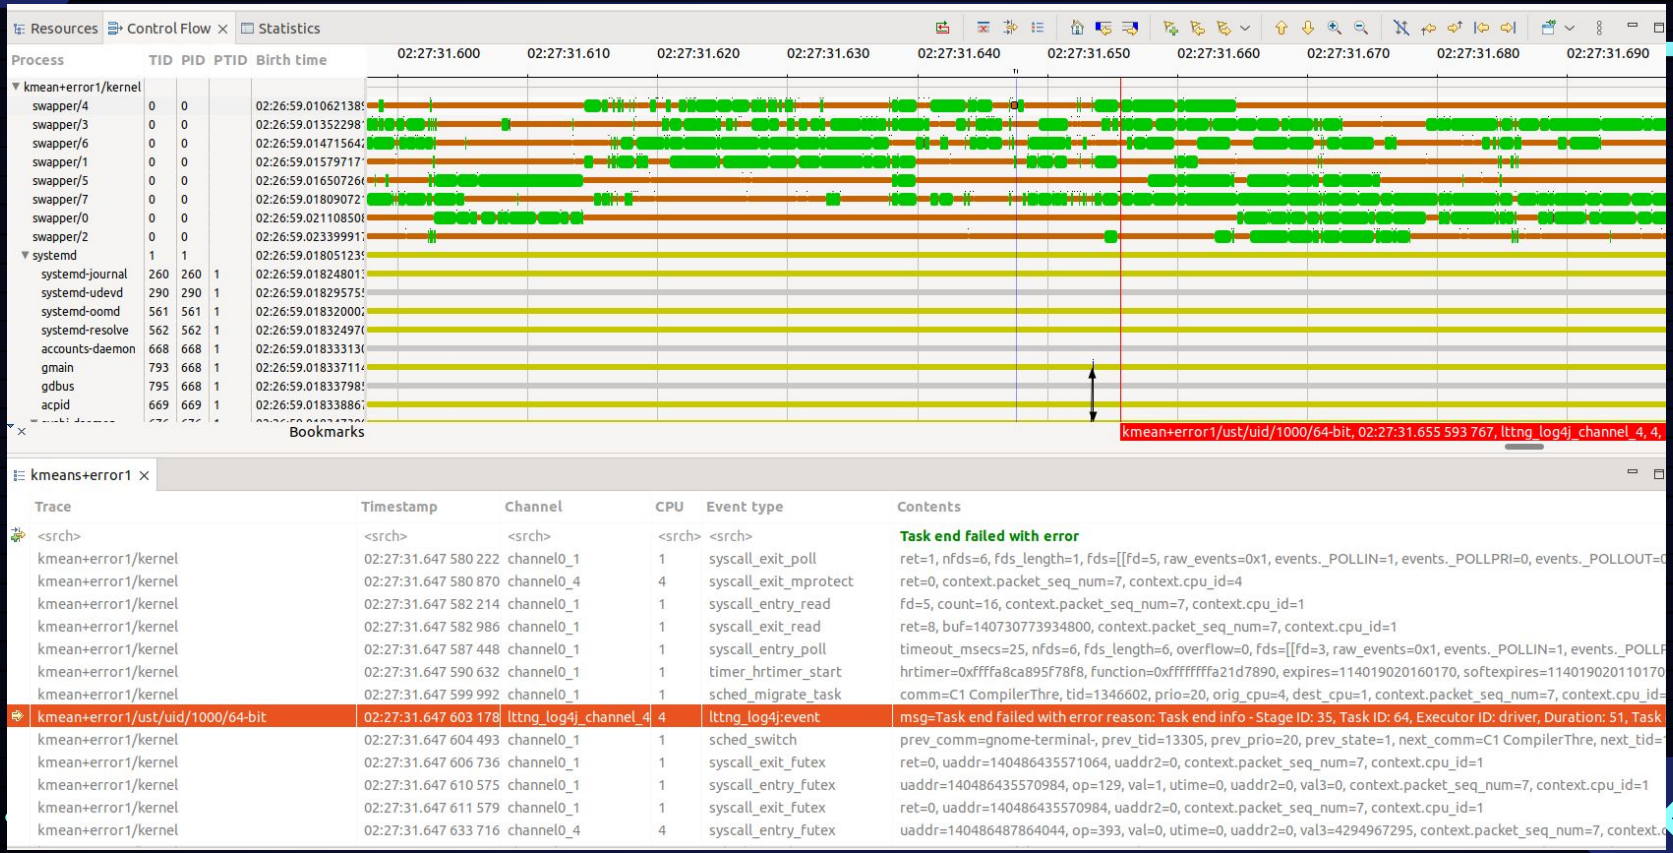

Conclusion and Future Work

1. I have all required logs from Apache spark logs with a little overhead in my lttng session
2. Adding more information from kernel events

1. Visualize the logs like Spark UI
2. Show the user which part of spark have issue related to their code

Thanks!

Do you have any questions?

Credits: This presentation template was created by **Slidesgo**, including icons by **Flaticon**, infographics & images by **Freepik**