

Debugging Spark Applications

Melike GEÇER

Supervised by


Dr. Haidar Osman

and

Prof. Oscar Nierstrasz

10.12.2019

Why is debugging hard in distributed systems?



Why is debugging hard?

Lack of tooling



About the Cluster

Logged in as: dr.w00t

Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Memory Reserved	VCores Used	VCores Total	VCores Reserved
0	0	0	0	0	0 B	6 GB	0 B	0	32	0

Cluster Nodes Metrics

Active Nodes	Decommissioning Nodes	Decommissioned Nodes	Lost Nodes	Unhealthy Nodes	Rebooted Nodes	Shutdown Nodes
4	0	0	0	0	0	0

Scheduler Metrics

Scheduler Type	Scheduling Resource Type	Minimum Allocation	Maximum Allocation	Maximum Cluster Application Priority
Capacity Scheduler	[memory-mb (unit=M), vcores]	<memory:128, vCores:1>	<memory:1536, vCores:4>	0

Cluster overview

Cluster ID: 157562052197

ResourceManager state: STARTED

ResourceManager HA state: active

ResourceManager HA zookeeper connection state: Could not find leader elector. Verify both HA and automatic failover are enabled.

ResourceManager RMStateStore: org.apache.hadoop.yarn.server.resourcemanager.recovery.NullRMStateStore

ResourceManager started on: Fri Dec 06 09:22:05 +0100 2019

ResourceManager version: 3.1.2 from 1019dde65bcf12e05ef48ac71e84550d589e5d9a by sunlig source checksum 7954337bd9688eca8aa32720d2c74 on 2019-01-29T02:04Z

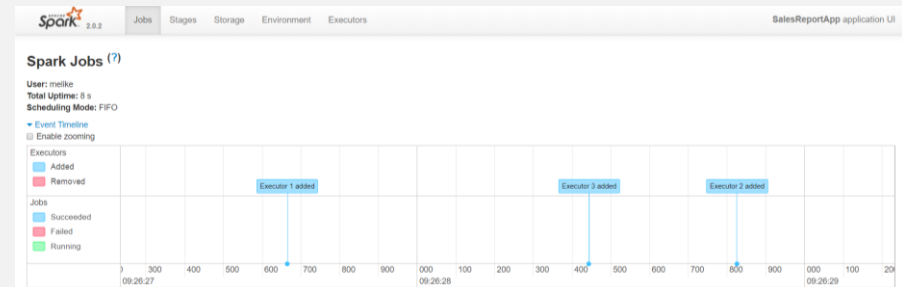
Hadoop version: 3.1.2 from 1019dde65bcf12e05ef48ac71e84550d589e5d9a by sunlig source checksum 64b8bdd4ca6e77cce75a93e09ab2a9 on 2019-01-29T01:39Z

```
melike@leela: ~$ cat /tmp/spark-66e379c7-9a68-45e1-bb29-df4ab63c75
019-12-06 09:26:30,440 INFO yarn.Client: Application report for application_157562052197_0001 (state: RUNNING)
019-12-06 09:26:31,446 INFO yarn.Client: Application report for application_157562052197_0001 (state: FINISHED)
019-12-06 09:26:31,449 INFO yarn.Client:
  client token: N/A
  diagnostics: User class threw exception: org.apache.spark.SparkException: Task not serializable
  ApplicationMaster host: 138.92.65.82
  ApplicationMaster RPC port: 0
  queue: default
  start time: 15756205753838
  final status: FAILED
  tracking URL: http://leela.unibe.ch:8088/proxy/application_157562052197_0001/
  user: melike
Exception in thread "main" org.apache.spark.SparkException: Application application_157562052197_0001 finished with failed status
    at org.apache.spark.deploy.yarn.Client.run(Client.scala:1132)
    at org.apache.spark.deploy.yarn.Client$.main(Client.scala:1178)
    at org.apache.spark.deploy.yarn.Client.main(Client.scala)
    at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
    at sun.reflect.NativeMethodAccessorImpl.invoke(Native Method)
    at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
    at java.lang.reflect.Method.invoke(Method.java:586)
    at org.apache.spark.deploy.SparkSubmit$.org$apache$spark$deploy$SparkSubmit$$runMain(SparkSubmit.scala:736)
    at org.apache.spark.deploy.SparkSubmit$.doRunMain$1(SparkSubmit.scala:185)
    at org.apache.spark.deploy.SparkSubmit$.submit(SparkSubmit.scala:210)
    at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala:124)
    at org.apache.spark.deploy.SparkSubmit.main(SparkSubmit.scala)
019-12-06 09:26:31,466 INFO util.ShutdownHookManager: Shutdown hook called
019-12-06 09:26:31,469 INFO util.ShutdownHookManager: Deleting directory /tmp/spark-66e379c7-9a68-45e1-bb29-df4ab63c75
```

Logs on the console

Logs on the Resource Manager

```
Java HotSpot(TM) 64-Bit Server VM warning: Ignoring option MaxPermSize=256m; support was removed in 8.0
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/tmp/hadoop-melike/nm-local-dir/usercache/melike/filecache/10/_spark_libs_7564521103123317951_zip/slf4j-log
SLF4J: Found binding in [jar:file:/opt/hadoop-3.1.2/share/hadoop/common/lib/slf4j-log4j12-1.7.25.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
2019-06-24 00:40:02,129 INFO util.SignalUtils: Registered signal handler for TERM
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for HUP
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for INT
2019-06-24 00:40:02,720 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where
2019-06-24 00:40:02,973 INFO yarn.ApplicationMaster: Preparing local resources
2019-06-24 00:40:03,935 INFO yarn.ApplicationMaster: Prepared local resources Map(__app__jar -> resource { scheme: "hdfs" host: "leela.unibe.c
2019-06-24 00:40:03,970 INFO yarn.ApplicationMaster: ApplicationAttemptId: appattemp1_1561329314867_0002_000002
2019-06-24 00:40:03,980 INFO spark.SecurityManager: Changing view acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing modify acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing view acls groups to:
2019-06-24 00:40:03,982 INFO spark.SecurityManager: Changing modify acls groups to:
2019-06-24 00:40:03,983 INFO spark.SecurityManager: SecurityManager: authentication disabled; ui acls disabled; users with view permissions: S
2019-06-24 00:40:04,005 INFO yarn.ApplicationMaster: Starting the user application in a separate Thread
2019-06-24 00:40:04,022 INFO yarn.ApplicationMaster: Waiting for spark context initialization
2019-06-24 00:40:04,023 INFO yarn.ApplicationMaster: Waiting for spark context initialization ...
2019-06-24 00:40:04,048 INFO spark.SparkContext: Running Spark version 2.0.2
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing view acls to: melike
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing modify acls to: melike
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing view acls groups to:
2019-06-24 00:40:04,099 INFO spark.SecurityManager: Changing modify acls groups to:
2019-06-24 00:40:04,099 INFO spark.SecurityManager: SecurityManager: authentication disabled; ui acls disabled; users with view permissions: S
2019-06-24 00:40:04,427 INFO util.Utils: Successfully started service "sparkDriver" on port 42097.
2019-06-24 00:40:04,457 INFO spark.SparkEnv: Registering MapOutputTracker
2019-06-24 00:40:04,484 INFO spark.SparkEnv: Registering BlockManagerMaster
2019-06-24 00:40:04,502 INFO storage.DiskBlockManager: Created local directory at /tmp/hadoop-melike/nm-local-dir/usercache/melike/appcache/app
2019-06-24 00:40:04,523 INFO memory.MemoryStore: MemoryStore started with capacity 93.3 MB
2019-06-24 00:40:04,641 INFO spark.SparkEnv: Registering OutputCommitCoordinator
2019-06-24 00:40:04,758 INFO util.log: Logging initialized @946ms
2019-06-24 00:40:04,813 INFO ui.JettyUtils: Adding filter: org.apache.hadoop.yarn.server.webproxy.amfilter.AmIpFilter
2019-06-24 00:40:04,880 INFO server.Server: jetty-9.2.z-SNAPSHOT
2019-06-24 00:40:04,908 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@557d5a49{/jobs,null,AVAILABLE}
```



History Server

Why is debugging hard?

Data is always dirty

employee_no	employee_name	designation	manager	hire_date	salary	Department_no
7000	SMITH	CLERK	7902	12/17/1980	800	70
7499	ALLEN	SALESMAN	7698	2/20/1981	1600	30
7521	WARD	SALESMAN	7698	2/22/1981	1250	30
7566	TURNER	MANAGER	7839	4/2/1981	2975	20
7654	MARTIN	SALESMAN	7698	9/28/1981	1250	30
7698	MILLER	MANAGER	7839	5/1/1981	thousand	30
7782	CLARK	MANAGER	7839	6/9/1981	2450	10
7788	SCOTT	ANALYST	7566	12/9/1982	3000	20
7839	KING	PRESIDENT	2000	11/17/1981	5000	10
7844	TURNER	SALESMAN	7698	9/8/1981	1500	30
7876	ADAMS	CLERK	7788	1/12/1983	1100	20
7900	JAMES	CLERK	7698	12/3/1981	950	30
7902	FORD	ANALYST	7566	12/3/1981	3000	20
7934	MILLER	CLERK	7782	1/23/1982	1300	10

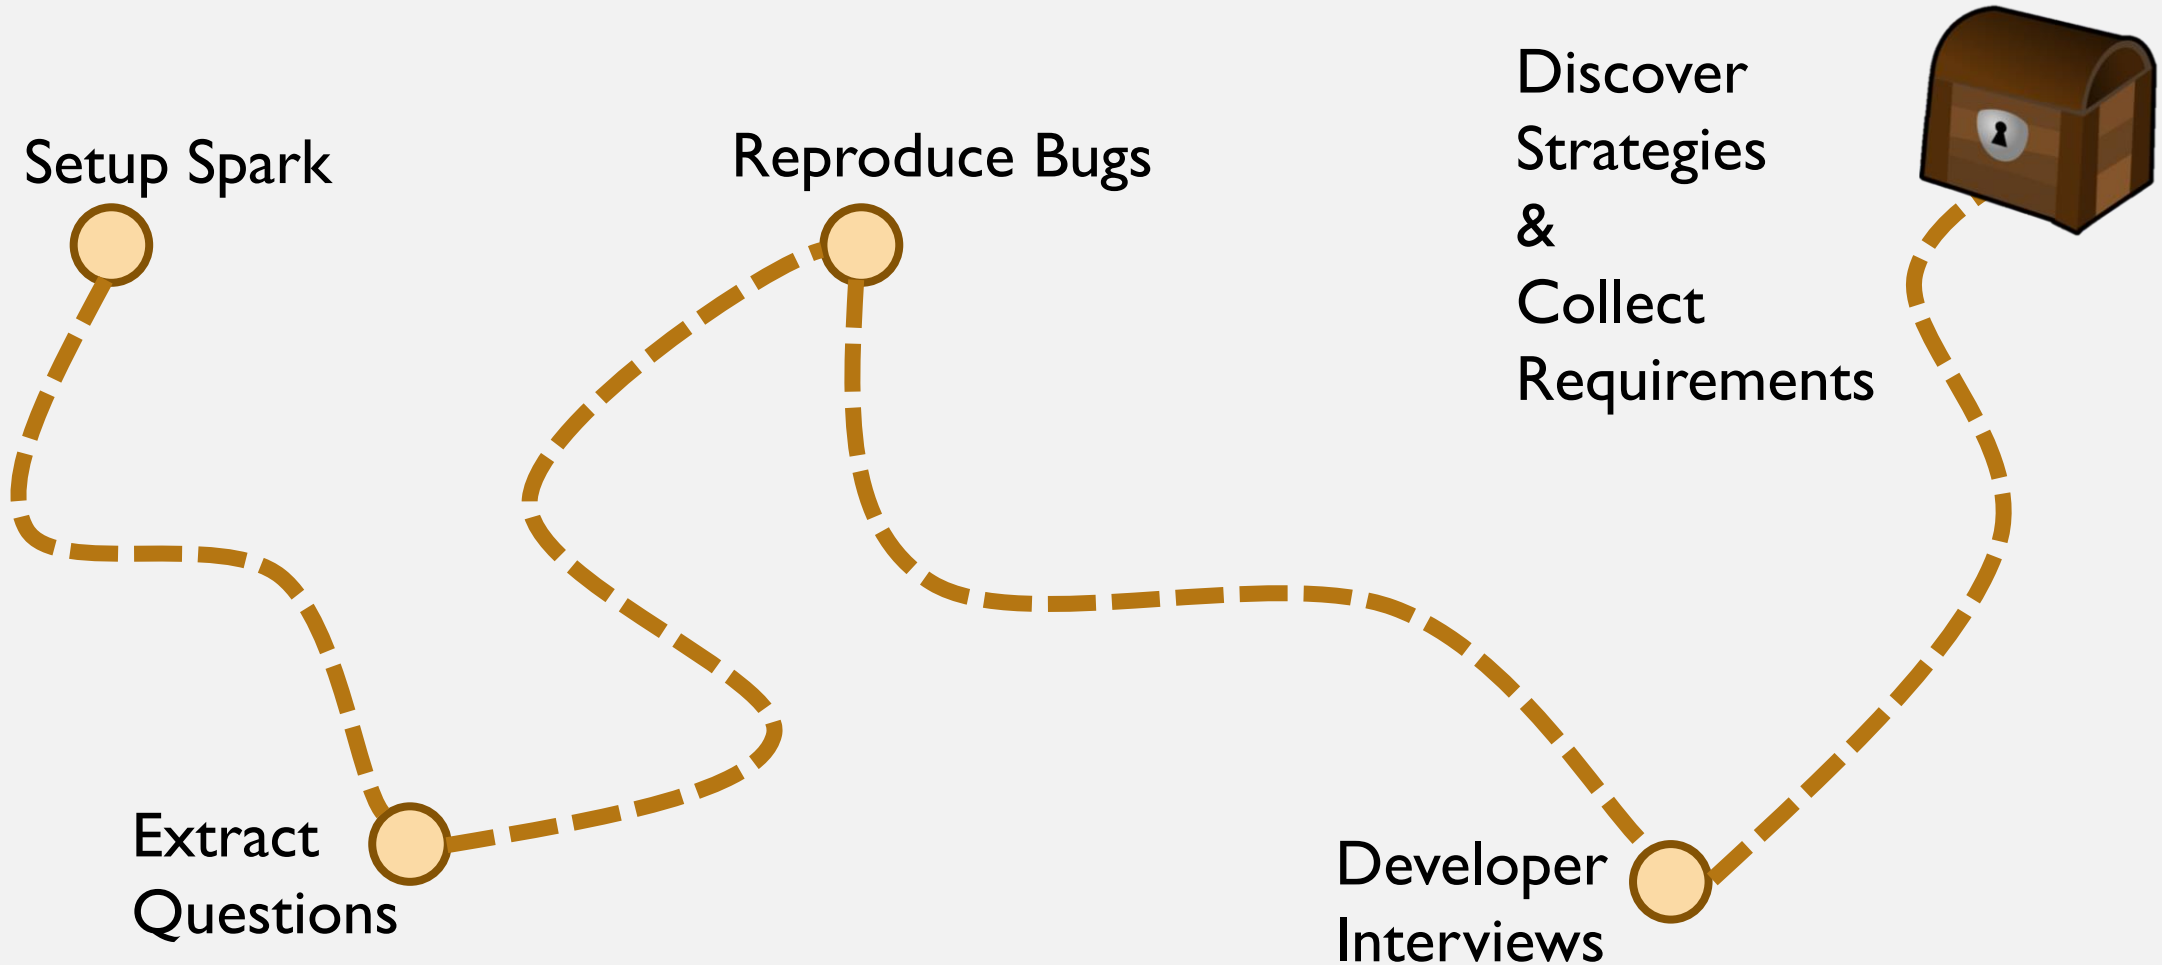
Research Questions

RQ1: How do developers debug distributed applications like Spark?

RQ2: What are the challenges that developers face when they debug such applications?

RQ3: How can we overcome these challenges?

Investigation Roadmap



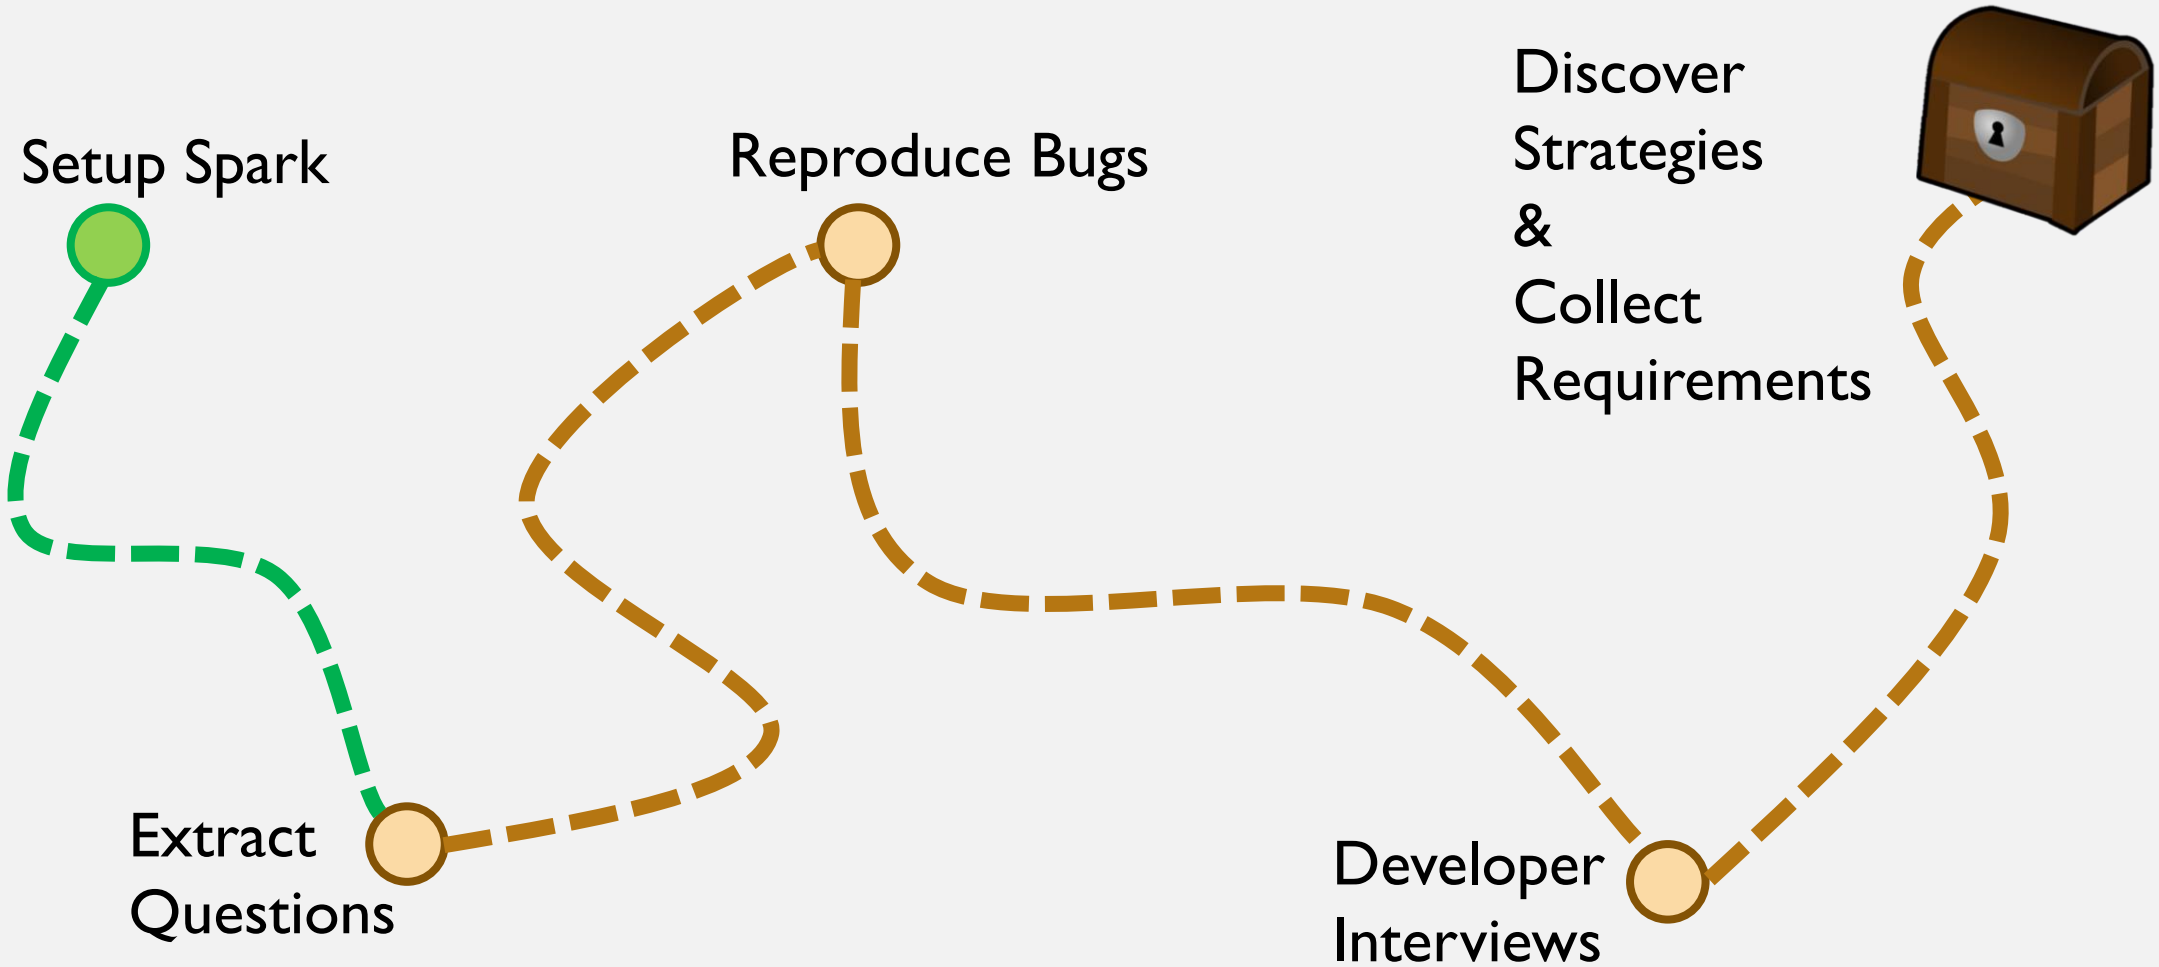
Setup Spark



ID	User	Name	Application Type	Queue
application_1561329314867_0002	melike	MainClass	SPARK	default

```
Java HotSpot(TM) 64-Bit Server VM warning: ignoring option MaxPermSize=256m; support was removed in 8.0
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/tmp/hadoop-melike/nm-local-dir/usercache/melike/filecache/10/_spark_libs_7564521103123317951.zip/slf4j-log4j12-1.7.16.jar!/org/slf4j/
SLF4J: Found binding in [jar:file:/opt/hadoop-3.1.2/share/hadoop/common/lib/slf4j-log4j12-1.7.25.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
2019-06-24 00:40:02,129 INFO util.SignalUtils: Registered signal handler for TERM
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for HUP
2019-06-24 00:40:02,131 INFO util.SignalUtils: Registered signal handler for INT
2019-06-24 00:40:02,720 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable
2019-06-24 00:40:02,973 INFO yarn.ApplicationMaster: Preparing Local resources
2019-06-24 00:40:03,935 INFO yarn.ApplicationMaster: Prepared Local resources Map(__app__.jar -> resource { scheme: "hdfs" host: "leela.unibe.ch" port: 9000 file: "/user/
2019-06-24 00:40:03,970 INFO yarn.ApplicationMaster: ApplicationAttemptId: appattempt_1561329314867_0002_000002
2019-06-24 00:40:03,980 INFO spark.SecurityManager: Changing view acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing modify acls to: melike
2019-06-24 00:40:03,981 INFO spark.SecurityManager: Changing view acls groups to:
2019-06-24 00:40:03,982 INFO spark.SecurityManager: Changing modify acls groups to:
```

Investigation Roadmap



Extract Questions



Extract question on
Stack Overflow

Have a pilot
interview

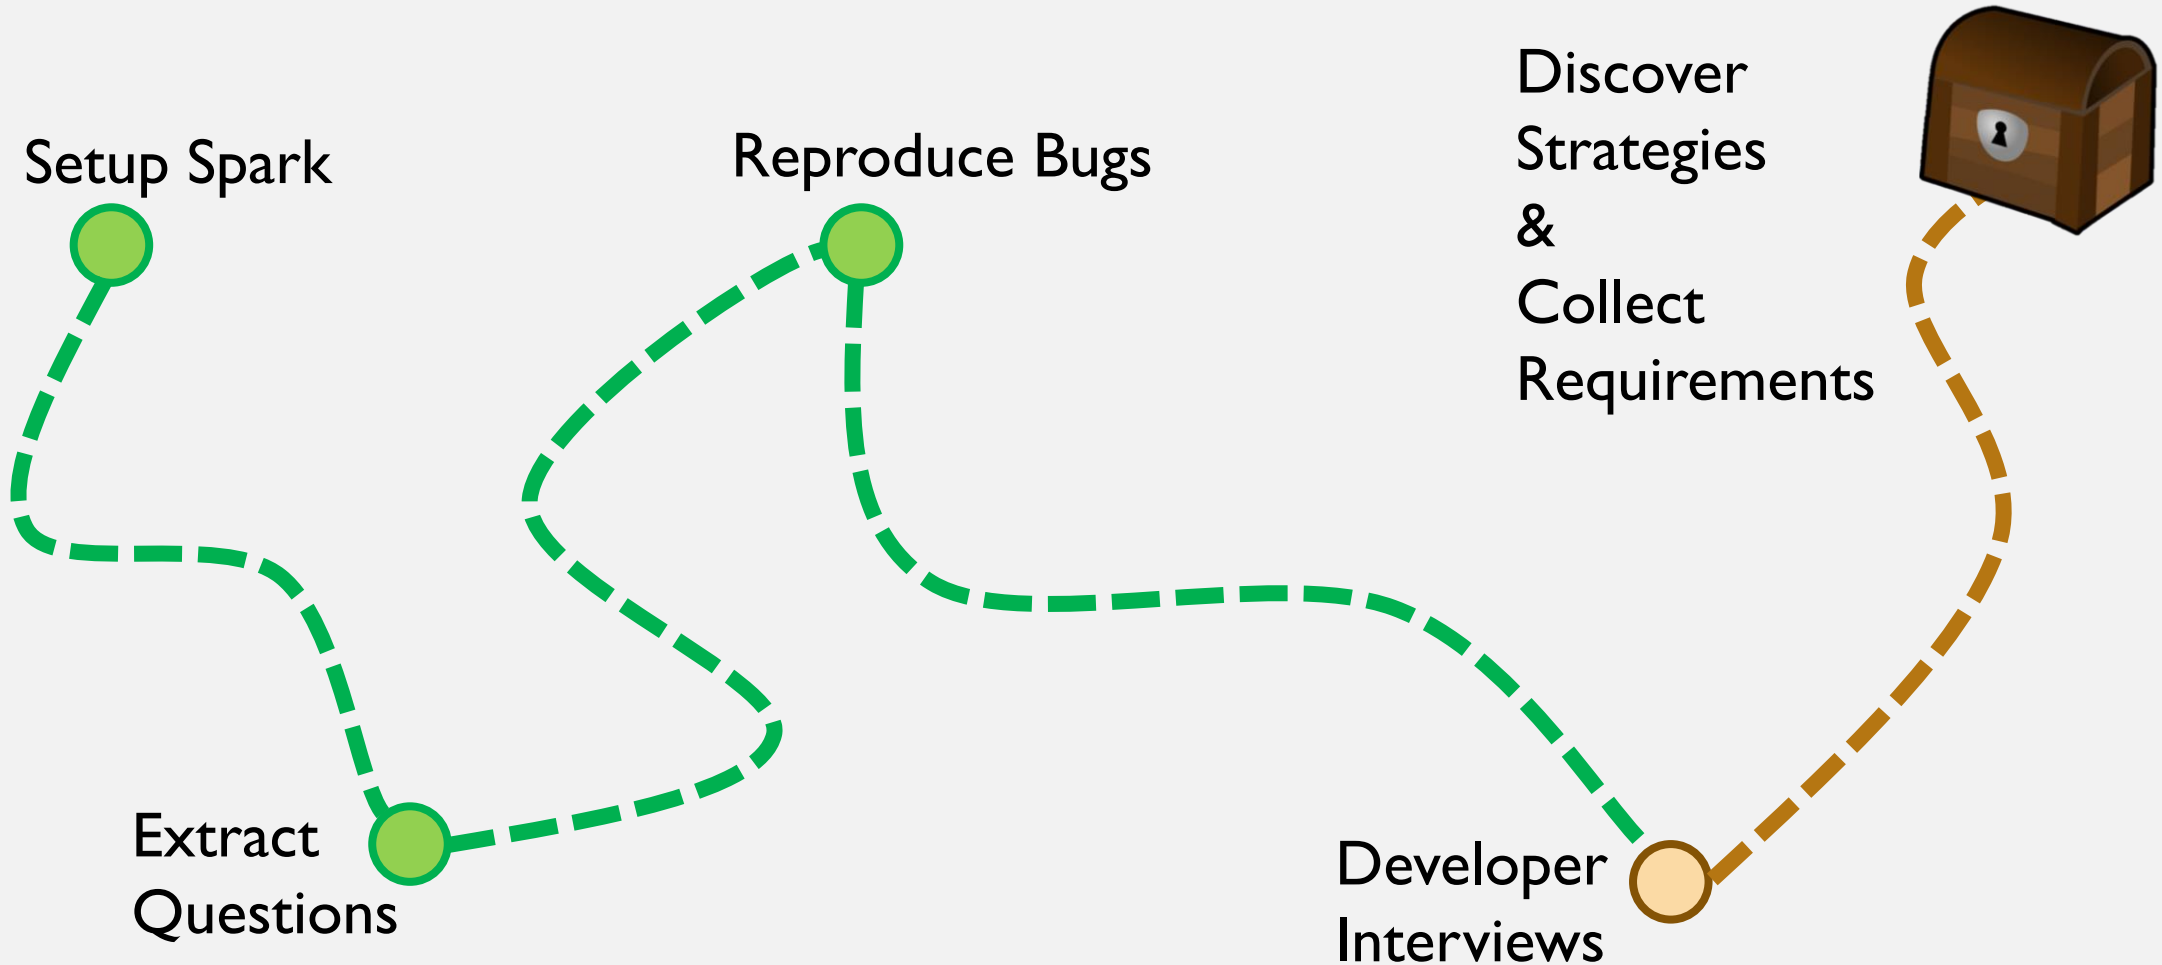
Talk to an expert
and ask a
recommendation

Review the
questions


Extract Questions

Exception	Reason	Solution
Task Not Serializable	The class is not serializable	<ul style="list-style-type: none">- Make the class Serializable- Change the method to a function- Copy the value to a local variable
Stack Overflow	Nested structures with many fields	<ul style="list-style-type: none">- Remove the nesting- Decrease the number of fields
Number Format Exception	A String value found where an integer expected in the production environment	<ul style="list-style-type: none">- Change the data- Discard the data
Stack Overflow	Making an action after 1000 transformations	<ul style="list-style-type: none">- Put checkpoints- Use caching- Decrease the number of transformations- Increase heap size of Java

Investigation Roadmap



Interview Setup



Semi-structured interview

An interviewer, an interviewee and an observer

With their consent, we record the screen and audio.

Interview Setup

Professional
Information

Coding
Interview


Post-mortem
Questionnaire

Professional Information

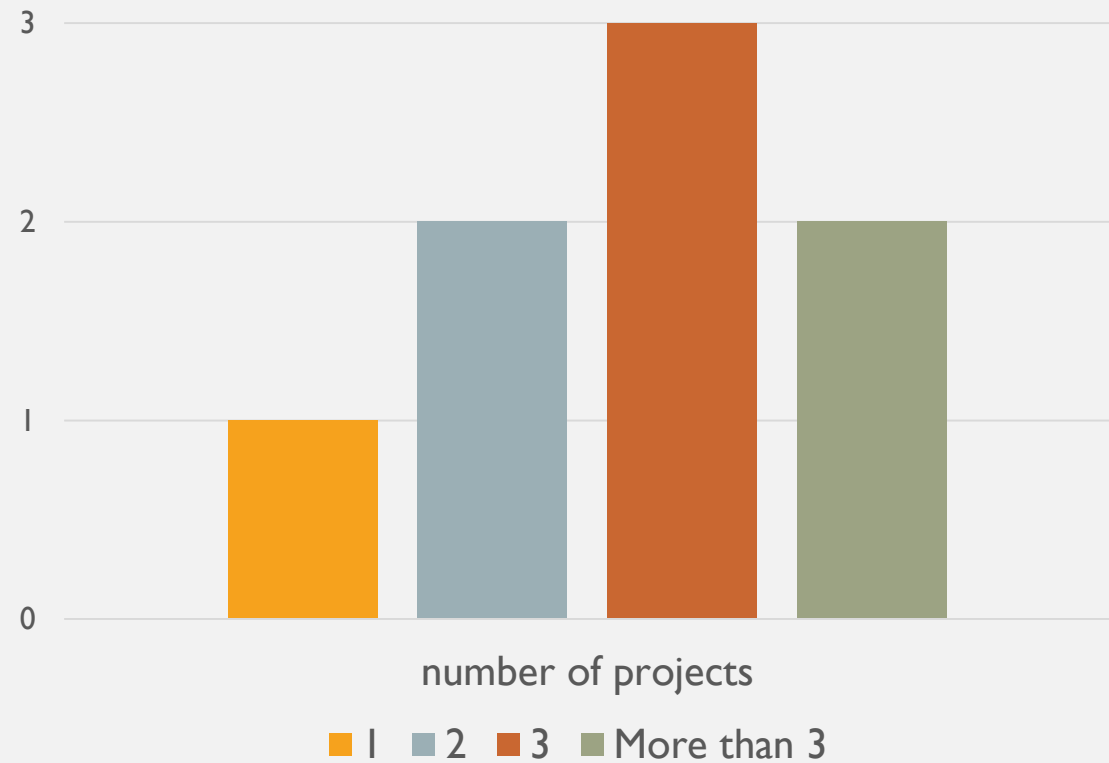
1. How long is your professional experience in Industry?
2. How long is your experience in Scala?
3. How long is your experience in Spark?
4. How many «different» projects that use Spark did you work on in your experience?
5. When you encounter a Spark job failure, what are your first steps to investigate?
6. What are your tools for debugging failed Spark jobs?

Interviewees

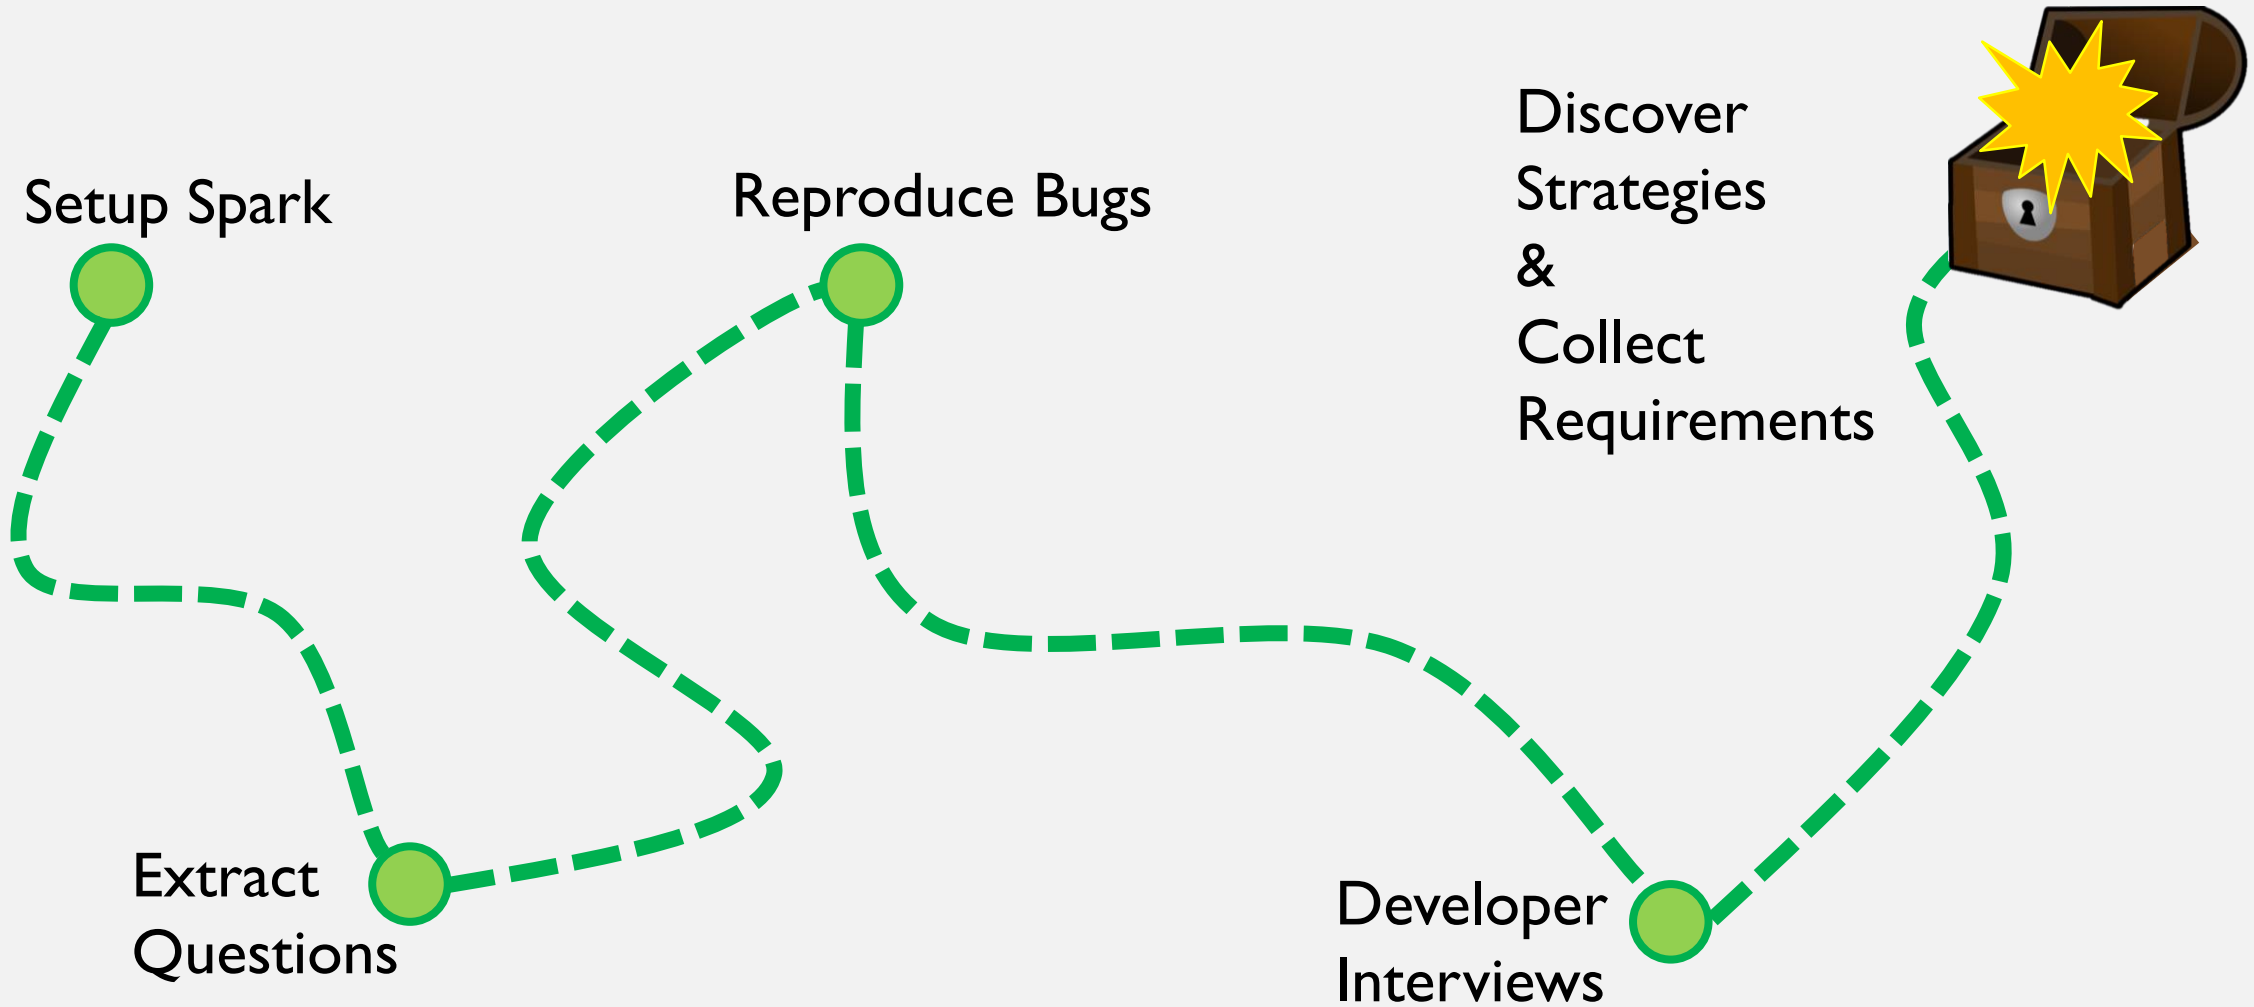
Experiences of interviewees



Number of «different» projects



Investigation Roadmap



Coding Interview - Patterns

A	Checks the logs on the console	Log Interactions
B	Sees the exception message	
C	Checks the logs on the Resource Manager	
E	Inspects the class that threw the exception	Code Inspection
N	Inspects the script	
D	Suspects something	Come Up w/ Hypothesis
F	Suggests to do something	
H	Makes changes in the code	Making Changes
P	Writes tests	
R	Changes the script	

Coding Interview - Patterns


Q	Runs their own tests	Test Execution
I	Runs the test with the new changes	
G	Runs the test without any changes in the code	
K	Searches the exception on the internet	Seek Help
J	Runs the code with the new changes on the cluster	Run On The Cluster
L	Checks the size of the file on HDFS	Checks The File's Size
M	Checks the file on cluster	Checks The File
O	Checks the file on local	

Coding Interview - Patterns

I #1	B #1	A	B	C	D	E	D	F	G	H	I			
	B #2	A	B	C	D	D	H	J						
	B #3	A	B	D										
	B #4	A	B	C	D	D								
I #2	B #1	A	R	J	A	B	E	K	H	I				
	B #2	C	B	D	L	D	D							
	B #3	A	G	B	D	M								
	B #4	A	N	C	D	D								
I #3	B #1	A	G	B	D	K	D	E	K	H	I			
	B #2	A	C	B	D	M	D	N	D					
	B #3	A	B	D	M									
	B #4	A	C	B	D	H	J							
I #4	B #1	M	A	C	B	D	H	A	C	H	G	C	E	H
	B #2	M	C	B	D	K	D	D	H					
	B #3	G	M	O	A	C	B	H	M					
	B #4	A	M	C	D	B	D	D						
I #5	B #1	A	B	G	H	I								
	B #2	C	B	D	H	D								
	B #3	G	A	B	F	O	A	D						
	B #4	A	B	M	D									
I #6	B #1	G	B	E	H	I								
	B #2	C	B	D	D									
	B #3	A	C	B	G									
	B #4	A	C	B	C	D								
I #7	B #1	A	D	C	B	E	C	G	D	H	I			
	B #2	A	C	B	D	D	P	Q						
	B #3	G	A	B	D									
	B #4	A	C	B	D	H								

A B C	Log Interactions
E N	Code Inspection
D F	Come Up w/ Hypothesis
H P R	Making Changes
Q I G	Test Execution
K	Seek Help
J	Run On The Cluster
L	Checks The File's Size
M O	Checks The File

Debugging Flow

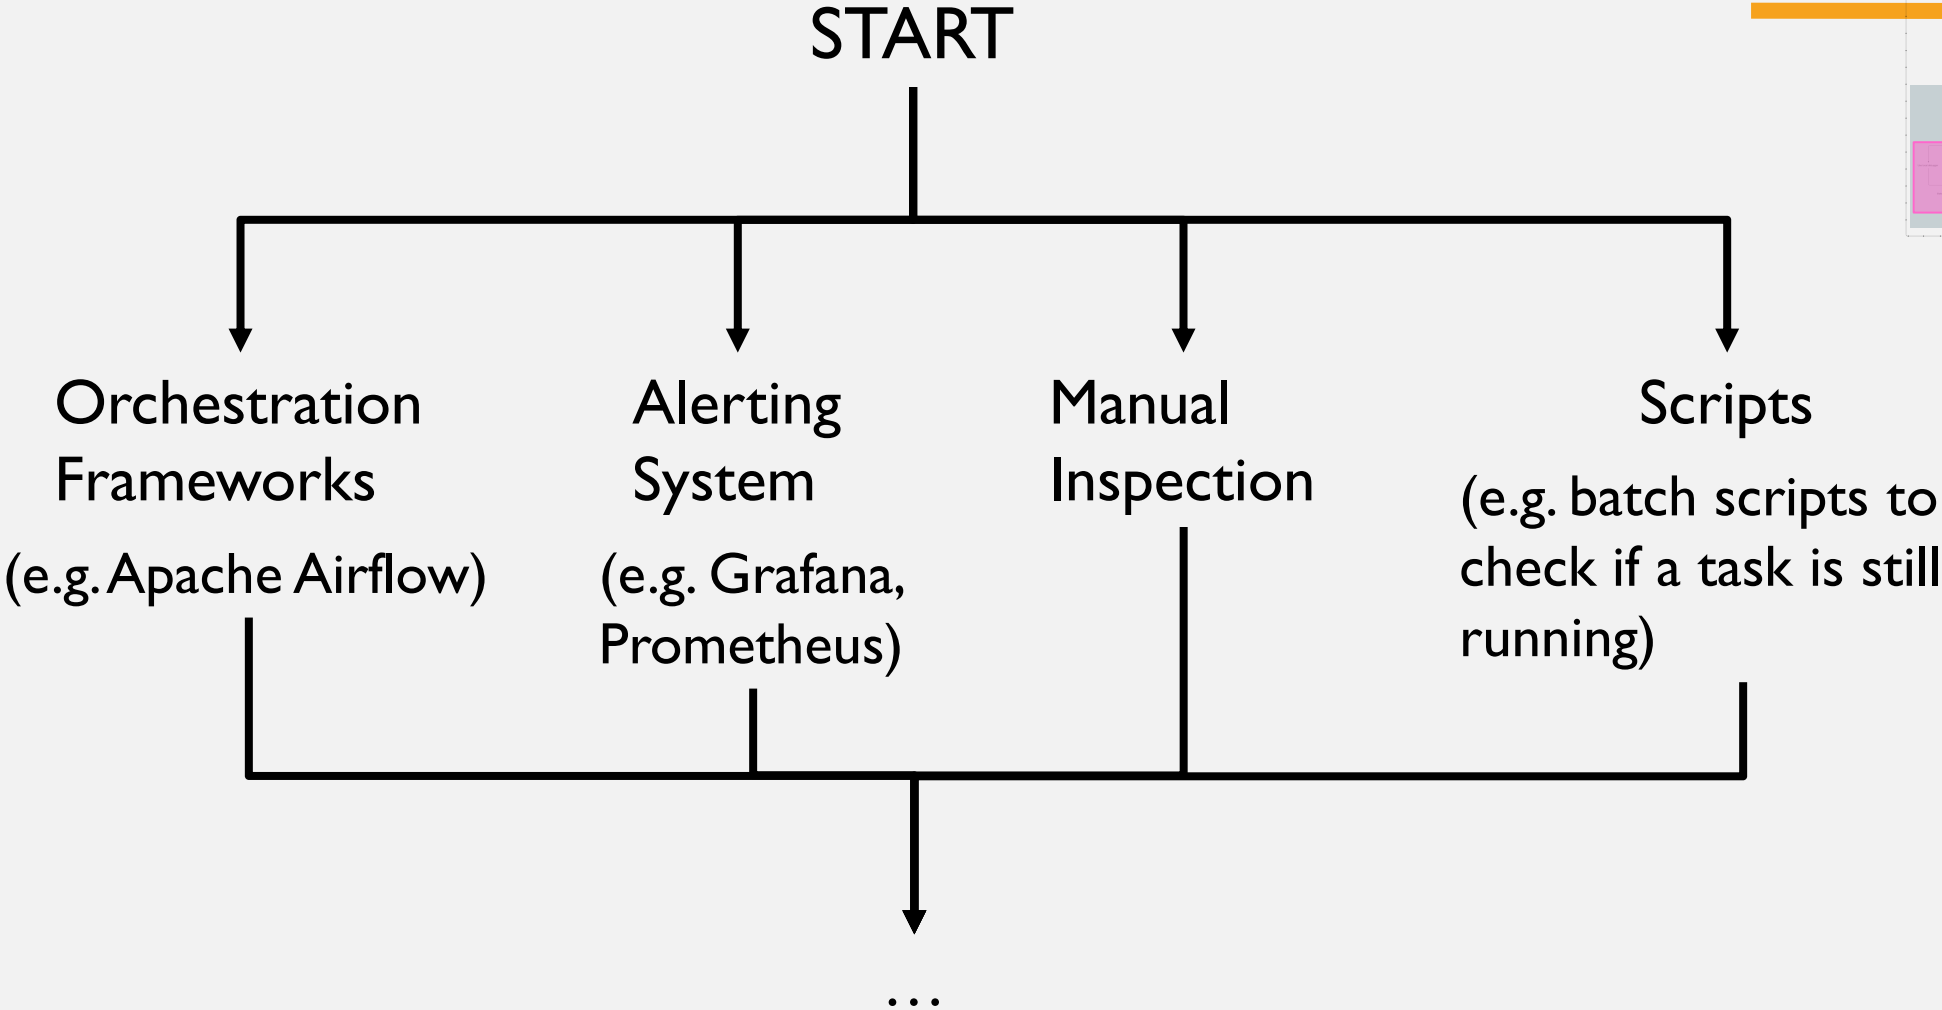


Learning the job has failed

Checking the logs

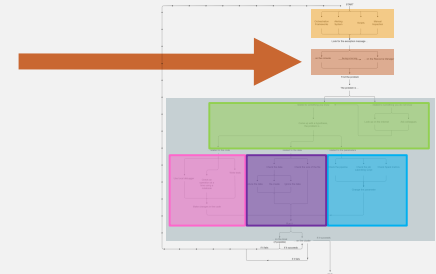
Solving the problem

Learning the job has failed



Checking the logs

```
2019-12-08 22:18:59,709 INFO yarn.Client: Application report for application_1575839673320_0002 (state: RUNNING)
2019-12-08 22:19:00,715 INFO yarn.Client: Application report for application_1575839673320_0002 (state: FINISHED)
2019-12-08 22:19:00,717 INFO yarn.Client:
  client token: N/A
  diagnostics: User class threw exception: org.apache.spark.SparkException: Task not serializable
  ApplicationMaster host: 130.92.65.82
  ApplicationMaster RPC port: 0
  name: default
  start time: 1575839983236
  Task status: FAILED
  Tracking URL: http://leela.unibe.ch:8088/proxy/application_1575839673320_0002/
  user: melike
Exception in thread "main" org.apache.spark.SparkException: Application application_1575839673320_0002 finished with failed status
at org.apache.spark.deploy.yarn.Client.run(Client.scala:1132)
at org.apache.spark.deploy.yarn.Client$.main(Client.scala:1173)
at org.apache.spark.deploy.yarn.Client.main(Client.scala)
at sun.reflect.NativeMethodAccessorImpl.invoke(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:498)
at org.apache.spark.deploy.SparkSubmit$.doRunMainIn(SparkSubmit.scala:185)
at org.apache.spark.deploy.SparkSubmit$.submit(SparkSubmit.scala:220)
at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala:124)
at org.apache.spark.deploy.SparkSubmit.main(SparkSubmit.scala)
2019-12-08 22:19:00,742 INFO util.ShutdownHookManager: Shutdown hook called
2019-12-08 22:19:00,745 INFO util.ShutdownHookManager: Deleting directory /tmp/spark-ba7ea2d-1919-486f-bff4-161c2c1b2
```



Look for the exception message ...

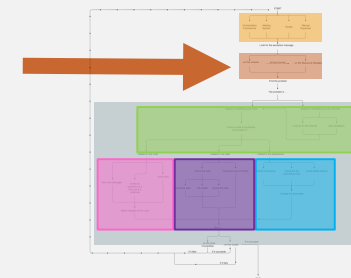
... on the console The log is too long ... on the Resource Manager

Find the problem

```
hadoop
  ResourceManager
  NodeManager
  Tools
  Query: 64-Bit Server VM warning: ignoring option MaxPermSize; support was removed in 8.0
  11:41:16: Error creating metrics: java.lang.NoClassDefFoundError: org.apache.hadoop.metrics.impl.MetricRegistryImpl
  11:41:16: Error loading in jar file (com.nimbusds.jose:apac...:1.11.2) (org.apache.hadoop.metrics.impl.MetricRegistryImpl)
  11:41:16: Error loading in jar file (com.nimbusds.jose:apac...:1.11.2) (org.apache.hadoop.metrics.impl.MetricRegistryImpl)
  11:41:16: Actual loading in jar file (org.apache.hadoop.metrics.impl.MetricRegistryImpl)
  2019-12-08 22:18:59,709 INFO yarn.Client: Application report for application_1575839673320_0002 (state: RUNNING)
  2019-12-08 22:19:00,715 INFO yarn.Client: Application report for application_1575839673320_0002 (state: FINISHED)
  2019-12-08 22:19:00,717 INFO yarn.Client:
    client token: N/A
    diagnostics: User class threw exception: org.apache.spark.SparkException: Task not serializable
    ApplicationMaster host: 130.92.65.82
    ApplicationMaster RPC port: 0
    name: default
    start time: 1575839983236
    Task status: FAILED
    Tracking URL: http://leela.unibe.ch:8088/proxy/application_1575839673320_0002/
    user: melike
  Exception in thread "main" org.apache.spark.SparkException: Application application_1575839673320_0002 finished with failed status
  at org.apache.spark.deploy.yarn.Client.run(Client.scala:1132)
  at org.apache.spark.deploy.yarn.Client$.main(Client.scala:1173)
  at org.apache.spark.deploy.yarn.Client.main(Client.scala)
  at sun.reflect.NativeMethodAccessorImpl.invoke(Native Method)
  at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57)
  at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
  at java.lang.reflect.Method.invoke(Method.java:498)
  at org.apache.spark.deploy.SparkSubmit$.doRunMainIn(SparkSubmit.scala:185)
  at org.apache.spark.deploy.SparkSubmit$.submit(SparkSubmit.scala:220)
  at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala:124)
  at org.apache.spark.deploy.SparkSubmit.main(SparkSubmit.scala)
  2019-12-08 22:19:00,742 INFO util.ShutdownHookManager: Shutdown hook called
  2019-12-08 22:19:00,745 INFO util.ShutdownHookManager: Deleting directory /tmp/spark-ba7ea2d-1919-486f-bff4-161c2c1b2
```

Checking the logs ... on the console

```
melike@leela: ~  
2019-12-08 22:18:59,709 INFO yarn.Client: Application report for application_1575839673320_0002 (state: RUNNING)  
2019-12-08 22:19:00,714 INFO yarn.Client: Application report for application_1575839673320_0002 (state: FINISHED)  
2019-12-08 22:19:00,717 INFO yarn.Client:  
  client token: N/A  
  diagnostics: User class threw exception: org.apache.spark.SparkException: Task not serializable  
  ApplicationMaster host: 130.92.65.82  
  ApplicationMaster RPC port: 0  
  queue: default  
  start time: 1575839903236  
  final status: FAILED  
  tracking URL: http://leela.unibe.ch:8088/proxy/application_1575839673320_0002/  
  user: melike  
Exception in thread "main" org.apache.spark.SparkException: Application application_1575839673320_0002 finished with failed status  
  at org.apache.spark.deploy.yarn.Client.run(Client.scala:1132)  
  at org.apache.spark.deploy.yarn.Client$.main(Client.scala:1178)  
  at org.apache.spark.deploy.yarn.Client.main(Client.scala)  
  at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)  
  at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57)  
  at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)  
  at java.lang.reflect.Method.invoke(Method.java:606)  
  at org.apache.spark.deploy.SparkSubmit$.org$apache$spark$deploy$SparkSubmit$$runMain(SparkSubmit.scala:736)  
  at org.apache.spark.deploy.SparkSubmit$.doRunMain$1(SparkSubmit.scala:185)  
  at org.apache.spark.deploy.SparkSubmit$.submit(SparkSubmit.scala:210)  
  at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala:124)  
  at org.apache.spark.deploy.SparkSubmit.main(SparkSubmit.scala)  
2019-12-08 22:19:00,742 INFO util.ShutdownHookManager: Shutdown hook called  
2019-12-08 22:19:00,745 INFO util.ShutdownHookManager: Deleting directory /tmp/spark-ba87ea2d-1919-486f-bff4-161c2c1be12b  
melike@leela:~$
```

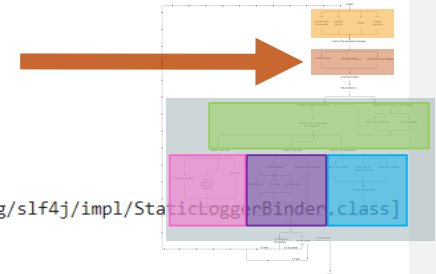


Checking the logs ... on the Resource Manager

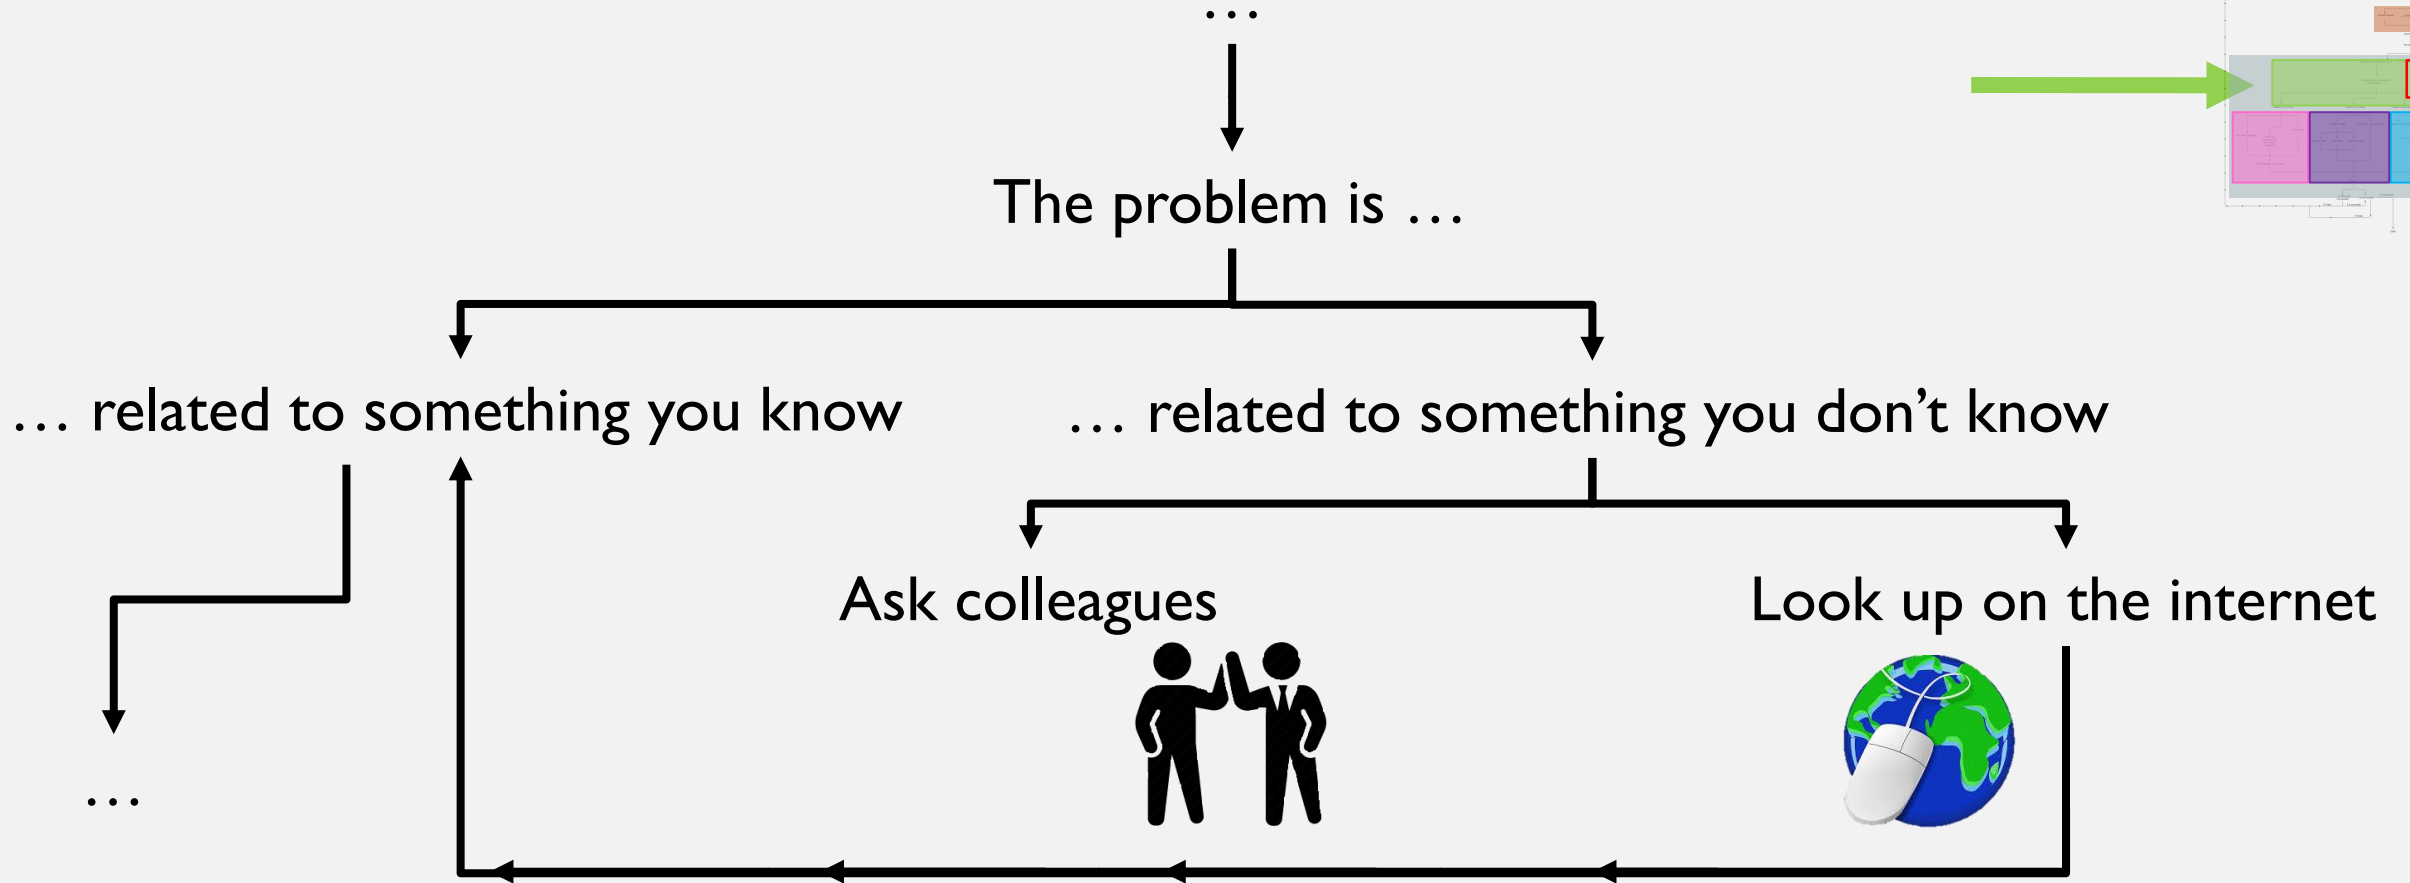


- ▼ ResourceManager
 - RM Home
- ▶ NodeManager
- ▶ Tools

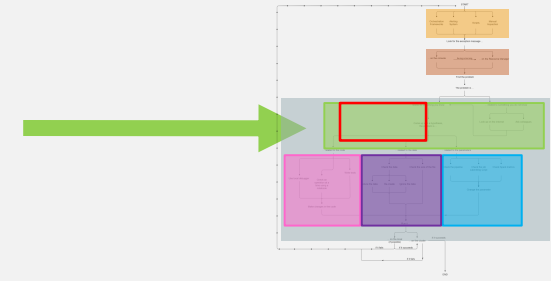
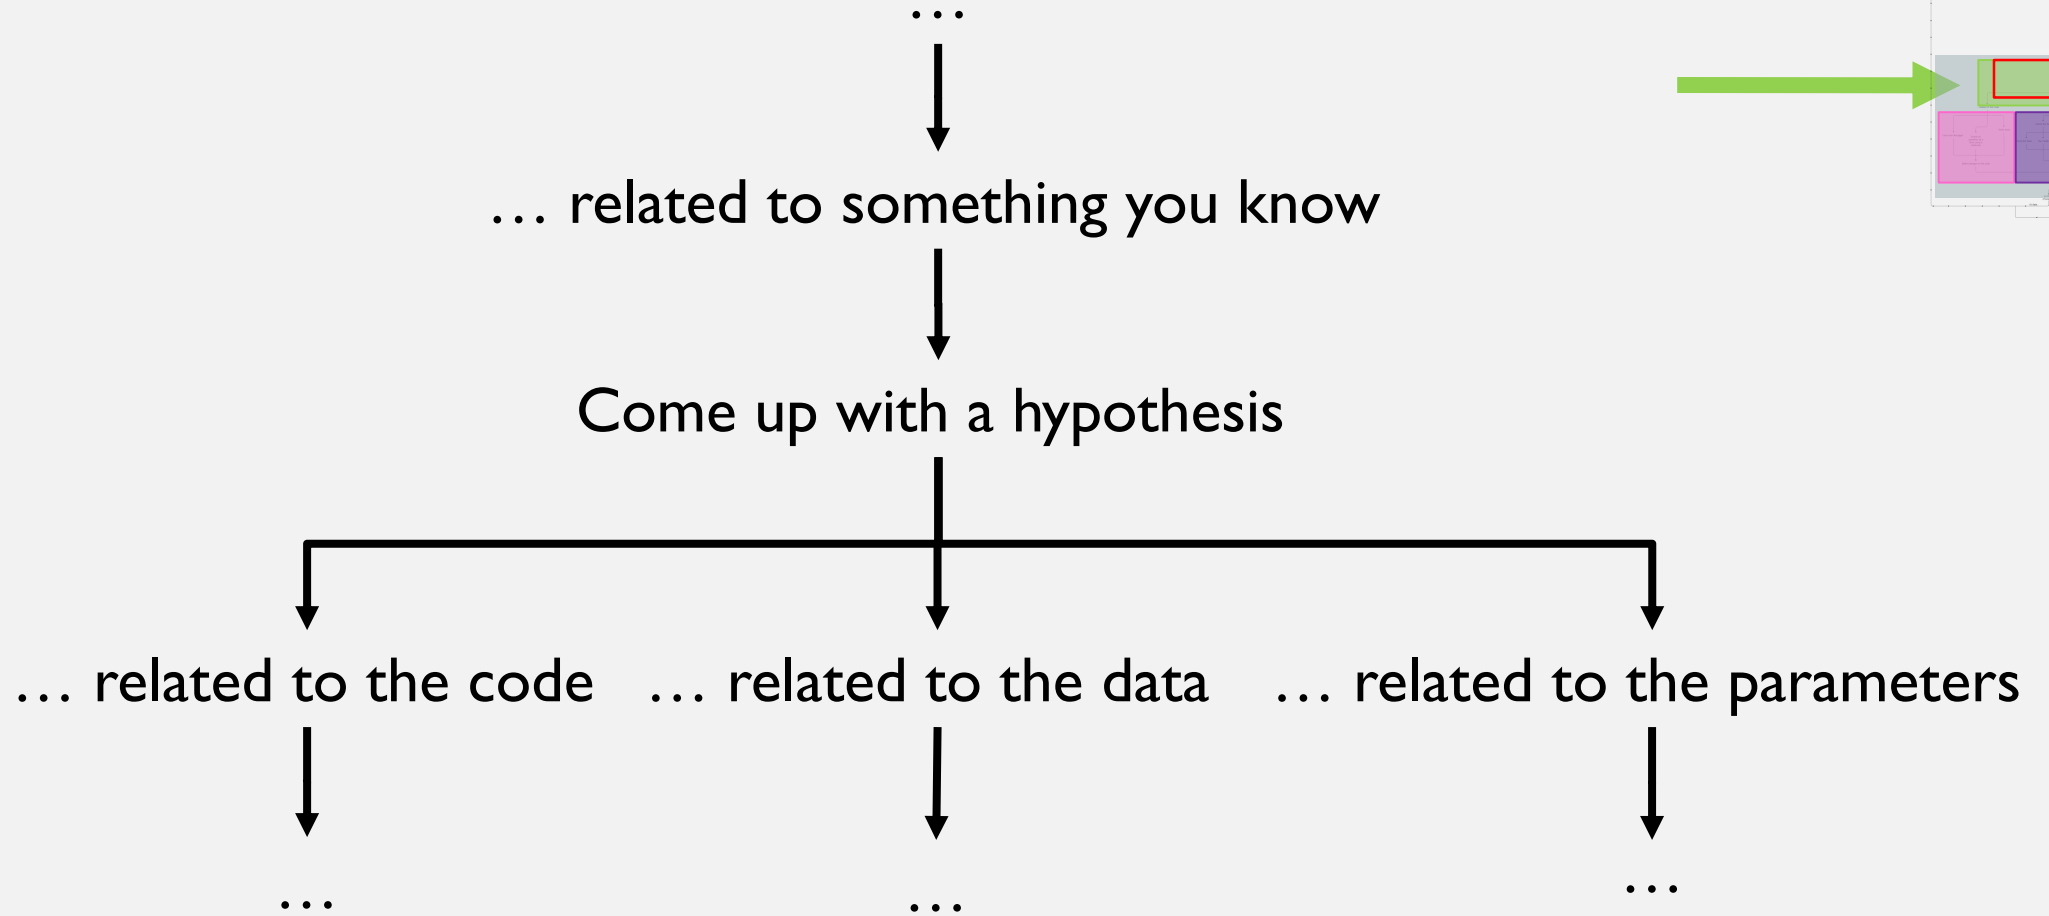
```
OpenJDK 64-Bit Server VM warning: ignoring option MaxPermSize=256m; support was removed in 8.0
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/tmp/hadoop-melike/nm-local-dir/usercache/melike/filecache/10/_spark_libs_887481056050304420.zip/slf4j-log4j12-1.7.16.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: Found binding in [jar:file:/opt/hadoop-3.1.2/share/hadoop/common/lib/slf4j-log4j12-1.7.25.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
2019-12-08 22:18:49,858 INFO util.SignalUtils: Registered signal handler for TERM
2019-12-08 22:18:49,860 INFO util.SignalUtils: Registered signal handler for HUP
2019-12-08 22:18:49,860 INFO util.SignalUtils: Registered signal handler for INT
2019-12-08 22:18:50,525 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using builtin-java classes where applicable
2019-12-08 22:18:50,823 INFO yarn.ApplicationMaster: Preparing Local resources
2019-12-08 22:18:52,149 INFO yarn.ApplicationMaster: Prepared Local resources Map(__app__.jar -> resource { scheme: "hdfs" host: "leela.unibe.ch" port: 9000 file: "/user/melike/.sparkStaging/applicati
2019-12-08 22:18:52,184 INFO yarn.ApplicationMaster: ApplicationAttemptId: appattempt_1575839673320_0002_000002
2019-12-08 22:18:52,196 INFO spark.SecurityManager: Changing view acls to: melike
2019-12-08 22:18:52,196 INFO spark.SecurityManager: Changing modify acls to: melike
2019-12-08 22:18:52,197 INFO spark.SecurityManager: Changing view acls groups to:
2019-12-08 22:18:52,198 INFO spark.SecurityManager: Changing modify acls groups to:
2019-12-08 22:18:52,198 INFO spark.SecurityManager: SecurityManager: authentication disabled; ui acls disabled; users with view permissions: Set(melike); groups with view permissions: Set(); users w
2019-12-08 22:18:52,224 INFO yarn.ApplicationMaster: Starting the user application in a separate Thread
2019-12-08 22:18:52,231 INFO yarn.ApplicationMaster: Waiting for spark context initialization
2019-12-08 22:18:52,232 INFO yarn.ApplicationMaster: Waiting for spark context initialization ...
2019-12-08 22:18:52,263 INFO spark.SparkContext: Running Spark version 2.0.2
2019-12-08 22:18:52,325 INFO spark.SecurityManager: Changing view acls to: melike
2019-12-08 22:18:52,325 INFO spark.SecurityManager: Changing modify acls to: melike
2019-12-08 22:18:52,325 INFO spark.SecurityManager: Changing view acls groups to:
2019-12-08 22:18:52,325 INFO spark.SecurityManager: Changing modify acls groups to:
2019-12-08 22:18:52,326 INFO spark.SecurityManager: SecurityManager: authentication disabled; ui acls disabled; users with view permissions: Set(melike); groups with view permissions: Set(); users w
2019-12-08 22:18:52,655 INFO util.Utils: Successfully started service 'sparkDriver' on port 41643.
2019-12-08 22:18:52,683 INFO spark.SparkEnv: Registering MapOutputTracker
2019-12-08 22:18:52,710 INFO spark.SparkEnv: Registering BlockManagerMaster
2019-12-08 22:18:52,731 INFO storage.DiskBlockManager: Created local directory at /tmp/hadoop-melike/nm-local-dir/usercache/melike/appcache/application_1575839673320_0002/blockmgr-ef493bb0-95ac-43e5-b
2019-12-08 22:18:52,759 INFO memory.MemoryStore: MemoryStore started with capacity 93.3 MB
2019-12-08 22:18:52,885 INFO spark.SparkEnv: Registering OutputCommitCoordinator
2019-12-08 22:18:53,017 INFO util.log: Logging initialized @4208ms
2019-12-08 22:18:53,079 INFO ui.JettyUtils: Adding filter: org.apache.hadoop.yarn.server.webproxy.amfilter.AmIpFilter
2019-12-08 22:18:53,153 INFO server.Server: jetty-9.2.z-SNAPSHOT
2019-12-08 22:18:53,182 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@5978551d{/jobs,null,AVAILABLE}
2019-12-08 22:18:53,182 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@57814a31{/jobs/json,null,AVAILABLE}
2019-12-08 22:18:53,182 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@5de2772d{/jobs/job,null,AVAILABLE}
2019-12-08 22:18:53,183 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@537a410d{/jobs/job/json,null,AVAILABLE}
2019-12-08 22:18:53,183 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@483136c2{/stages,null,AVAILABLE}
2019-12-08 22:18:53,183 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@15bc2d37{/stages/json,null,AVAILABLE}
2019-12-08 22:18:53,184 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@5b296bb2{/stages/stage,null,AVAILABLE}
2019-12-08 22:18:53,184 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@5a5a2bf3{/stages/stage/json,null,AVAILABLE}
2019-12-08 22:18:53,184 INFO handler.ContextHandler: Started o.s.j.s.ServletContextHandler@1h1c915c{/stages/pool,null,AVAILABLE}
```



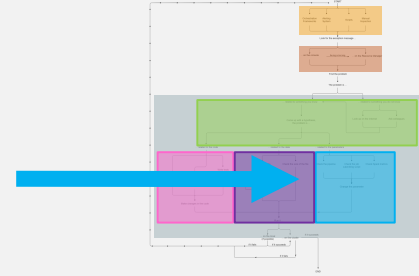
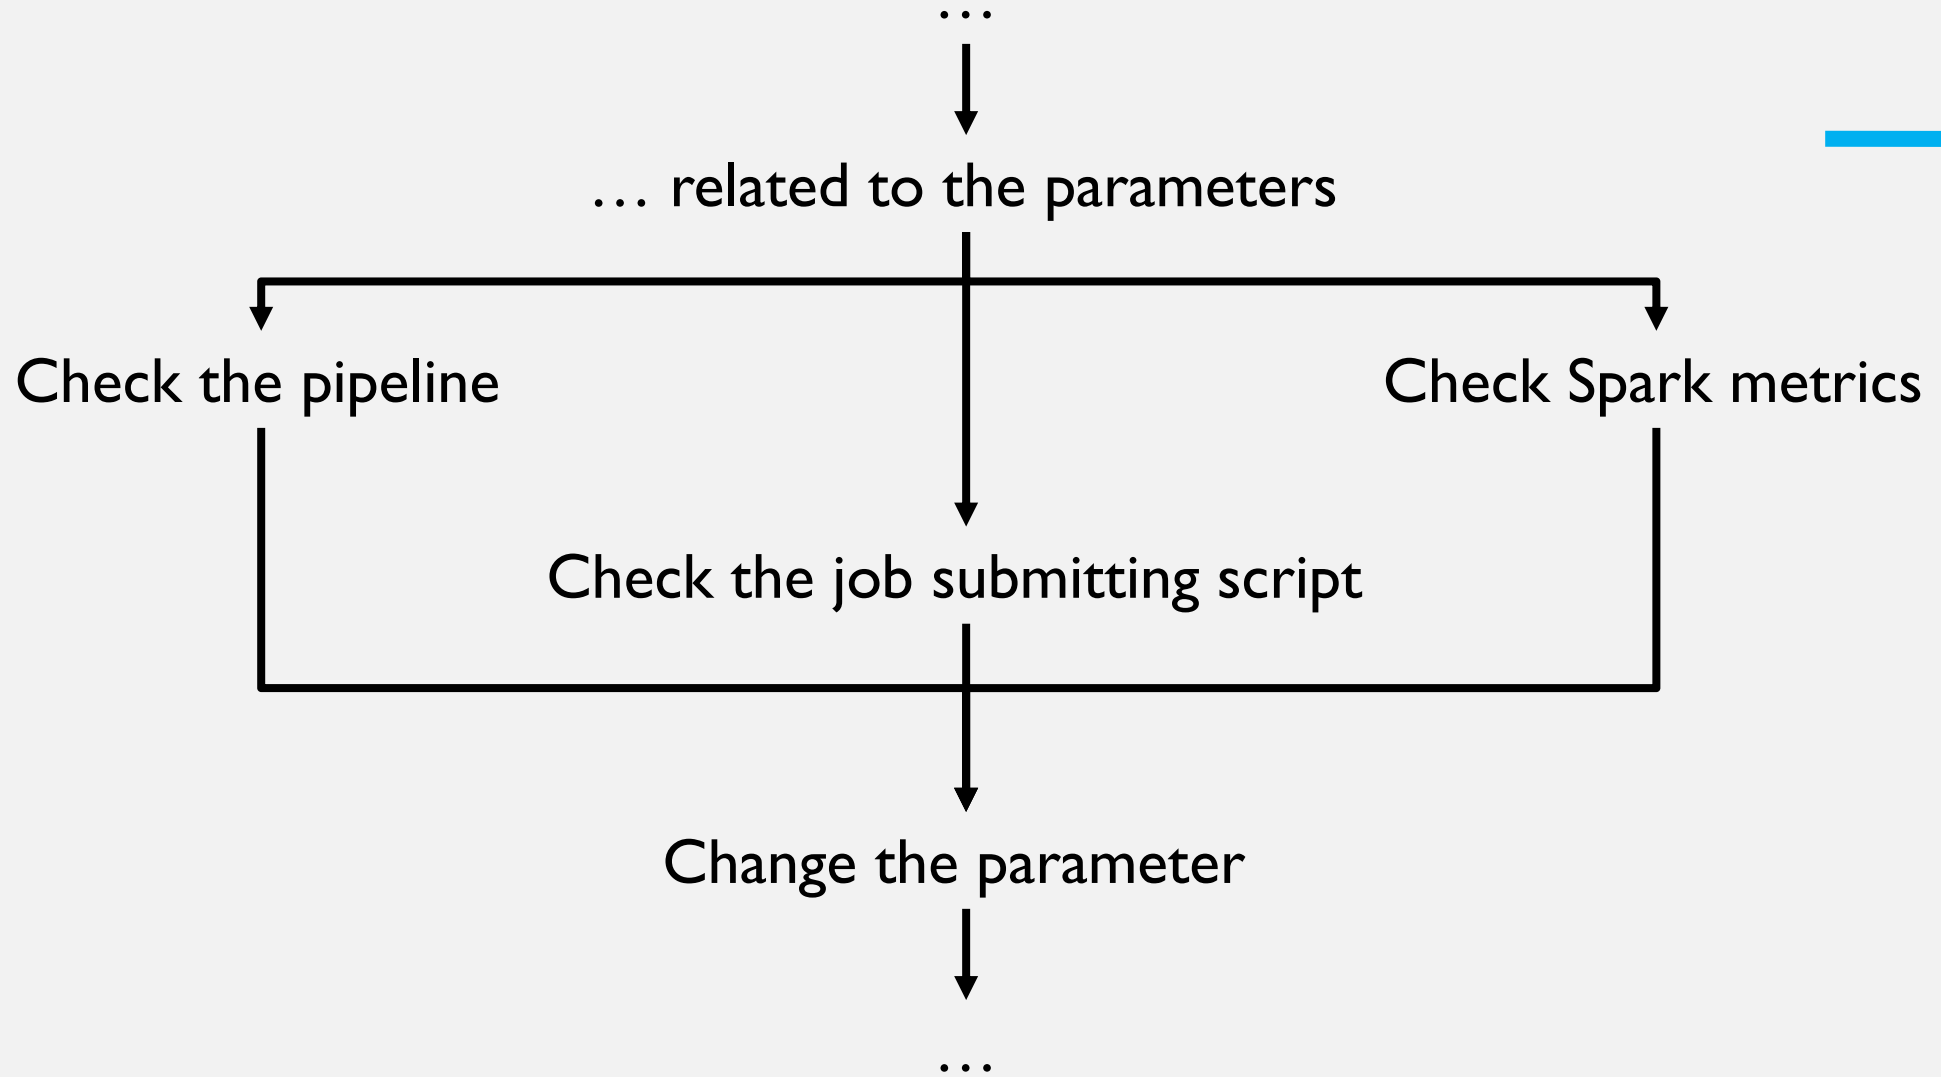
Solving the problem



Solving the problem



Solving the problem



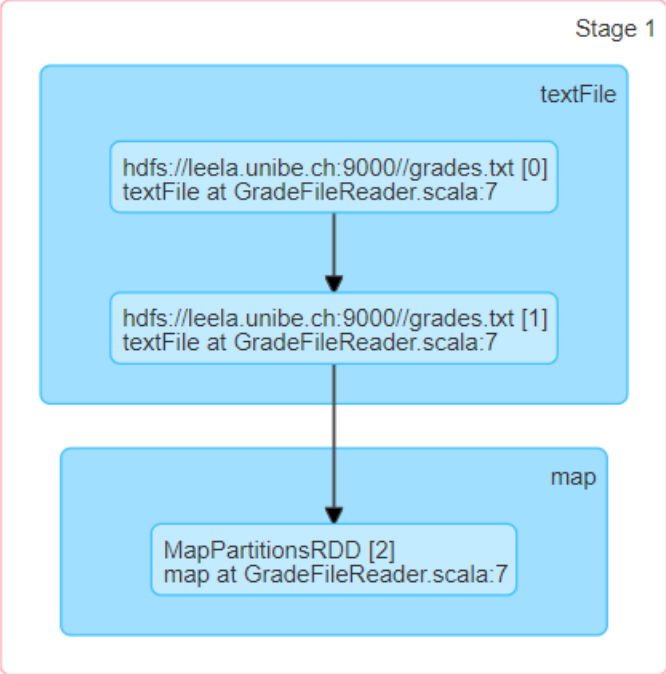
Solving the problem - Check the pipeline

APACHE Spark 2.0.2 Jobs Stages Storage

Details for Stage 1 (Attempt 0)

Total Time Across All Tasks: 1 s
Locality Level Summary: Rack local: 1
Input Size / Records: 24.0 B / 11

▼ DAG Visualization



Stage 1

textFile

hdfs://leela.unibe.ch:9000//grades.txt [0]
textFile at GradeFileReader.scala:7

hdfs://leela.unibe.ch:9000//grades.txt [1]
textFile at GradeFileReader.scala:7

map

MapPartitionsRDD [2]
map at GradeFileReader.scala:7

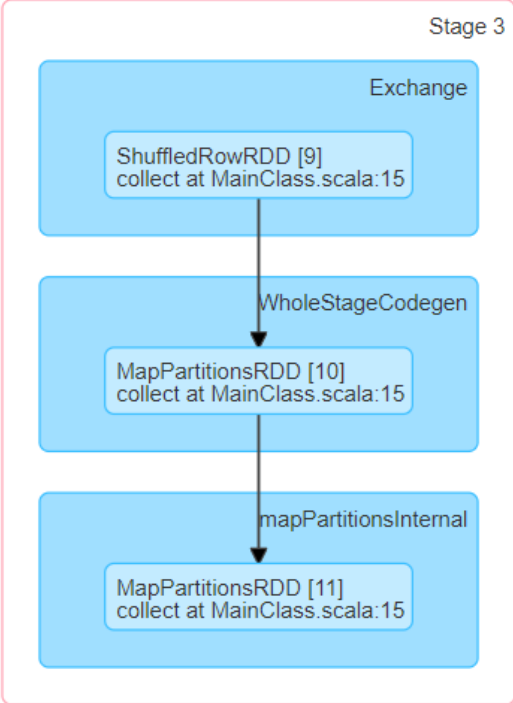
Detailed description: This diagram shows the execution plan for Stage 1. It starts with a 'textFile' action that reads from 'hdfs://leela.unibe.ch:9000//grades.txt' at index 0. This leads to another 'textFile' action at index 1. Finally, a 'map' action is performed on the resulting RDD, creating a 'MapPartitionsRDD [2]'.

APACHE Spark 2.0.2 Jobs Stages Storage

Details for Stage 3 (Attempt 0)

Total Time Across All Tasks: 5 s
Locality Level Summary: Node local: 3; Process local: 201
Shuffle Read: 616.0 B / 3

▼ DAG Visualization



Stage 3

Exchange

ShuffledRowRDD [9]
collect at MainClass.scala:15

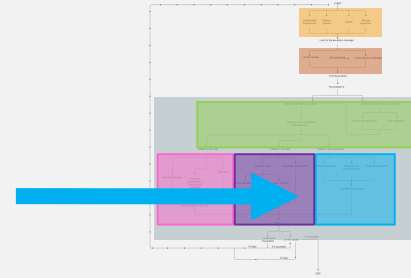
WholeStageCodegen

MapPartitionsRDD [10]
collect at MainClass.scala:15

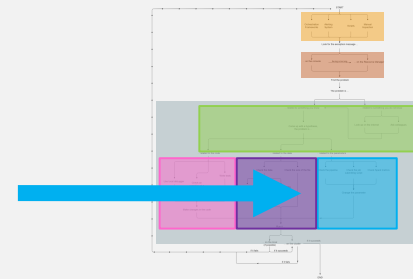
mapPartitionsInternal

MapPartitionsRDD [11]
collect at MainClass.scala:15

Detailed description: This diagram shows the execution plan for Stage 3. It begins with an 'Exchange' action that creates a 'ShuffledRowRDD [9]' from a 'collect' operation. This is followed by 'WholeStageCodegen' which produces a 'MapPartitionsRDD [10]'. Finally, a 'mapPartitionsInternal' action is performed, resulting in a 'MapPartitionsRDD [11]'.



Solving the problem - Check the job submitting script

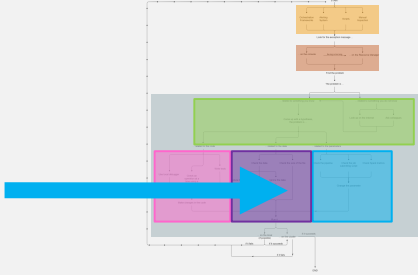
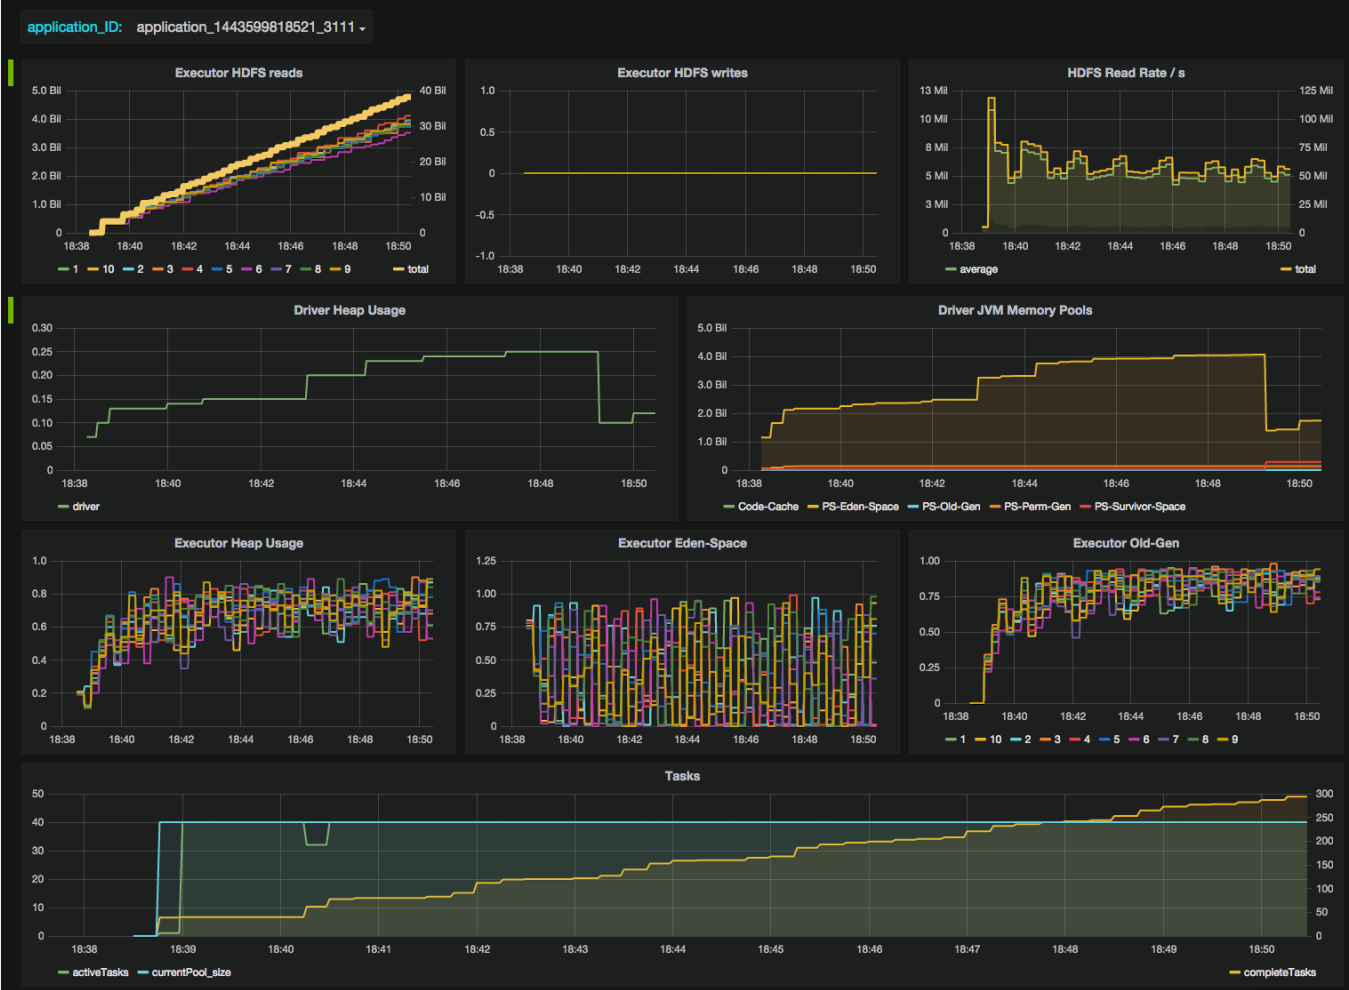


melike@leela: ~


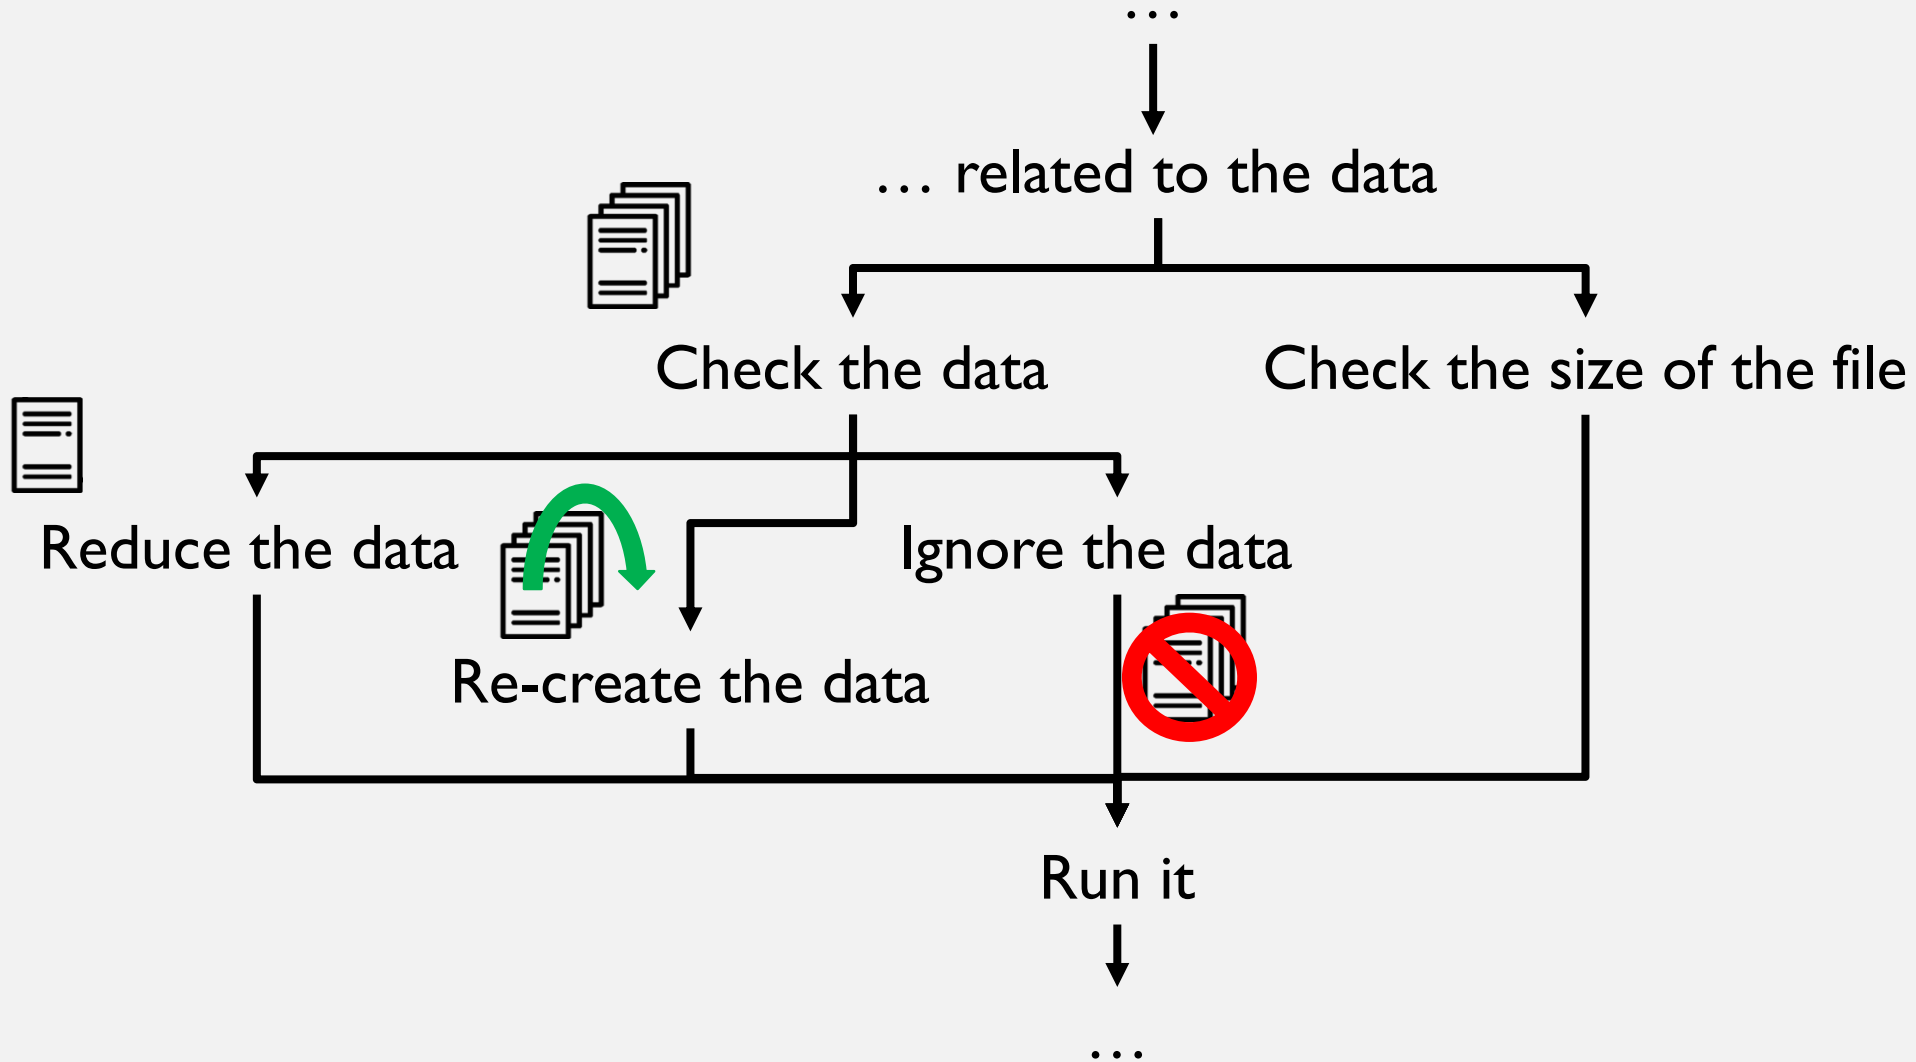
```
melike@leela:~$ spark-submit --num-executors 3 --master yarn --deploy-mode cluster --class ch.unibe.scg.MainClass CodeSa  
mpleOne-assembly-0.1.jar hdfs://leela.unibe.ch:9000//sales.txt
```

Solving the problem - Check Spark metrics

e.g. Grafana dashboard



Solving the problem



Solving the problem

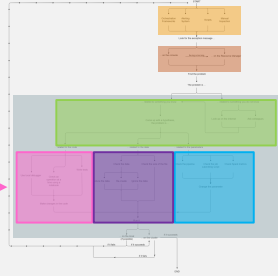


...
↓
... related to the code

Use local debugger

Write tests


Check an operation at a time using a notebook



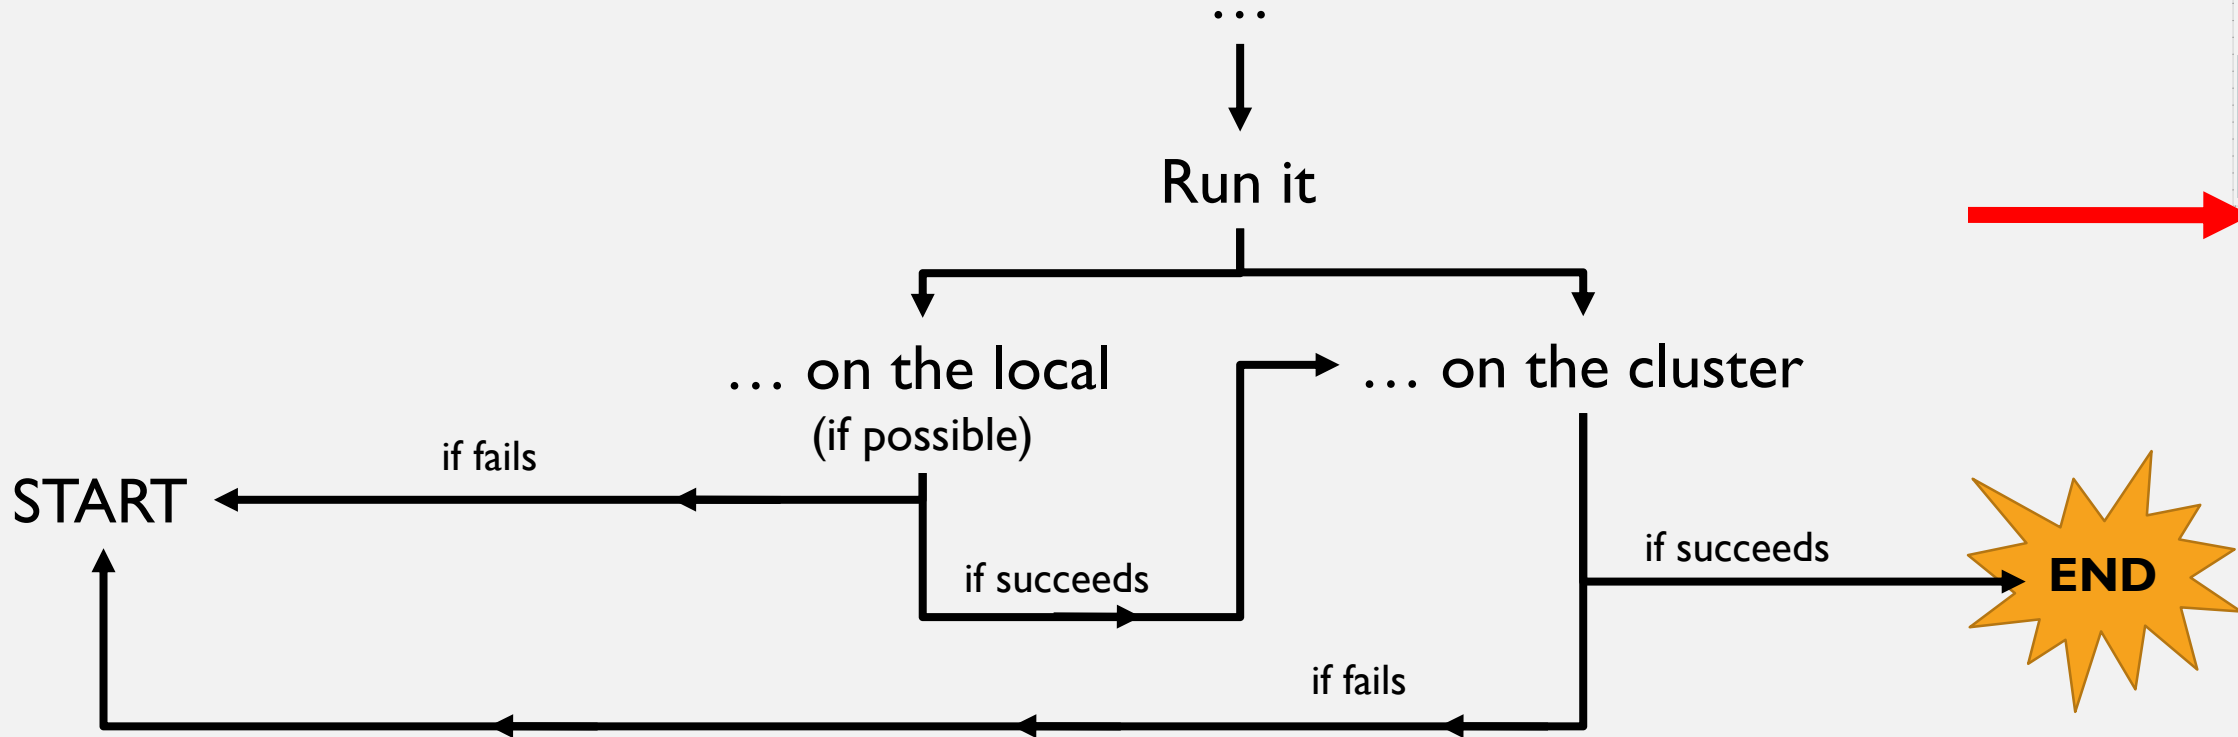
Make changes in the code

Run it

...



Epilogue



Postmortem Questionnaire


1. Have you seen these bugs before?
2. What do you do, after you see the exception message?
3. How do you know if there is a job failure?
4. What is the nastiest type of bugs you encountered in Spark applications?
5. Would you describe a tool that you miss when debugging Spark applications? What does it do? Why do you think it's not there?

Have you seen these bugs before?

Bug #1
Task Not Serializable



Bug #2
Stack Overflow



Bug #3
Number Format



Bug #4
Stack Overflow



What is the nastiest type of bugs you encountered in Spark applications?

Spark is difficult for **new learners**.

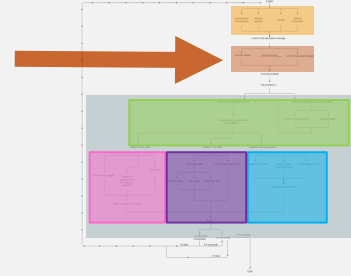
Assigning executors to partitions.

A job fails with **Memory Overhead**.

«Other components are nastier.»

What are the challenges that developers face when they debug such applications?

Logs don't give much insight about the exception
«Logs have messages that come directly from JVM»



Look for the exception message ...

... on the console The log is too long ... on the Resource Manager

Find the problem

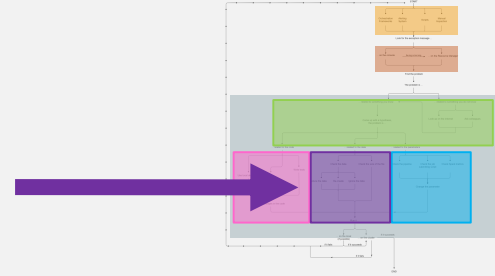
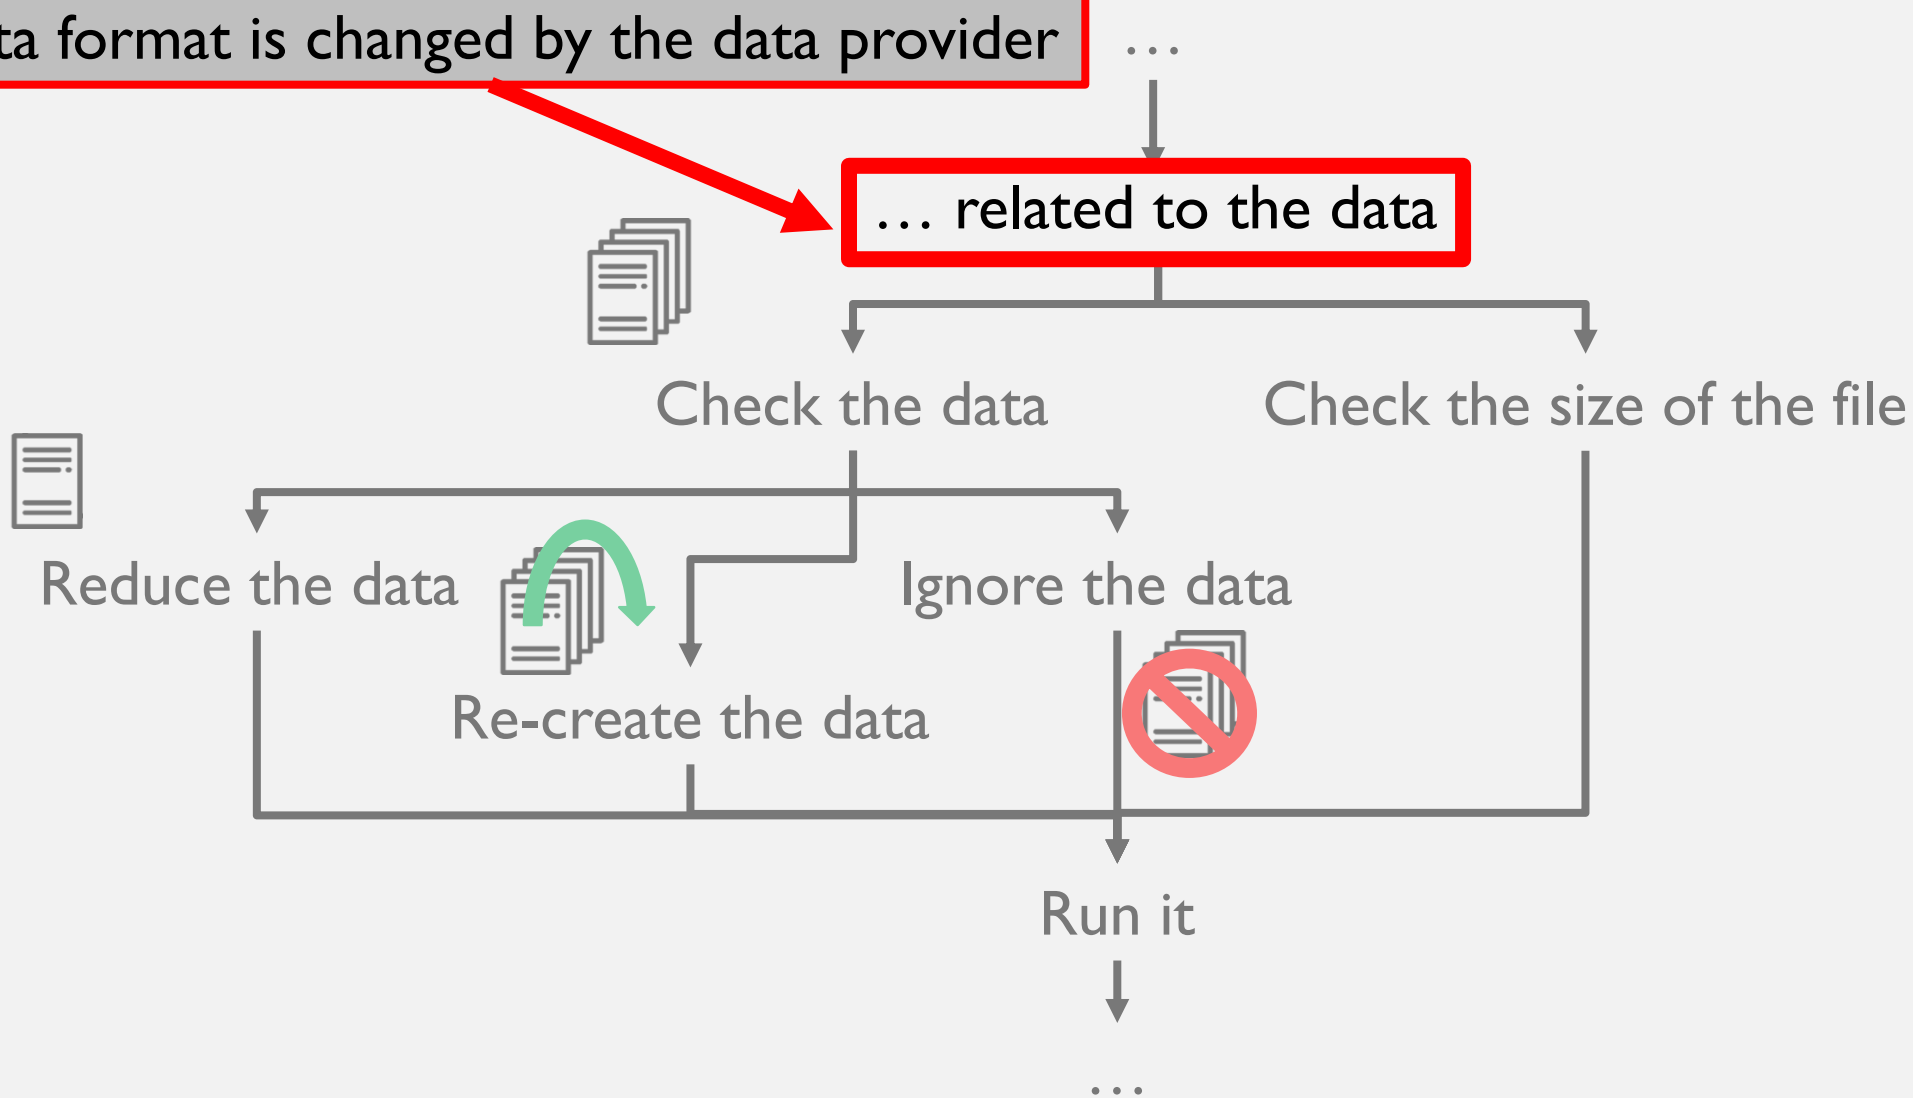
```
meike@meike: ~$ cat /tmp/spark-bad7ead2-1919-486f-bff4-161c2c1621e2/2019-12-08_22:19:00_745.INFO
2019-12-08 22:18:59,769 INFO yarn.Client: Application report for application_157583967320_0002 (state: RUNNING)
2019-12-08 22:19:00,745 INFO yarn.Client: Application report for application_157583967320_0002 (state: FINISHED)
2019-12-08 22:19:00,747 INFO yarn.Client:
client token: N/A
diagnostics: User class threw exception: org.apache.spark.SparkException: Task not serializable
ApplicationMaster host: 130.92.65.82
ApplicationMaster RPC port: 0
queue: default
start time: 1575839903236
final status: FAILED
tracking URL: http://leela.unibe.ch:8088/proxy/application_157583967320_0002/
user: meike
Exception in thread "main" org.apache.spark.SparkException: Application application_157583967320_0002 finished with failed status
at org.apache.spark.deploy.yarn.Client.run(Client.scala:1132)
at org.apache.spark.deploy.yarn.Client$.main(Client.scala:1174)
at org.apache.spark.deploy.yarn.Client.main(Client.scala)
at sun.reflect.NativeMethodAccessorImpl.invoke(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:606)
at org.apache.spark.deploy.SparkSubmit$.org$apache$spark$deploy$SparkSubmit$$runMain(SparkSubmit.scala:736)
at org.apache.spark.deploy.SparkSubmit$.doRunMain$1(SparkSubmit.scala:185)
at org.apache.spark.deploy.SparkSubmit$.submit(SparkSubmit.scala:210)
at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala:126)
at org.apache.spark.deploy.SparkSubmit$.main(SparkSubmit.scala)
2019-12-08 22:19:00,745 INFO util.ShutdownHookManager: Shutdown hook called
2019-12-08 22:19:00,745 INFO util.ShutdownHookManager: Deleting directory /tmp/spark-bad7ead2-1919-486f-bff4-161c2c1621e2
```

```
hadoop
- ResourceManager
- BitMatrix
- NodeManager
- Tools
2019-12-08 22:19:00,745 INFO util.ShutdownHookManager: Deleting directory /tmp/spark-bad7ead2-1919-486f-bff4-161c2c1621e2
```

What are the challenges that developers face when they debug such applications?

Data format is changed by the data provider

... related to the data



What are the challenges that developers face when they debug such applications?

Debugging on local doesn't always help
Some problems occur only on the cluster

Use local debugger

... related to the code

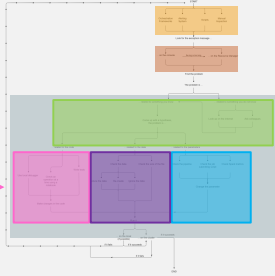
Write tests

Check an operation at a time using a notebook

Make changes in the code

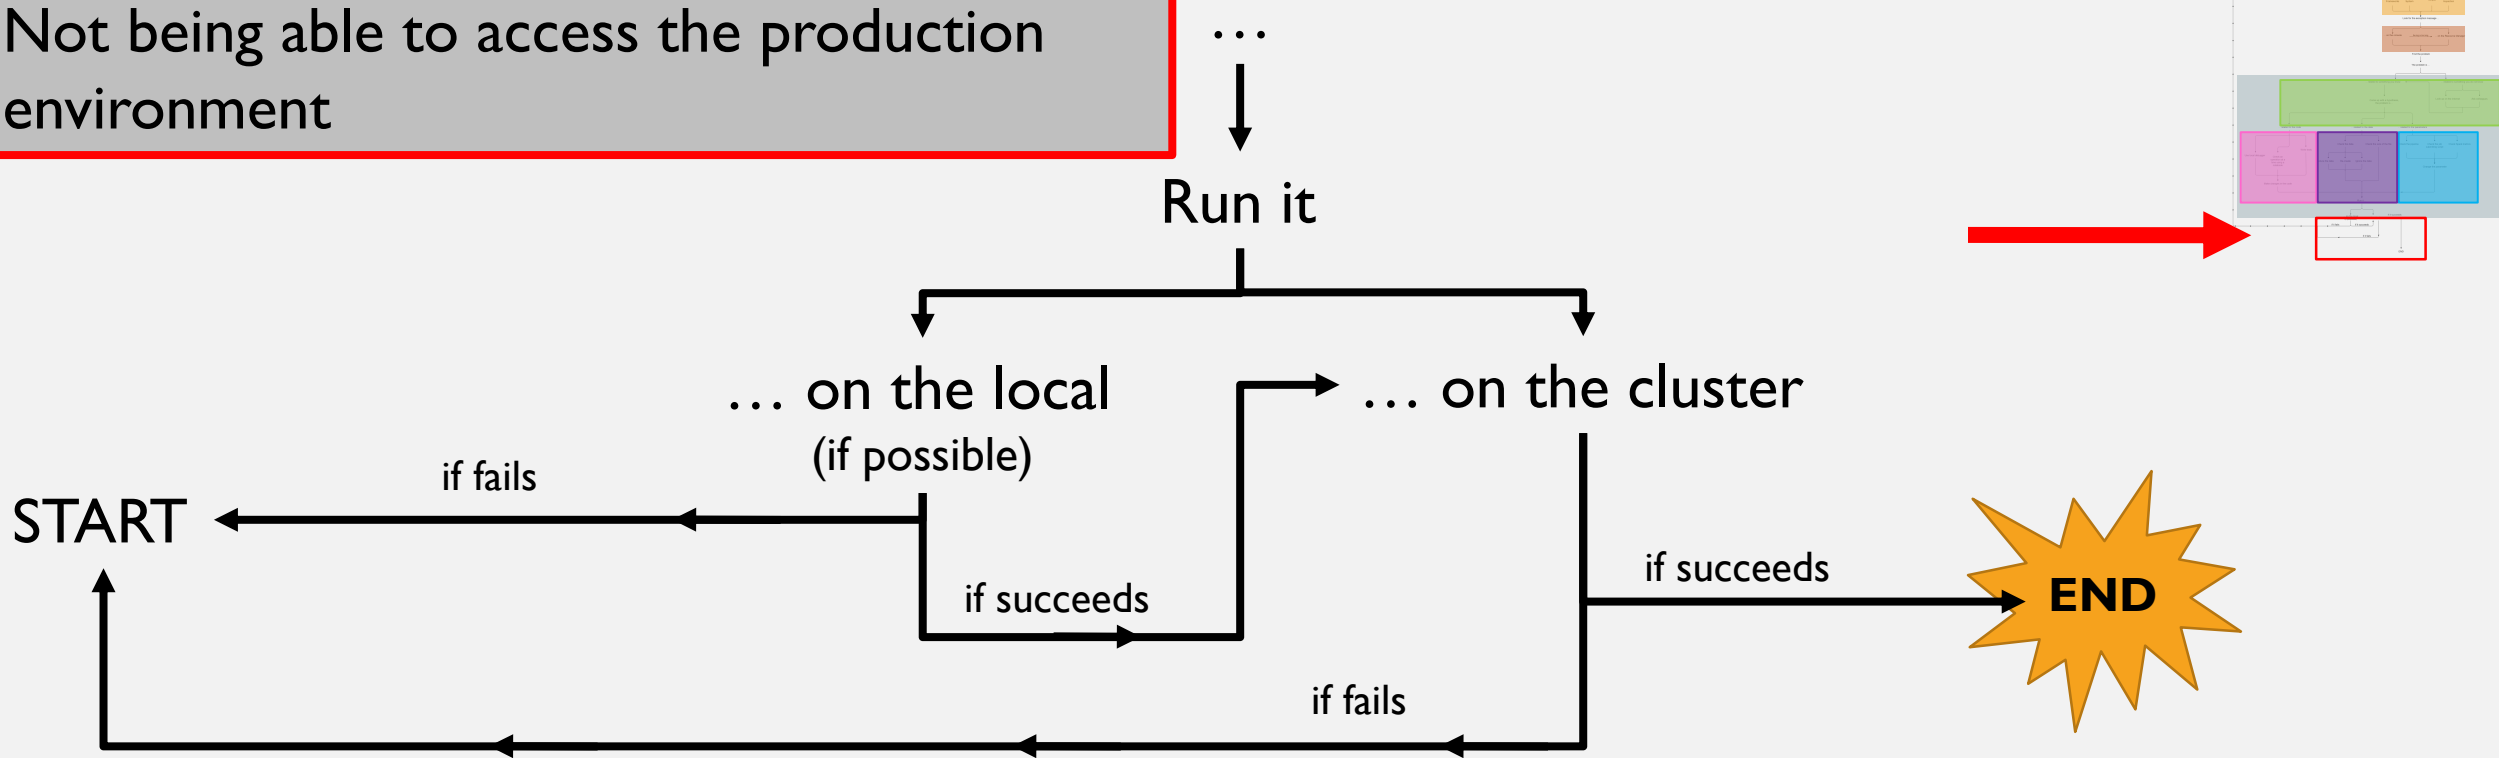
Run it

...




What are the challenges that developers face when they debug such applications?

Not being able to access the production environment



How can we overcome these challenges?

3 categories of enhancements were proposed:



Code

```
org.apache.spark.SparkException: Task not serializable
    at org.apache.spark.util.ClosureCleaner$.ensureSerializable(ClosureCleaner.scala:298)
    at org.apache.spark.util.ClosureCleaner$.org$spark$util$ClosureCleaner$$clean(ClosureCleaner.scala:288)
    at org.apache.spark.util.ClosureCleaner$.clean(ClosureCleaner.scala:108)
    at org.apache.spark.SparkContext.clean(SparkContext.scala:2839)
    at org.apache.spark.rdd.RDD$$anonfun$map$1.apply(RDD.scala:366)
    at org.apache.spark.rdd.RDD$$anonfun$map$1.apply(RDD.scala:365)
    at org.apache.spark.rdd.RDDOperationScope$.withScope(RDDOperationScope.scala:151)
    at org.apache.spark.rdd.RDDOperationScope$.withScope(RDDOperationScope.scala:112)
    at org.apache.spark.rdd.RDD.withScope(RDD.scala:358)
    at org.apache.spark.rdd.RDD.map(RDD.scala:365)
    at MathOperations.divisibleBy3(MainClass.scala:27)
    at MainClass$.main(MainClass.scala:13)
    at MainClass.main(MainClass.scala)
    at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
    at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
    at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
    at java.lang.reflect.Method.invoke(Method.java:498)
    at org.apache.spark.deploy.yarn.ApplicationMaster$$anon$2.run(ApplicationMaster.scala:627)
Caused by: java.io.NotSerializableException: MathOperations
Serialization stack:
  - object not serializable (class: MathOperations, value: MathOperations@64b55111)
  - field (class: MathOperations$$anonfun$1, name: $outer, type: class MathOperations)
  - object (class: MathOperations$$anonfun$1, $function)
    at org.apache.spark.serializer.SerializationDebugger$.improveException(SerializationDebugger.scala:40)
    at org.apache.spark.serializer.JavaSerializationStream.writeObject(JavaSerializer.scala:46)
    at org.apache.spark.serializer.JavaSerializerInstance.serialize(JavaSerializer.scala:108)
    at org.apache.spark.util.ClosureCleaner$.ensureSerializable(ClosureCleaner.scala:295)
    ... 17 more
```

Log



Dashboard

Tool Recommendations to Enhance Coding

- Monitor data after each transformation
- Find the line in a data file, which causes a problem
- A distributed debugger
 - A unified debugger interface (Leske, Chiş and Nierstrasz, 2016)
 - A remote debugging model based on reflection (Papoulias et al., 2015)


Tool Recommendations to Enhance Logs

- A Log parser
 - A tool which separates important and useless parts of logs
 - Better exception messages, instead of taking the message from JVM
 - Using ML techniques to learn and advice solutions to certain problems
- See in which line the application dies like in Java
- See logs per stage on Spark UI

Tool Recommendation: Dashboard


- A tool which displays configuration of the cluster
 - Memory, GC, Network Bandwidth, and other parameters
- A tool which can compare different environments
 - when there's no access to the other environment

Investigation Roadmap




Duration vs. Number of Steps


B#1



B#2



B#3



B#4

