



KubeCon



CloudNativeCon



redhat.

Europe 2018

SCALABLE MONITORING USING PROMETHEUS WITH APACHE SPARK

DIANE FEDDEMA
PRINCIPAL SOFTWARE ENGINEER
ZAK HASSAN
SOFTWARE ENGINEER



YOUR SPEAKERS



KubeCon



CloudNativeCon

Europe 2018

DIANE FEDDEMA

PRINCIPAL SOFTWARE ENGINEER - EMERGING TECHNOLOGY, DATA ANALYTICS

- Currently focused on developing and applying Data Science and Machine Learning techniques for performance analysis, automating these analyses and displaying data in novel ways.
- Previously worked as a performance engineer at the National Center for Atmospheric Research, NCAR, working on optimizations and tuning in parallel global climate models.

ZAK HASSAN

SOFTWARE ENGINEER - EMERGING TECHNOLOGY, DATA ANALYTICS

- Currently focused on developing analytics platform on OpenShift and leveraging Apache Spark as the analytics engine. Also, developing data science apps and working on making metrics observable through cloud-native technology.
- Previously worked as a Software Consultant in the financial services and insurance industry, building end-to-end software solutions for clients.



redhat®

OVERVIEW



KubeCon



CloudNativeCon

Europe 2018

OBSERVABILITY

- Motivation
- What Is Spark?
- What Is Prometheus?
- Our Story
- Spark Cluster JVM Instrumentation

PERFORMANCE TUNING

- Tuning Spark jobs
- Spark Memory Model
- Prometheus as a performance tool
- Comparing cached vs non-cached dataframes
- Demo



redhat.

MOTIVATION



KubeCon



CloudNativeCon

Europe 2018

- Rapid experimentation of data science apps
- Identify bottlenecks
- Improve performance
- Resolve incidents quicker
- Improving memory usage to tune spark jobs



OUR STORY



- Instrumented spark jvm to expose metrics in a kubernetes pod.
- Added ability to monitor spark with prometheus
- Experimented with using Grafana with Prometheus to provide more insight
- Sharing our experiments and experience with using this to do performance analysis of spark jobs.
- Demo at the very end

June 1, 2017 - <https://github.com/radanalyticsio/openshift-spark/pull/28>

- Added agent to report jolokia metrics endpoint in kubernetes pod

Nov 7, 2017 - <https://github.com/radanalyticsio/openshift-spark/pull/35>

- Added agent to report prometheus metrics endpoint in kubernetes pod



WHAT IS PROMETHEUS



KubeCon



CloudNativeCon

Europe 2018

- Open source monitoring
- in 2016 prometheus become the 2nd member of the CNCF
- scrapes metrics from a endpoint.
- Client libraries in **Go, Java, Python,C#/.Net, Node.JS, Haskell, Erlang, Rust, Ruby.**
- Kubernetes comes instrumented out of the box with prometheus endpoints.
- If you don't have native integration with prometheus there are lots of community exporters that allow lots of things to expose metrics in your infrastructure to get monitored.



redhat.

WHAT IS SPARK



KubeCon



CloudNativeCon

Europe 2018

Spark is an in demand data processing engine with a thriving community and steadily growing install base

- Supports interactive data exploration in addition to apps
- Batch and stream processing
- Machine learning libraries
- Distributed
- Separate storage and compute (in memory processing)
- new external scheduler kubernetes



SPARK FEATURES




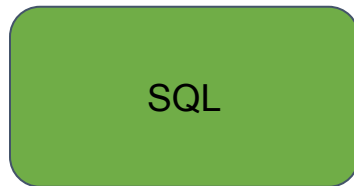
KubeCon



CloudNativeCon

Europe 2018

- Can run standalone, with yarn, mesos or **Kubernetes** as the cluster manager
- Has language bindings for Java, Scala, Python, and R
- Access data from JDBC, HDFS, S3 or regular filesystem
- Can persist data in different data formats: parquet, avro, json, csv, etc.



redhat.

SPARK APPLICATION



KubeCon



CloudNativeCon

Europe 2018



redhat.

SPARK IN CONTAINERS

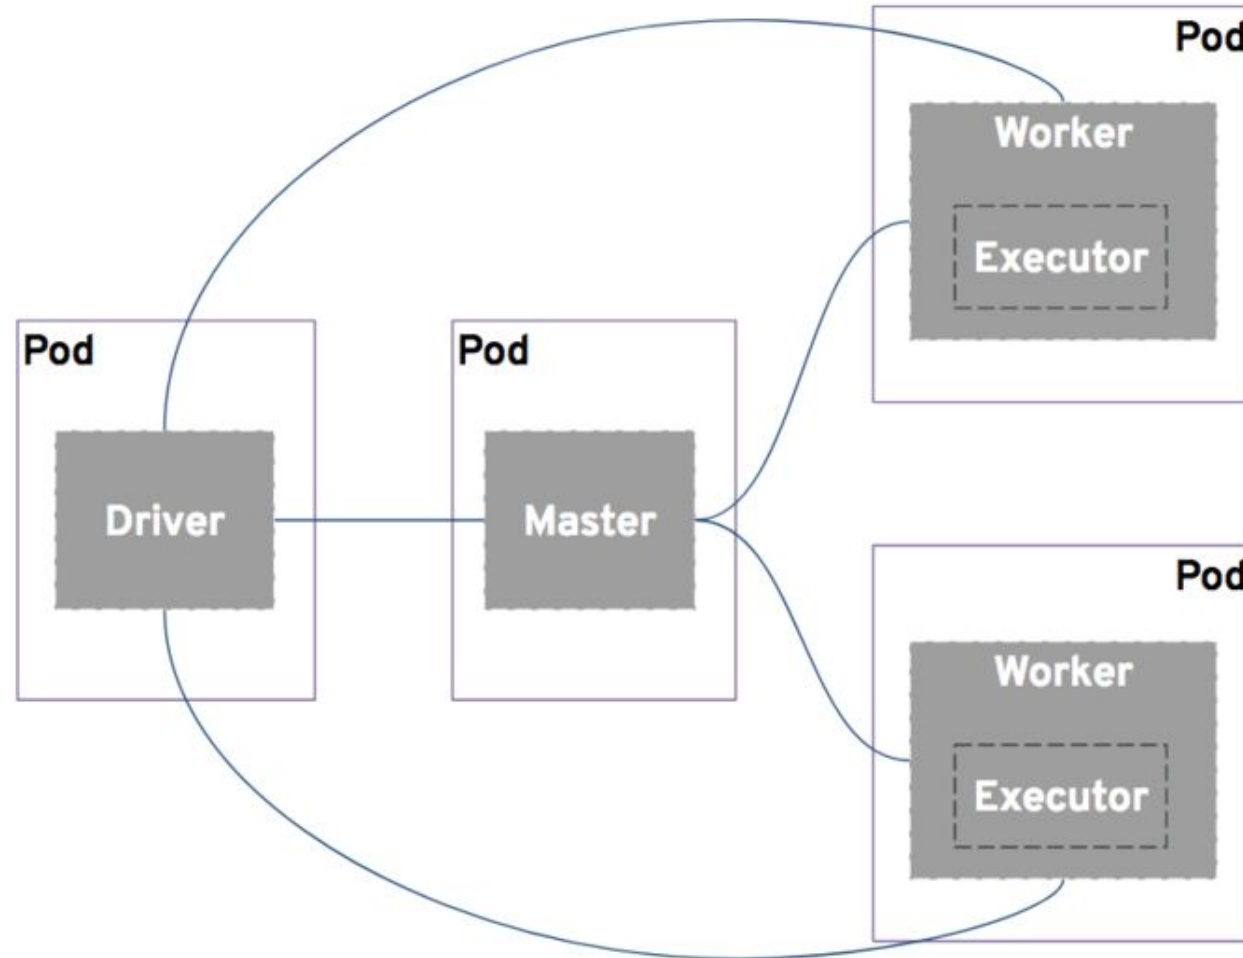


KubeCon



CloudNativeCon

Europe 2018



redhat.

SPARK CLUSTER INSTRUMENT

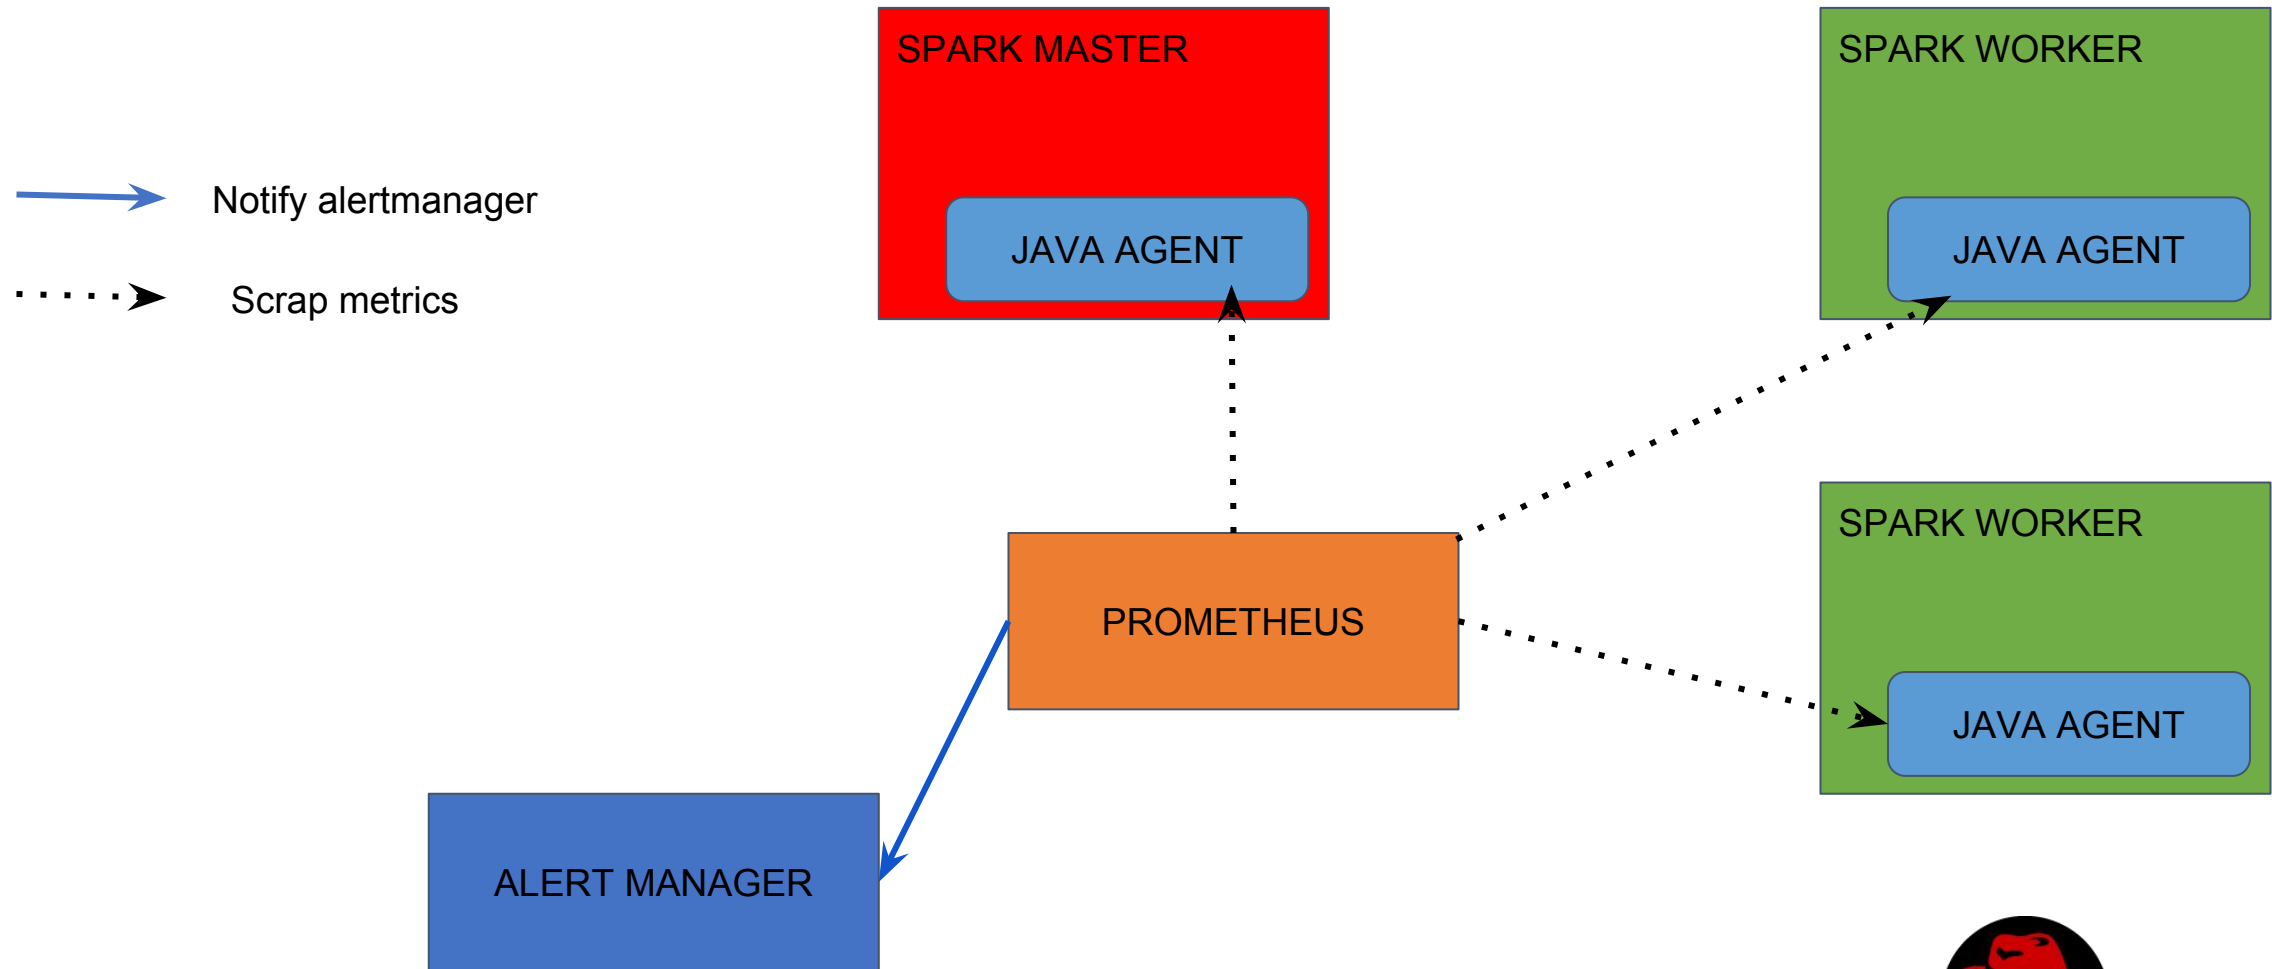


KubeCon



CloudNativeCon

Europe 2018



redhat.

INSTRUMENT JAVA AGENT



KubeCon



CloudNativeCon

Europe 2018

```
30 elif [ ${SPARK_METRICS_ON} == "prometheus" ]; then
31     JAVA_AGENT=" -javaagent:${SPARK_HOME}/agent-bond.jar=${SPARK_HOME}/conf/agent.properties"
32     metrics=" with prometheus metrics enabled"
33 else
34     JAVA_AGENT=" -javaagent:${SPARK_HOME}/jolokia-jvm-1.3.6-agent.jar=port=7777,host=0.0.0.0"
35     metrics=" with jolokia metrics enabled (deprecated, set SPARK_METRICS_ON to 'prometheus')"
36 fi
37
38 if [ -z ${SPARK_MASTER_ADDRESS+_} ]; then
39     echo "Starting master$metrics"
40     exec $SPARK_HOME/bin/spark-class$JAVA_AGENT org.apache.spark.deploy.master.Master
41 else
42     echo "Starting worker$metrics, will connect to: $SPARK_MASTER_ADDRESS"
43     while true; do
44         echo "Waiting for spark master to be available ..."
45         curl --connect-timeout 1 -s -X GET $SPARK_MASTER_UI_ADDRESS > /dev/null
46         if [ $? -eq 0 ]; then
47             break
48         fi
49         sleep 1
50     done
51     exec $SPARK_HOME/bin/spark-class$JAVA_AGENT org.apache.spark.deploy.worker.Worker $SPARK_MASTER_ADDRESS
```



redhat.

PROMETHEUS TARGETS



KubeCon



CloudNativeCon

Europe 2018

Prometheus Alerts Graph Status Help

Targets

kubernetes-apiservers (1/1 up)

Endpoint	State	Labels	Last Scrape	Error
https://10.19.47.23:8443/metrics	UP	instance="10.19.47.23:8443"	47.748s ago	

kubernetes-cadvisor (2/2 up)

Endpoint	State	Labels	Last Scrape	Error
https://10.19.47.25:10250/metrics/cadvisor	UP	beta_kubernetes_io_arch="amd64" beta_kubernetes_io_os="linux" instance="et10.et.eng.bos.redhat.com" kubernetes_io_hostname="et10.et.eng.bos.redhat.com" region="infra" zone="default"	1.713s ago	
https://10.19.47.23:10250/metrics/cadvisor	UP	beta_kubernetes_io_arch="amd64" beta_kubernetes_io_os="linux" instance="et9.et.eng.bos.redhat.com" kubernetes_io_hostname="et9.et.eng.bos.redhat.com" region="primary" zone="default"	30.001s ago	

kubernetes-controllers (1/1 up)

Endpoint	State	Labels	Last Scrape	Error
https://10.19.47.23:8444/metrics	UP	instance="10.19.47.23:8444"	35.983s ago	

kubernetes-nodes (2/2 up)

Endpoint	State	Labels	Last Scrape	Error
https://10.19.47.25:10250/metrics	UP	beta_kubernetes_io_arch="amd64" beta_kubernetes_io_os="linux" instance="et10.et.eng.bos.redhat.com" kubernetes_io_hostname="et10.et.eng.bos.redhat.com" region="infra" zone="default"	33.888s ago	
https://10.19.47.23:10250/metrics	UP	beta_kubernetes_io_arch="amd64" beta_kubernetes_io_os="linux" instance="et9.et.eng.bos.redhat.com" kubernetes_io_hostname="et9.et.eng.bos.redhat.com" region="primary" zone="default"	44.336s ago	

spark-cluster-m-1-fq2dj (1/1 up)

Endpoint	State	Labels	Last Scrape	Error
http://10.128.0.141:7777/metrics	UP	instance="10.128.0.141:7777"	16.304s ago	

spark-cluster-w-1-b55mg (1/1 up)



redhat.

PULL METRICS



KubeCon



CloudNativeCon

Europe 2018

- Prometheus lets you configure how often to scrap and which endpoints to scrap. The prometheus server will pull in the metrics that are configured.

 `prometheus.yaml`

```
1  global:
2    scrape_interval:    15s
3    evaluation_interval: 15s
4  alerting:
5    alertmanagers:
6      - static_configs:
7        - targets:
8          - alertmanager:9093
9  rule_files:
10   - "simple_rule.yml"
11  scrape_configs:
12   - job_name: 'prometheus'
13     static_configs:
14       - targets: ['localhost:9090']
```



redhat.

ALERTMANAGER



KubeCon



CloudNativeCon

Europe 2018

- PromQL query is used to create rules to notify you if the rule is triggered.
- Currently alertmanager will receive the notification and is able to notify you via email, slack or other options (see docs for details) .

 `simple_rule.yml`

```
1  groups:
2  - name: spark.rules
3    rules:
4    - alert: SparkOutage
5      expr: up == 0
6      for: 5s
7      labels:
8        severity: critical
9      annotations:
10     description: erik spark cluster is down and out
11     summary: erik spark Instance down
```



redhat.

PROMQL



KubeCon



CloudNativeCon

Europe 2018

- Powerful query language to get metrics on kubernetes cluster along with spark clusters.
- What are gauges and counters?

Gauges: Latest value of metric

Counters: Total number of event occurrences. Might be suffix “***total**”.

You can use this format to get the last minute **prom_metric_total[1m]**

PART 2: Tuning Spark Jobs with Prometheus



KubeCon



CloudNativeCon

Europe 2018

Things we would like to know when tuning Spark programs:

- How much memory is the driver using?
- How much memory are the workers using?
- How is the JVM begin utilized by spark?
- Is my spark job saturating the network?
- What is the cluster view of network, cpu and memory utilization?

We will demonstrate how **Prometheus** coupled with **Grafana** on **Kubernetes** can help answer these types of questions.



Our Example Application



KubeCon



CloudNativeCon

Europe 2018

Focus on Memory:

Efficient memory use is key to good performance in Spark jobs.

How:

We will create Prometheus + Grafana dashboards to evaluate memory usage under different conditions?

Example:

Our Spark Python example will compare memory usage with and without caching to illustrate how memory usage and timing change for a PySpark program performing a cartesian product followed by a groupby operation



A Little Background



KubeCon



CloudNativeCon

Europe 2018

Memory allocation in Spark

- Spark is an "in-memory" computing framework
- Memory is a limited resource!
- There is competition for memory
- Caching reusable results can save overall memory usage under certain conditions
- Memory runs out in many large jobs forcing spills to disk



redhat.

Spark Unified Memory Model

LRU eviction and user defined memory configuration options




KubeCon



CloudNativeCon

Europe 2018

Total JVM Heap Memory allocated to SPARK JOB



redhat.



KubeCon



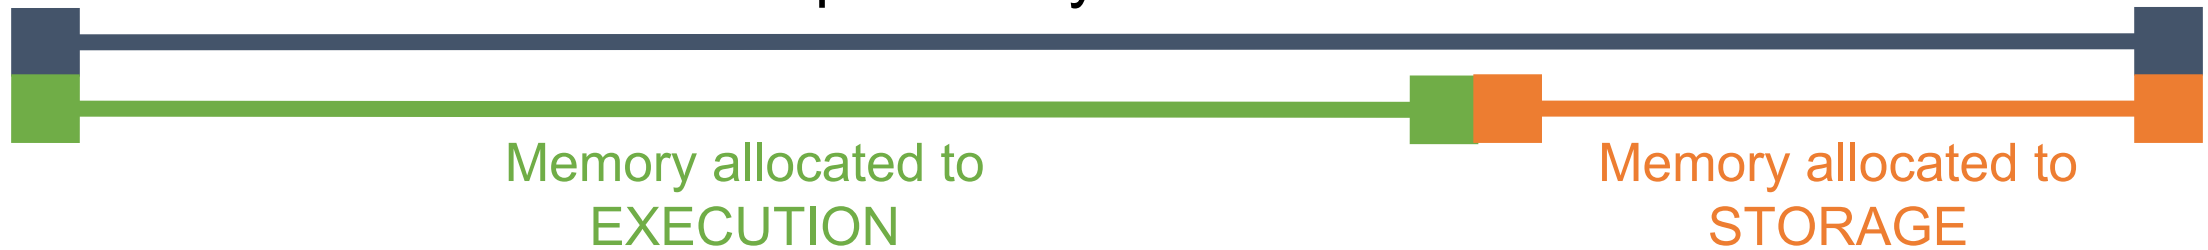
CloudNativeCon

Europe 2018

Spark Unified Memory Model

LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



redhat.

Spark Unified Memory Model

LRU eviction and user defined memory configuration options



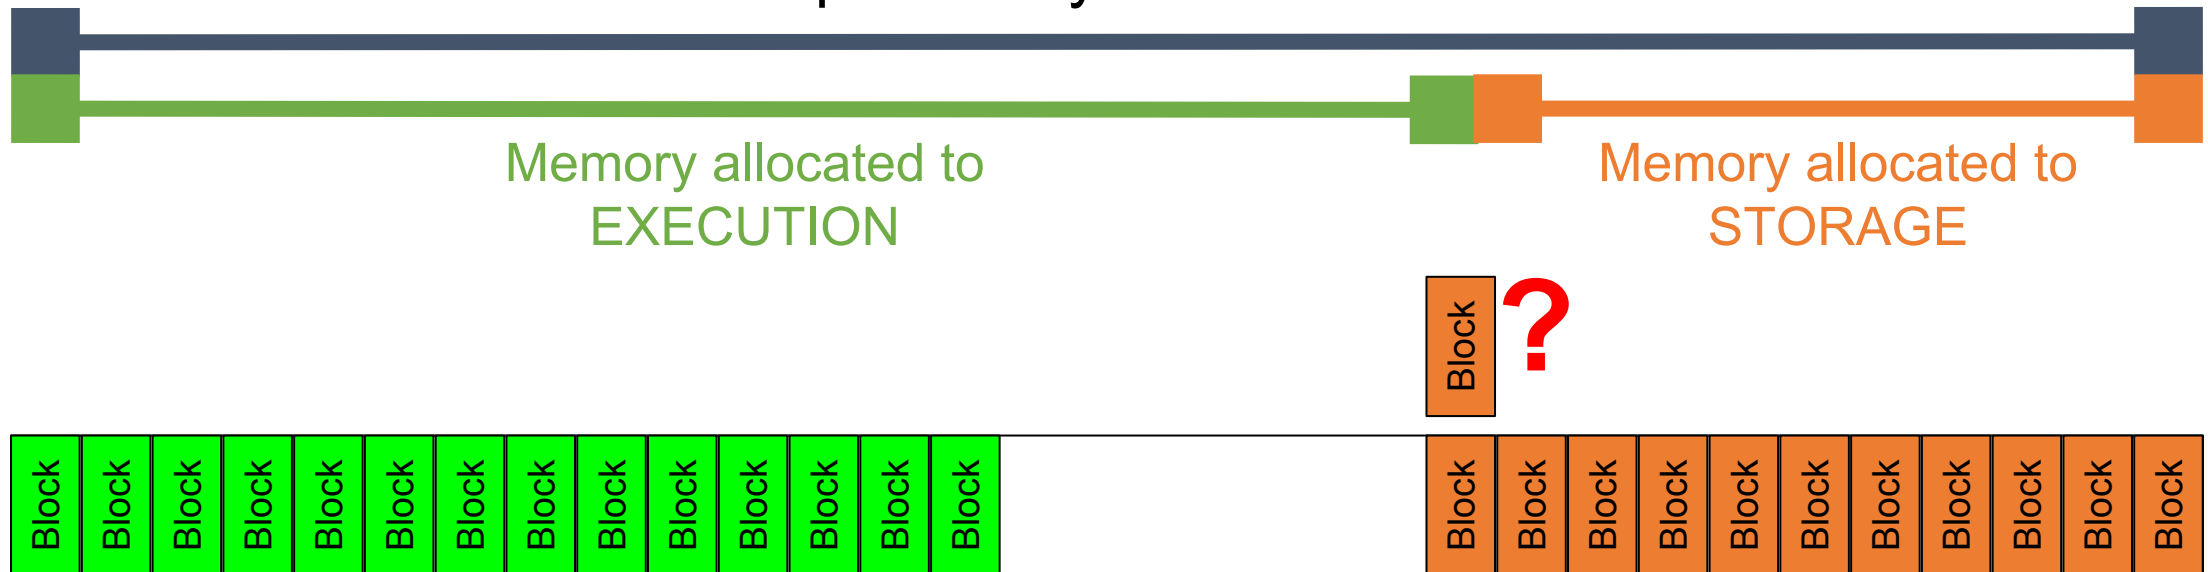
KubeCon



CloudNativeCon

Europe 2018

Total JVM Heap Memory allocated to SPARK JOB



redhat.



KubeCon



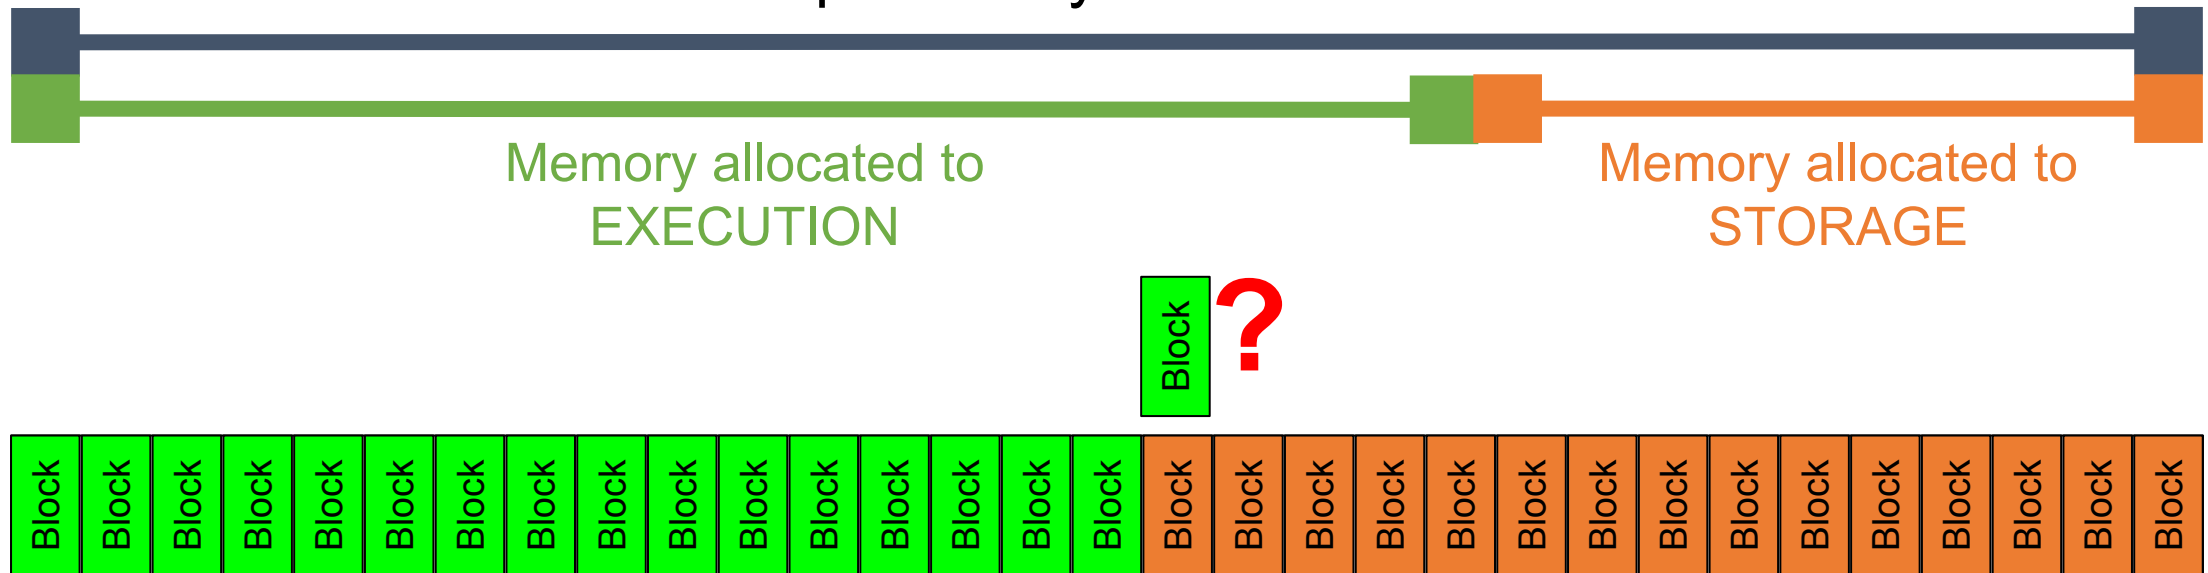
CloudNativeCon

Europe 2018

Spark Unified Memory Model

LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



redhat.



KubeCon



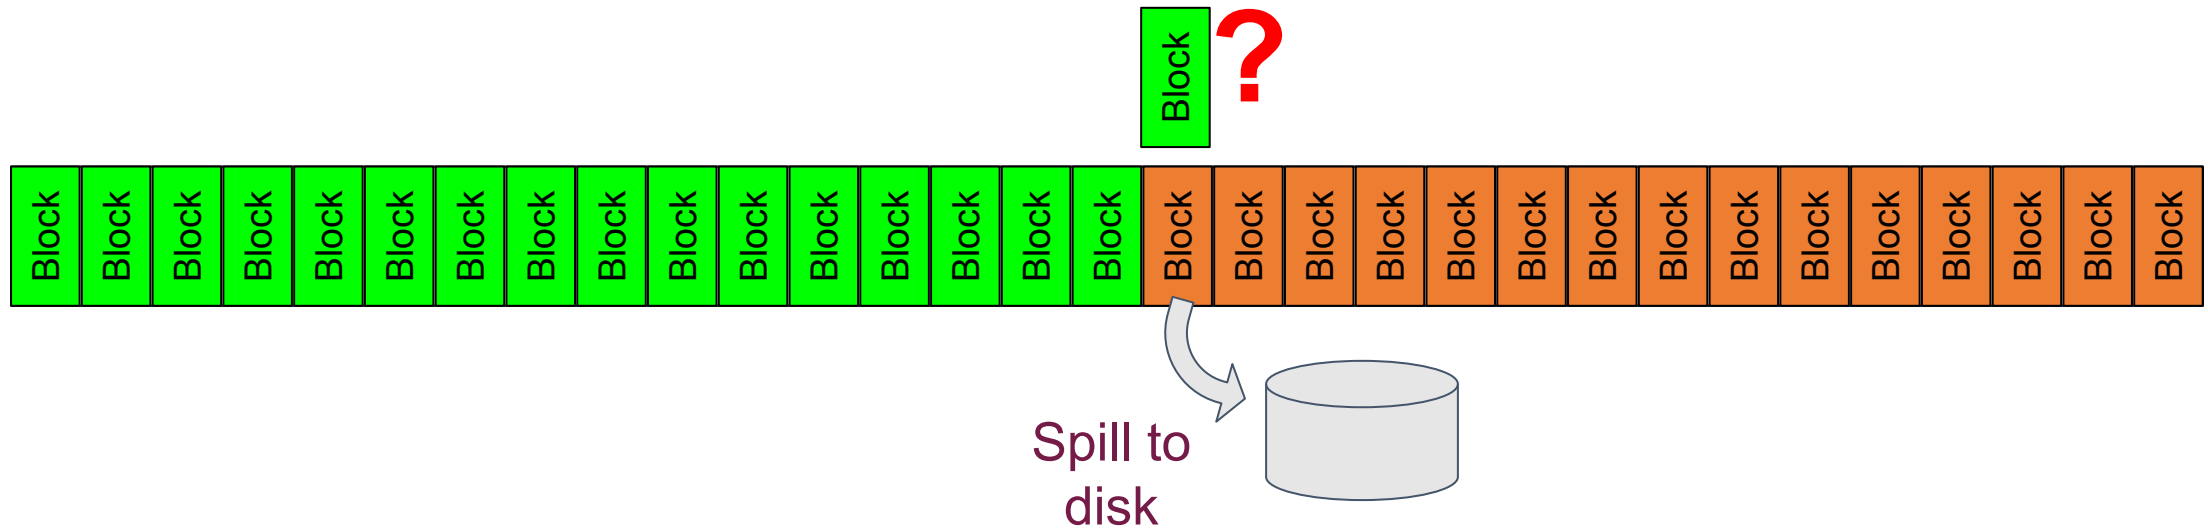
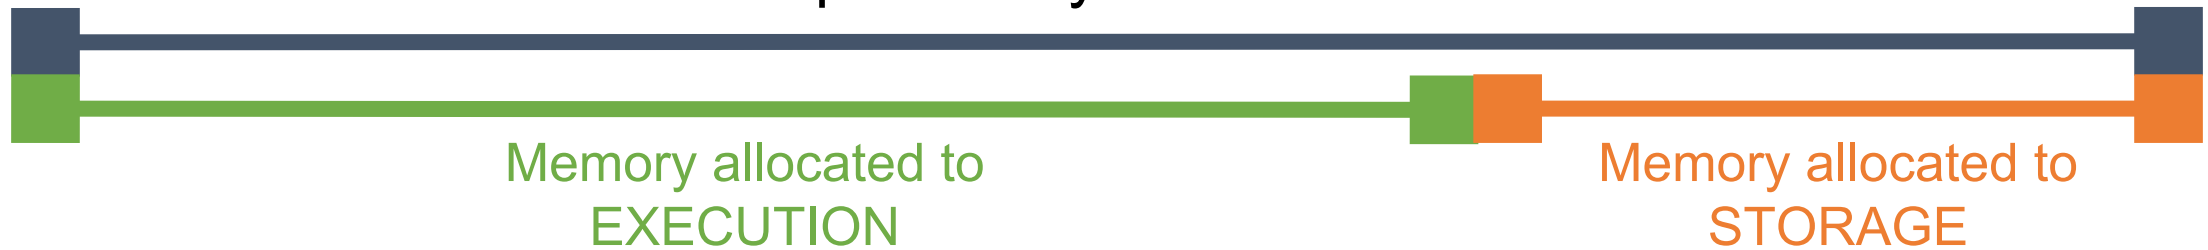
CloudNativeCon

Europe 2018

Spark Unified Memory Model

LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



Spark Unified Memory Model

LRU eviction and user defined memory configuration options



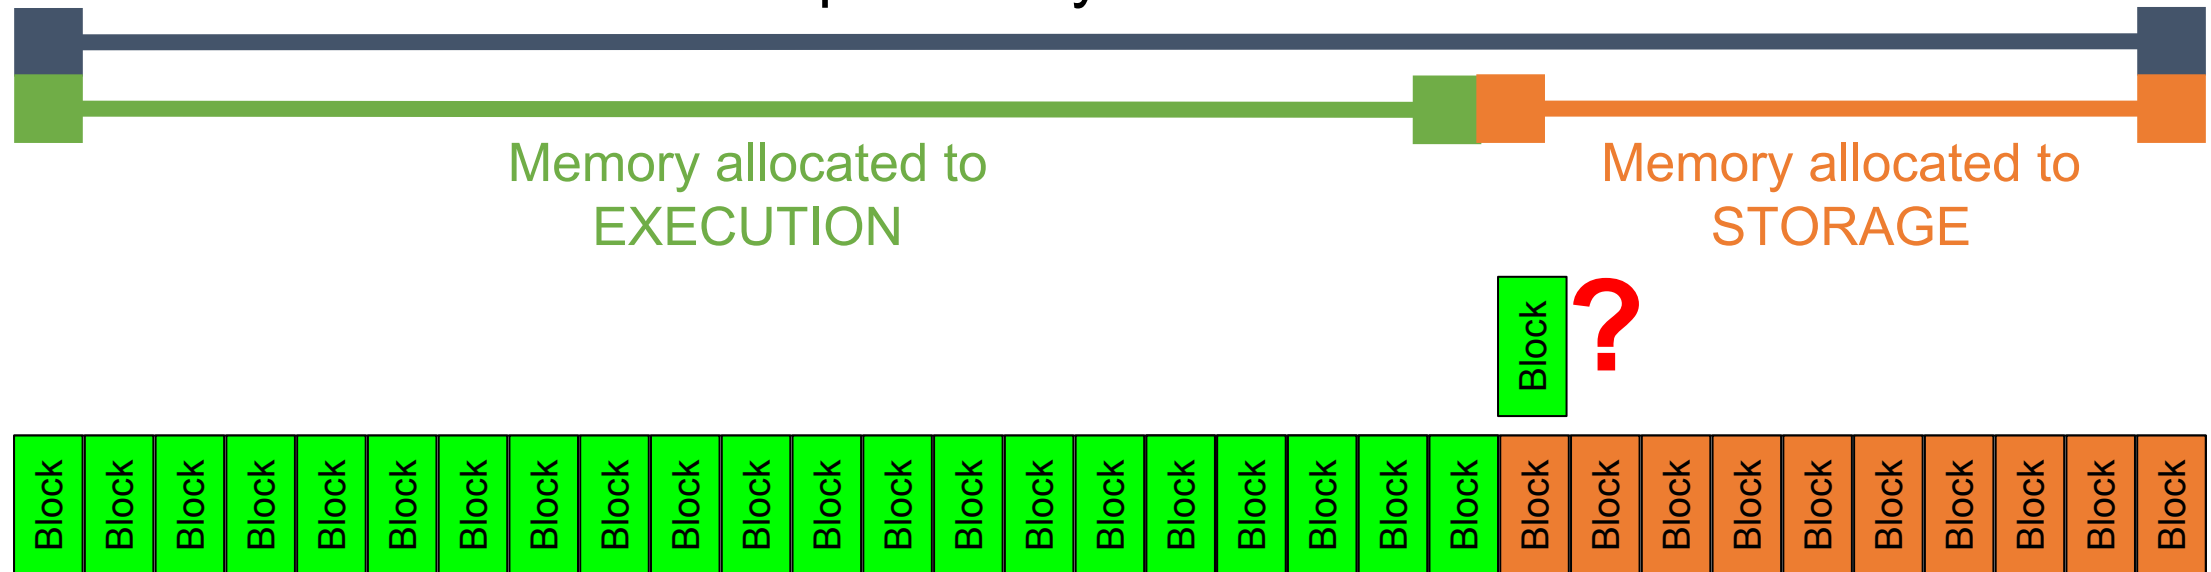
KubeCon



CloudNativeCon

Europe 2018

Total JVM Heap Memory allocated to SPARK JOB



redhat.



KubeCon



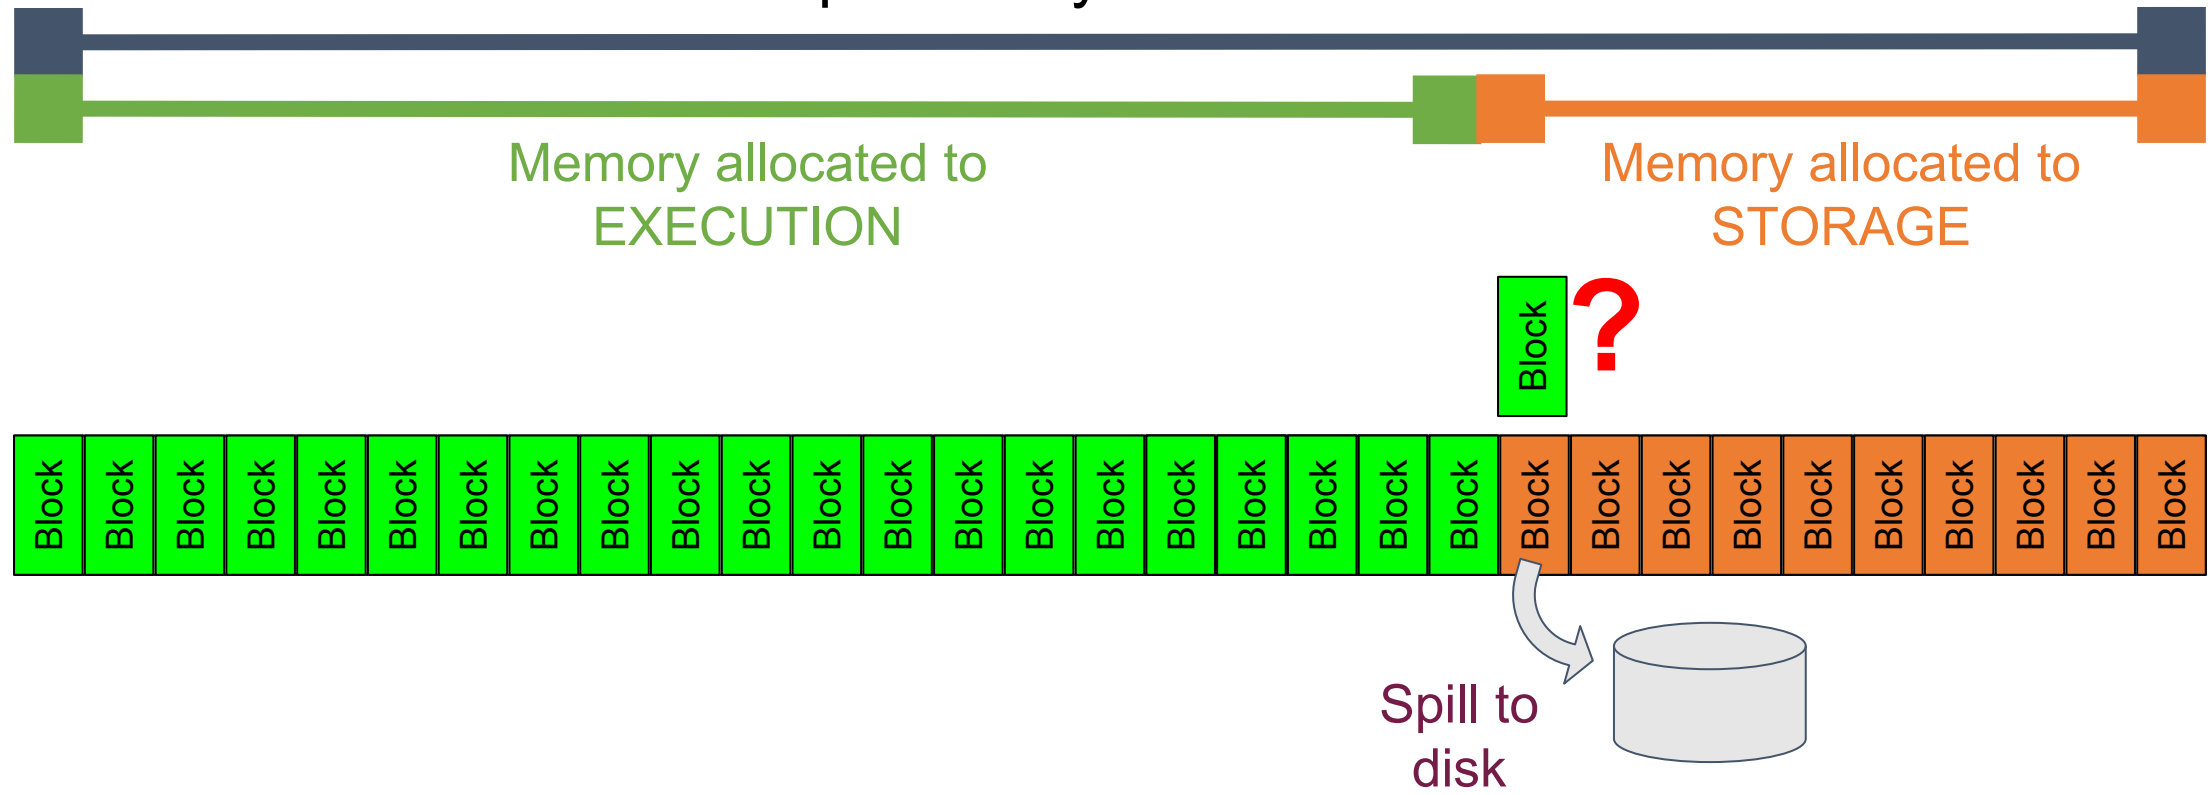
CloudNativeCon

Europe 2018

Spark Unified Memory Model

LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



redhat.



KubeCon



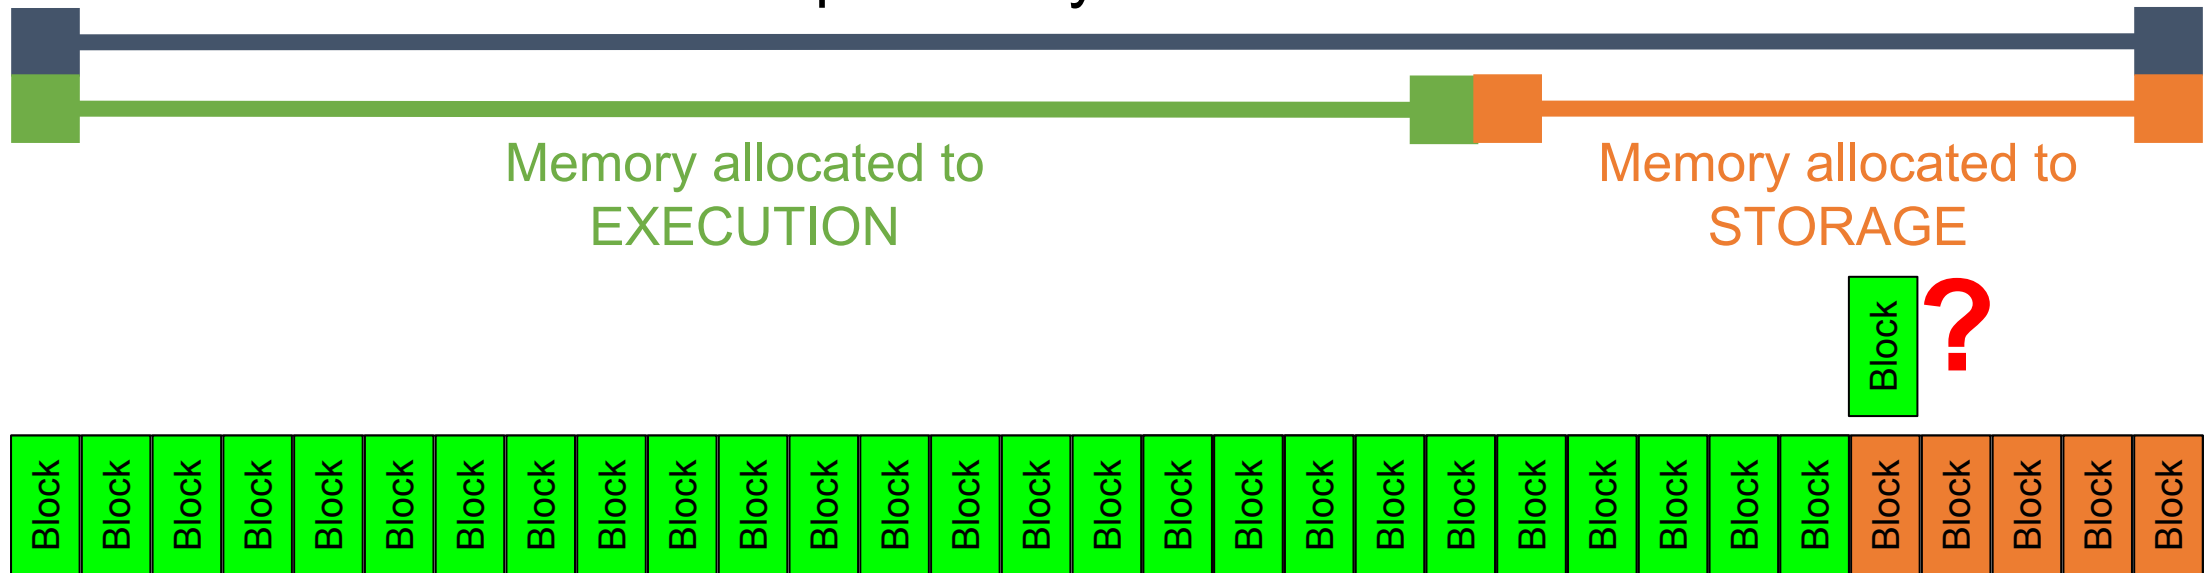
CloudNativeCon

Europe 2018

Spark Unified Memory Model

LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



redhat.



KubeCon



CloudNativeCon

Europe 2018

Spark Unified Memory Model

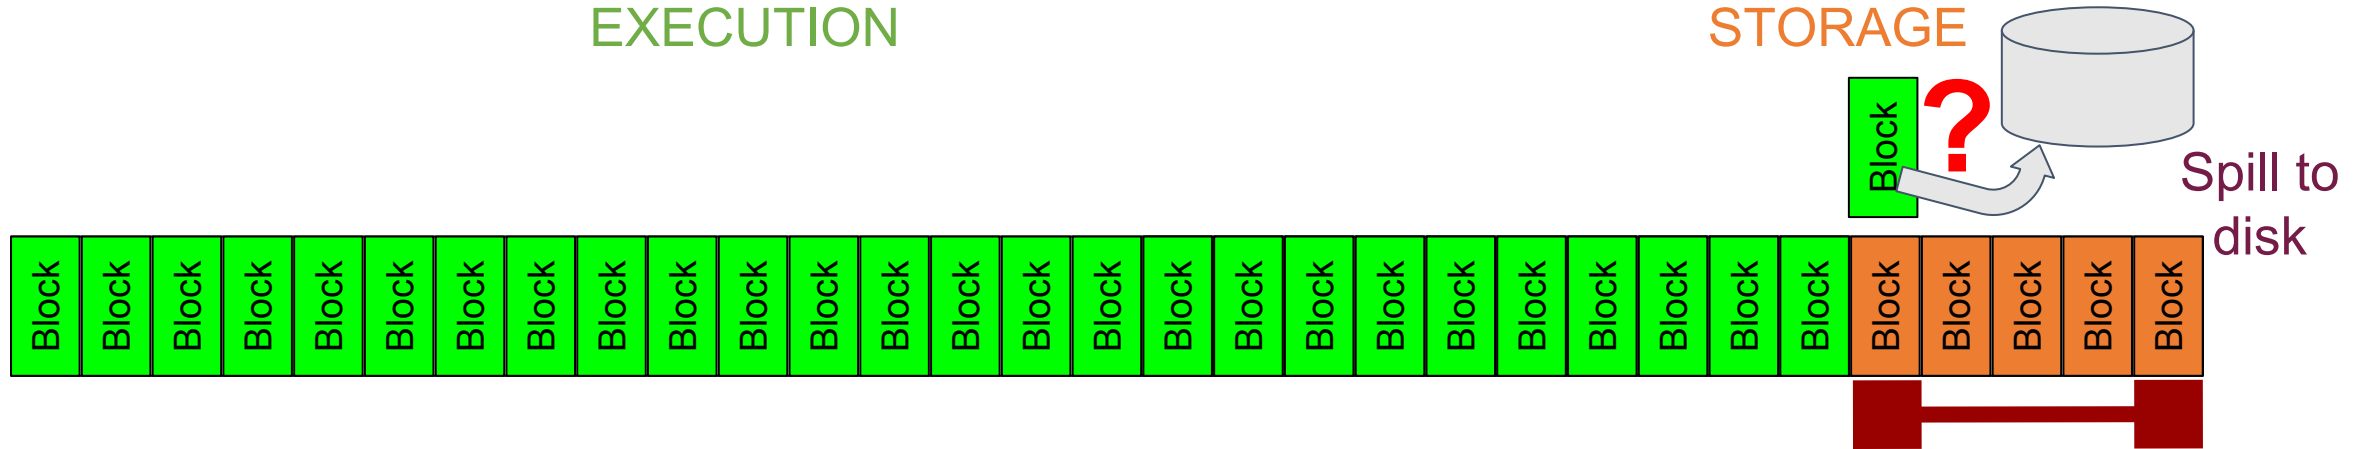
LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



Memory allocated to EXECUTION

Memory allocated to STORAGE



Spill to disk

Spark.memory.storageFraction

User specified unevictable amount



redhat



KubeCon



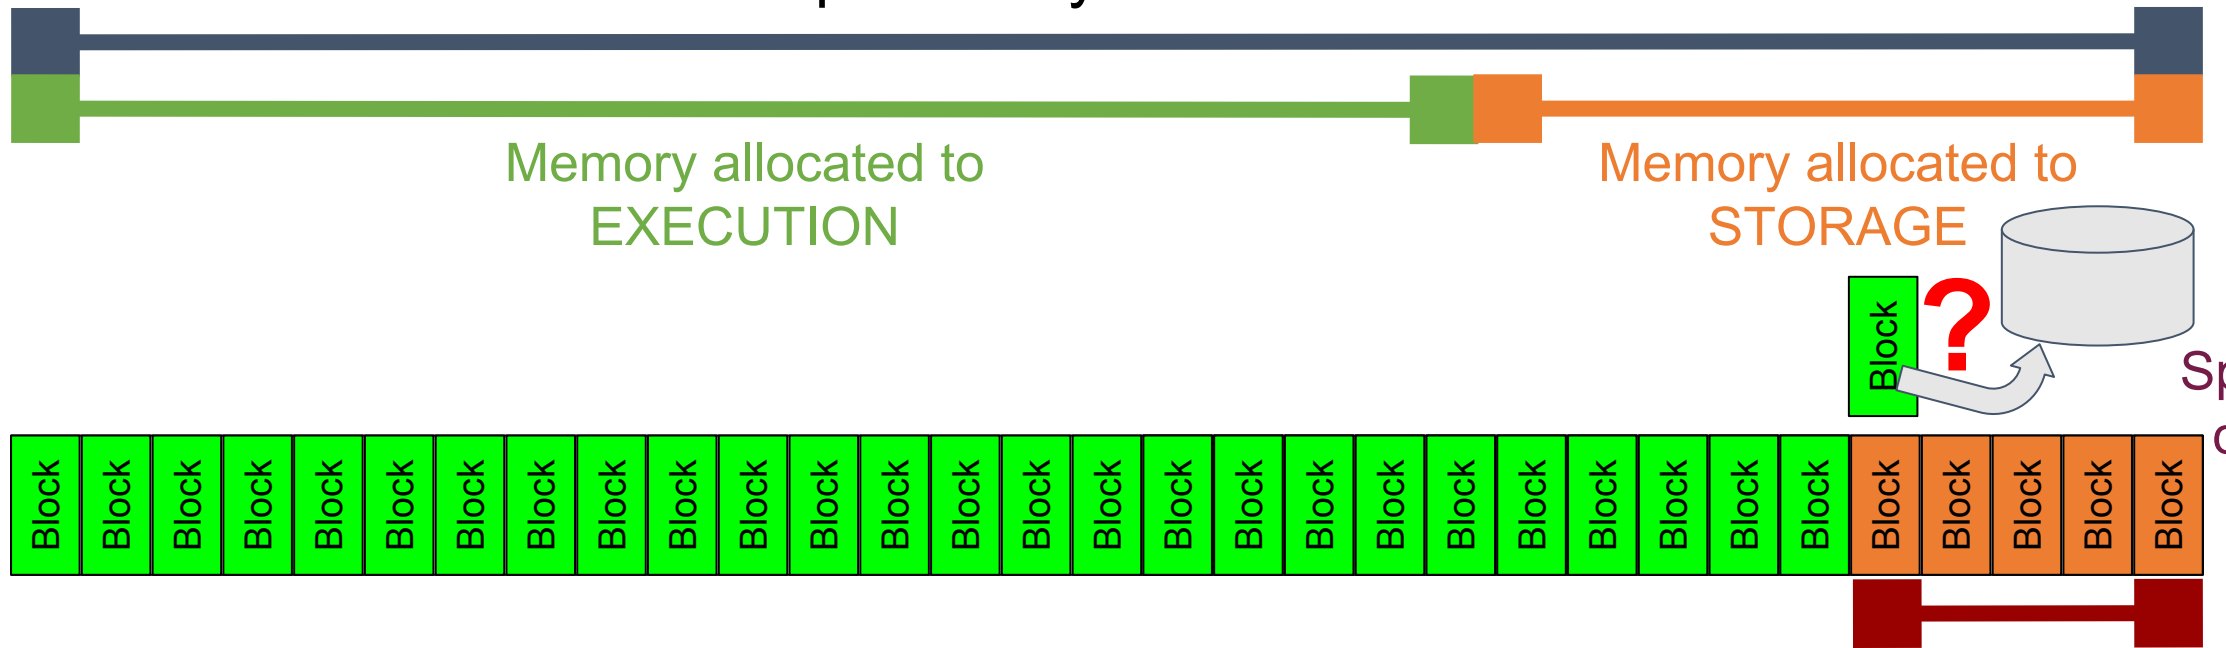
CloudNativeCon

Europe 2018

Spark Unified Memory Model

LRU eviction and user defined memory configuration options

Total JVM Heap Memory allocated to SPARK JOB



EXECUTION takes precedence over STORAGE up to user defined unevictable amount

User specified unevictable amount

Spark.memory.storageFraction



redhat

Using Spark SQL and Spark RDD API together in a tuning exercise



KubeCon



CloudNativeCon

Europe 2018

We want to use Spark SQL to manipulate dataframes

Spark SQL is a component of Spark

- it provides structured data processing
- it is implemented as a library on top of Spark

Three main APIs:

- SQL syntax
- Dataframes
- Datasets

Two backend components:

- Catalyst - query optimizer
- Tungsten - off-heap memory management eliminates overhead of Java Objects



redhat.

Performance Optimizations with Spark SQL

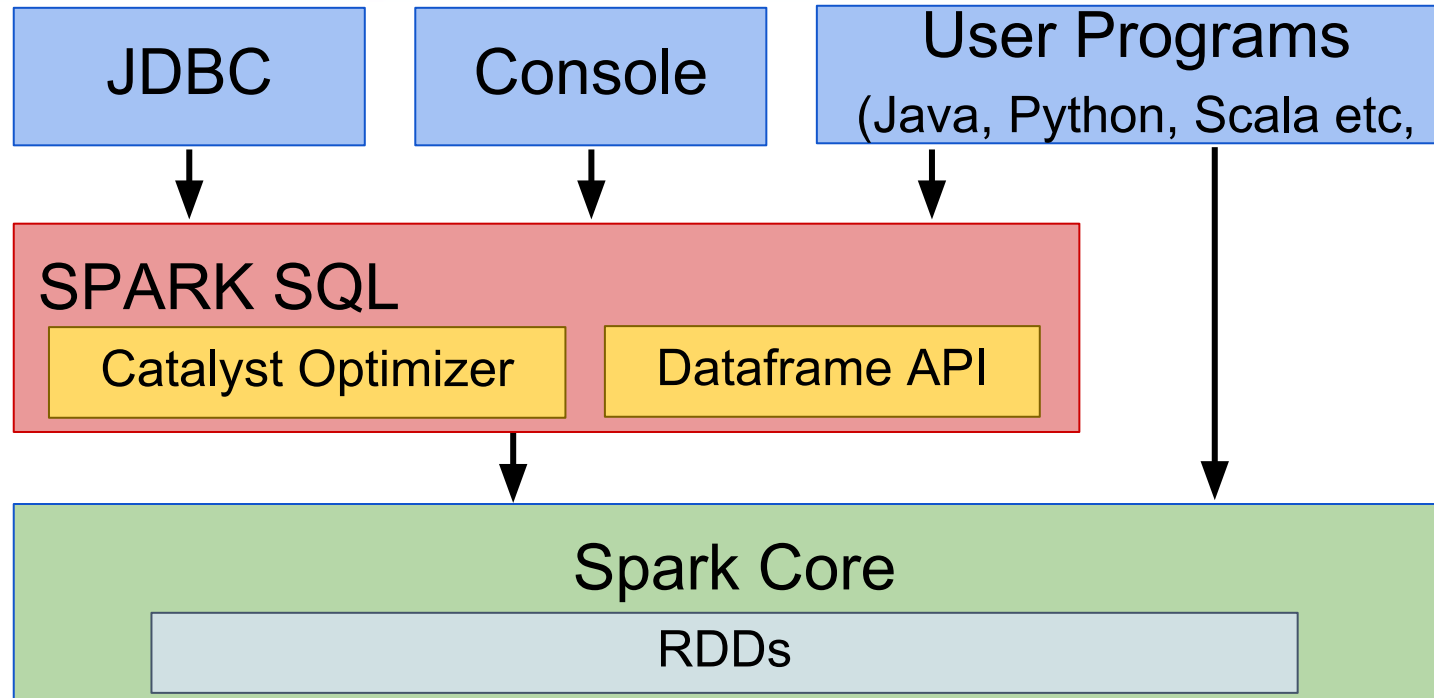


KubeCon



CloudNativeCon

Europe 2018



Spark SQL performance benefits:

- Catalyst compiles Spark SQL programs down to an RDD
- Tungsten provides more efficient data storage compared to Java objects on the heap
- Dataframe API and RDD API can be freely intermixed



redhat.

Using Prometheus + Grafana for performance optimization



KubeCon



CloudNativeCon

Europe 2018

Specific code example:

Compare **non-cached** and **cached** dataframes that are reused in a groupBy transformation

When is good idea to use cache in a dataframe?

- when a result of a computation is going to be reused later
- when it is costly to recompute that result
- in cases where algorithms make several passes over the data



redhat.

Determining memory consumption for dataframes you want to cache



KubeCon



CloudNativeCon

Europe 2018



lg_200k_cartprod_cache2.py application UI

Jobs Stages **Storage** Environment Executors SQL

Storage

RDDs

RDD Name	Storage Level	Cached Partitions	Fraction Cached	Size in Memory	Size on Disk
Scan ExistingRDD[E#9,F#10,G#11,H#12]	Memory Deserialized 1x Replicated	32	100%	6.1 MB	0.0 B
Scan ExistingRDD[A#0,B#1,C#2,D#3]	Memory Deserialized 1x Replicated	32	100%	6.1 MB	0.0 B



redhat.

Example: Code for non-cached run



KubeCon



CloudNativeCon

Europe 2018

```
rdd1 = RandomRDDs.normalVectorRDD(spark, nRow, nCol, numPartitions, seed)
seed = 3
rdd2 = RandomRDDs.normalVectorRDD(spark, nRow, nCol, numPartitions, seed)
sc = spark.sparkContext
# convert each tuple in the rdd to a row
randomNumberRdd1 = rdd1.map(lambda x: Row(A=float(x[0]), B=float(x[1]), C=float(x[2]), D=float(x[3])))
randomNumberRdd2 = rdd2.map(lambda x: Row(E=float(x[0]), F=float(x[1]), G=float(x[2]), H=float(x[3])))
# create dataframe from rdd
schemaRandomNumberDF1 = spark.createDataFrame(randomNumberRdd1)
schemaRandomNumberDF2 = spark.createDataFrame(randomNumberRdd2)
cross_df = schemaRandomNumberDF1.crossJoin(schemaRandomNumberDF2)
# aggregate
results = schemaRandomNumberDF1.groupBy("A").agg(func.max("B"),func.sum("C"))
results.show(n=100)
print "-----Count in cross-join----- {0}".format(cross_df.count())
```



redhat.

Example: Code for **cached** run



KubeCon



CloudNativeCon

Europe 2018

```
rdd1 = RandomRDDs.normalVectorRDD(spark, nRow, nCol, numPartitions, seed)
seed = 3
rdd2 = RandomRDDs.normalVectorRDD(spark, nRow, nCol, numPartitions, seed)
sc = spark.sparkContext
# convert each tuple in the rdd to a row
randomNumberRdd1 = rdd1.map(lambda x: Row(A=float(x[0]), B=float(x[1]), C=float(x[2]), D=float(x[3])))
randomNumberRdd2 = rdd2.map(lambda x: Row(E=float(x[0]), F=float(x[1]), G=float(x[2]), H=float(x[3])))
# create dataframe from rdd
schemaRandomNumberDF1 = spark.createDataFrame(randomNumberRdd1)
schemaRandomNumberDF2 = spark.createDataFrame(randomNumberRdd2)
# cache the dataframe
schemaRandomNumberDF1.cache()
schemaRandomNumberDF2.cache()
cross_df = schemaRandomNumberDF1.crossJoin(schemaRandomNumberDF2)
# aggregate
results = schemaRandomNumberDF1.groupBy("A").agg(func.max("B"),func.sum("C"))
results.show(n=100)
print "-----Count in cross-join----- {0}".format(cross_df.count())
```



redhat.

Query plan comparison



KubeCon



CloudNativeCon

Europe 2018

Non-Cached

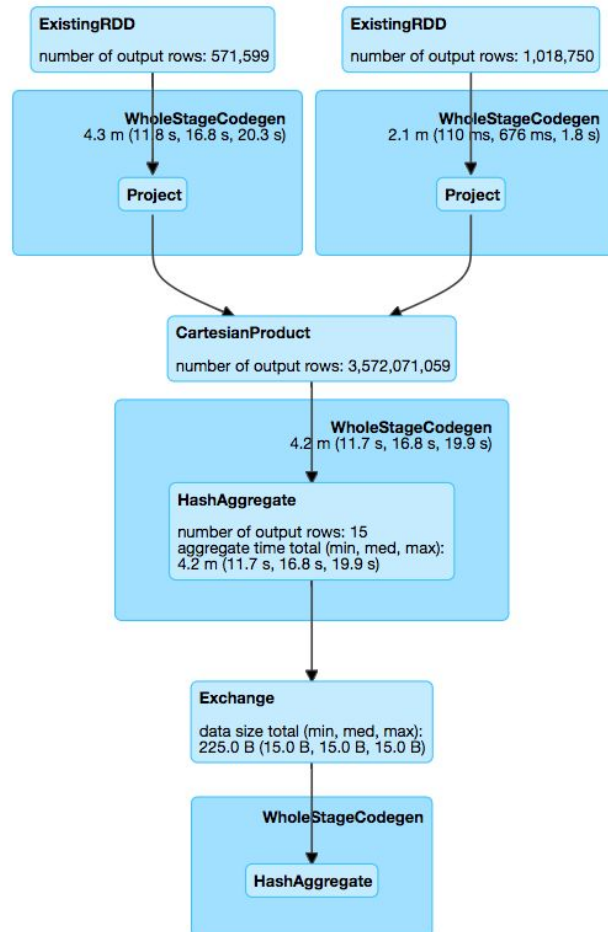
Cached

Details for Query 1

Submitted Time: 2018/04/12 14:29:04

Duration: 22 s

Running Jobs: 3



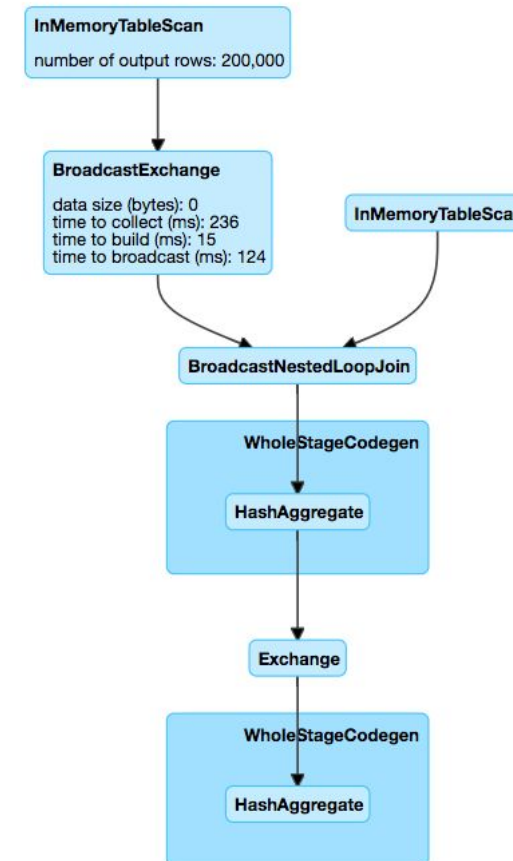
Details for Query 1

Submitted Time: 2018/04/13 04:11:24

Duration: 4 s

Running Jobs: 4

Succeeded Jobs: 3



redhat.

Example: Comparing cached vs non-cached runs



KubeCon

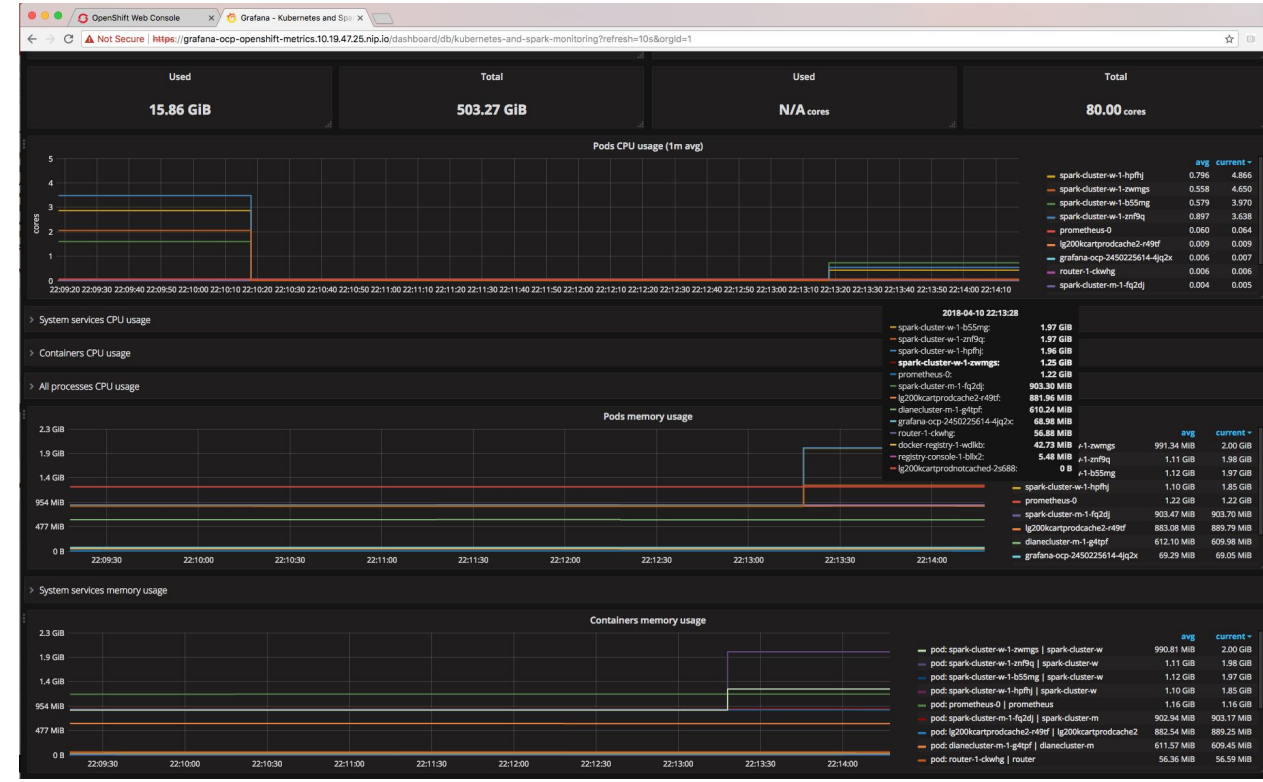
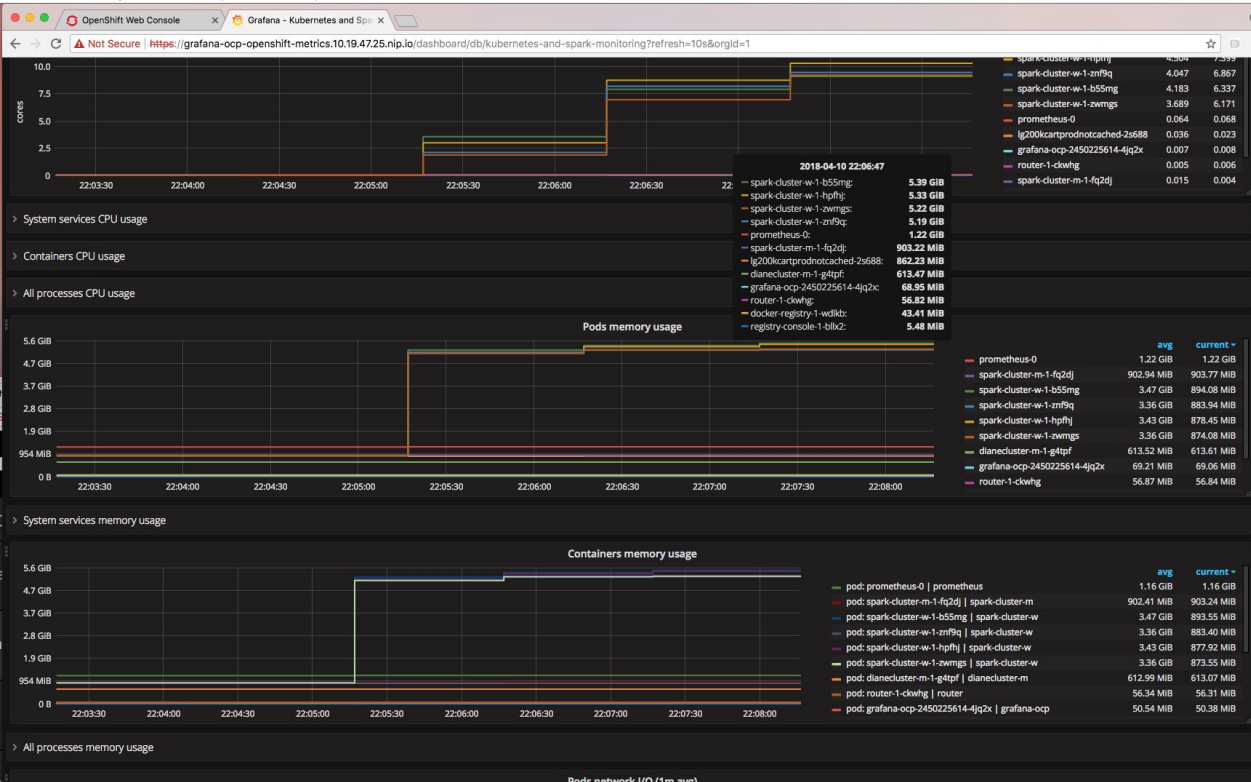


CloudNativeCon

Europe 2018

Prometheus dashboard: non-cached

Prometheus dashboard: cached



Example: Comparing cached vs non-cached runs



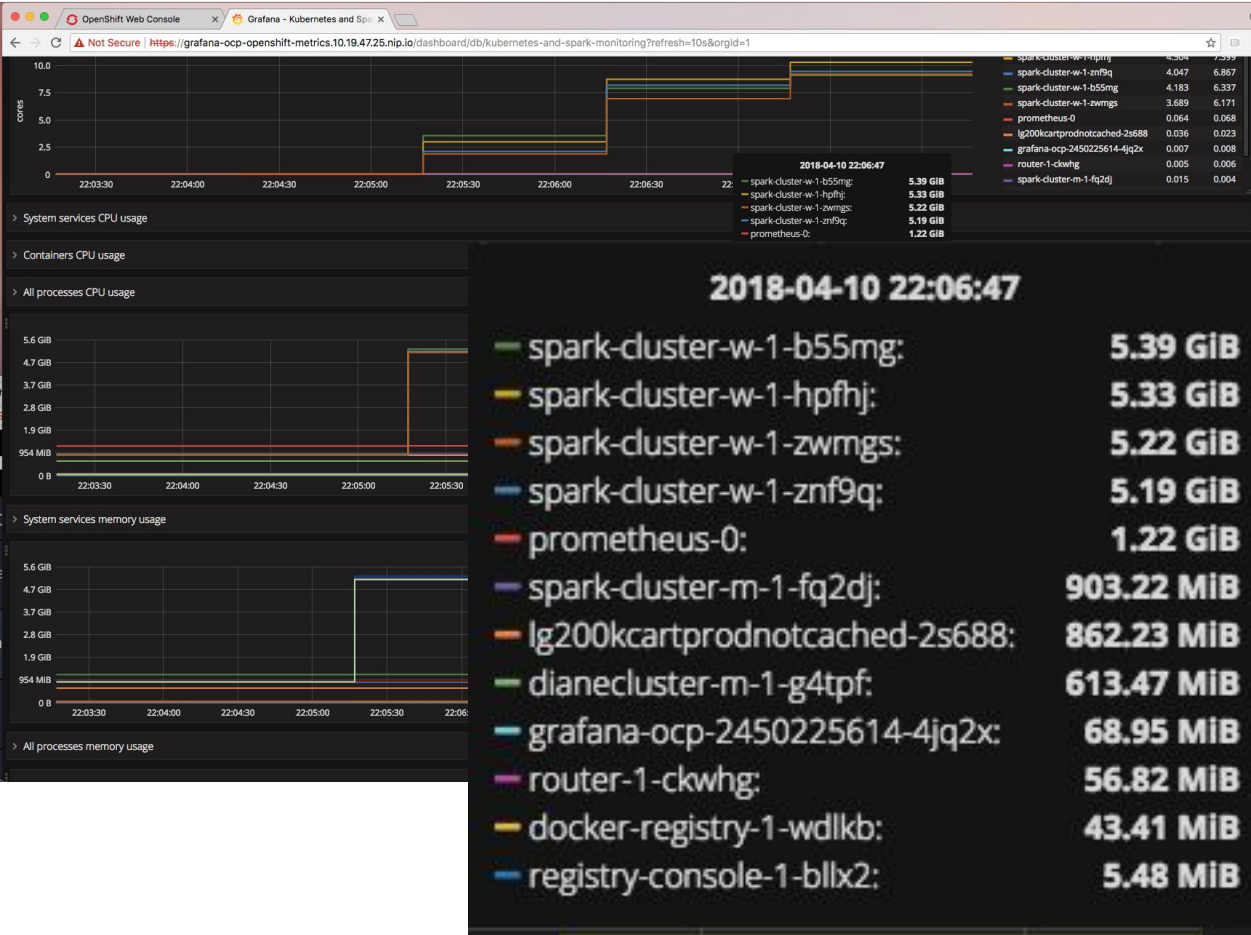
KubeCon



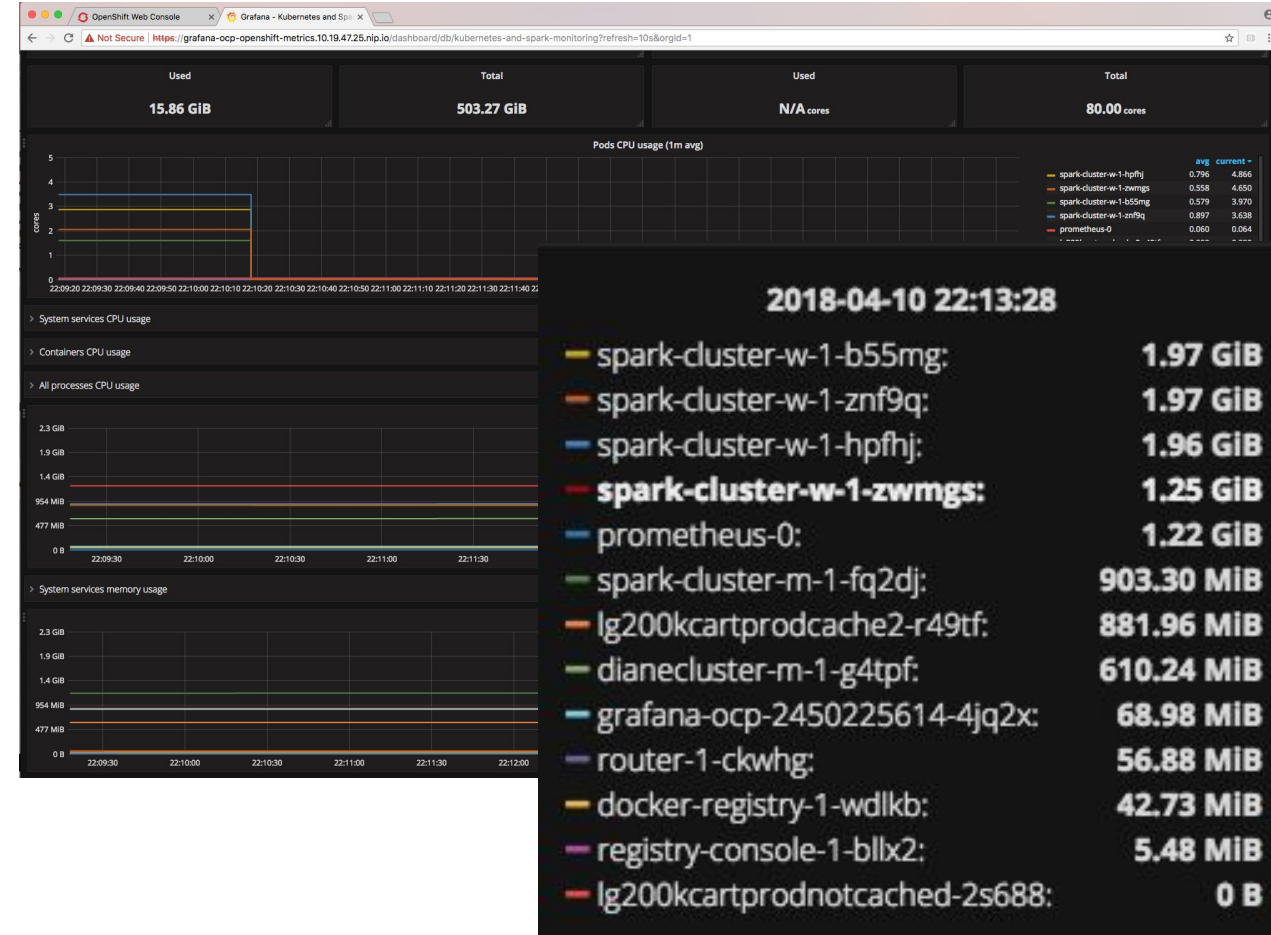
CloudNativeCon

Europe 2018

Prometheus dashboard: non-cached



Prometheus dashboard: cached



Comparing non-cached vs cached runs



KubeCon



CloudNativeCon

Europe 2018

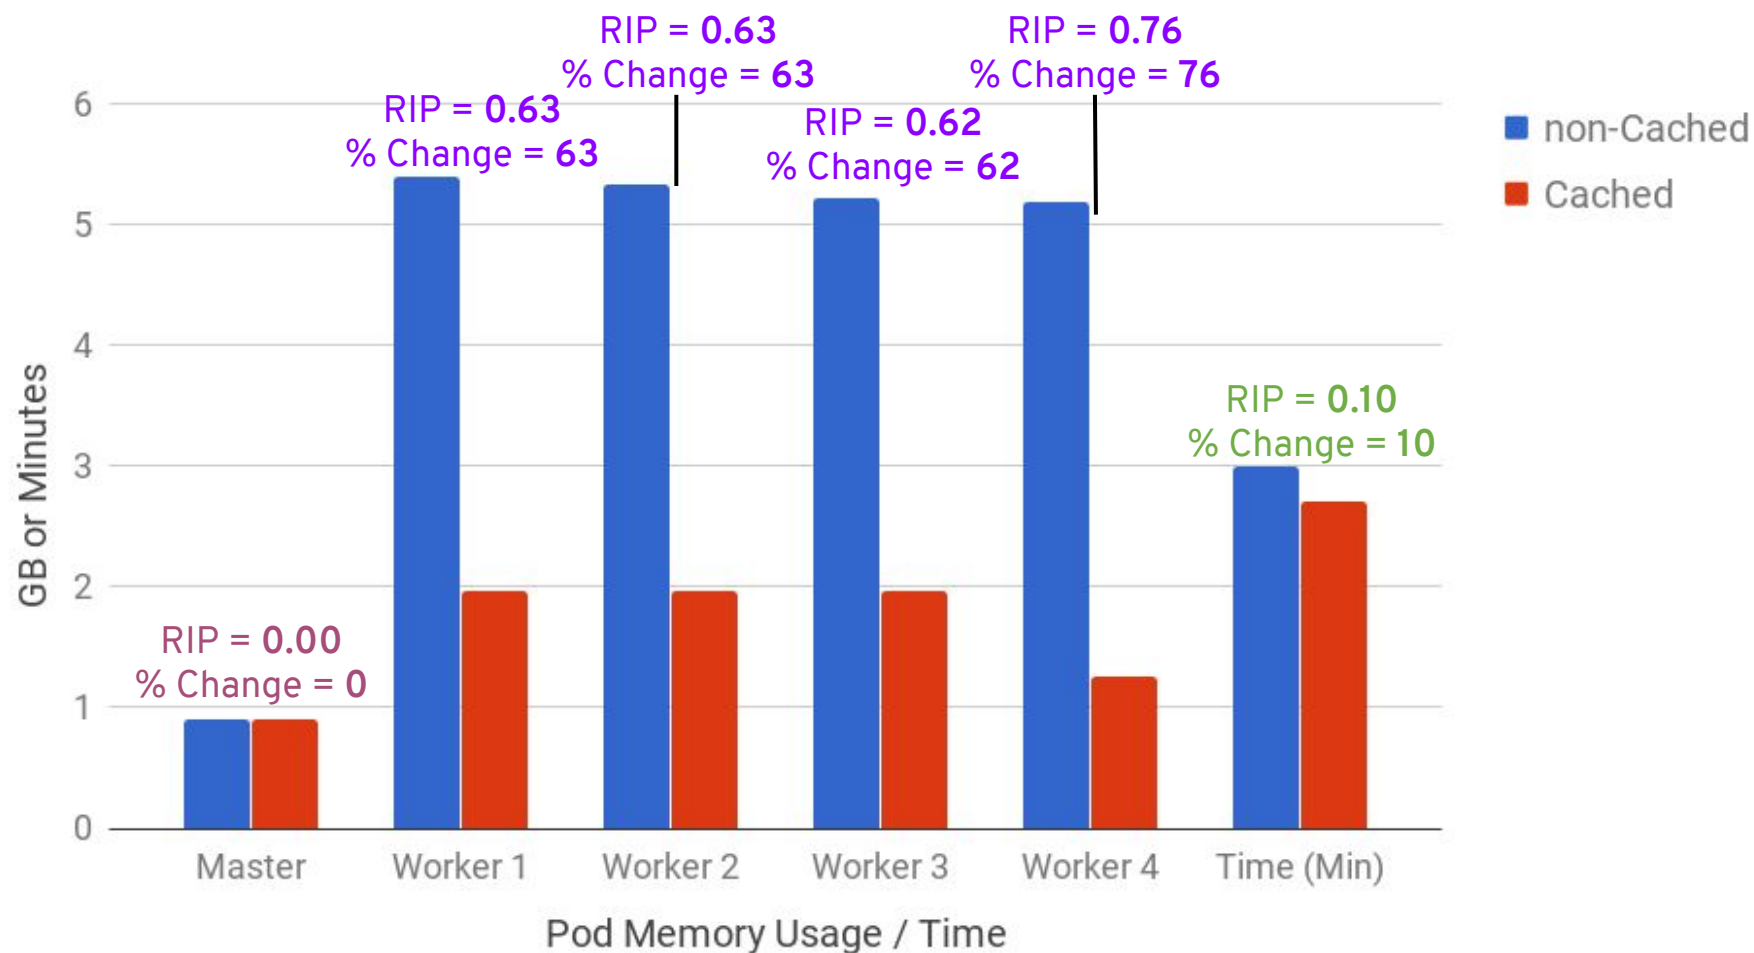
Highwater mark memory usage for master and worker pods (GB) and timing (min)

RIP (Relative Index of Performance)

RIP: 0 to 1 = Improvement

0 to -1 = Degradation

% Change: negative values = Improvement



redhat.

Demo Time!



KubeCon



CloudNativeCon

Europe 2018

SPARK JOB + PROMETHEUS + GRAFANA DEMO



TRY THIS AT HOME



KubeCon



CloudNativeCon

Europe 2018

- Download prometheus here: <https://prometheus.io/download/>
- GET SAMPLE PROM CONFIG:
 - sample_prom_config.yml
<https://gist.github.com/zmhassan/7fdc763095ebe09d5516c8c395fa163e>
 - sample_alertmanager_simple_rule.yml
 - <https://gist.github.com/zmhassan/6dc27c4238fbd253df9c061df7dfe208>



redhat.

Recap



KubeCon



CloudNativeCon

Europe 2018

You learned:

- About our story of how we added prometheus to monitor our spark cluster metrics.
- Spark Features?
- What is Prometheus?
- How to Create Custom Instrumentation?
- Spark Applications and how memory works
- Spark Cluster JVM Instrumentation
- Monitoring tips and tricks.
- How to deploy a spark job and monitor it via grafana dashboard.
- cache vs non-cached dataframes



Thank You!



KubeCon



CloudNativeCon

Europe 2018

Questions?