

Apache Spark : Fast and Easy Data Processing

Sujee Maniyam

Elephant Scale LLC

sujee@elephantscale.com

<http://elephantscale.com>

- ❑ Fast & Expressive Cluster computing engine
- ❑ Compatible with Hadoop
- ❑ Came out of Berkeley AMP Lab
- ❑ Now Apache project
- ❑ Version 1.1 just released (Sep 2014)

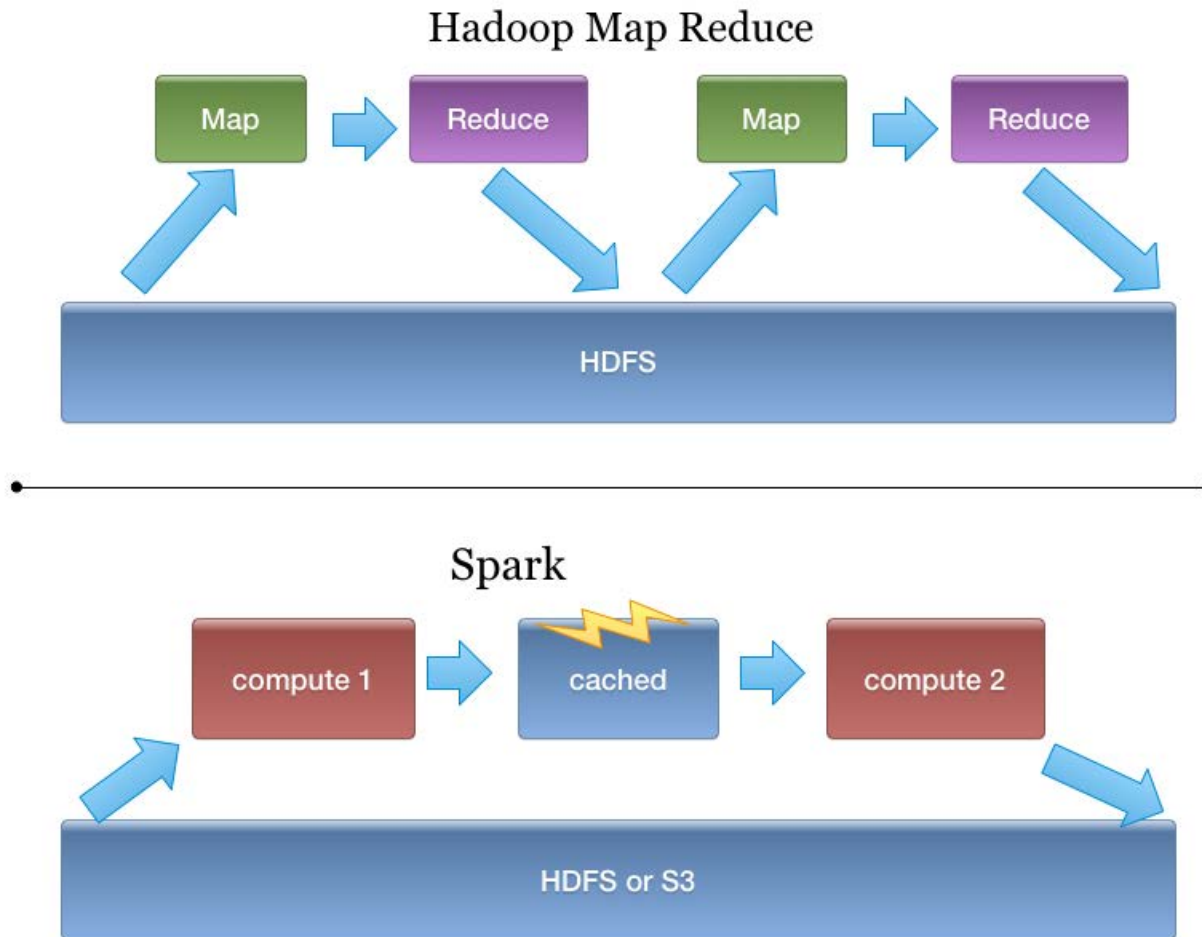
Comparison With Hadoop

| Hadoop | Spark |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Distributed Storage + Distributed Compute | Distributed Compute Only |
| MapReduce framework | Generalized computation |
| Usually data on disk (HDFS) | On disk / in memory |
| Not ideal for iterative work | Great at Iterative workloads (machine learning ..etc) |
| Batch process | <ul style="list-style-type: none">- Upto 10x faster for data on disk- Upto 100x faster for data in memory |
| | Compact code Java, Python, Scala supported |
| | Shell for ad-hoc exploration |

Spark Vs Hadoop

- ❑ Spark is 'easier' than Hadoop
- ❑ 'friendlier' for data scientists
- ❑ Interactive shell (adhoc exploration)
- ❑ API supports multiple languages
 - ❑ Java, Scala, Python
- ❑ Great for small (Gigs) to medium (100s of Gigs) data

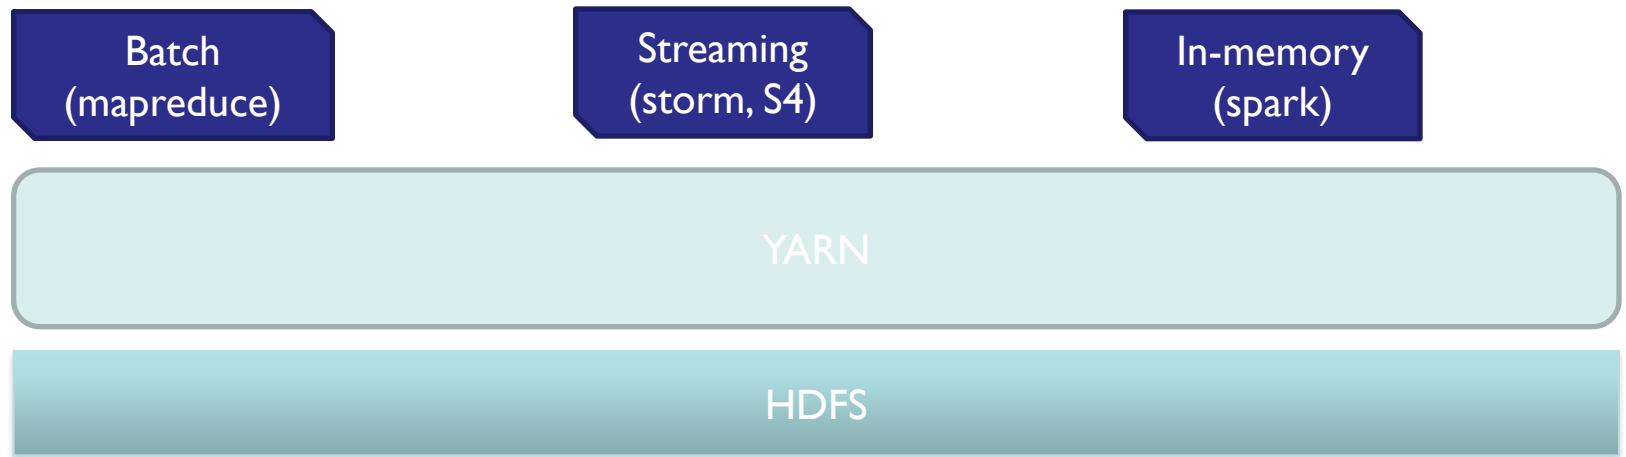
Spark Vs. Hadoop



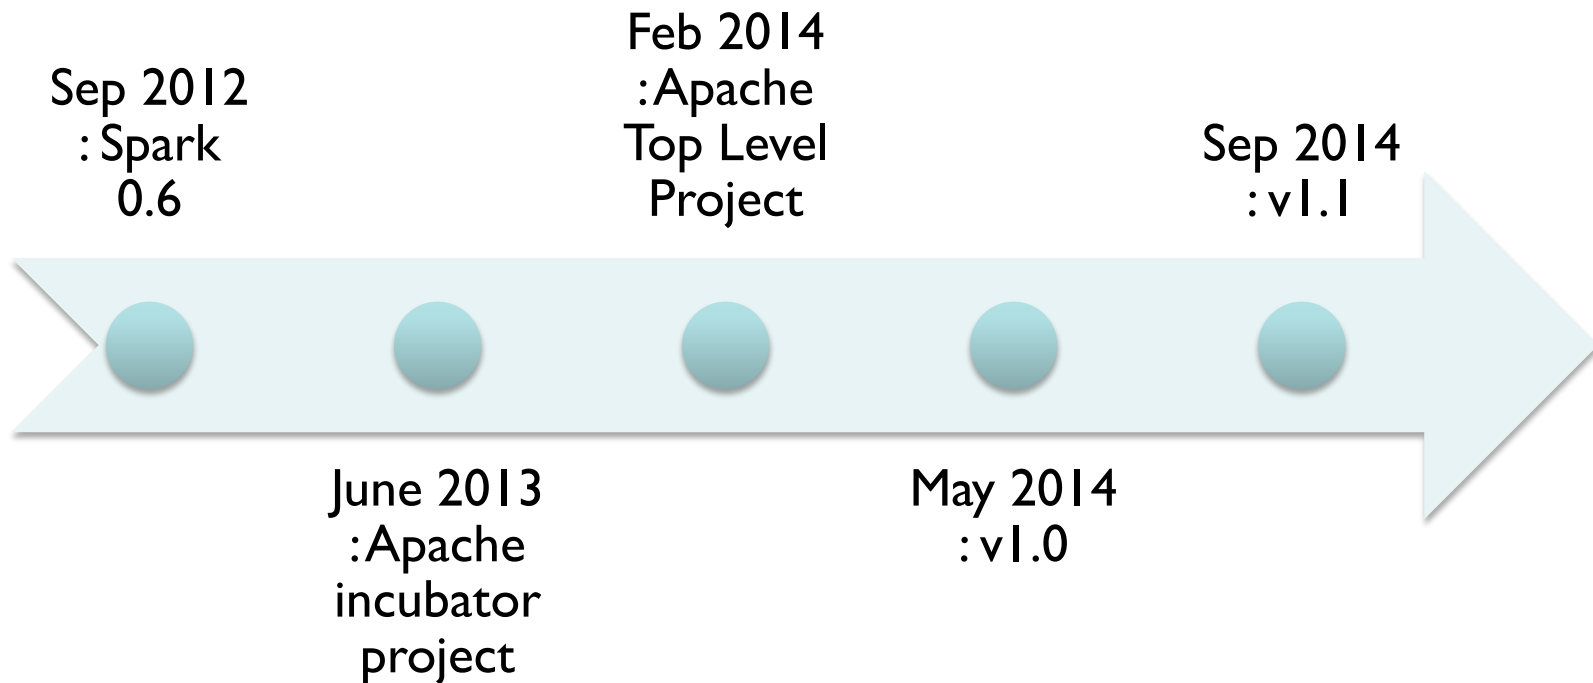
Is Spark Replacing Hadoop?


- ❑ Right now, Spark runs on Hadoop / YARN
 - ❑ Complimentary
- ❑ Can be seen as generic MapReduce
- ❑ Spark is really great if data fits in memory (few hundred gigs),
- ❑ People are starting to use Spark as their only compute platform
- ❑ Future ???

Hadoop + Yarn : Universal OS for Cluster Computing



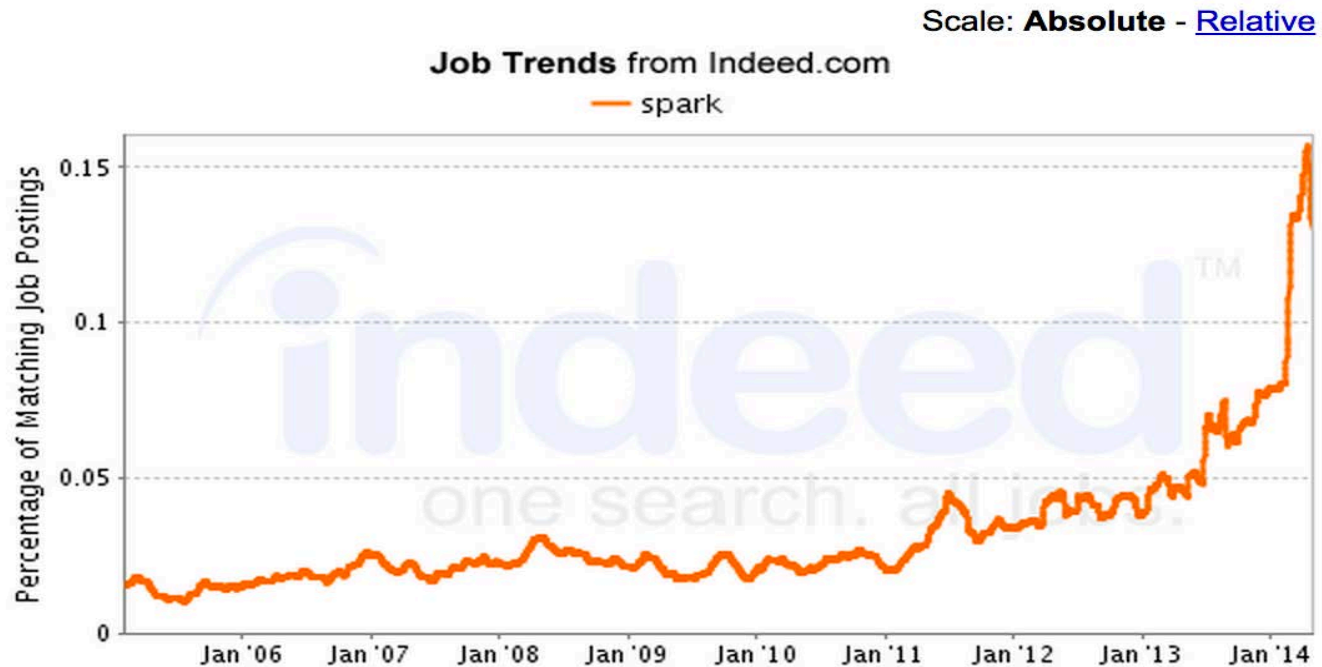
Bit of History of Spark



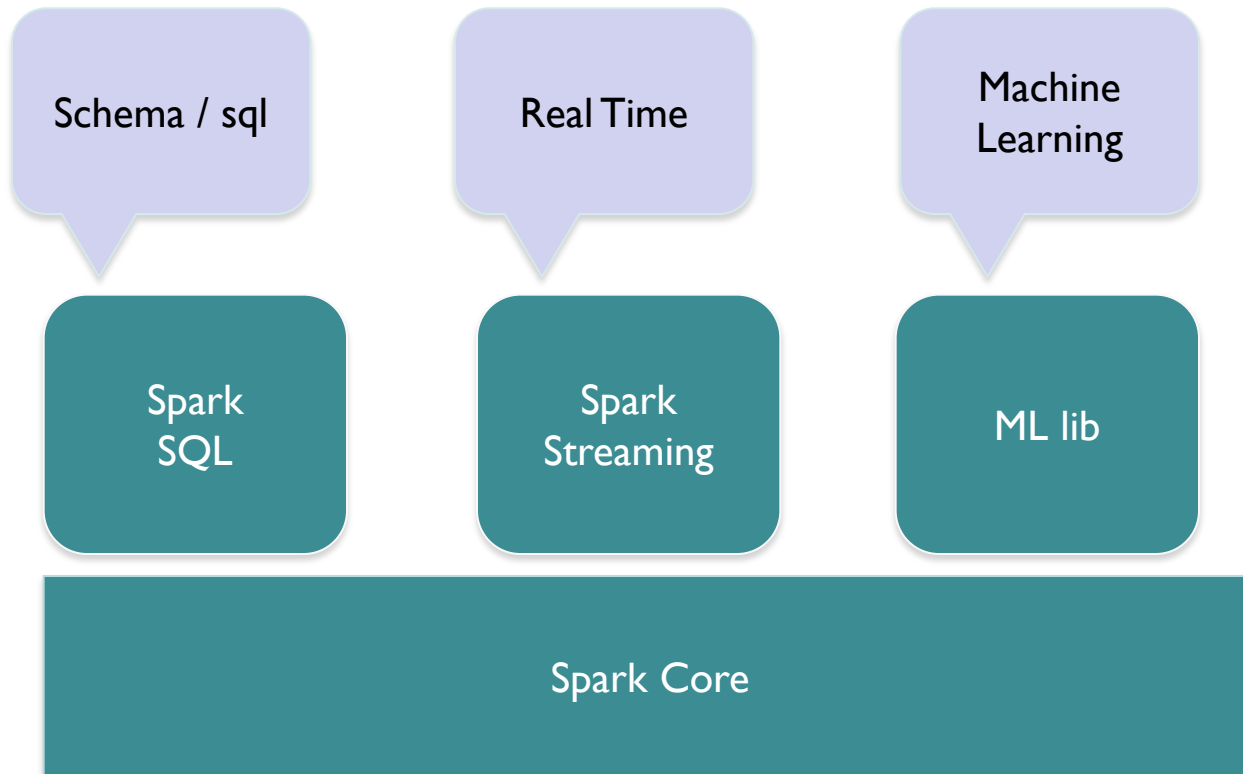


Spark Job Trends

spark Job Trends



Spark Eco-System



- ❑ Distributed compute engine
- ❑ Sort / shuffle algorithms
- ❑ In case of node failures -> re-computes missing pieces

Spark Streaming

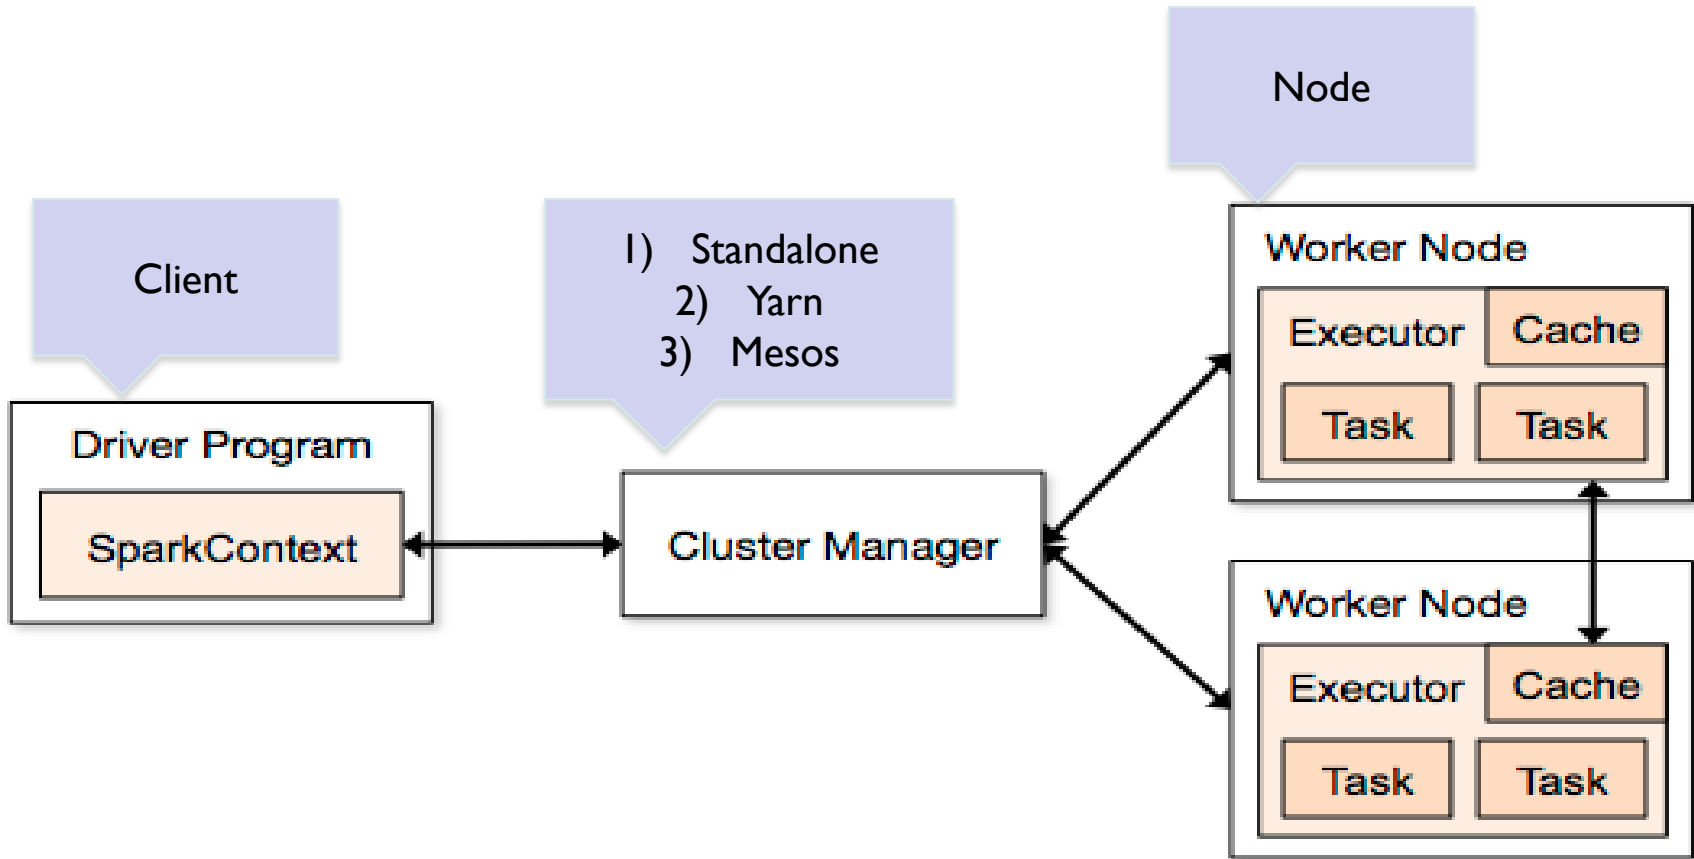
- ❑ Process data streams in **real time**
- ❑ Stock ticks / click streams ...etc



Machine Learning (ML Lib)

- ❑ Out of the box ML capabilities !
- ❑ Lots of common algorithms are supported
- ❑ Classification / Regressions
 - ❑ Linear models (linear R, logistic regression, SVM)
 - ❑ Decision trees
- ❑ Collaborative filtering (recommendations)
- ❑ K-Means clustering
- ❑ ...
- ❑ More to come

Spark Architecture



Spark Architecture

- ❑ Multiple 'applications' can run at the same time
- ❑ Driver (or 'main') launches an application
- ❑ Each application gets its own 'executor'
 - ❑ Isolated (runs in different JVMs)
 - ❑ Also means data can not be shared across applications
- ❑ Cluster Managers:
 - ❑ multiple cluster managers are supported
 - ❑ 1) Standalone : simple to setup
 - ❑ 2) YARN : on top of Hadoop
 - ❑ 3) Mesos : General cluster manager (AMP lab)

Spark Data Model : RDD

- ❑ Resilient Distributed Dataset (RDD)
- ❑ Can live in
 - ❑ Memory (best case scenario)
 - ❑ Or on disk (FS, HDFS, S3 ...etc)
- ❑ Each RDD is split into multiple partitions
- ❑ Partitions may live on different nodes
- ❑ Partitions can be computed in parallel on different nodes

RDD : Loading

- Use Spark context to load RDDs from disk / external storage

```
val sc = new SparkContext(...)
```

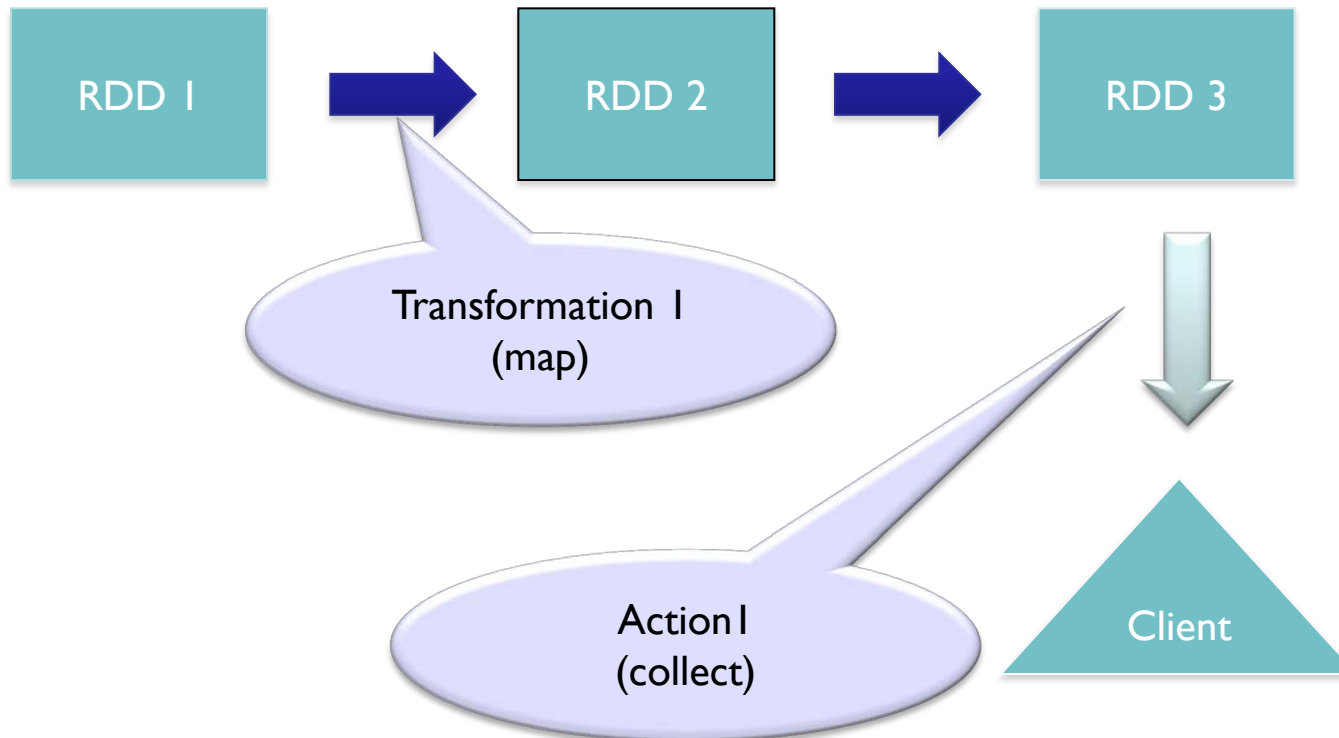
```
val f = sc.textFile("/data/input1.txt") // single file
```

```
sc.textFile("/data/") // load all files under dir
```

```
sc.textFile("/data/*.log") // wild card matching
```

- ❑ Two kinds of operations on RDDs
 - ❑ 1) Transformations
 - ❑ Create a new RDD from existing ones (e.g. Map)
 - ❑ 2) Actions
 - ❑ E.g. Returns the results to clients (e.g. Reduce)
- ❑ Transformations are **lazy**.. Actions **force** transformations

Transformations / Actions



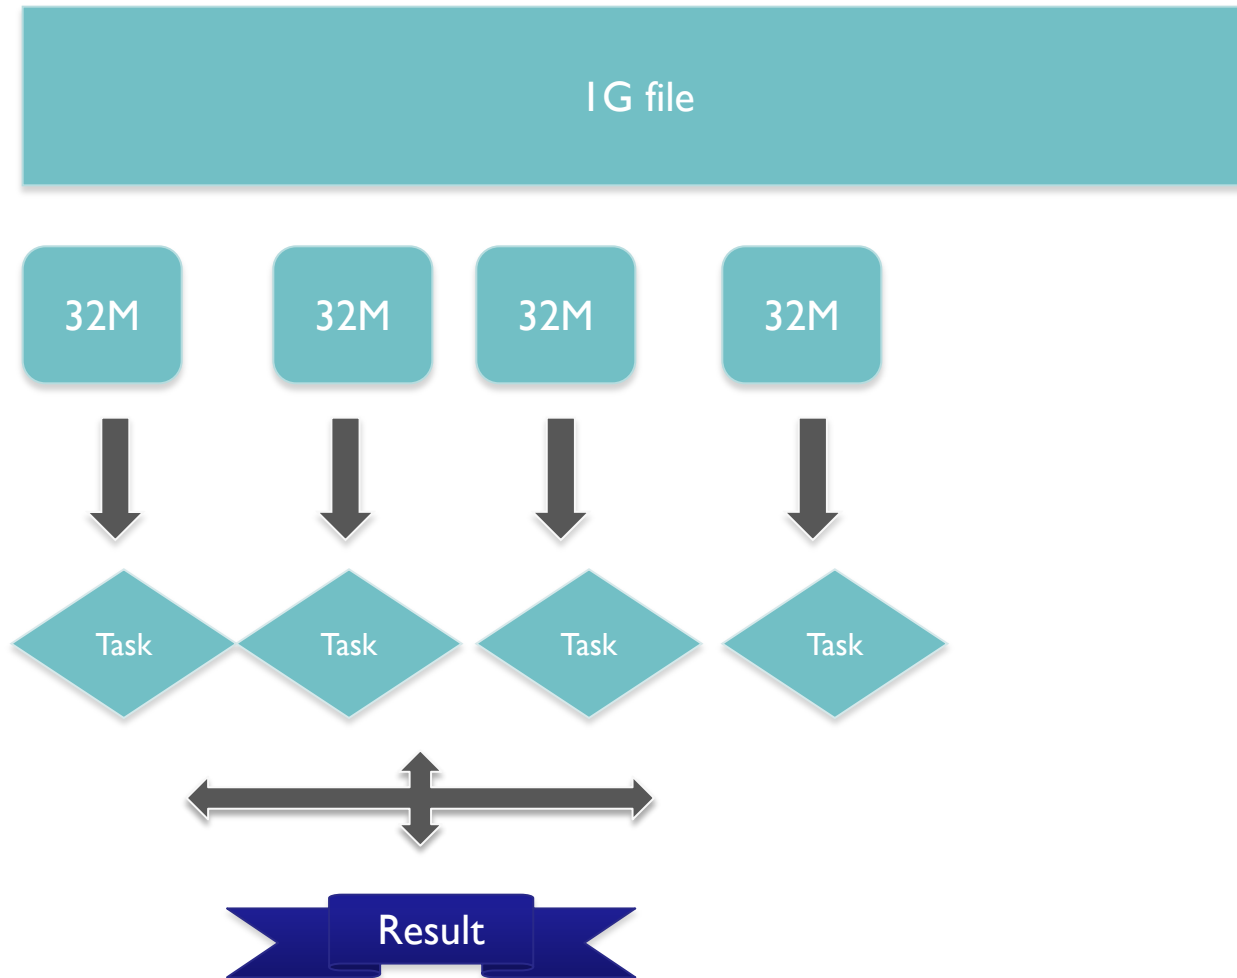
RDD Transformations

| Transformation | Description | Example |
|-----------------|----------------------------------------|-----------------------------------------------------------|
| filter | Filters through each record (aka grep) | <code>f.filter(line => line.contains("ERROR"))</code> |
| union | Merges two RDDs | <code>rdd1.union(rdd2)</code> |
| ...see docs ... | | |

RDD Actions

| Action | Description | Example |
|---------------------------------|------------------------------------------------------------------------------------------------------------------|-------------|
| count() | Counts all records in an rdd | f.count() |
| first() | Extract the first record | f.first () |
| take(n) | Take first N lines | f.take(10) |
| collect() | Gathers all records for RDD. All data has to fit in memory of ONE machine (don't use for big data sets) | f.collect() |
| See documentation .. | | |
| | | |

Partitions Explained

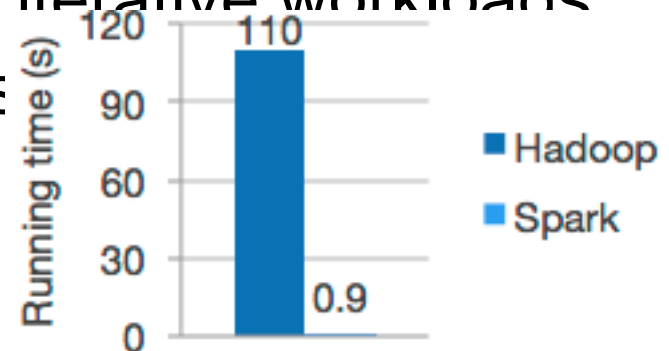


RDD : Saving

- ❑ `saveAsTextFile ()` and `saveAsSequenceFile()`
`f.saveAsTextFile("/output/directory")` // a directory
- ❑ Output usually is a directory
 - ❑ RDDs will be saved as multiple files in the dir
 - ❑ Each partition → one output file

Caching of RDDs

- ❑ RDDs can be loaded from disk and computed
 - ❑ Hadoop mapreduce model
- ❑ Also RDDs can be cached in memory
- ❑ Subsequent operations are much faster
- f.persist() // on disk or memory
- f.cache() // memory only
- ❑ In memory RDDs are great for iterative workloads
 - ❑ Machine learning algorithms



Demo Time !



Demo : Spark-shell

- ❑ Invoke spark shell
- ❑ Load a data set
- ❑ Do basic operations (count / filter)

Demo: RDD Caching

- ❑ From Spark shell
- ❑ Load an RDD
- ❑ Demonstrate the difference between cached and non-cached

Demo : Map Reduce

□ Quick Word count

```
val input = sc.textFile("...")  
  val counts = input.flatMap(  
    line => line.split(" ")).  
    map(word => (word, 1)).  
    reduceByKey(_+_)
```

Thanks !

Sujee Maniyam

sujee@elephantscale.com

<http://elephantscale.com>

Expert consulting & training in Big Data



Credits

- ❑ <http://spark.apache.org/>
- ❑ <http://www.strategictechplanning.com>