outline

- ***basis & internals***
- ecosystem

火龙果•整理
uml.org.cn

Current Major Release

- **Just released Spark 1.1**

火龙果•整理
uml.org.cn

Spark : What & Why

- Apache Spark is a fast and general engine for large-scale data processing.
- Speed
- Ease of Use
- Generality
- Integrated with Hadoop

火龙果•整理
uml.org.cn

BDAS

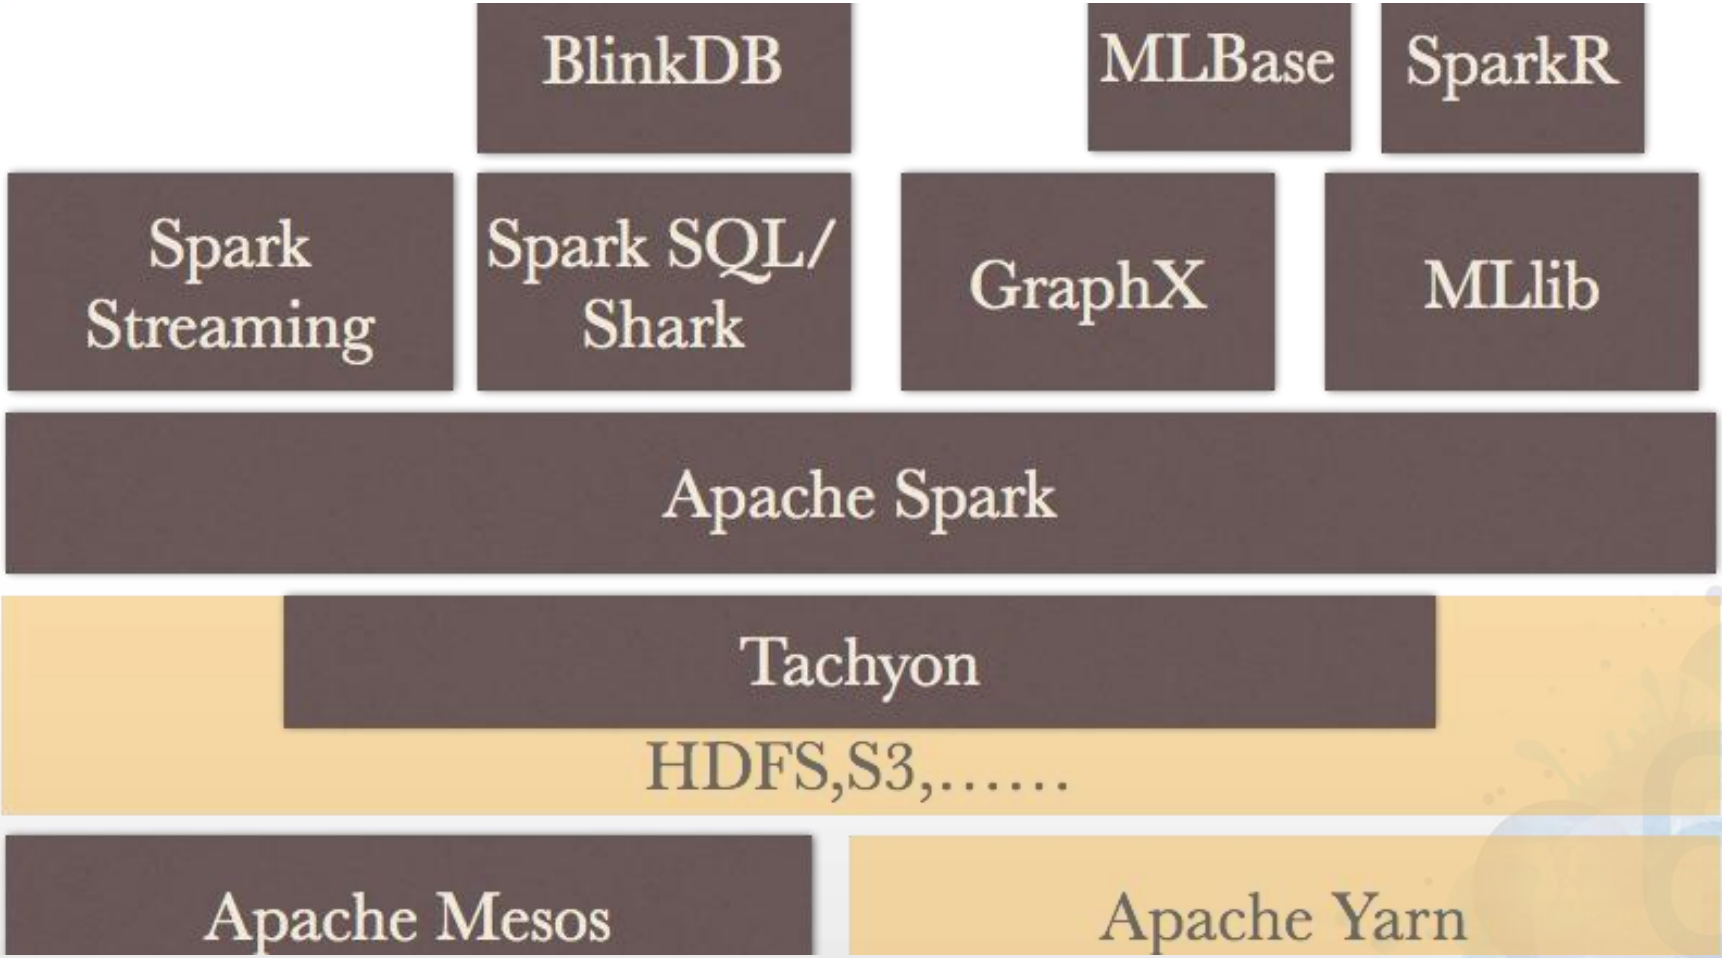

火龙果•整理
uml.org.cn

one stack to rule them all

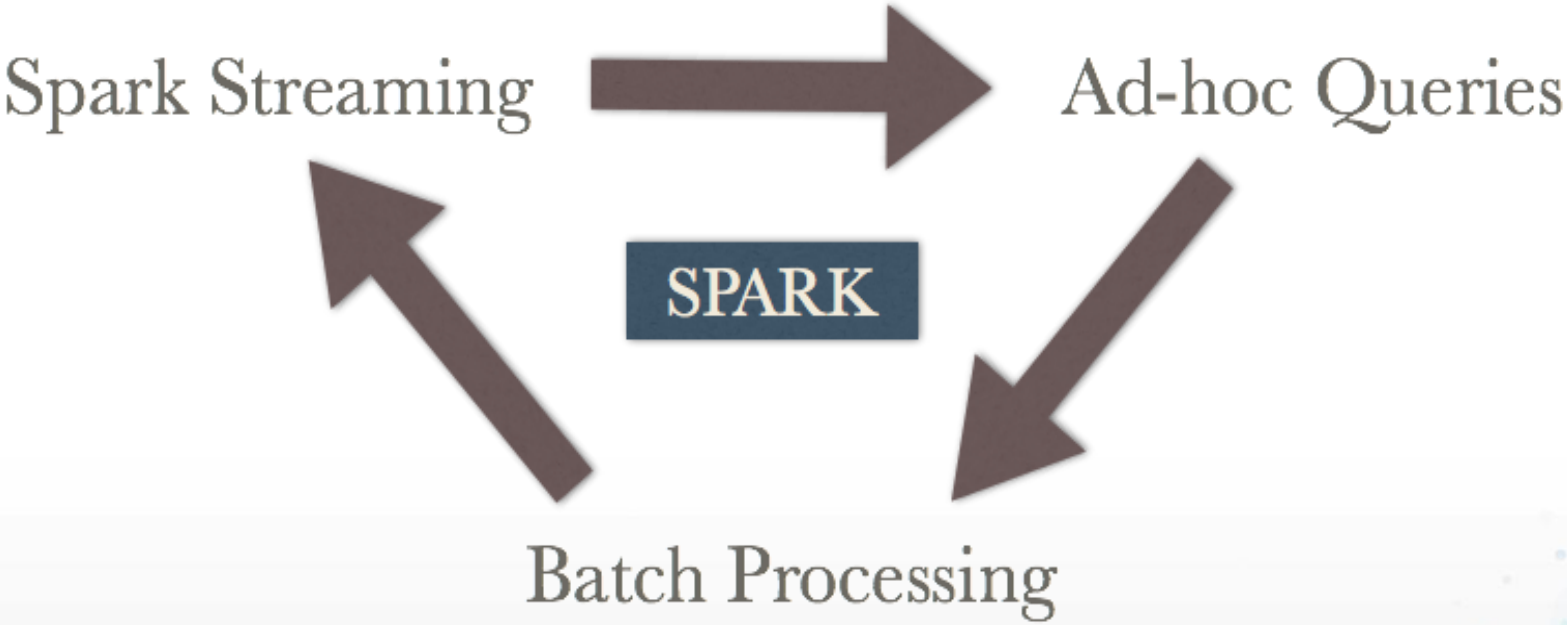

火龙果•整理
uml.org.cn

Key Concept-RDD

- A list of partitions
- A function for computing each split
- A list of dependencies on other RDDs
- Optionally, a Partitioner for key-value RDDs
- Optionally, a list of preferred locations to compute each split on

火龙果•整理
uml.org.cn

Key Concept-Lineage

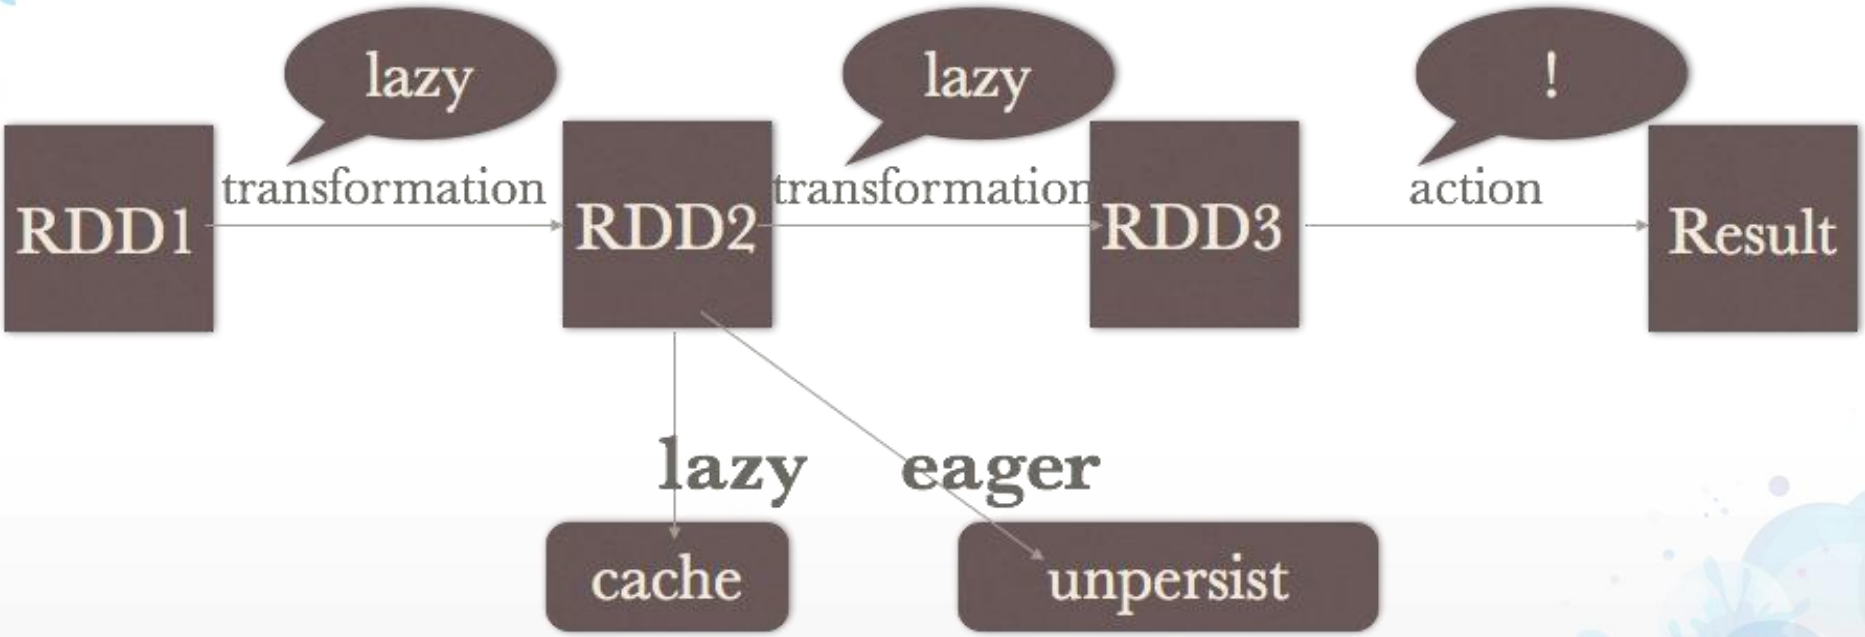

unroll partition safely when caching

火龙果•整理
uml.org.cn

Key Concept-Dependency

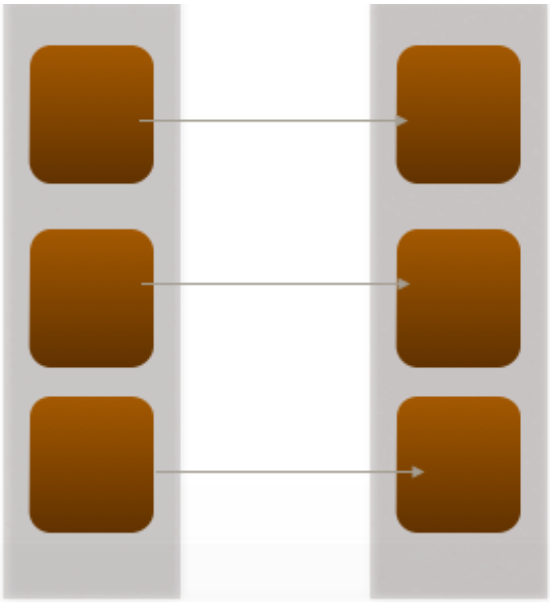

Narrow Dependency

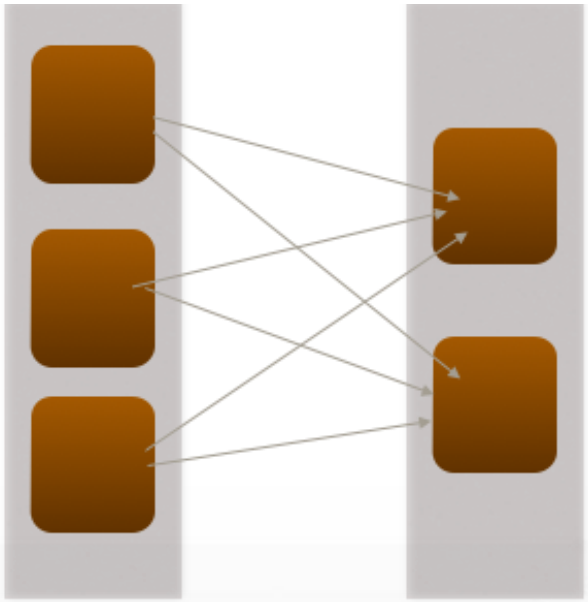

Wide Dependency

火龙果•整理
uml.org.cn

Key Concept-ClusterManager

- Local
- Standalone
- Yarn
- Mesos

火龙果•整理
uml.org.cn

Cluster Overview

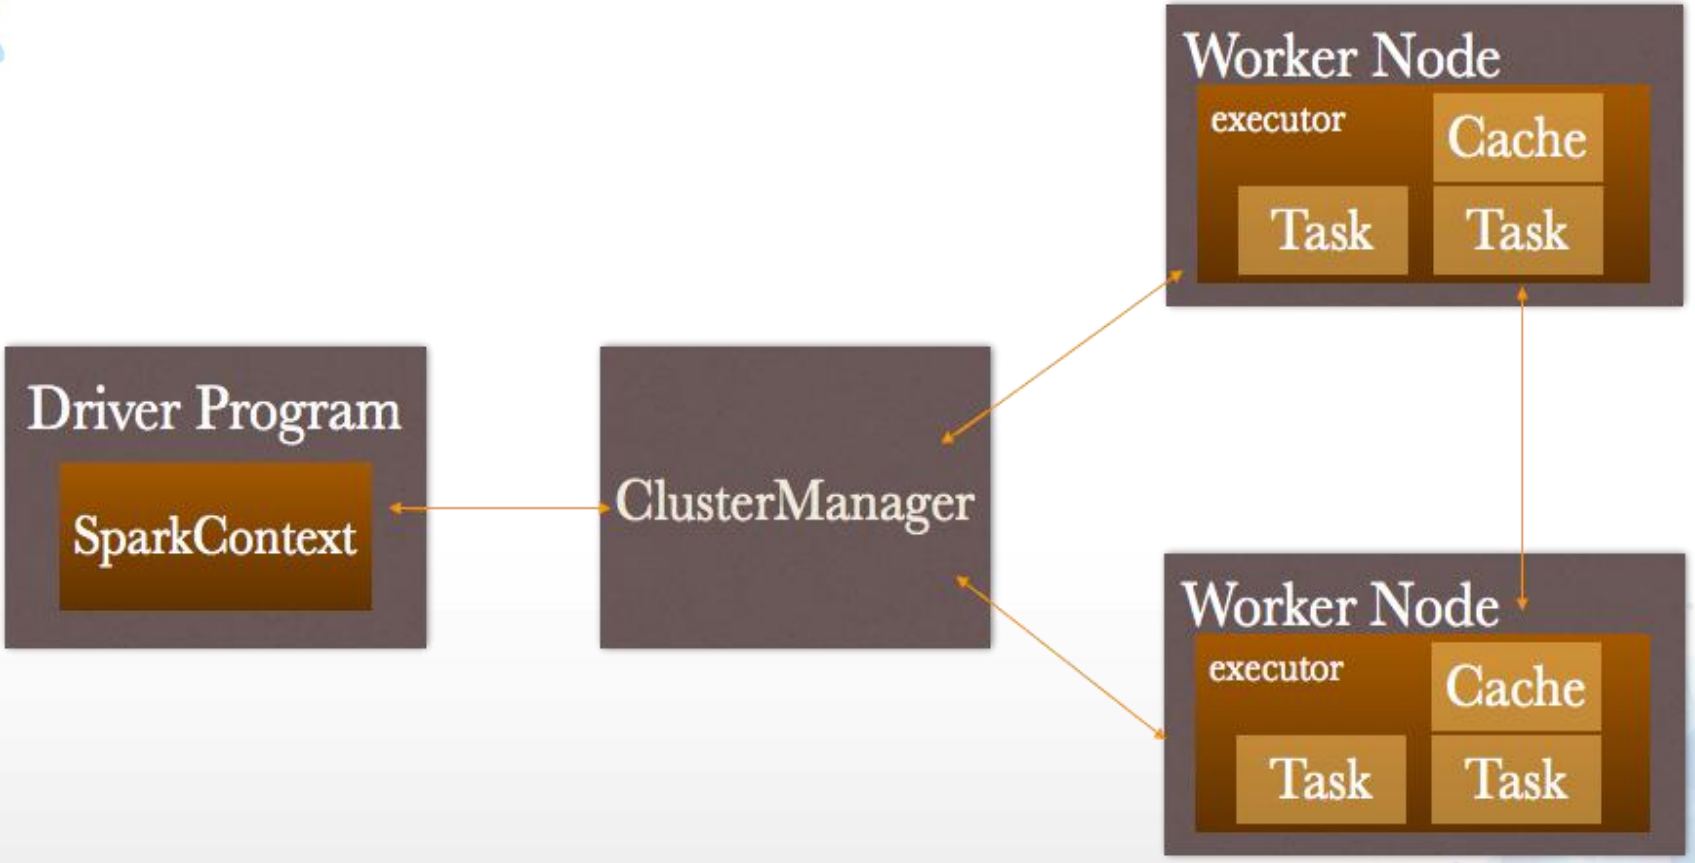

火龙果•整理
uml.org.cn

schedule

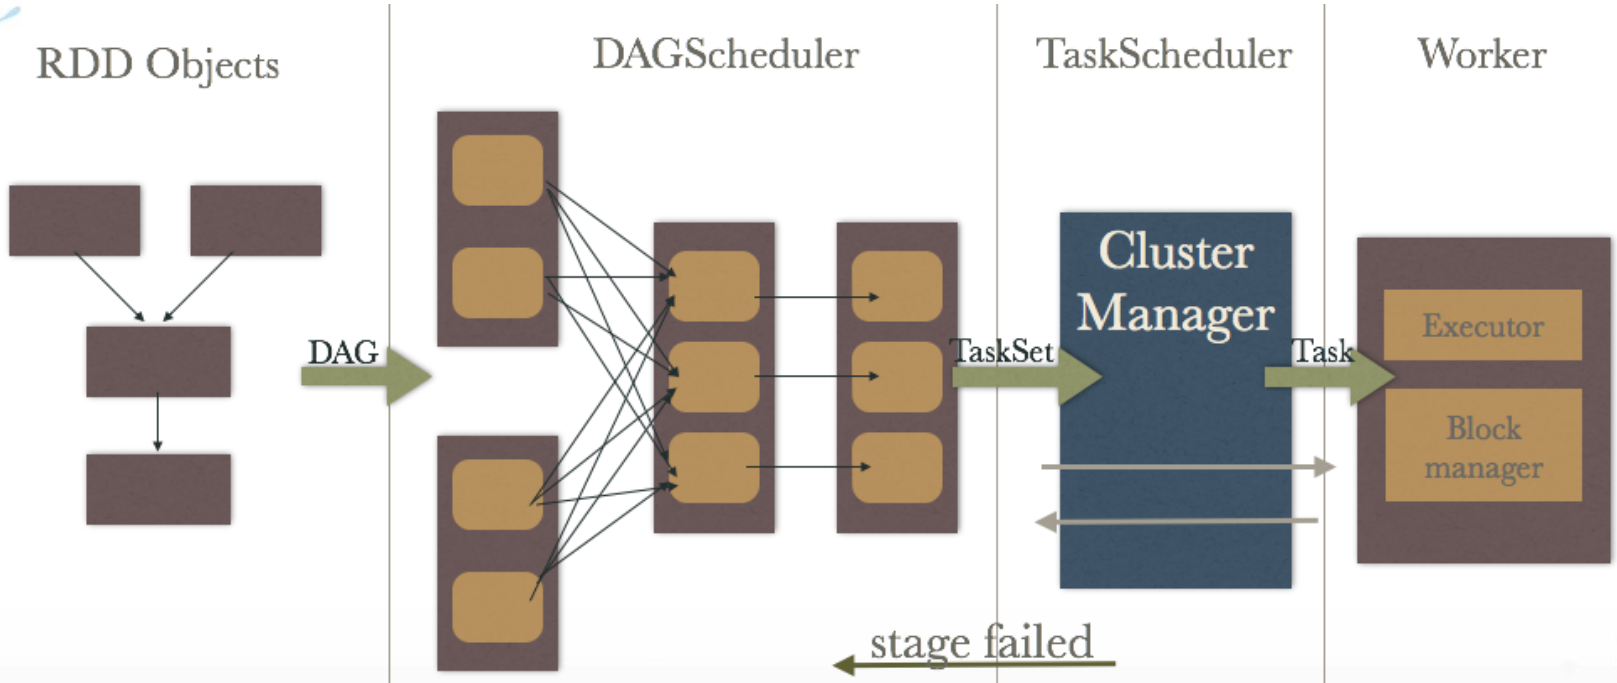

Build operator DAG

split graph into stages of tasks

submit each stage as ready

launch tasks via cluster manager

retry failed or straggling tasks

execute tasks

store and serve blocks

火龙果•整理
uml.org.cn

Executor

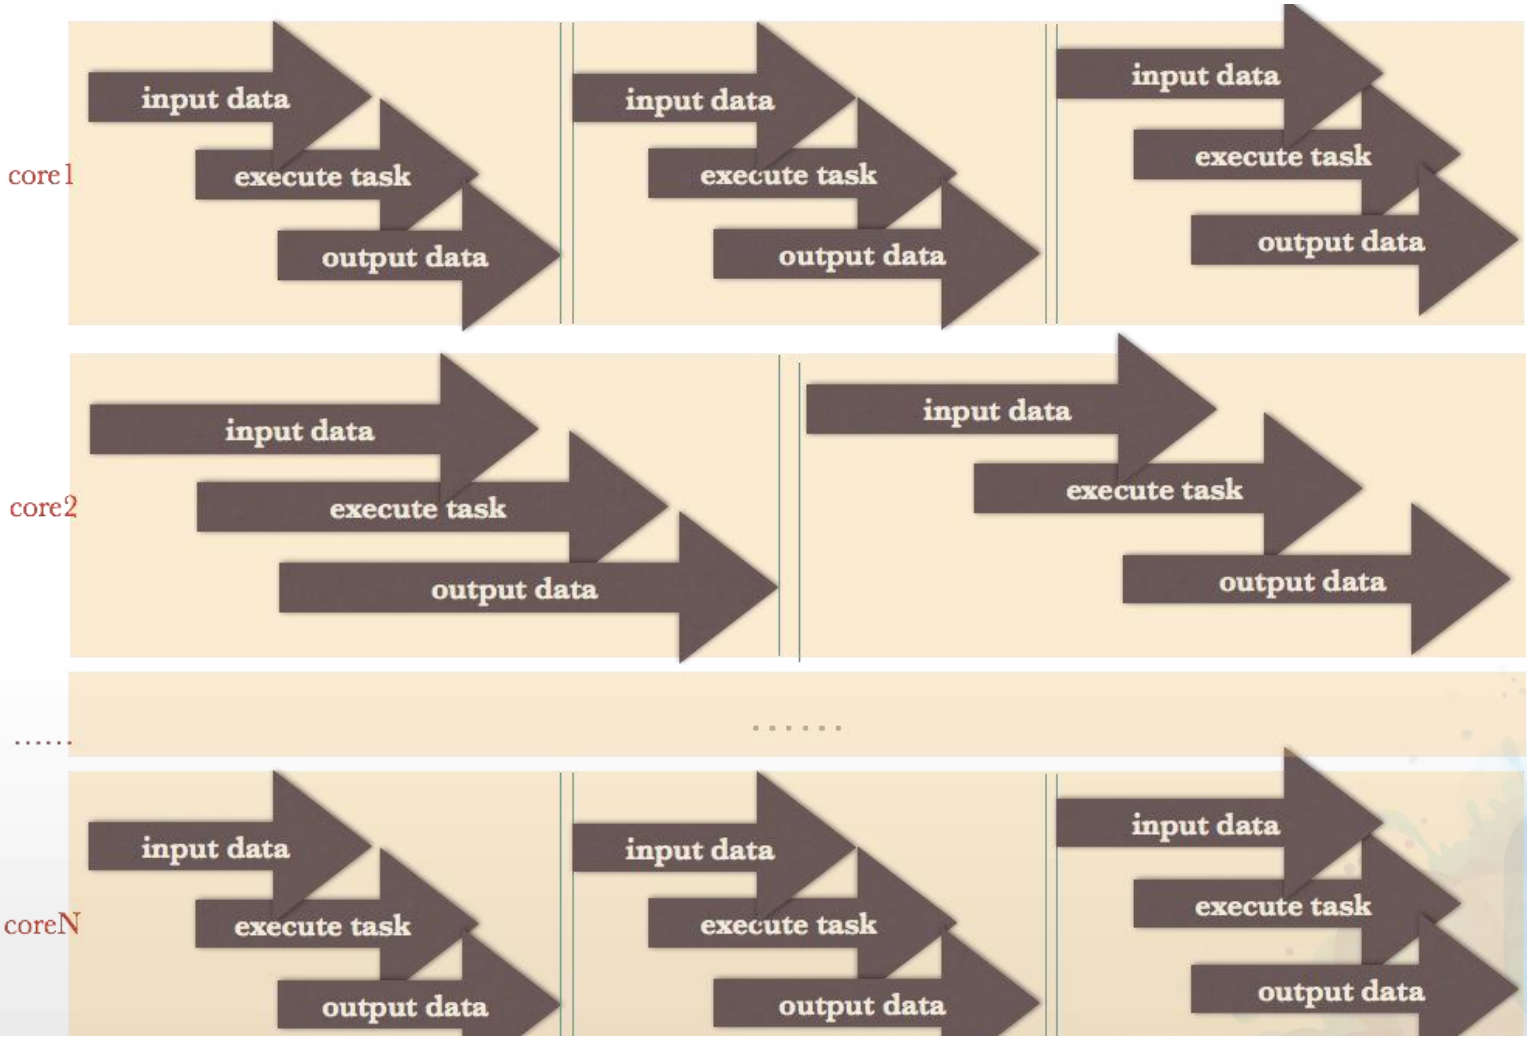

火龙果•整理
uml.org.cn

Shuffle

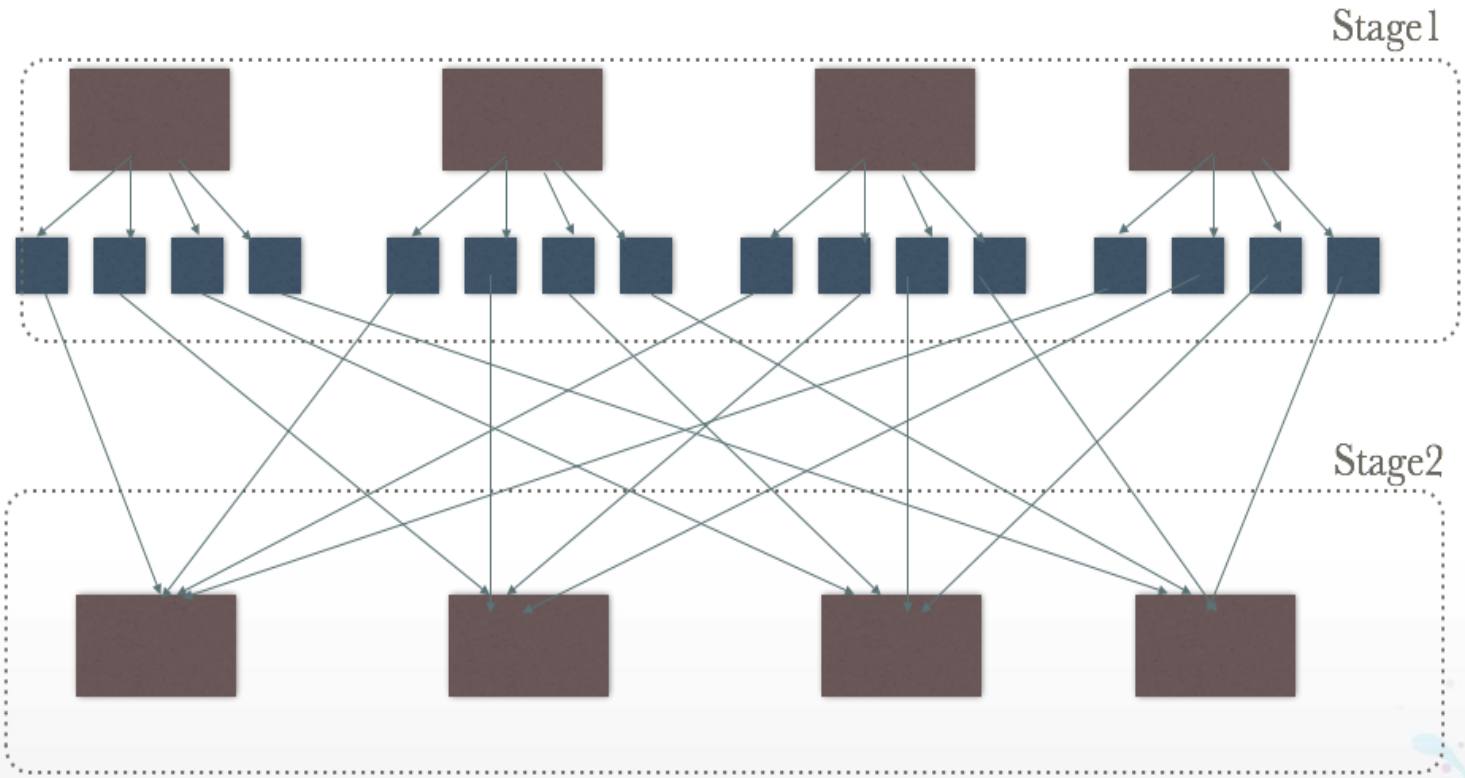

Sort-based shuffle supported

火龙果•整理
uml.org.cn

Shuffle

- Pull-based (not push-based)
- Write intermediate files to disk
- Build hash map within each partition
- Can spill across keys
- A single key-value pair must fit in memory

火龙果•整理
uml.org.cn

Better Metrics System

- Previously: only collect after task completed
- Now : report when task is still running

火龙果•整理
uml.org.cn

outline

- basis & internals
- ***ecosystem***

火龙果•整理
uml.org.cn

Spark Streaming

- Mini-batch

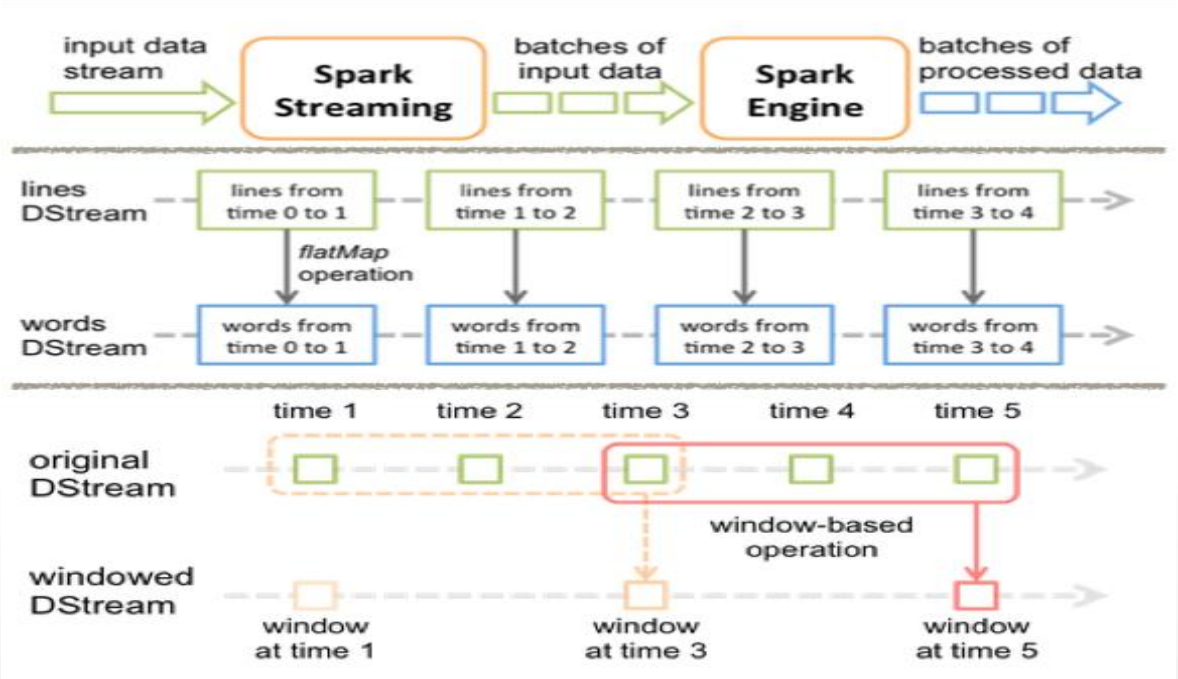

rate limiting supported!

Streaming + MLlib

火龙果•整理
uml.org.cn

MLlib

- Spark implementation of some common machine learning algorithms and utilities
- classification
- regression
- clustering
- collaborative filtering
- dimensionality reduction

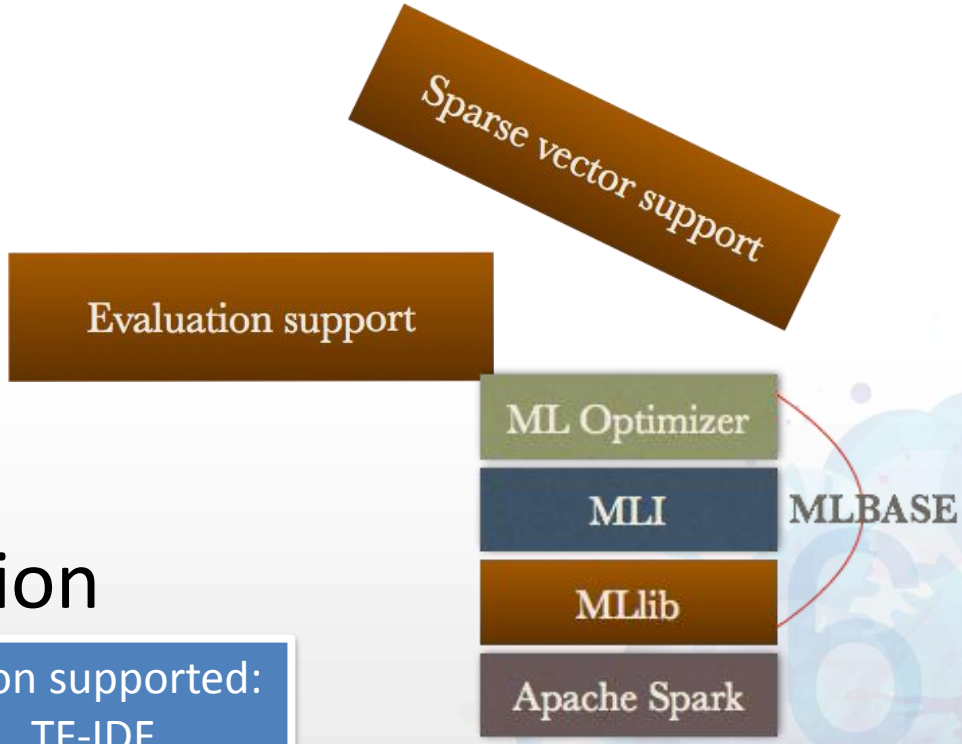

feature extraction supported:
Word2Vec , TF-IDF

火龙果•整理
uml.org.cn

GraphX

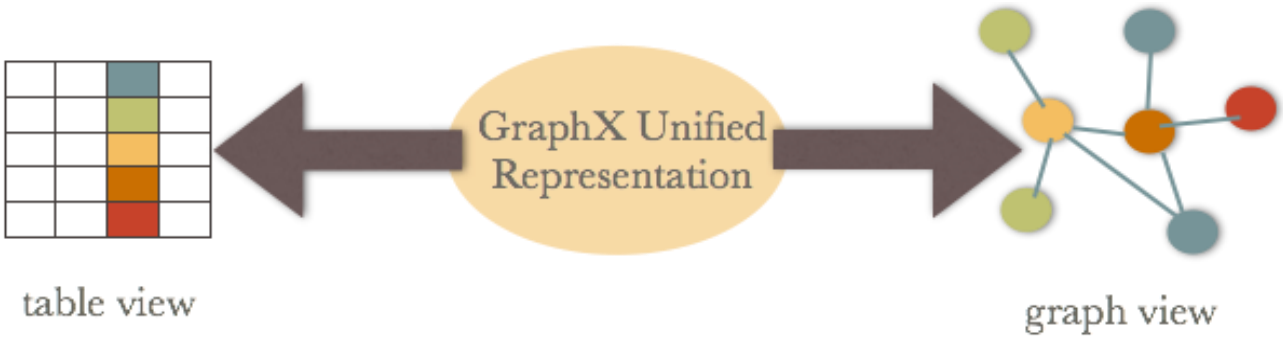

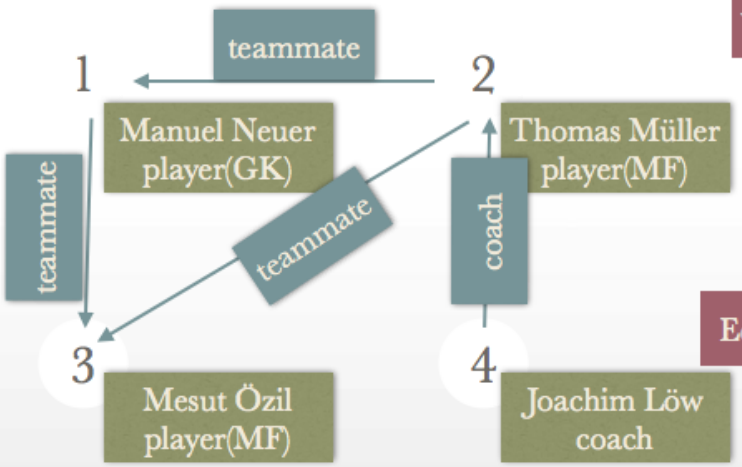

Vertex Table

id	
1	(Manuel Neuer,player)
2	(Thomas Müller,player)
3	(Mesut Özil,player)
4	(Joachim Löw,coach)

Edge Table

SrcId	DstId	Property(E)
2	1	teammate
2	3	teammate
1	3	teammate
4	2	coach

火龙果•整理
uml.org.cn

GraphX

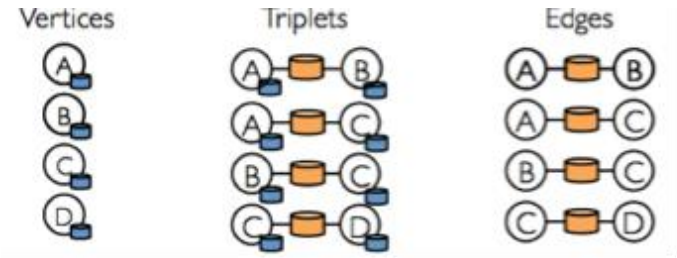

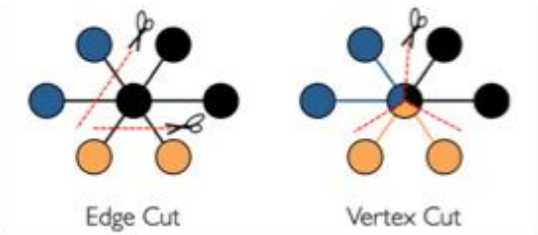

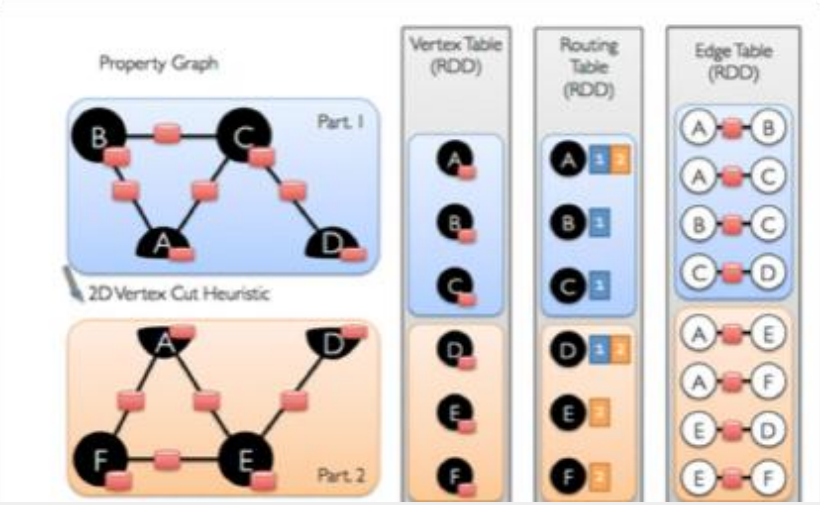

火龙果•整理
uml.org.cn

Spark SQL

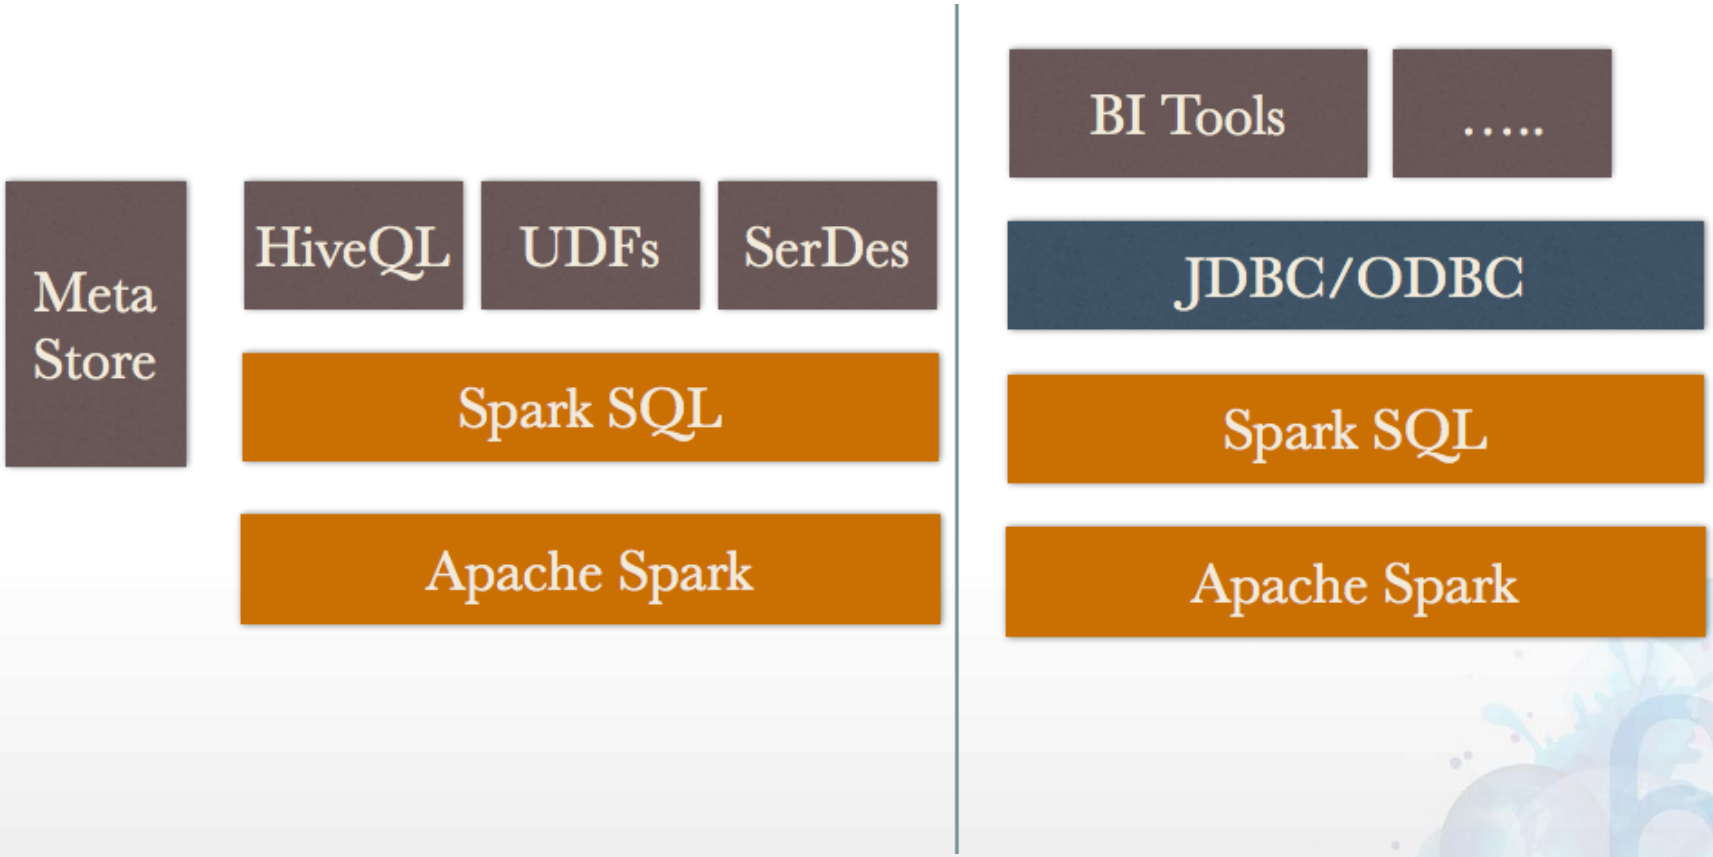

火龙果•整理
uml.org.cn

Spark SQL

- Data Sources
- RDDs/Parquet Files/JSON Datasets/Hive Table
- DSL
- JDBC Server

Programmatically Specifying the Schema

火龙果•整理
uml.org.cn

Shark

- Development in Shark has been ended and subsumed by Spark SQL
- Mission completed !!!

火龙果•整理
uml.org.cn

Tachyon

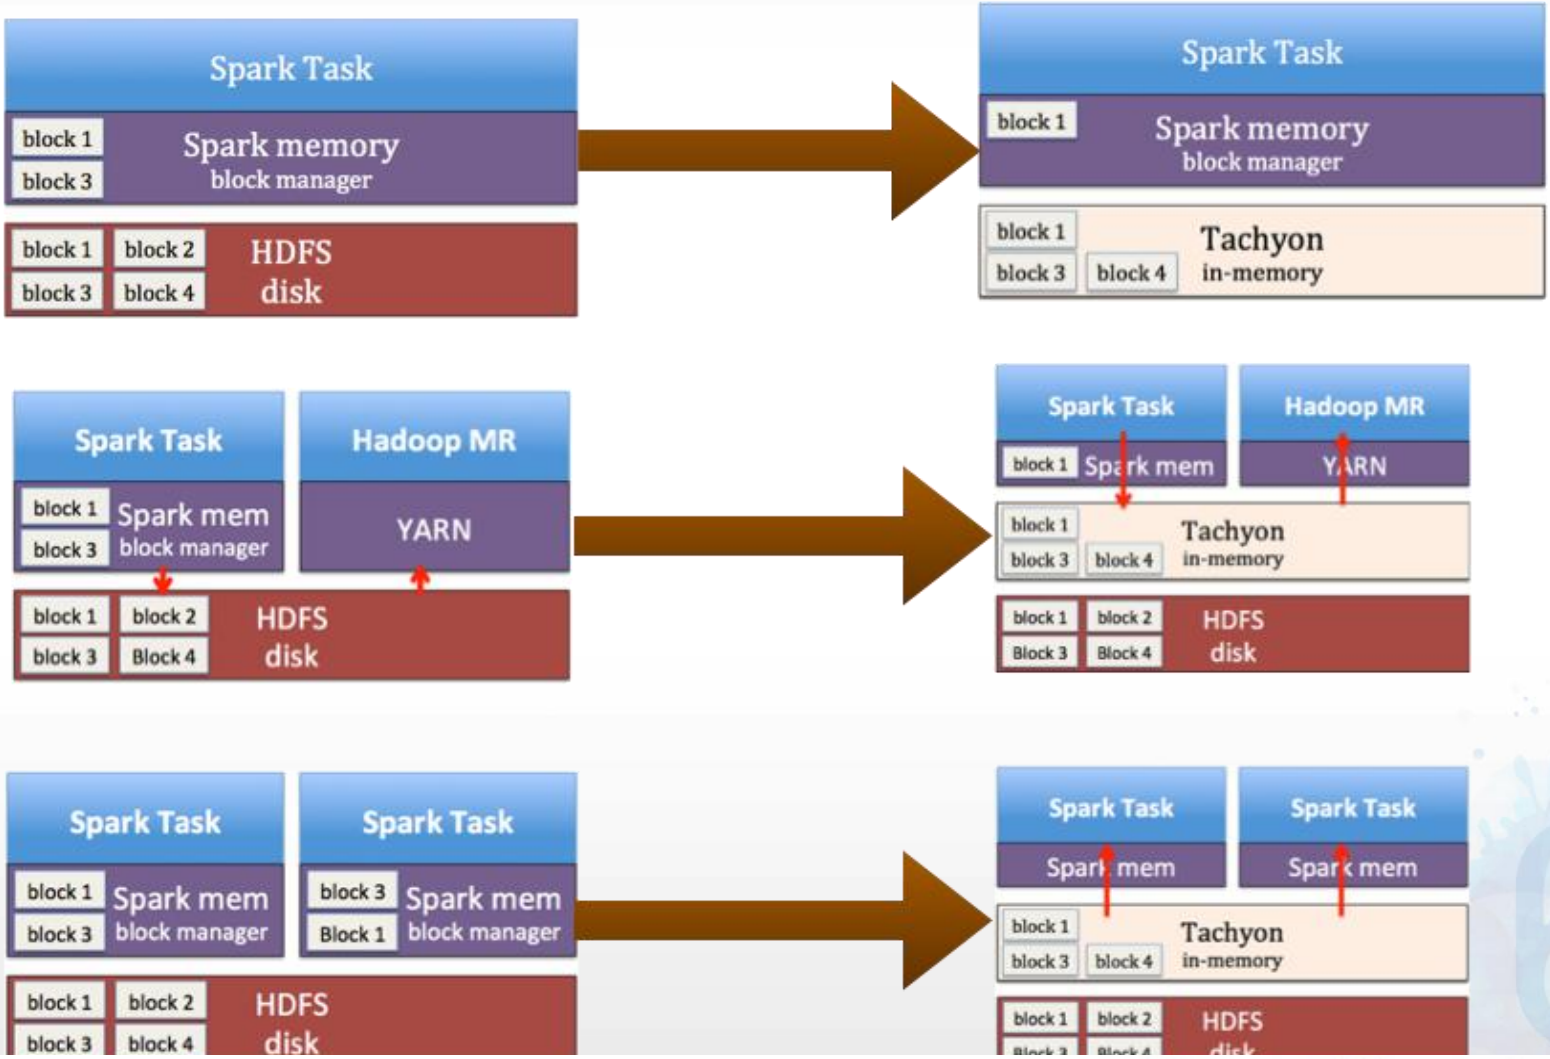

火龙果•整理
uml.org.cn

Tachyon

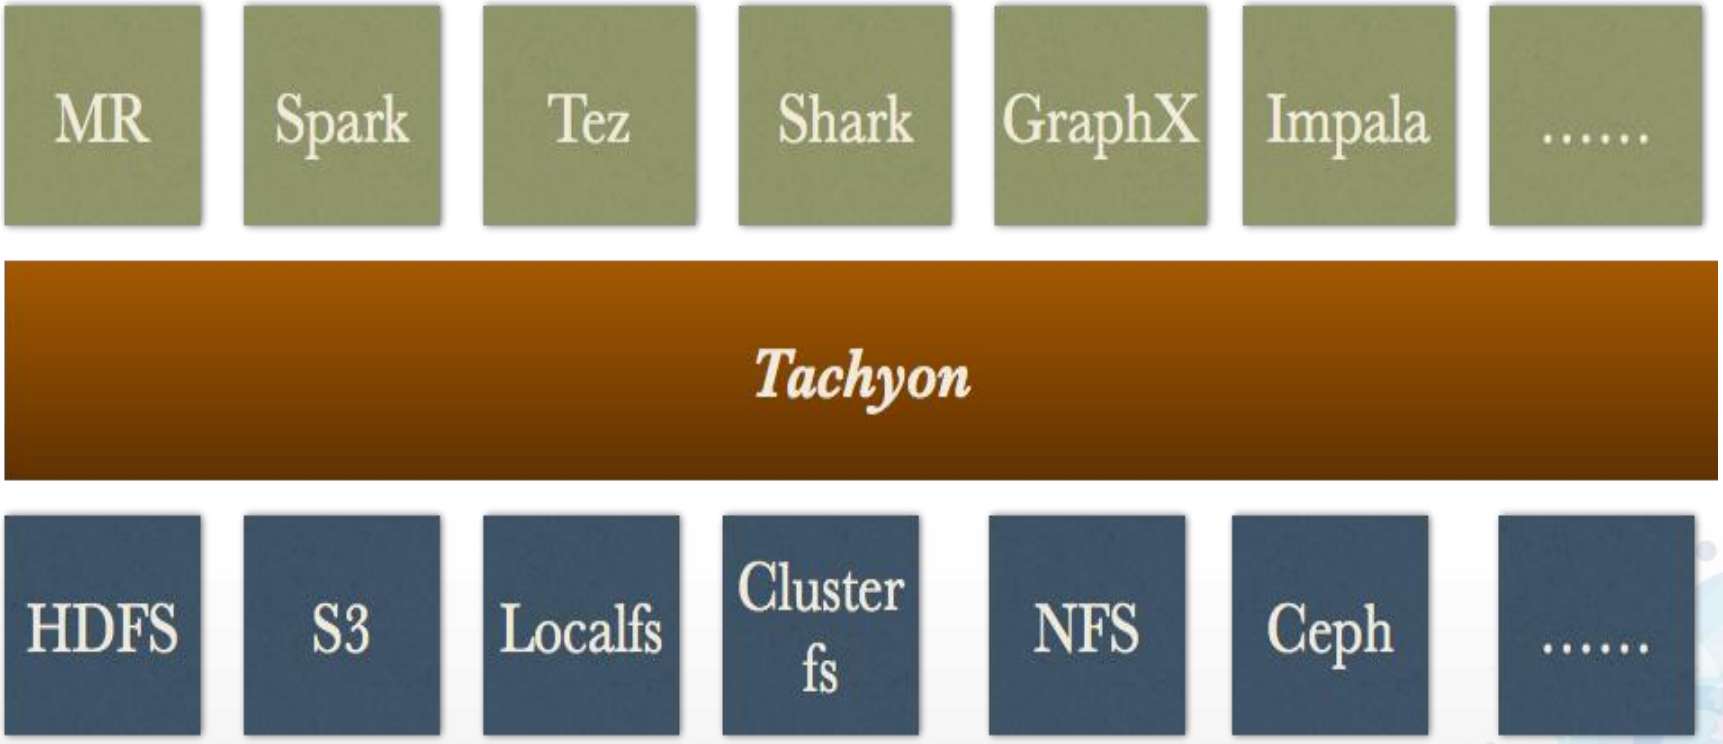

火龙果•整理
uml.org.cn

SparkR

R + RDD = RRDD

RDDs as Distributed Lists

```
sc <- sparkR.init("local")
lines <- textFile(sc, "hdfs://data.txt")
wordsPerLine <- lapply(lines, function(line) { length(unlist(strsplit(line, " "))) })
```

火龙果•整理
uml.org.cn

BlinkDB

- Queries with Bounded Errors and Bounded Response Times on Very Large Data

```
SELECT avg(sessionTime)
FROM Table
WHERE city='San Francisco'
WITHIN 2 SECONDS
```

Queries with Time Bounds

```
SELECT avg(sessionTime)
FROM Table
WHERE city='San Francisco'
ERROR 0.1 CONFIDENCE 95.0%
```

Queries with Error Bounds

火龙果•整理
uml.org.cn