دیدنش با چشم چون ممکن نبود - اندر آن تاریکیش کف می‌بسود

مولانا

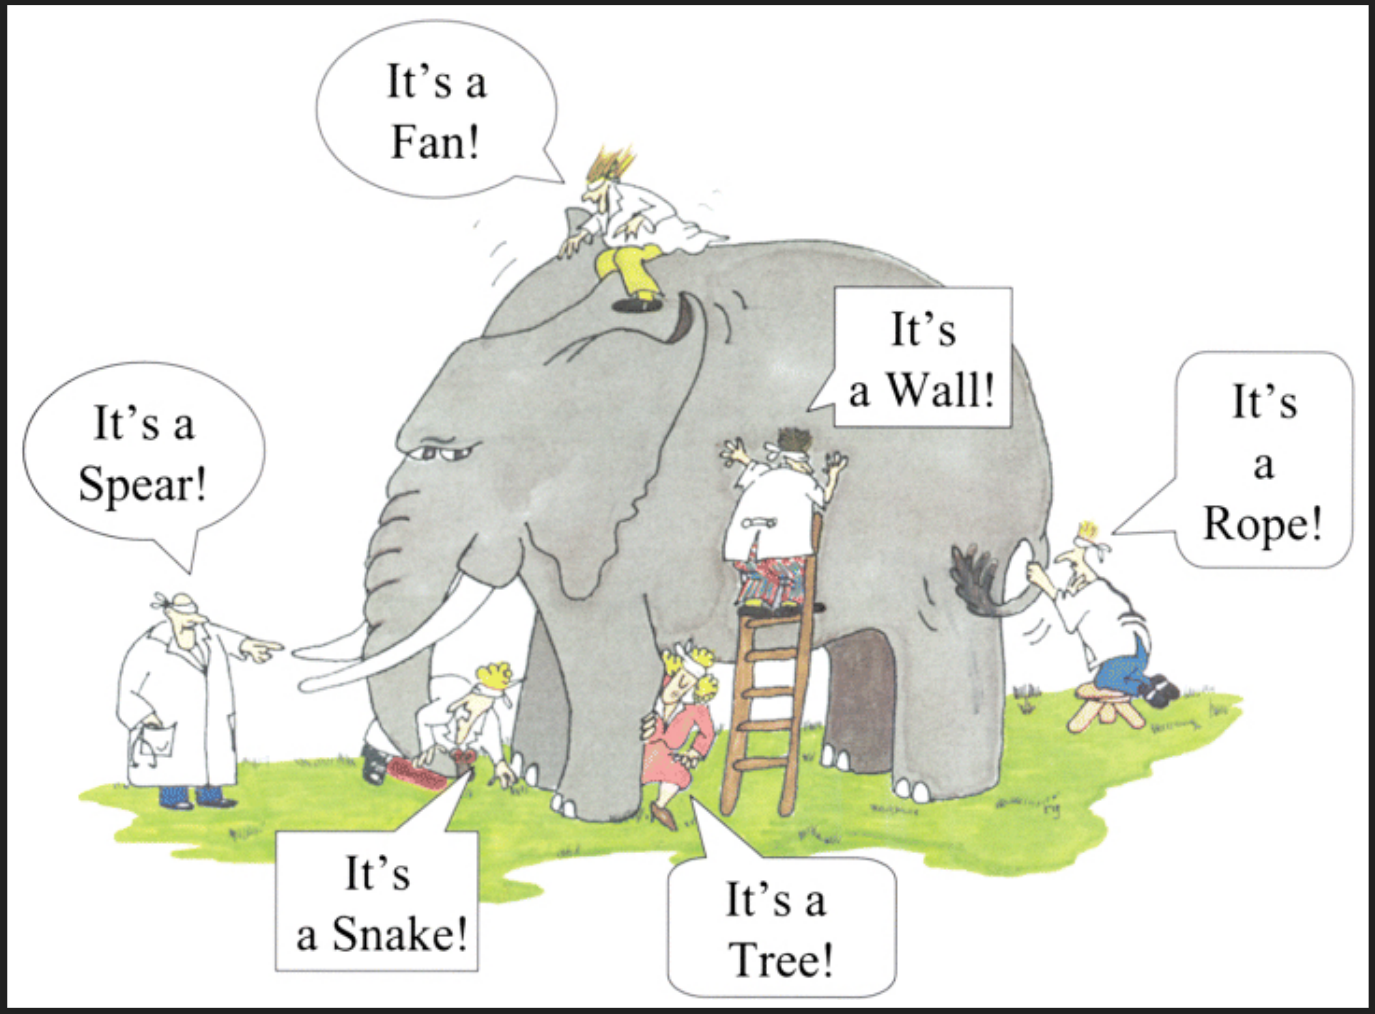
It's a Fan!
It's a Spear!
It's a Wall!
It's a Rope!
It's a Snake!
It's a Tree!

INTRODUCTION TO HADOOP

Big Data, Hadoop & more

Milad As (Ravexina)

milad (𝔸) tuxgeek (dot) ir

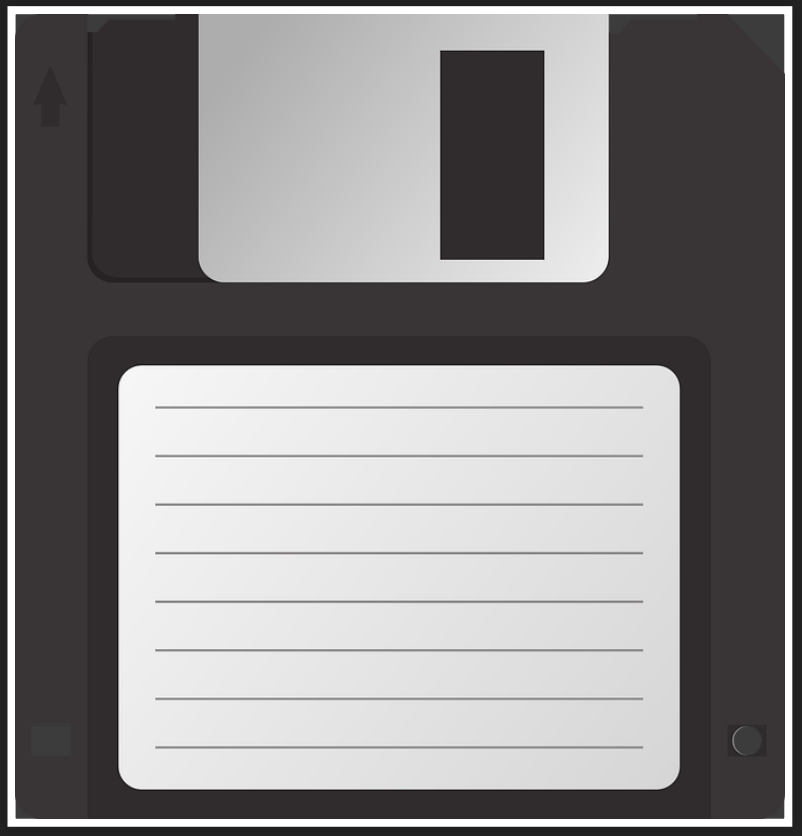

SAVE BUTTON?

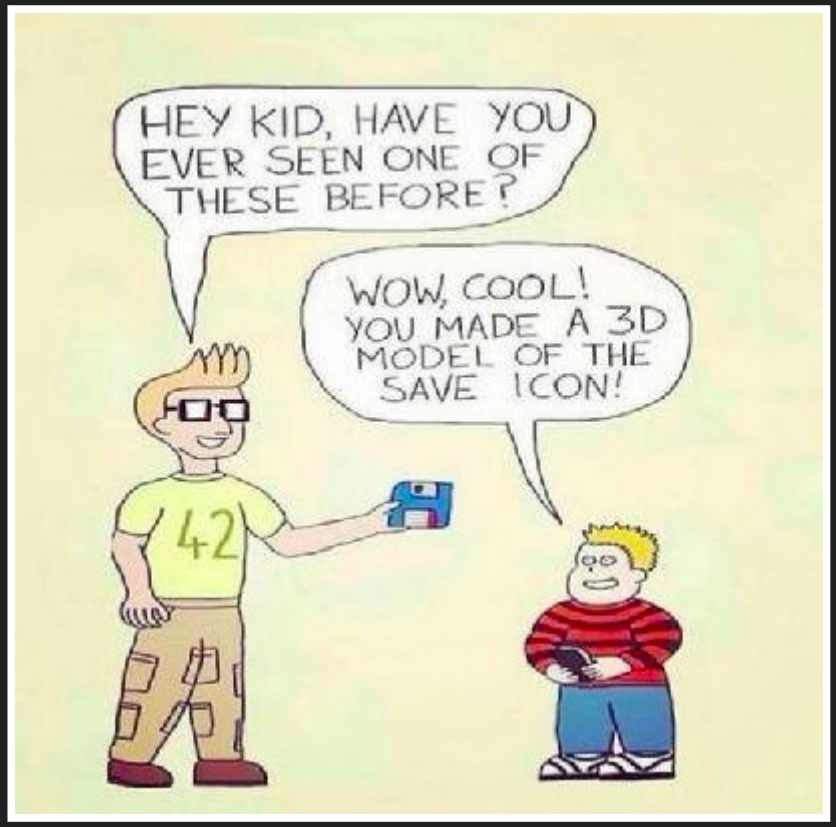
HEY KID, HAVE YOU EVER SEEN ONE OF THESE BEFORE?
WOW, COOL! YOU MADE A 3D MODEL OF THE SAVE ICON!
42

FLOPPY DISK

BUT THINGS CHANGED!

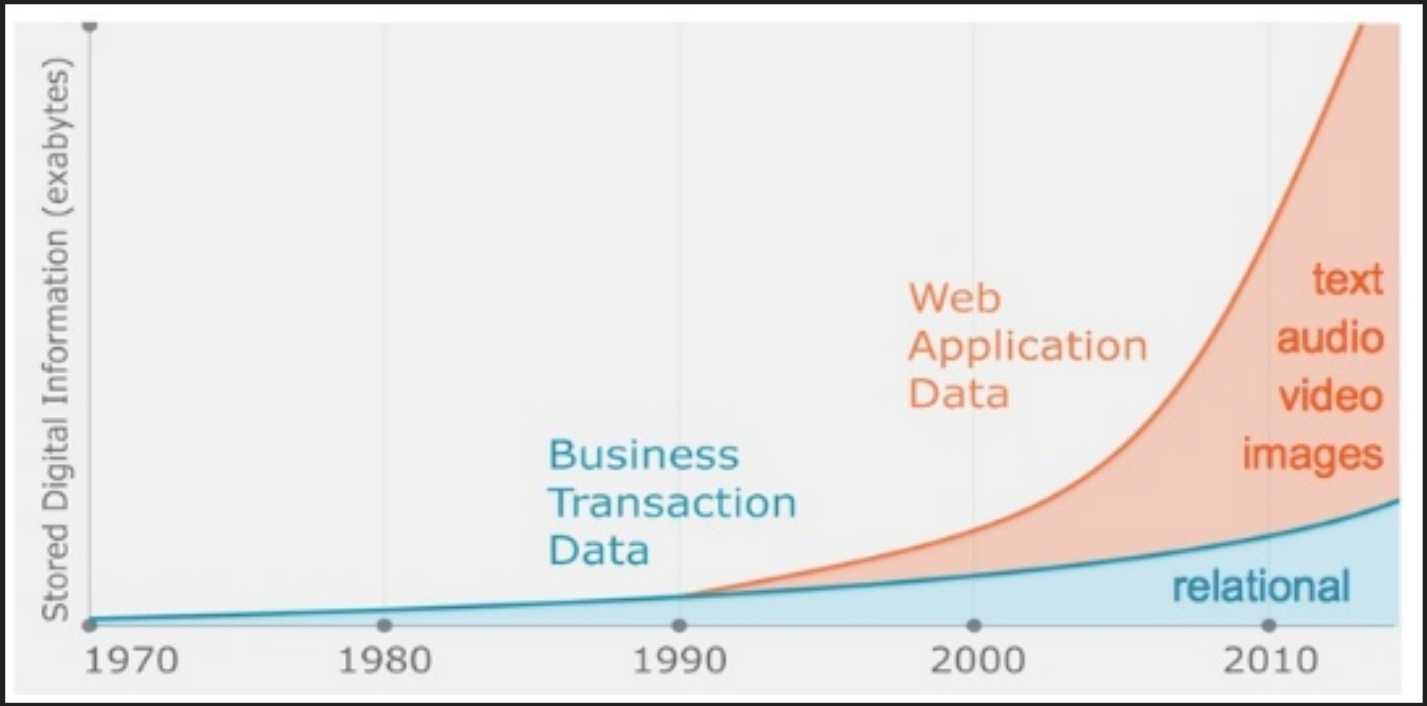

INCREASE IN

Sources (Users, Clients), Forms

FLOOD IS COMING

MySQL, MariaDB, Access, SQL Server, Oracle

Time out on expensive servers...

OVERLOAD

You must change your technologies

WHAT MAKES BIG DATA

Implicit and Explicit

Forums, Messengers, Comments, Q/A, Reviews

Share ...!?

Did you just pause that?

Share, view, follow

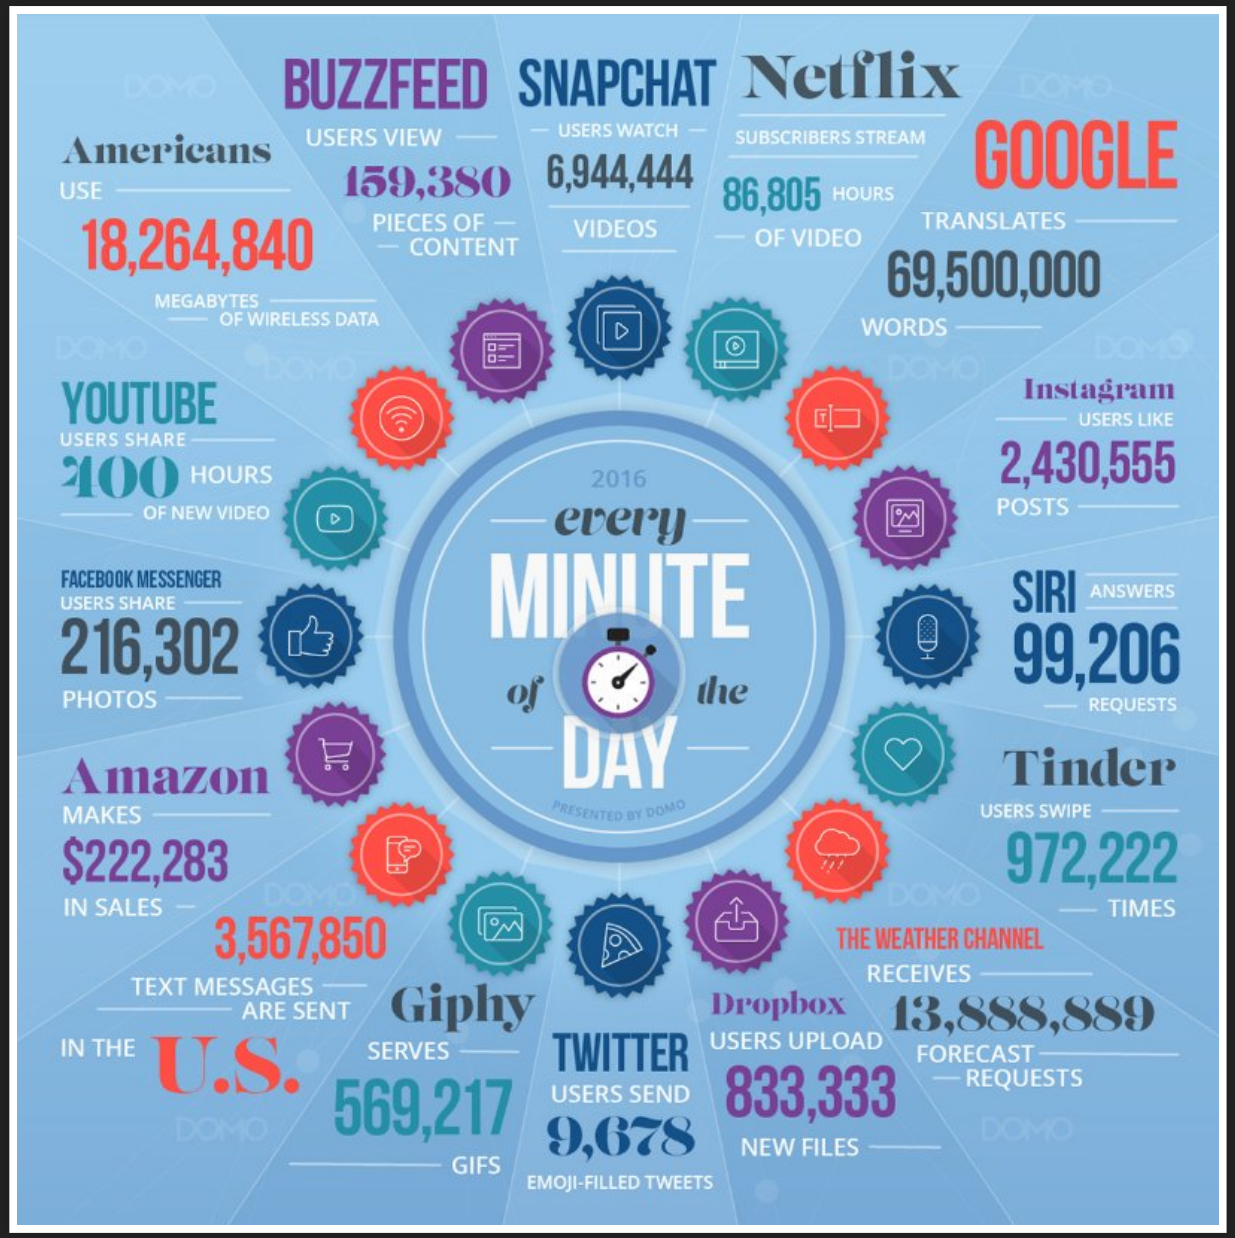
2016
every
MINUTE
of the
DAY
PRESENTED BY DOMO
BUZZFEED
USERS VIEW
159,380
PIECES OF CONTENT
SNAPCHAT
USERS WATCH
6,944,444
VIDEOS
Netflix
SUBSCRIBERS STREAM
86,805 HOURS
OF VIDEO
GOOGLE
TRANSLATES
69,500,000
WORDS
Americans
USE
18,264,840
MEGABYTES
OF WIRELESS DATA
YOUTUBE
USERS SHARE
100 HOURS
OF NEW VIDEO
Instagram
USERS LIKE
2,430,555
POSTS
FACEBOOK MESSENGER
USERS SHARE
216,302
PHOTOS
SIRI ANSWERS
99,206
REQUESTS
Amazon
MAKES
$222,283
IN SALES
Tinder
USERS SWIPE
972,222
TIMES
3,567,850
TEXT MESSAGES
ARE SENT
IN THE U.S.
Giphy
SERVES
569,217
GIFS
TWITTER
USERS SEND
9,678
EMOJI-FILLED TWEETS
Dropbox
USERS UPLOAD
833,333
NEW FILES
THE WEATHER CHANNEL
RECEIVES
13,888,889
FORECAST
REQUESTS

THESE ARE ONLY THE VISIBLE PARTS OF BIGDATA

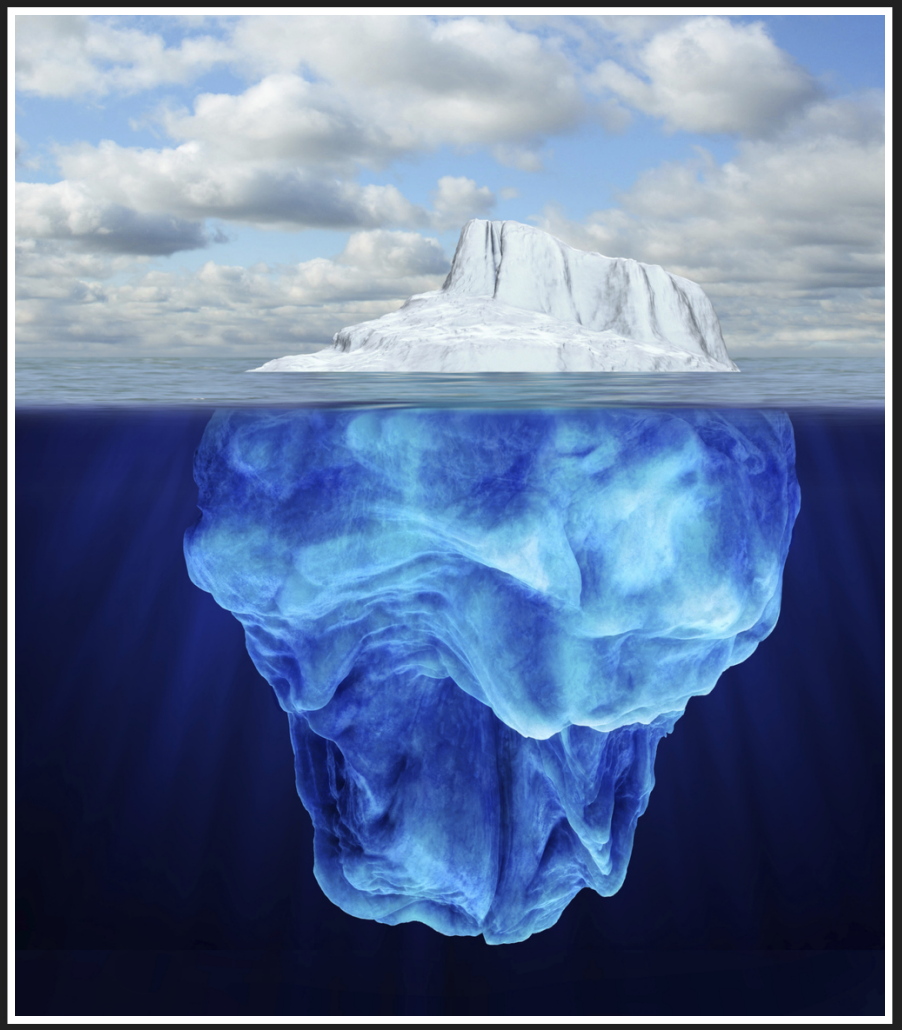

Applications

Applications

Servers (Web, Email, Proxy, ...)

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Hardwares Access point

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Hardwares Access point

Life

LOGS

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Hardwares Access point

Life

LOGS

Clicks

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Hardwares Access point

Life

LOGS

Clicks

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Hardwares Access point

Life

LOGS

Clicks

Applications

Servers (Web, Email, Proxy, ...)

Systems journalct

Hardwares Access point

Life

MAIN SOURCE

Mysterious and valuable part!

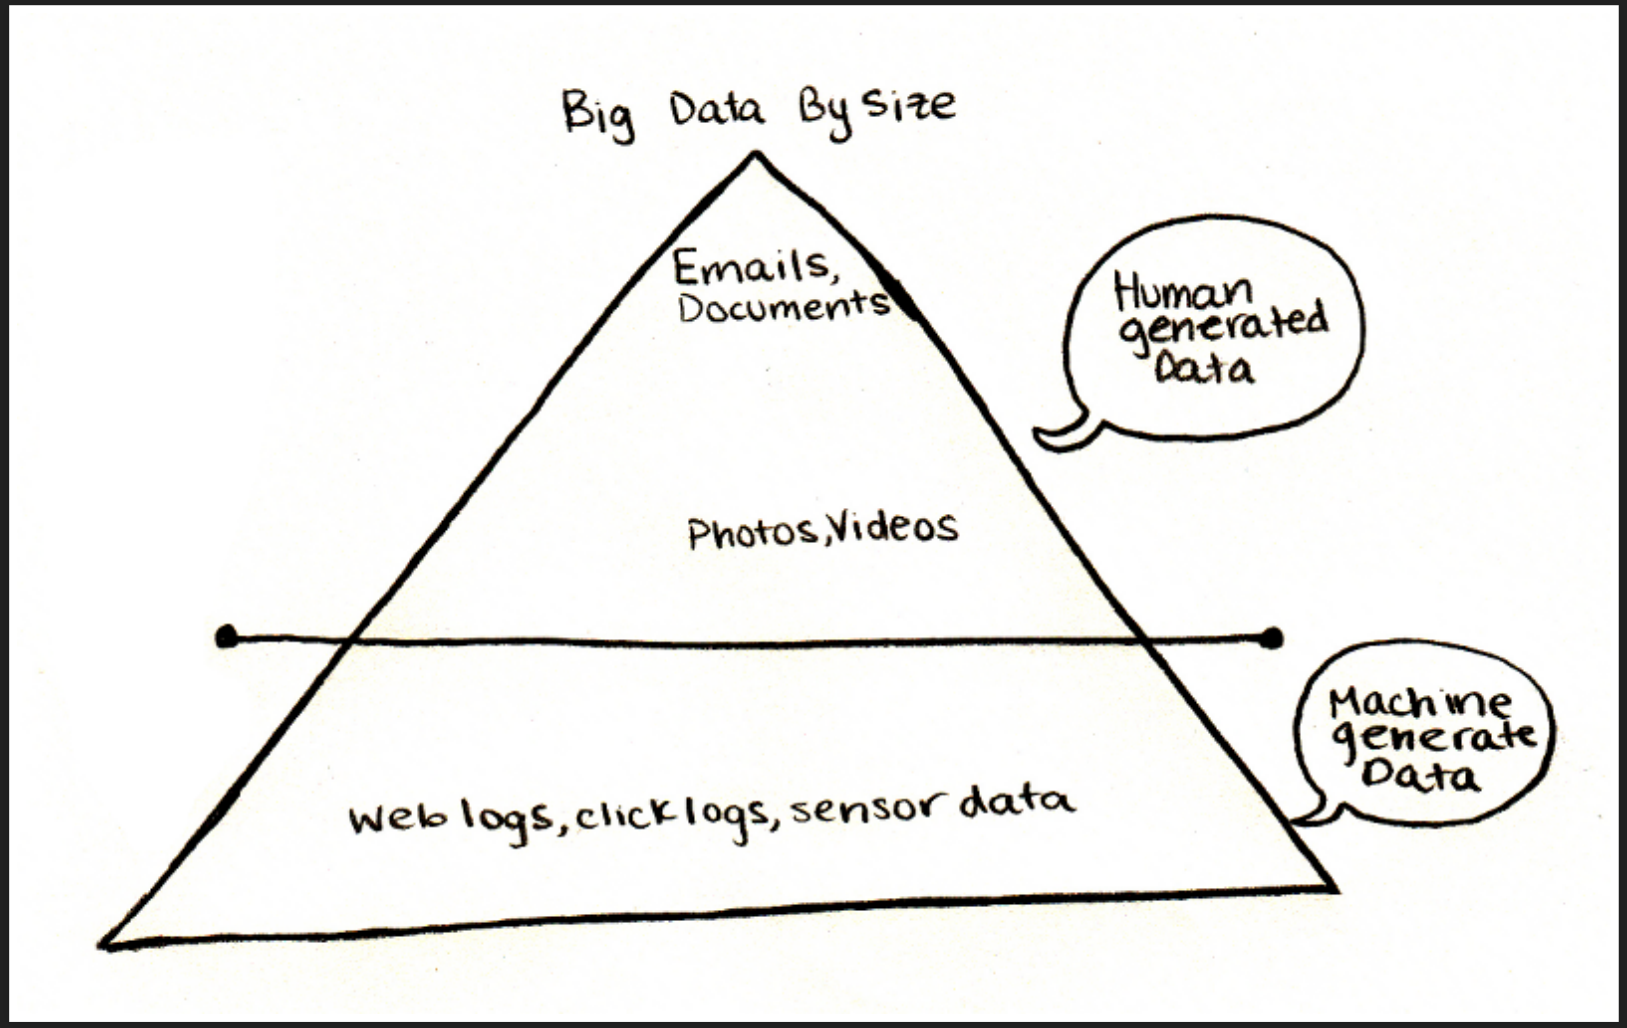

DIGITAL FOOTPRINT

A simple tweet

HOW CAN WE PROCESS THESE DATA?

SCALABILITY

Low resistance but inexpensive and infinite

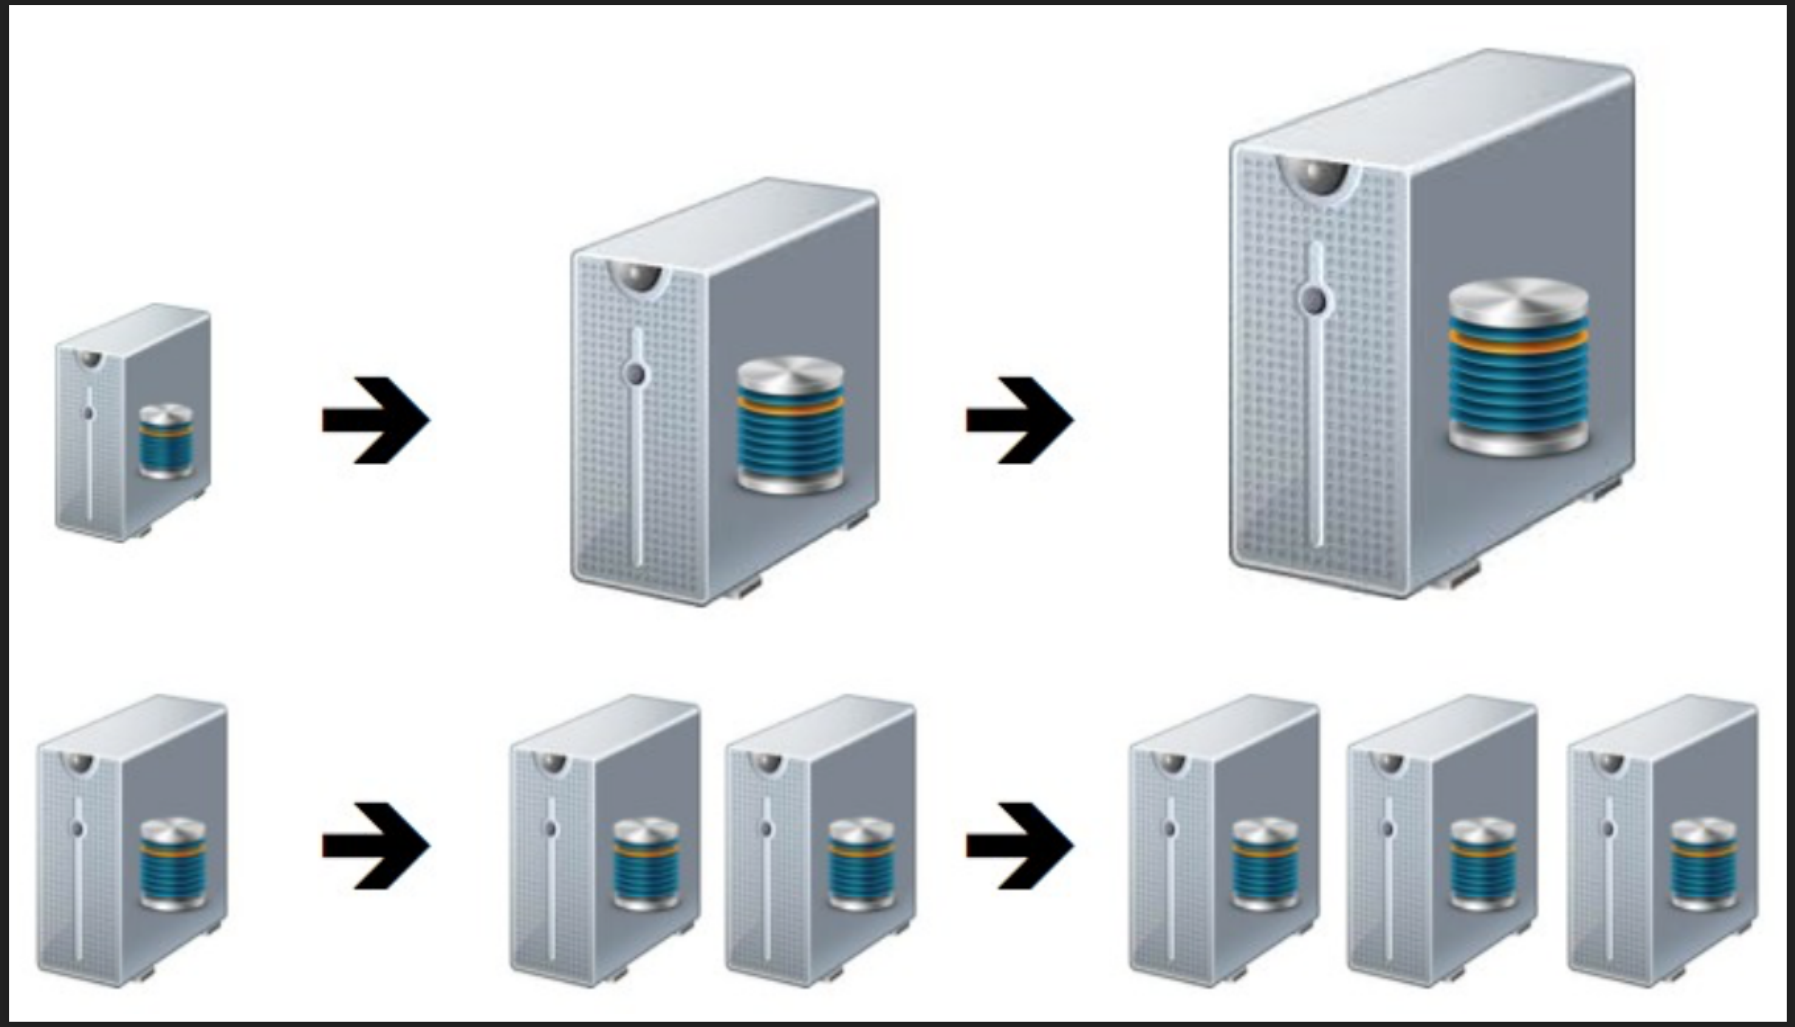

Organizations

SATURATION

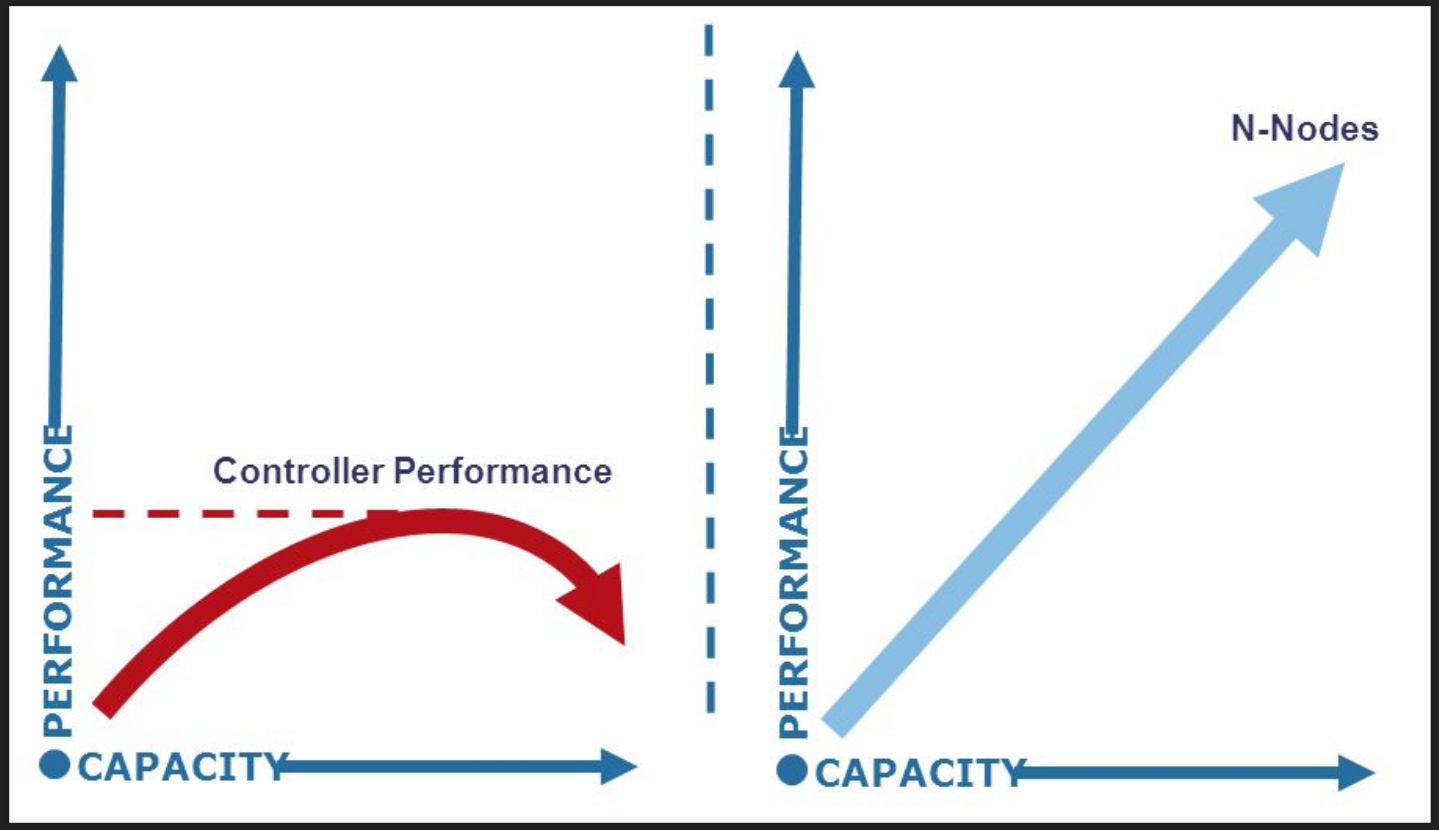

Bottlenecks like: Disk IO, load avg

No matter how much memory I have...

SATURATION

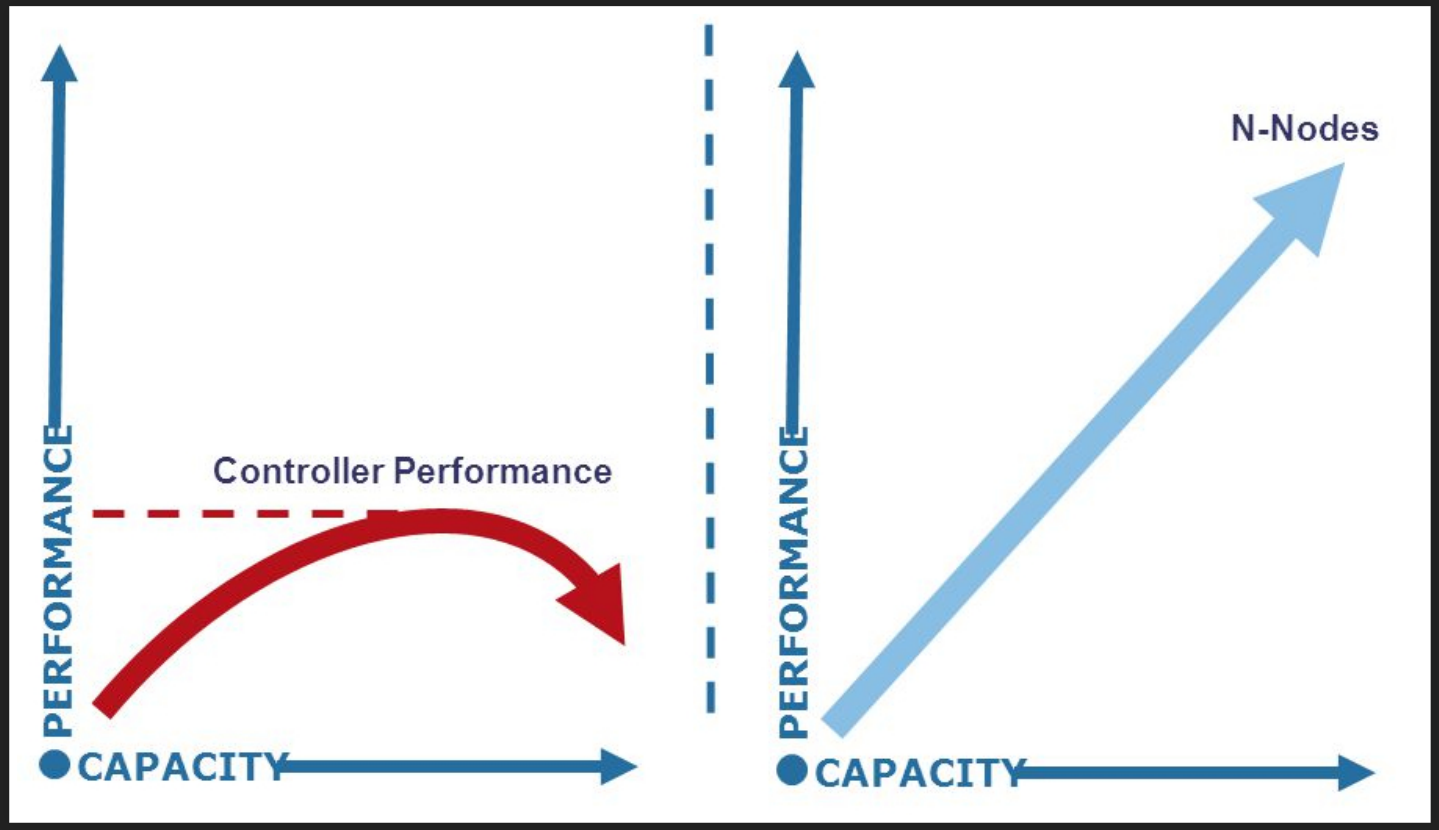

Bottlenecks like: Disk IO, load avg

No matter how much memory I have...

Now we know what we are looking for!

CLIENT SERVER ARCHITECTURE

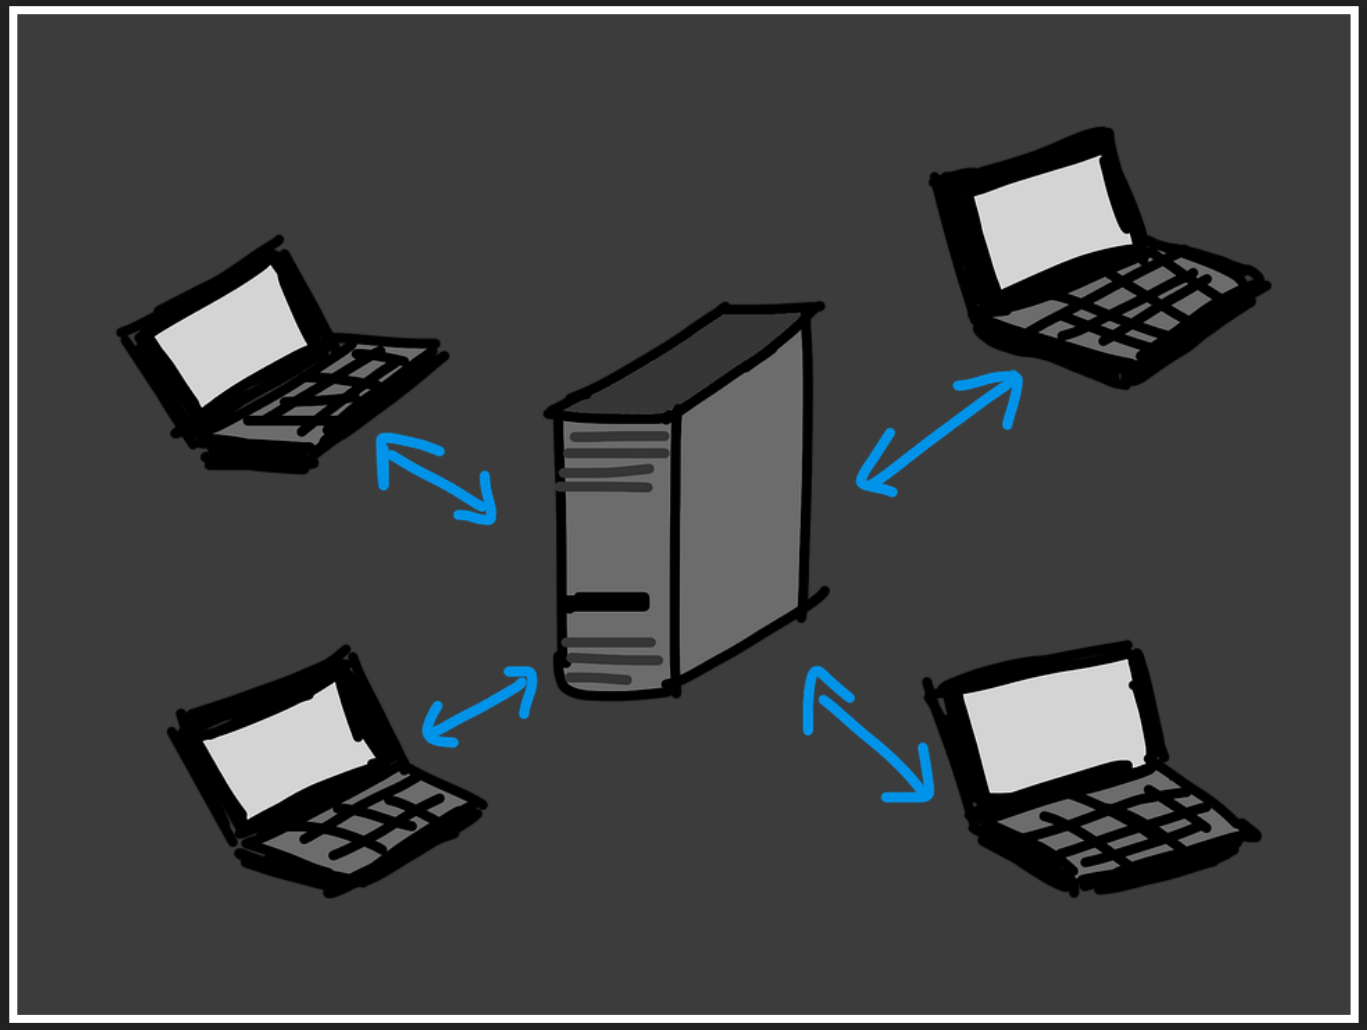

RAID

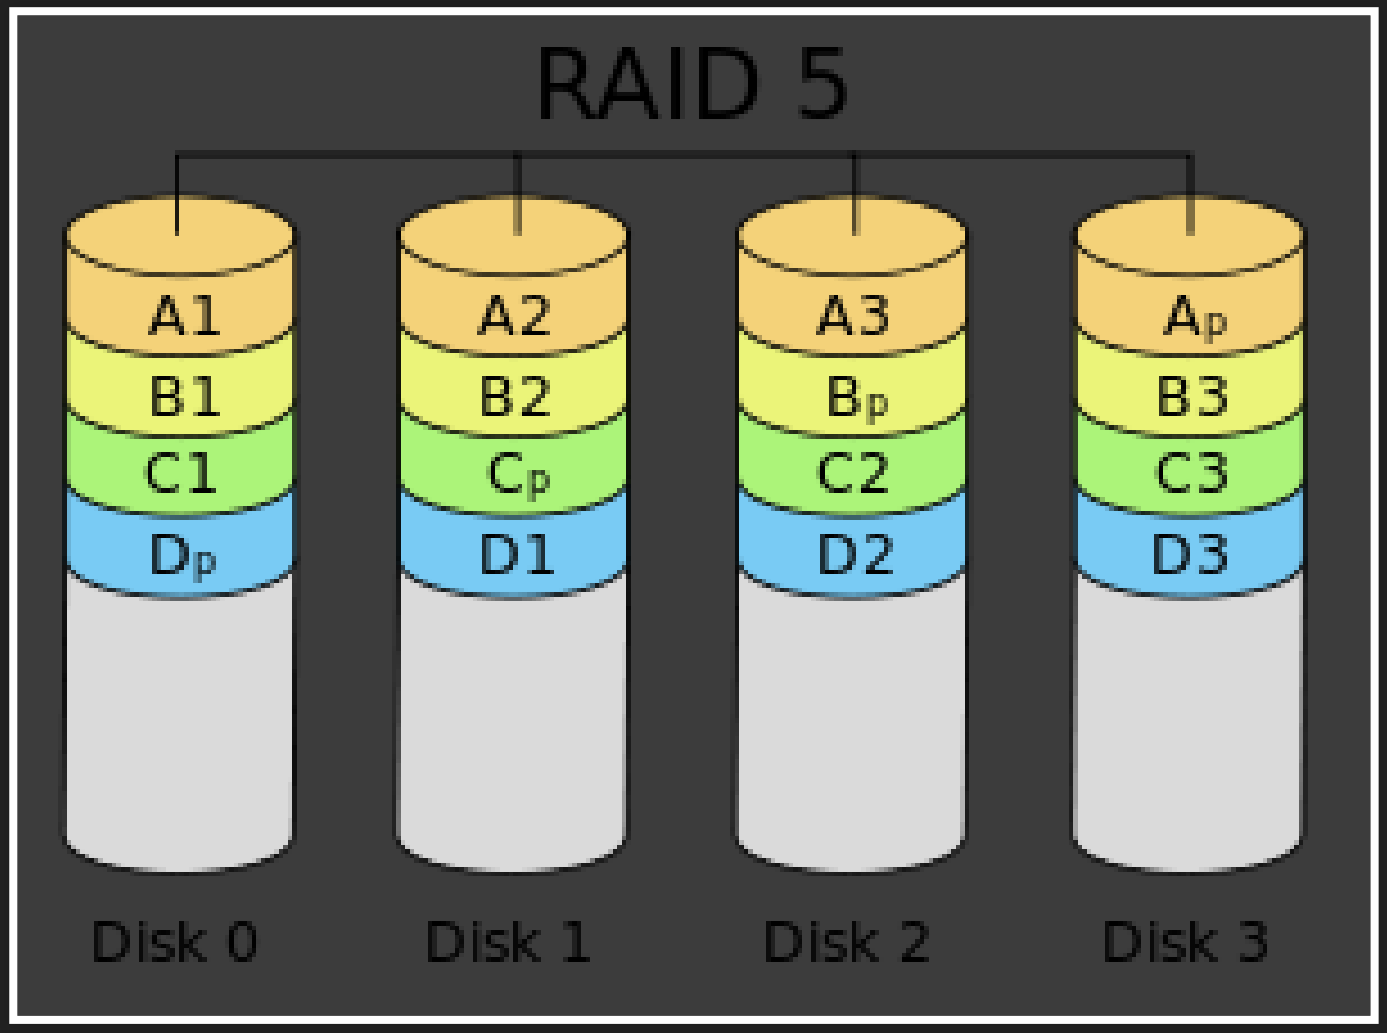
RAID 5
A1
B1
C1
Dp
A2
B2
Cp
D1
A3
Bp
C2
D2
Ap
B3
C3
D3
Disk 0
Disk 1
Disk 2
Disk 3

GOOGLE

GFS, ???

DOUG CUTTING

Apache Lucene (IR)

Apache Nutch (Web Crawler)

WHITE PAPER

MAPREDUCE

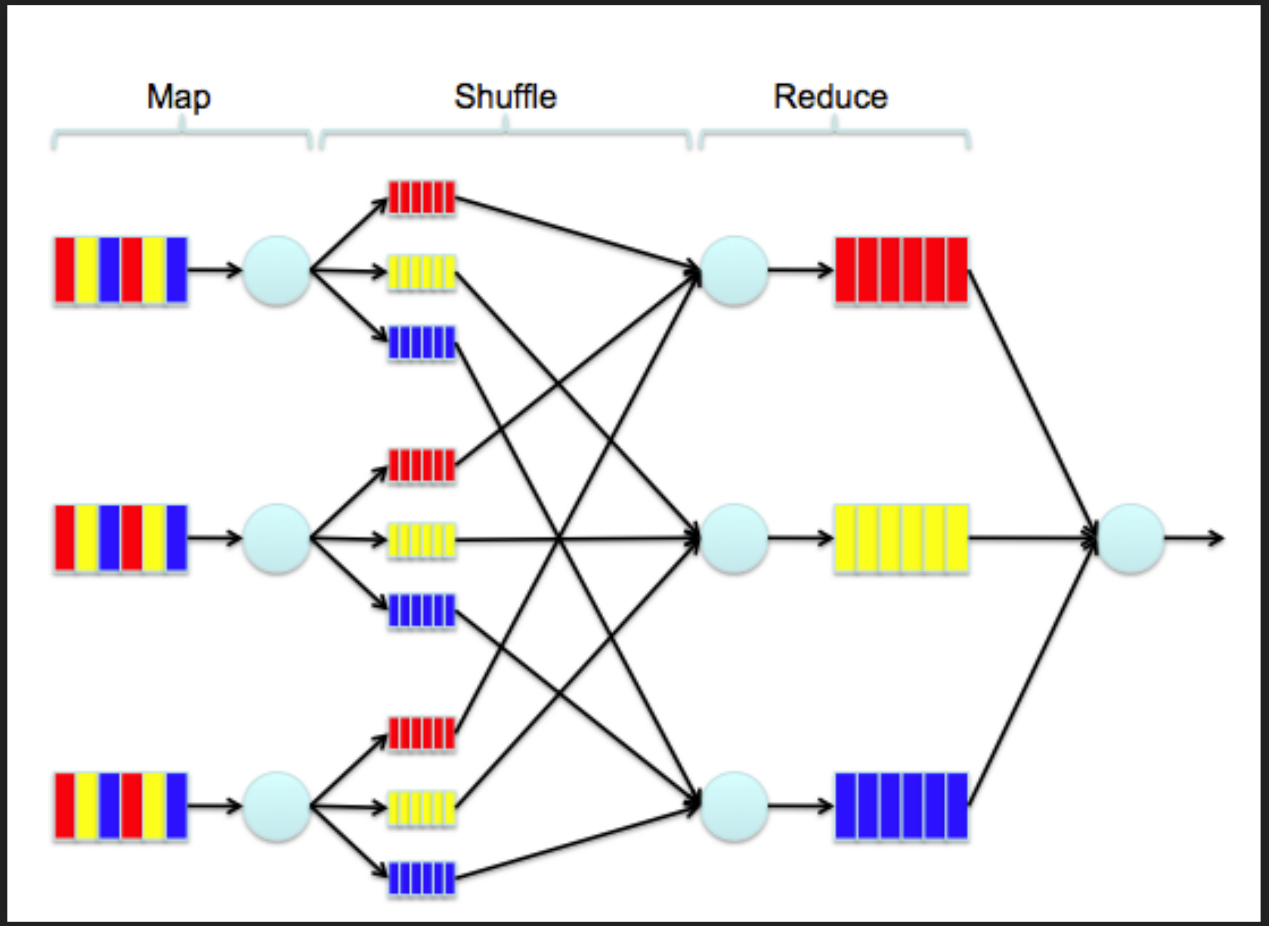

It's just a framework, C4.5

HADOOP COMES INTO PLAY!

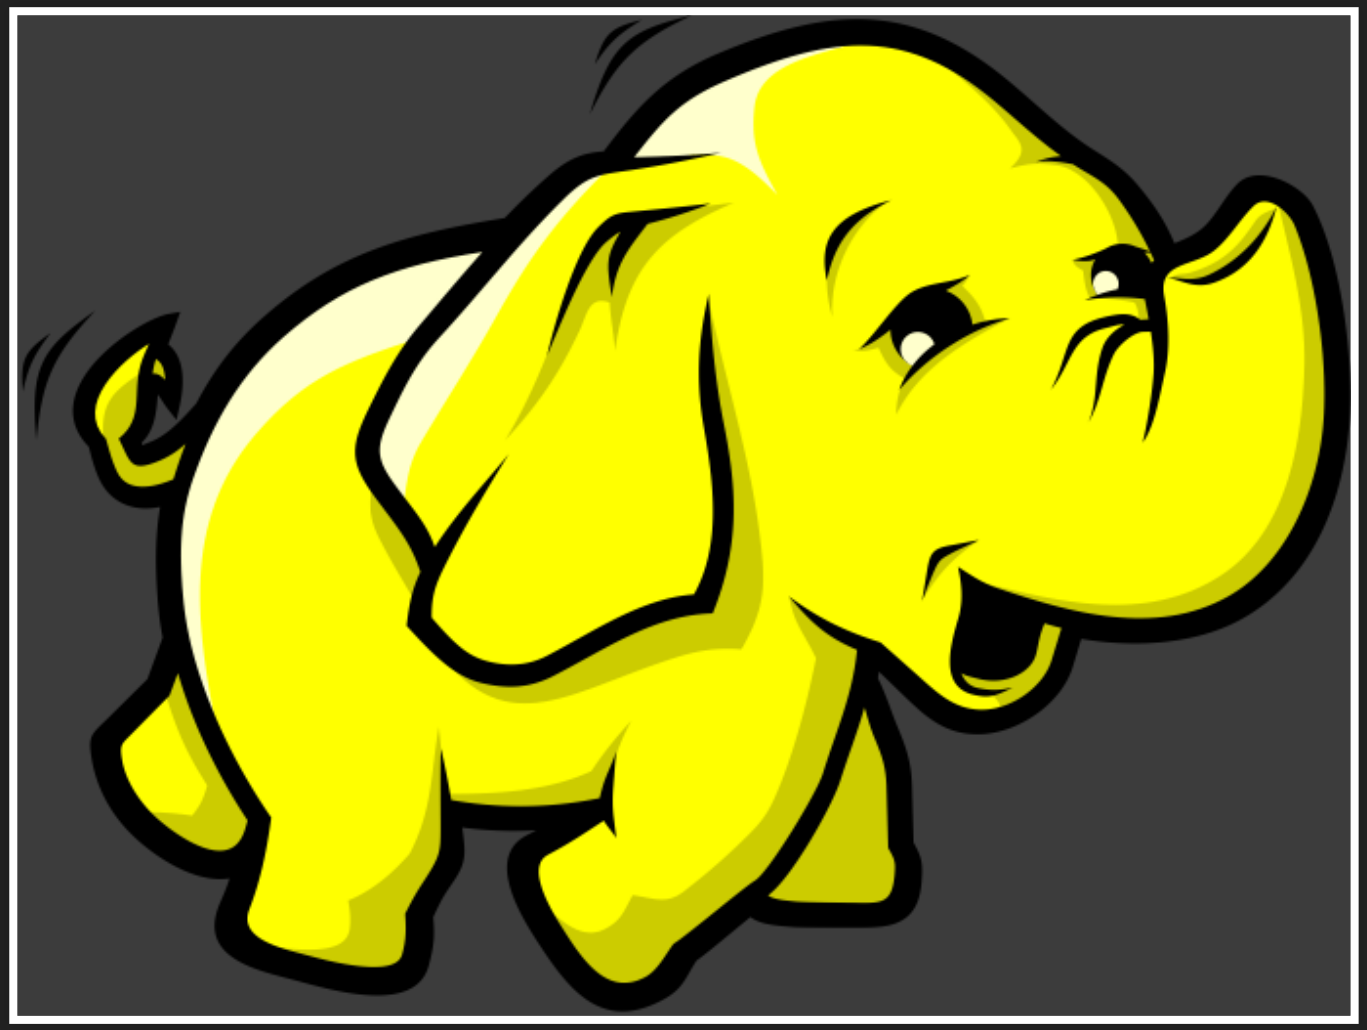

NAMING

HORTONWORKS DEFINATION

An opensource software platform for distributed storage and distributed processing of very large data sets on computer clusters built from commodity hardware

TWO SIDES OF HADOOP

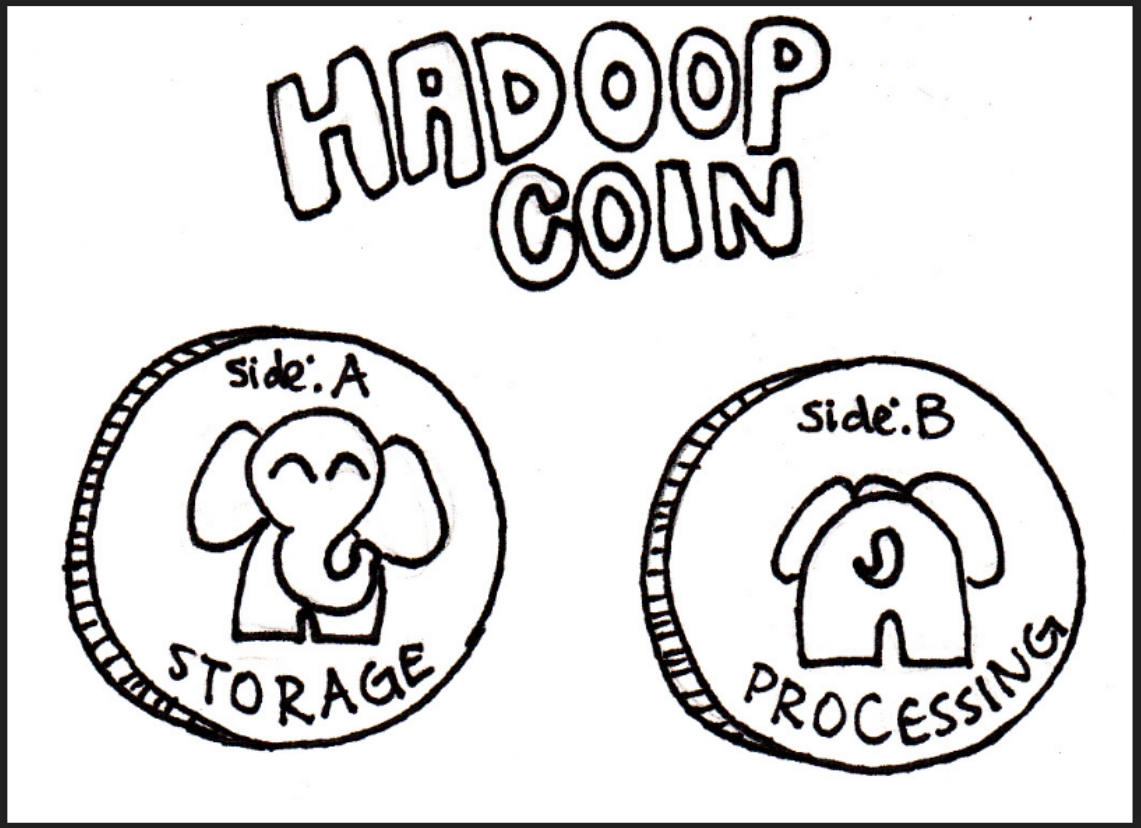

TWO SIDES OF HADOOP

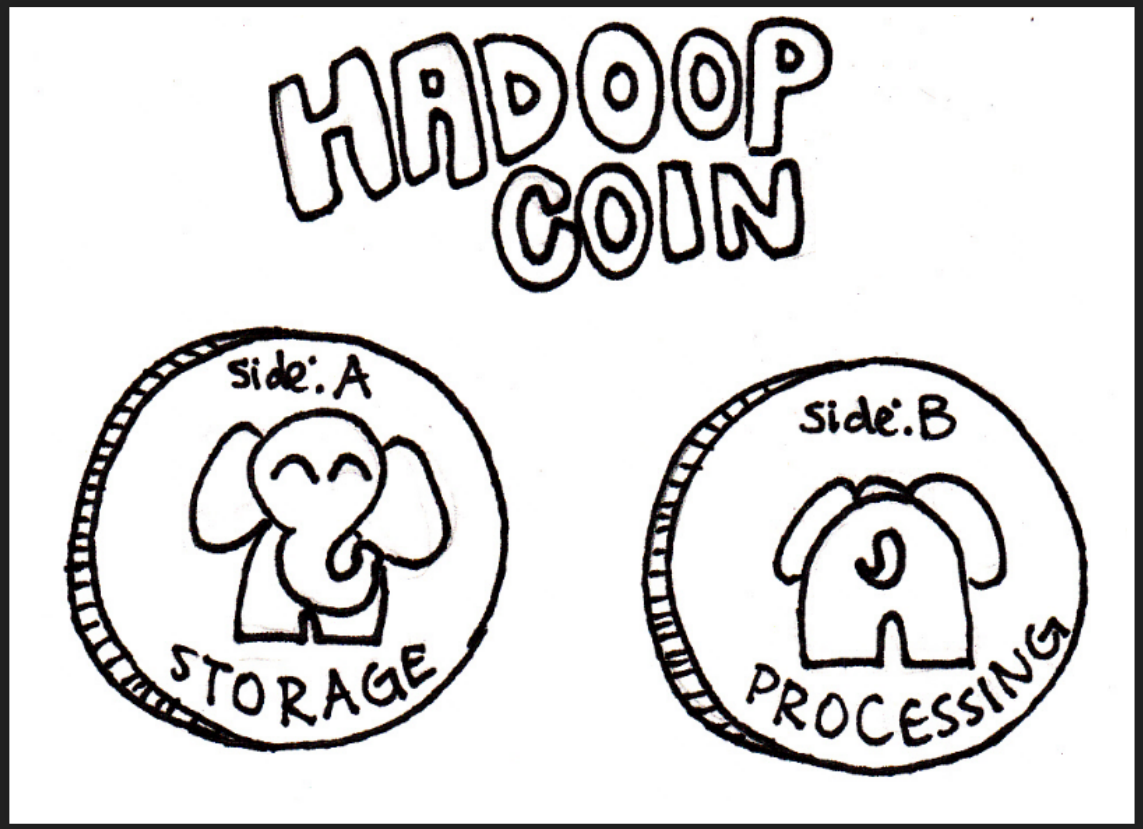

GFS, MapReduce

TWO SIDES OF HADOOP

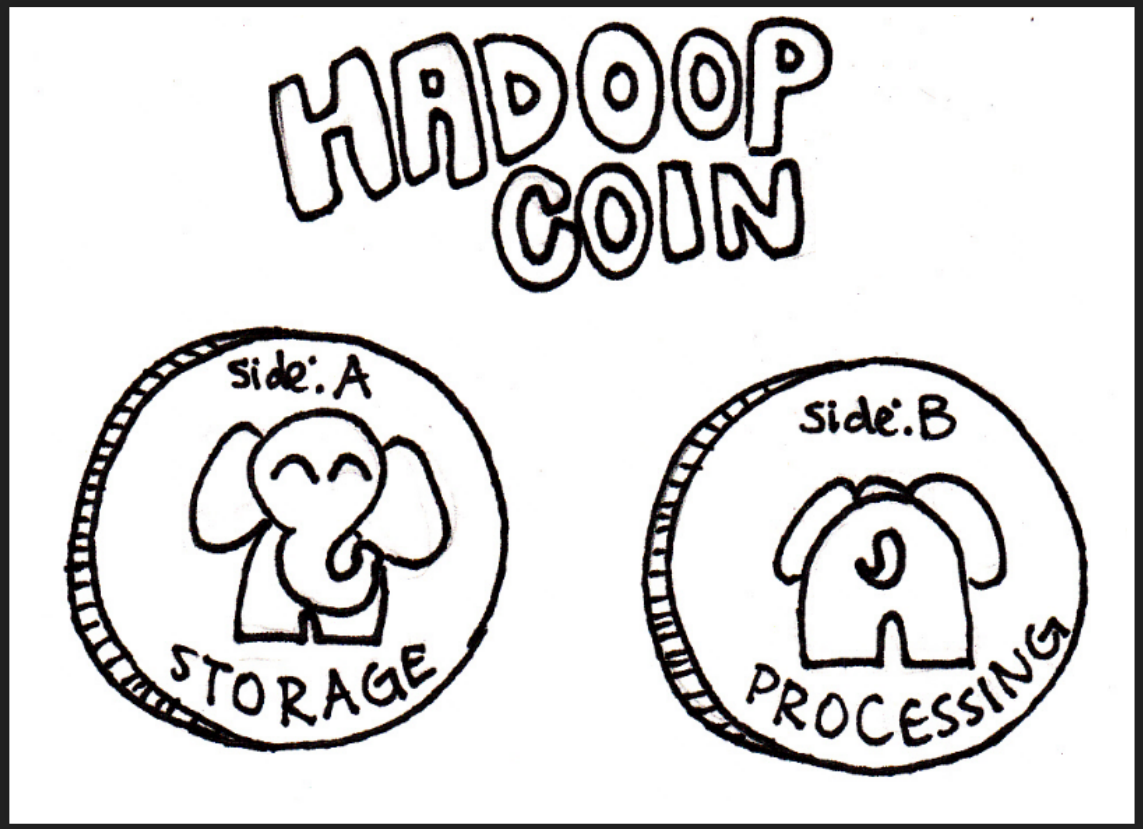

GFS, MapReduce

HDFS, Hadoop MapReduce

WHAT WE DO WITH DATA

WHAT WE DO WITH DATA

- Save
- Transfer
- Join
- Index
- Analytics
- Aggregate
- Visualize

ECOSYSTEM

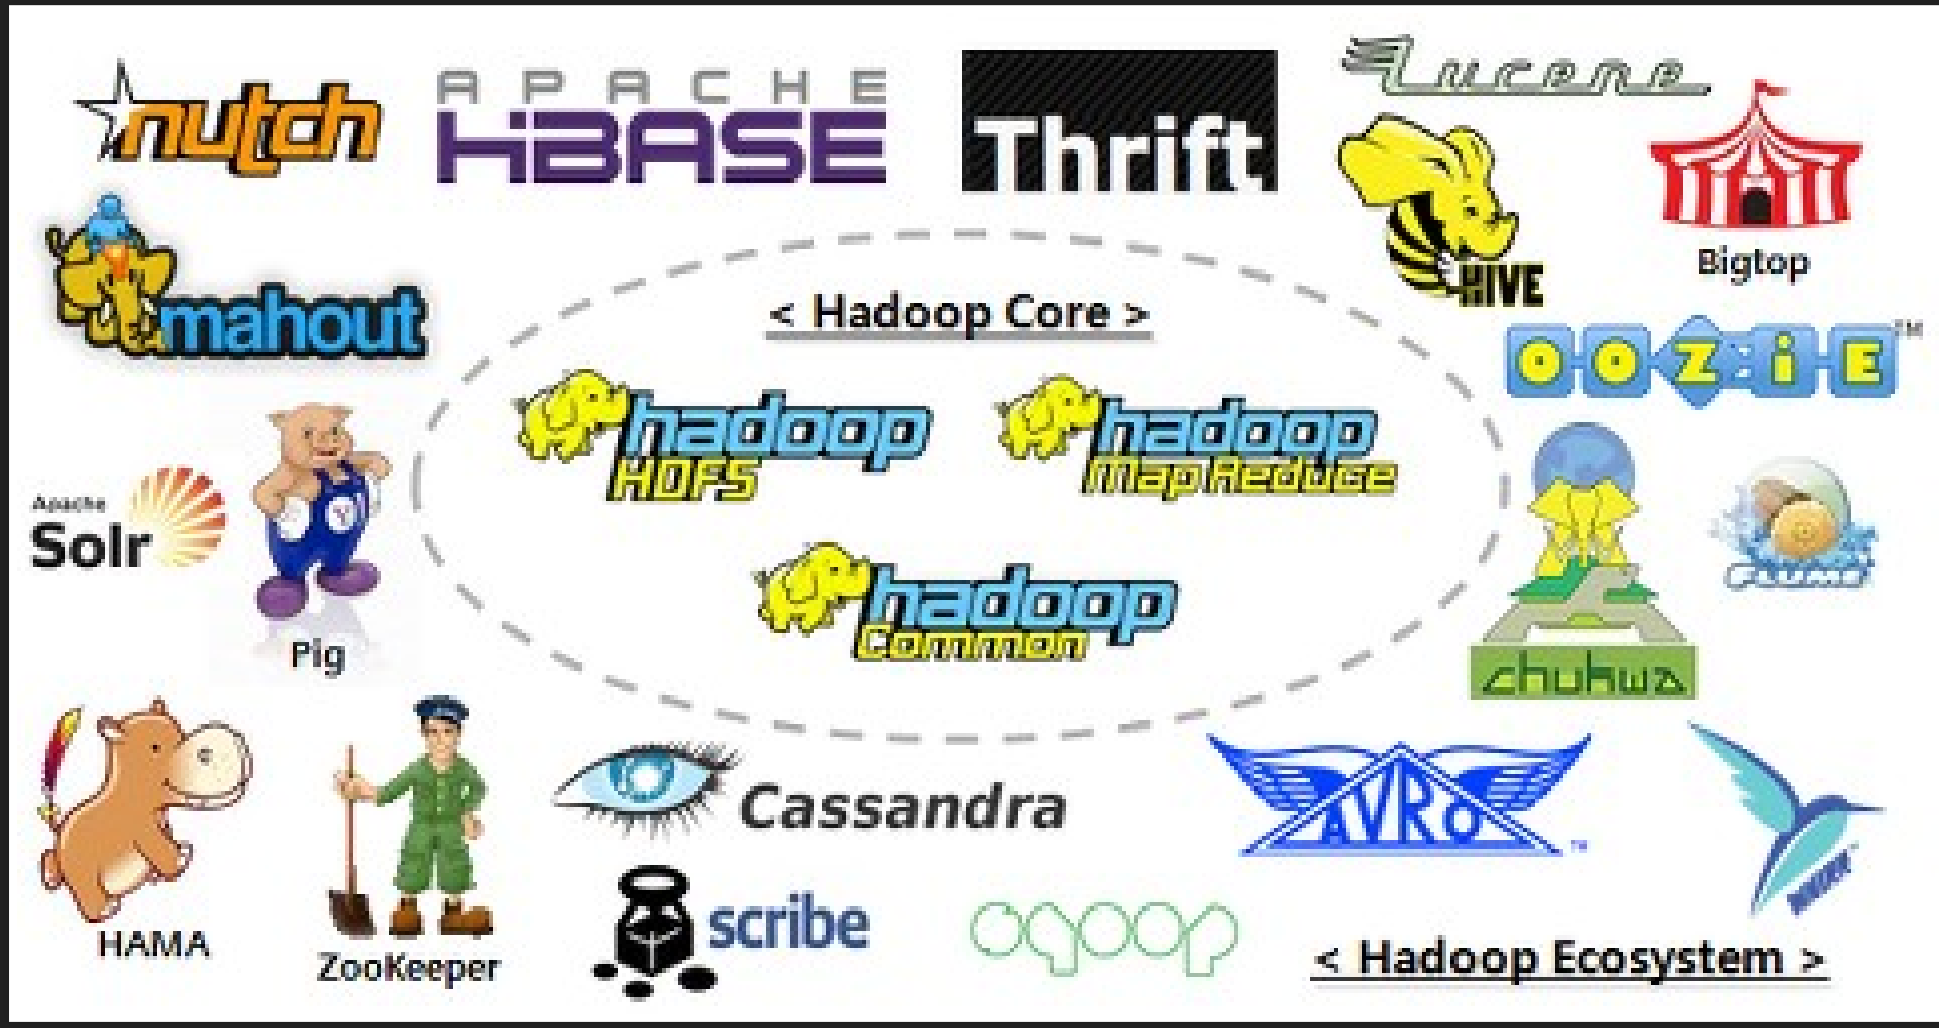
nutch
APACHE HBASE
Thrift
Lucene
HIVE
Bigtop
mahout
< Hadoop Core >
OOZIE
hadoop HDFS
hadoop Map Reduce
hadoop Common
Apache Solr
Pig
chukwa
Flume
Cassandra
AVRO
HAMA
ZooKeeper
scribe
< Hadoop Ecosystem >

ECOSYSTEM

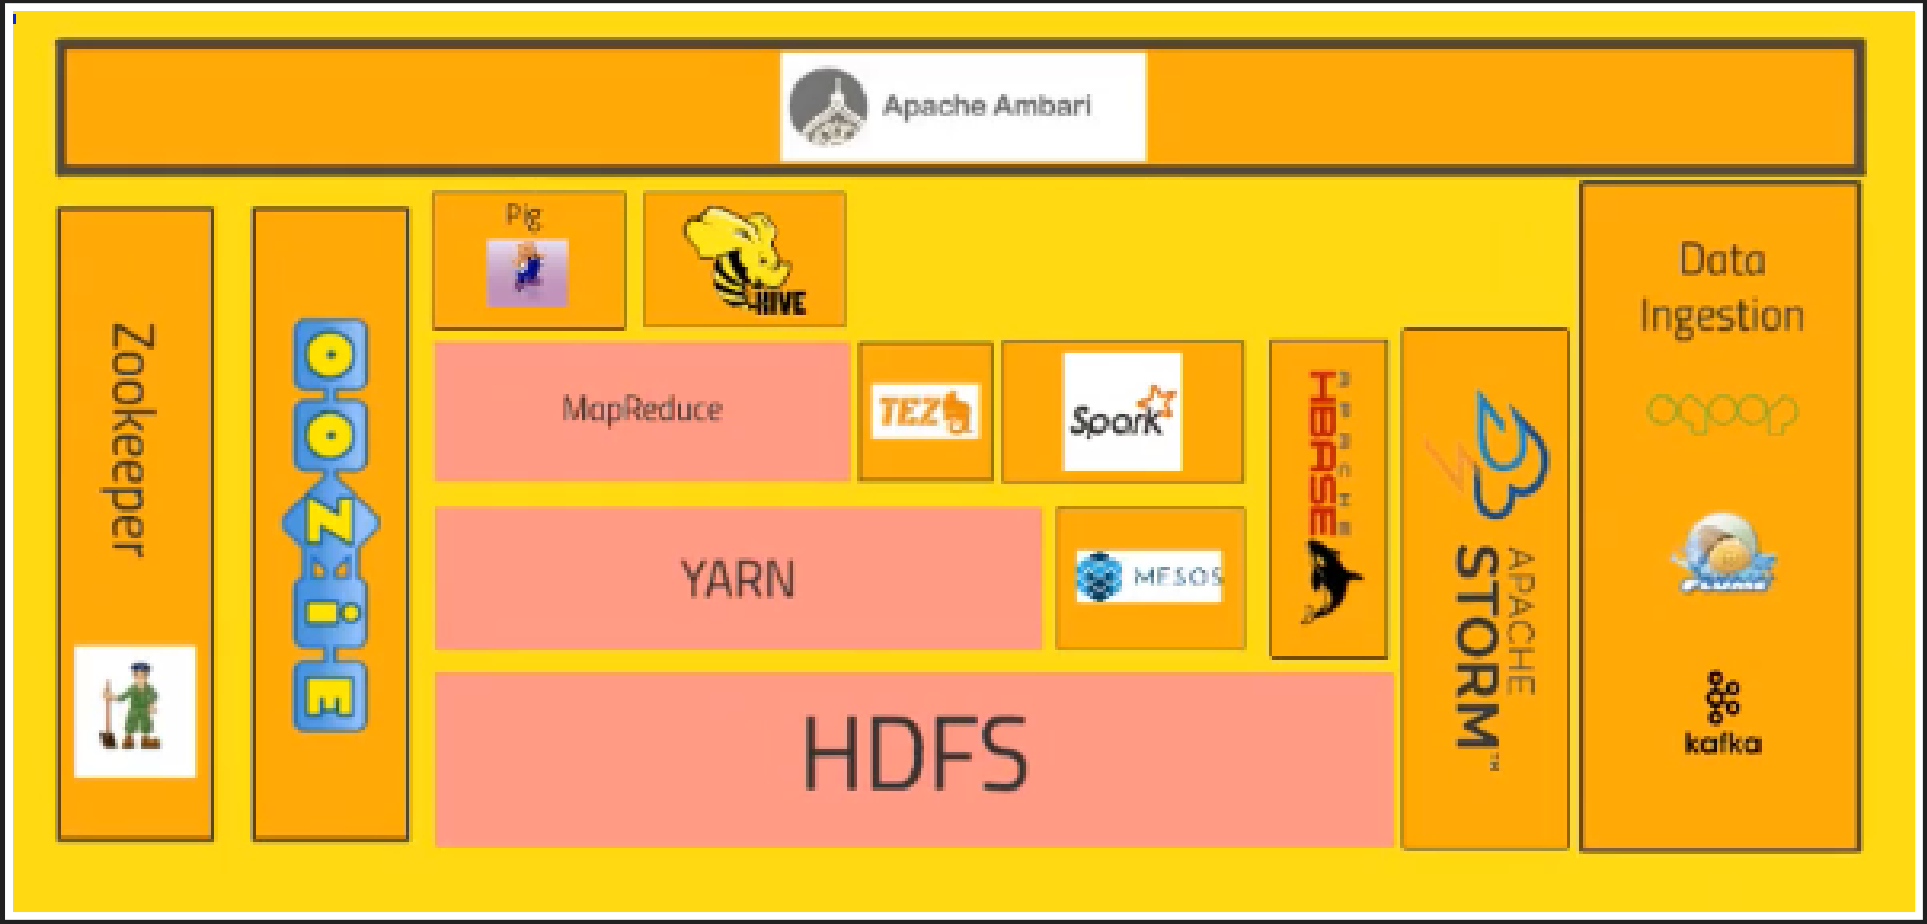

HDFS | YARN | MapReduce | Pig | Hive | TEZ | Storm | Spark | oozie | Zookeeper | Others

HDFS

Allows us to distribute the storage

All hard drives look likes a single huge hard disk

Keeps copy of data

No single point of failure

Back

YARN

Manage the resouces

What get to run tasks and when

Which node is available

We build applications on top of it

Back

MAPREDUCE

Got remvoed from YARN

Back

PIG

No java, Python? Scripting language like SQL
Transforms the script to something than can be run on MapReduce

Back

HIVE

SQL

Makes the data to look like a RDBMS

Back

TEZ

Hive on TEZ is faster than MapReduce

Back

SPARK

Sitting at same level of MapReduce on top of YARN Or MESOS

Python, Java, Scala

Fast

Active development

Handle SQL Query

Machine learning

Handle Stream data

Back

STORM

Processing streaming data

Sensors, Logs

Spark streaming does the same thing

Update machine learning model

Update data as it comes

Back

OOZIE

SCHEDULE OF JOBS

Complicated steps
Load to hive, Query using spark then transform to HBASE.

Back

ZOOKEEPER

Cordinates everything on clusters

Which node is up or down

Many of these apps relay on zookeeper

Back

DATA INGESTION

SQOOP

Turn into hadoop, talks to ODBS JDBC

FLUME

Transport Web logs into hadoop (spark, storm)

KAFKA

Like flume but more general, cluster of PC or webservers or whatever to broadcast into Hadoop cluster.

DISTRIBUTIONS

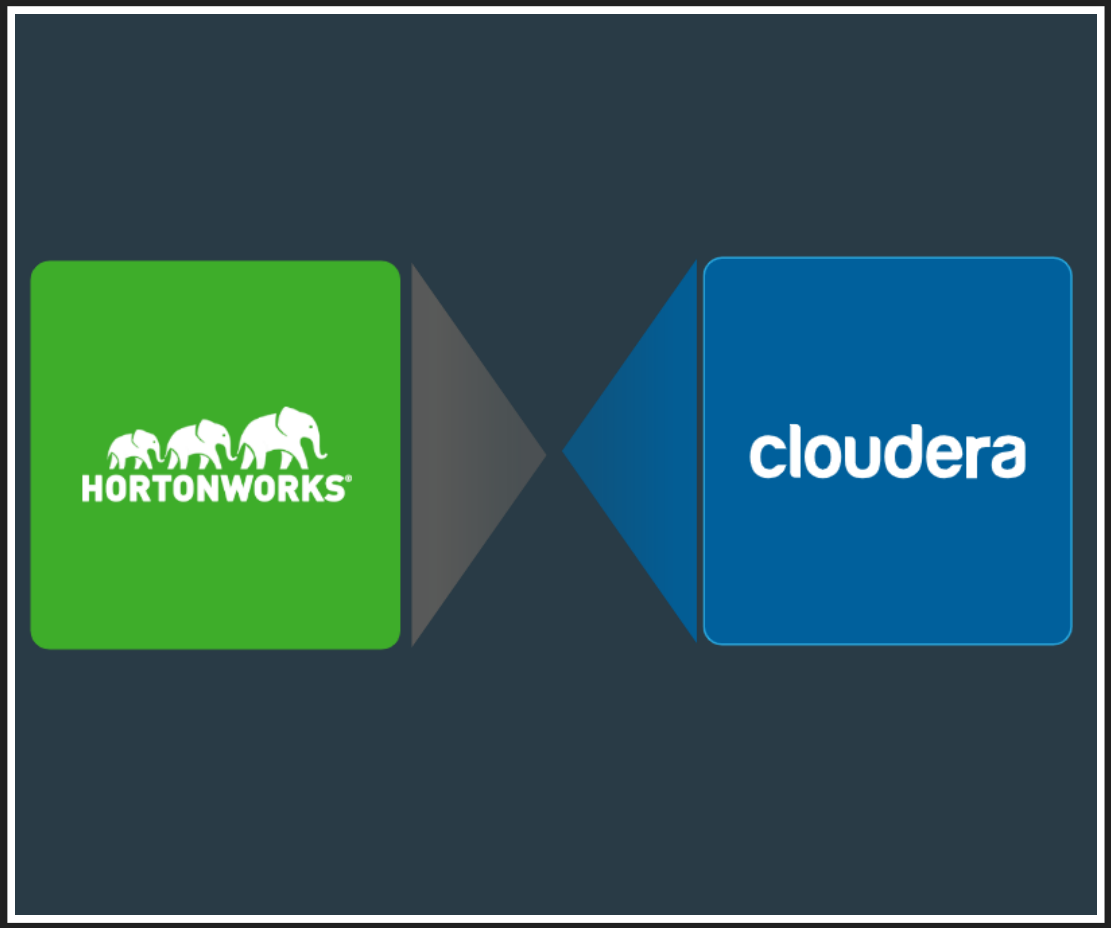

HORTONWORKS SANDBOX

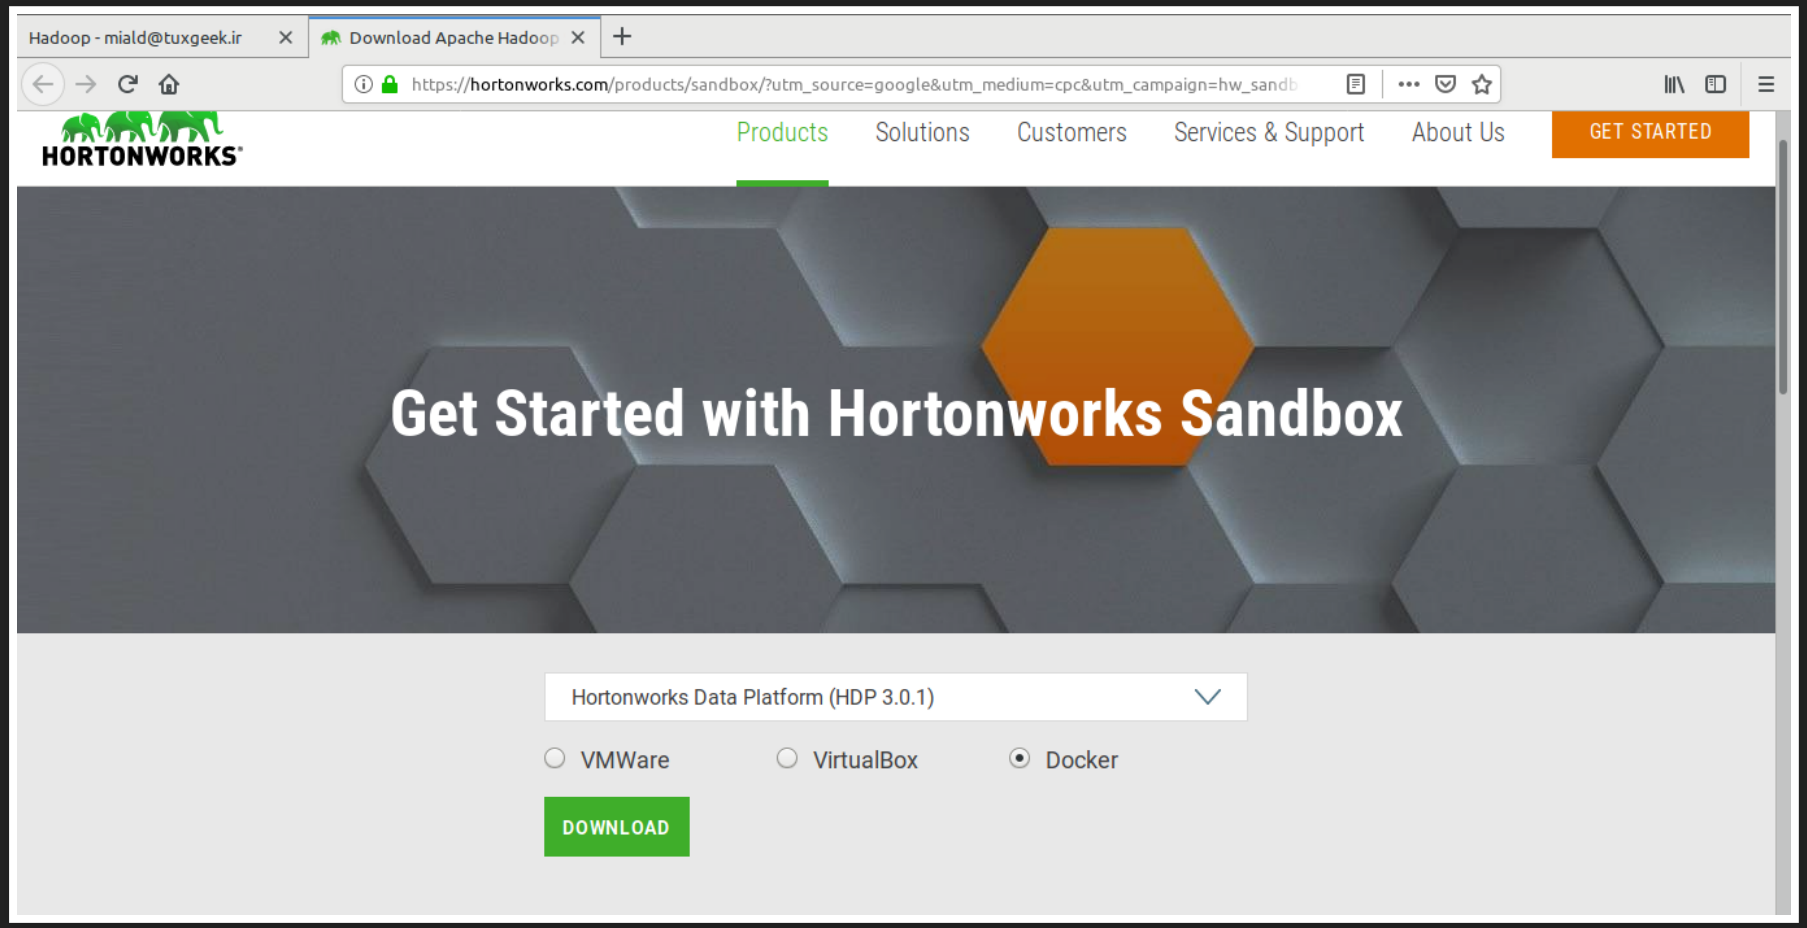

AN HIVE EXAMPLE

TTY

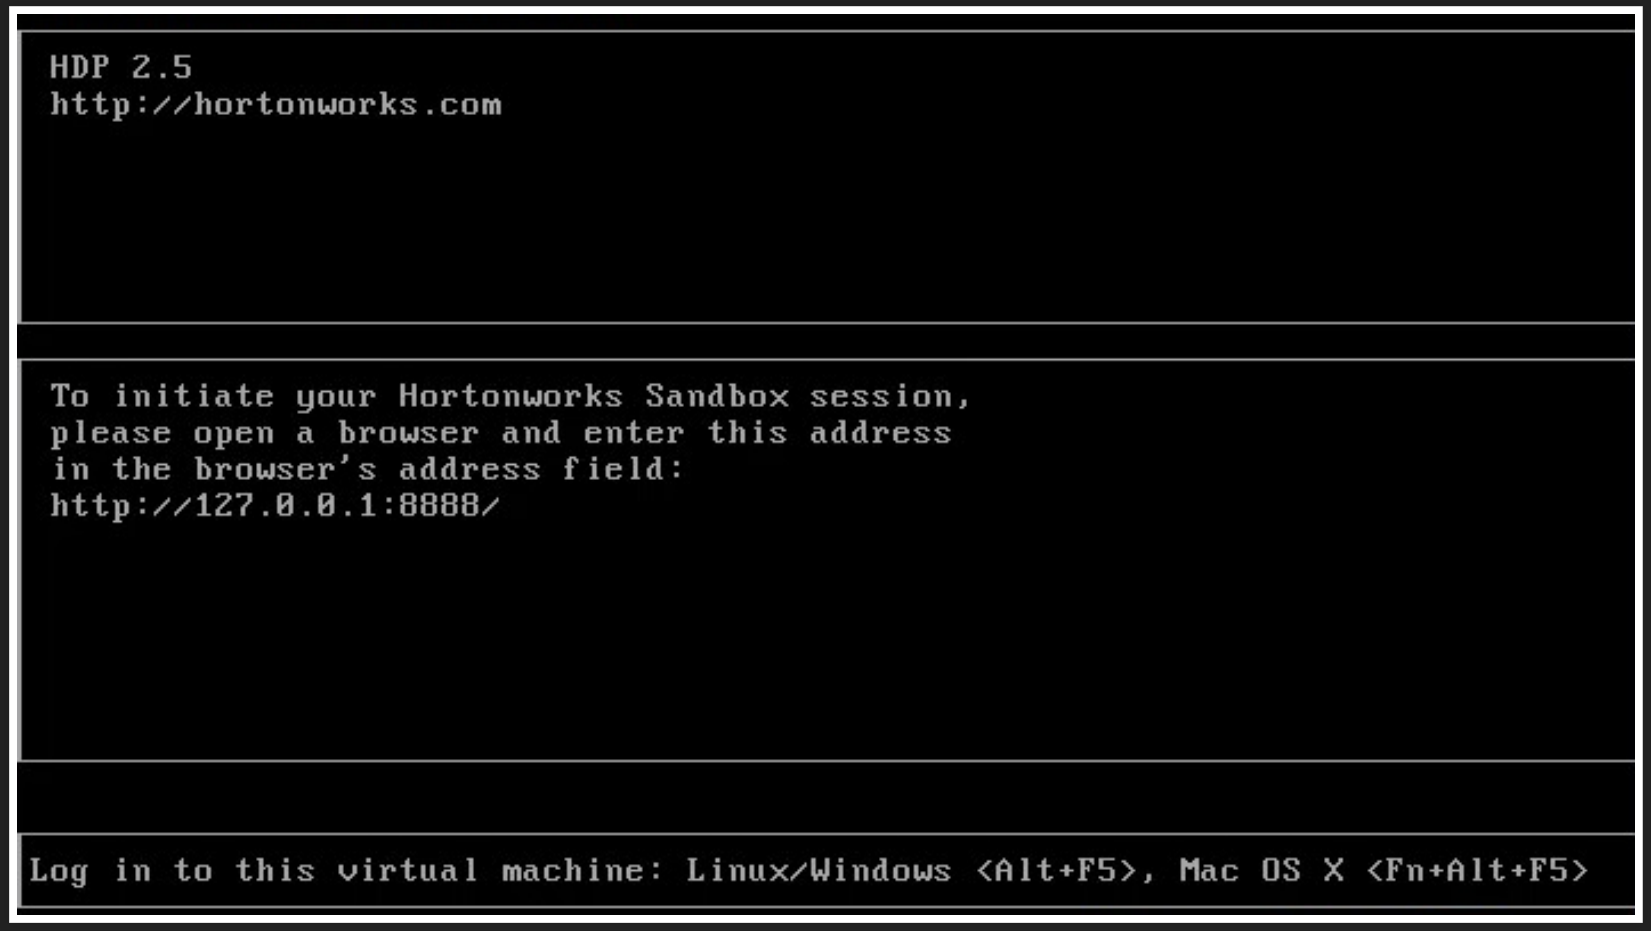
HDP 2.5
http://hortonworks.com
To initiate your Hortonworks Sandbox session,
please open a browser and enter this address
in the browser's address field:
http://127.0.0.1:8888/
Log in to this virtual machine: Linux/Windows <Alt+F5>, Mac OS X <Fn+Alt+F5>

APACHE AMBARI

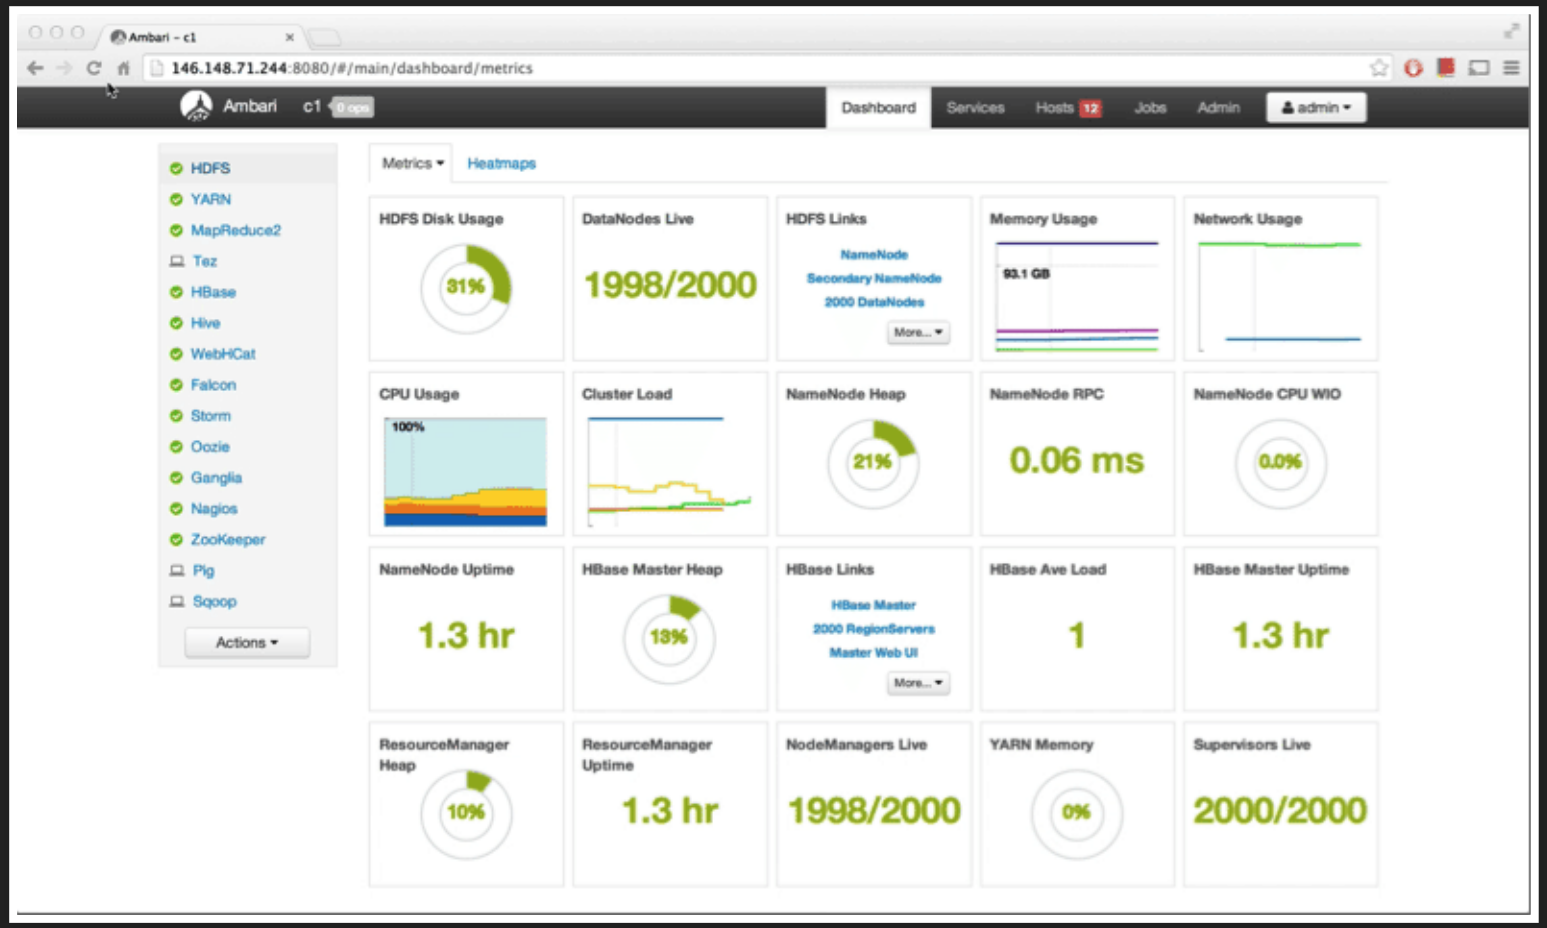

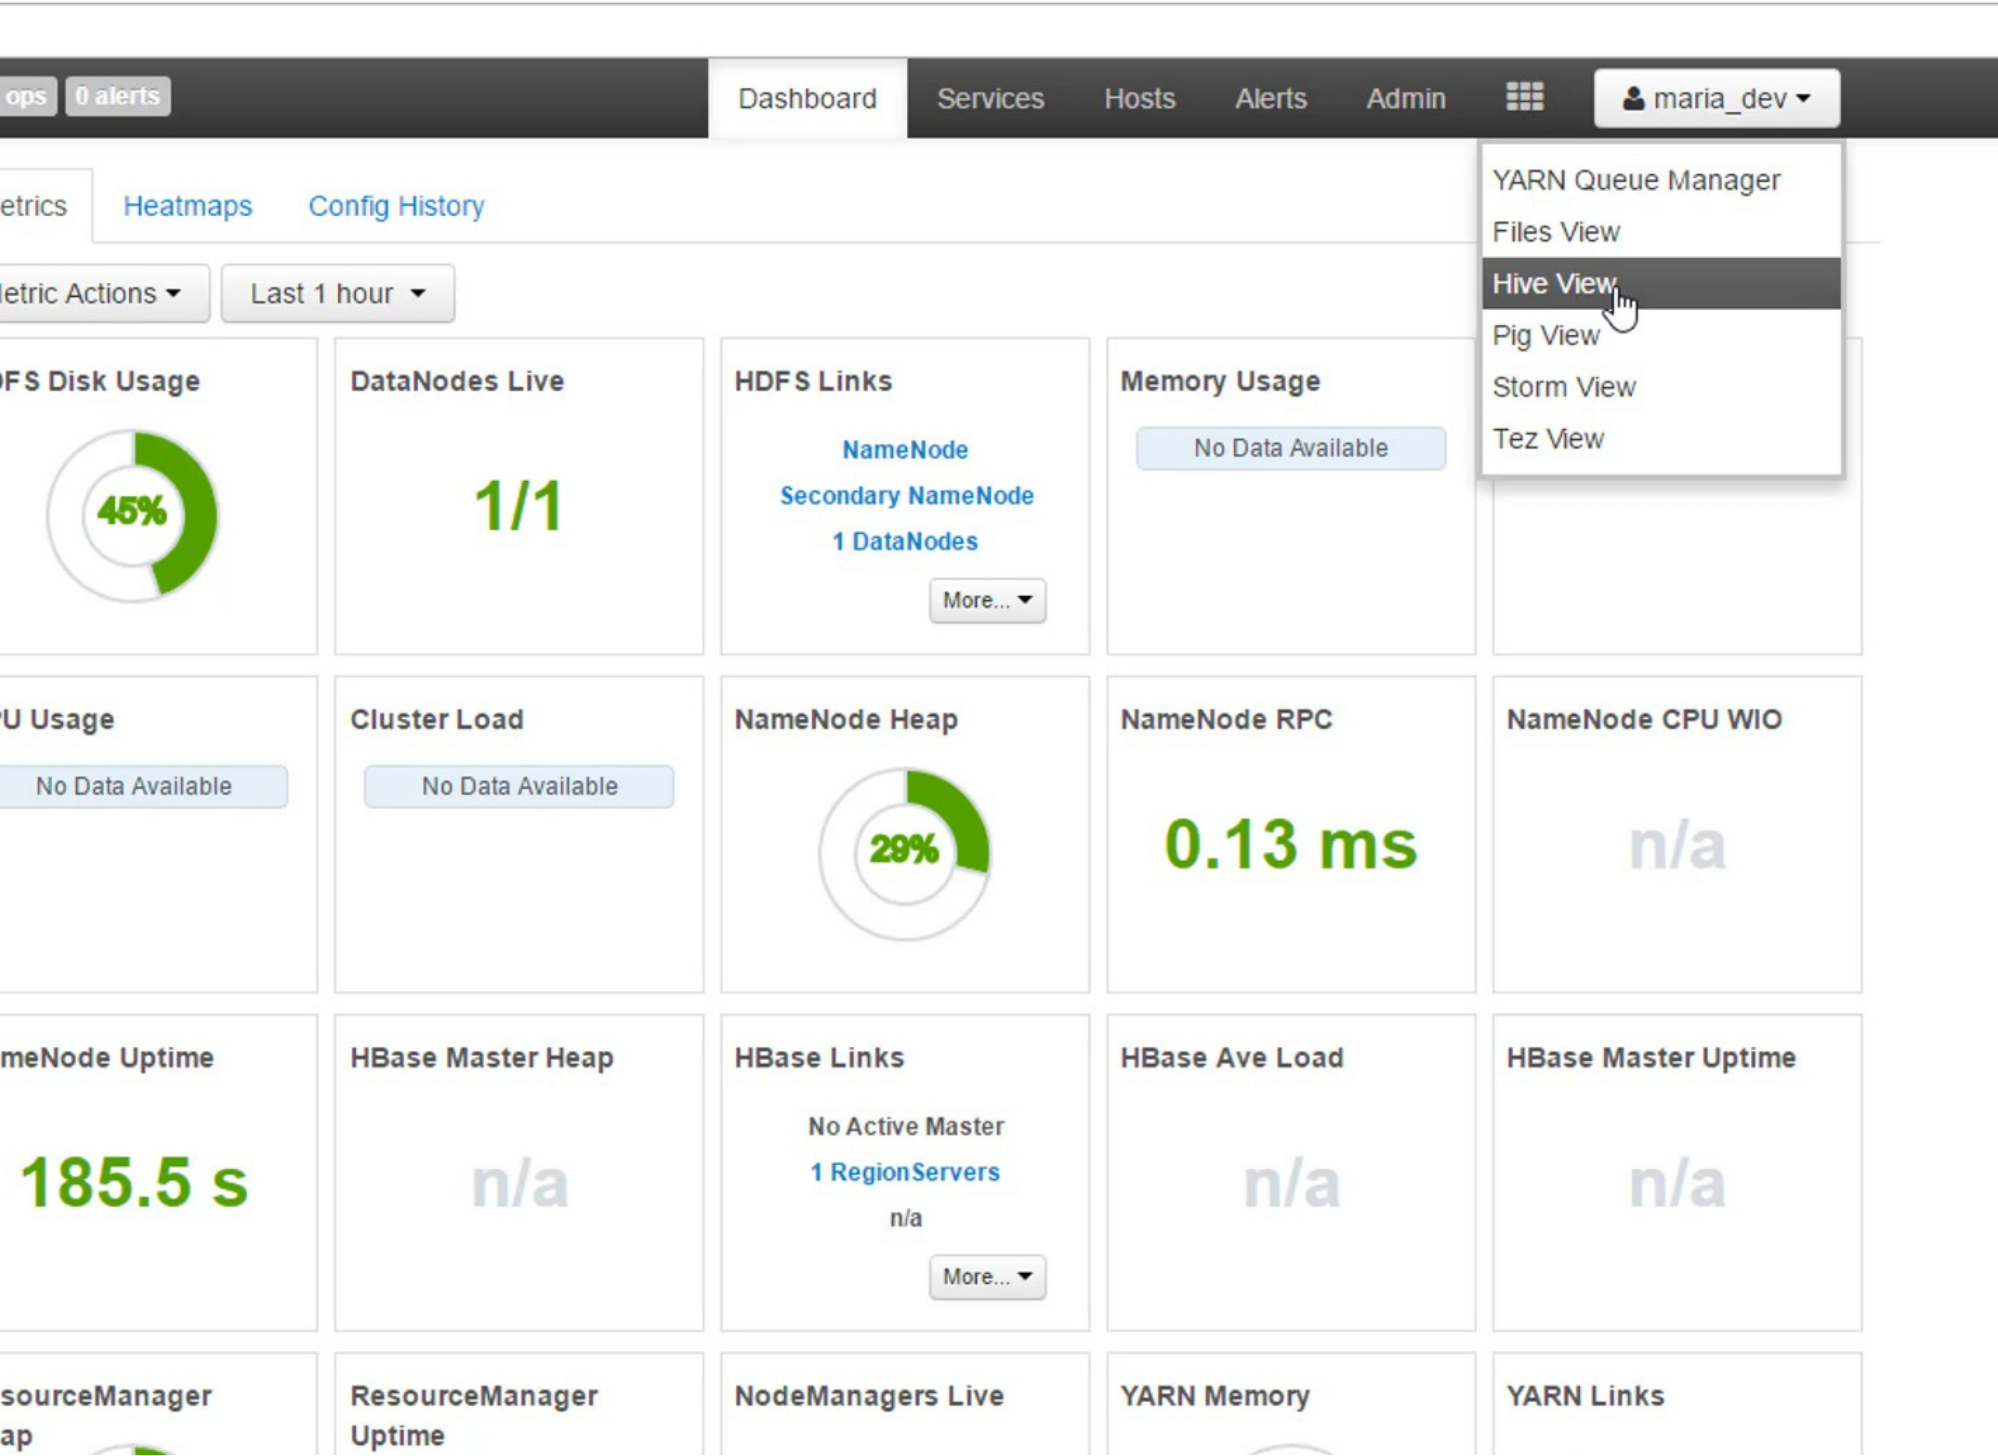
ops
0 alerts
Dashboard
Services
Hosts
Alerts
Admin
maria_dev
YARN Queue Manager
Files View
Hive View
Pig View
Storm View
Tez View
etrics
Heatmaps
Config History
letric Actions
Last 1 hour
DFS Disk Usage
45%
DataNodes Live
1/1
HDFS Links
NameNode
Secondary NameNode
1 DataNodes
More...
Memory Usage
No Data Available
PU Usage
No Data Available
Cluster Load
No Data Available
NameNode Heap
29%
NameNode RPC
0.13 ms
NameNode CPU WIO
n/a
meNode Uptime
185.5 s
HBase Master Heap
n/a
HBase Links
No Active Master
1 RegionServers
n/a
More...
HBase Ave Load
n/a
HBase Master Uptime
n/a
sourceManager
ap
ResourceManager
Uptime
NodeManagers Live
YARN Memory
YARN Links

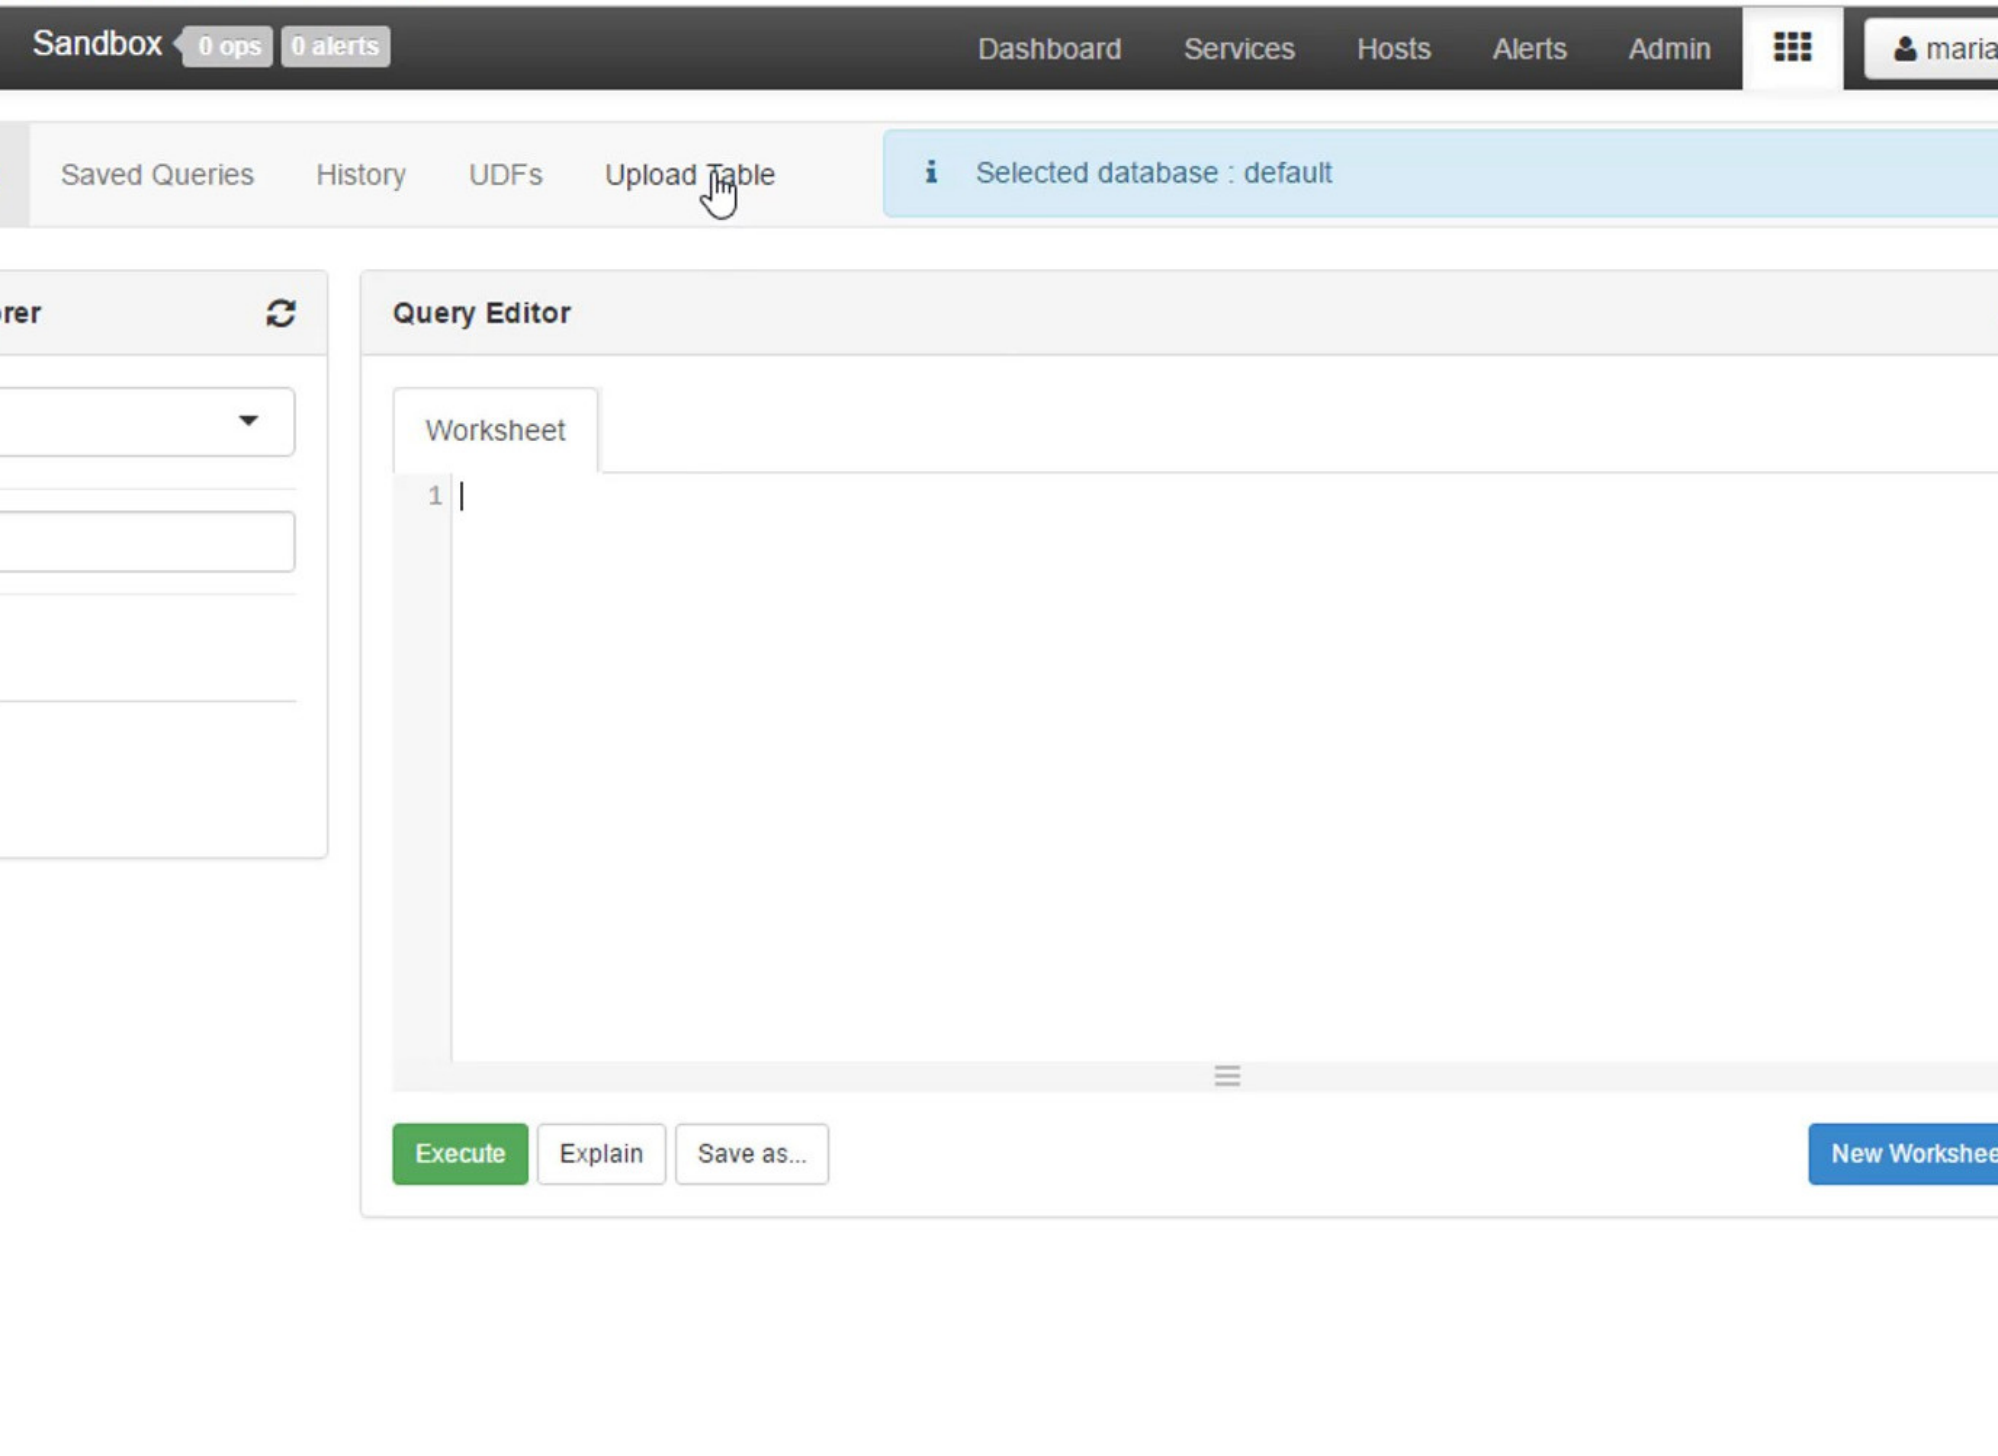
Sandbox
0 ops
0 alerts
Dashboard
Services
Hosts
Alerts
Admin
maria
Saved Queries
History
UDFs
Upload Table
Selected database : default
Query Editor
Worksheet
1
Execute
Explain
Save as...
New Worksheet

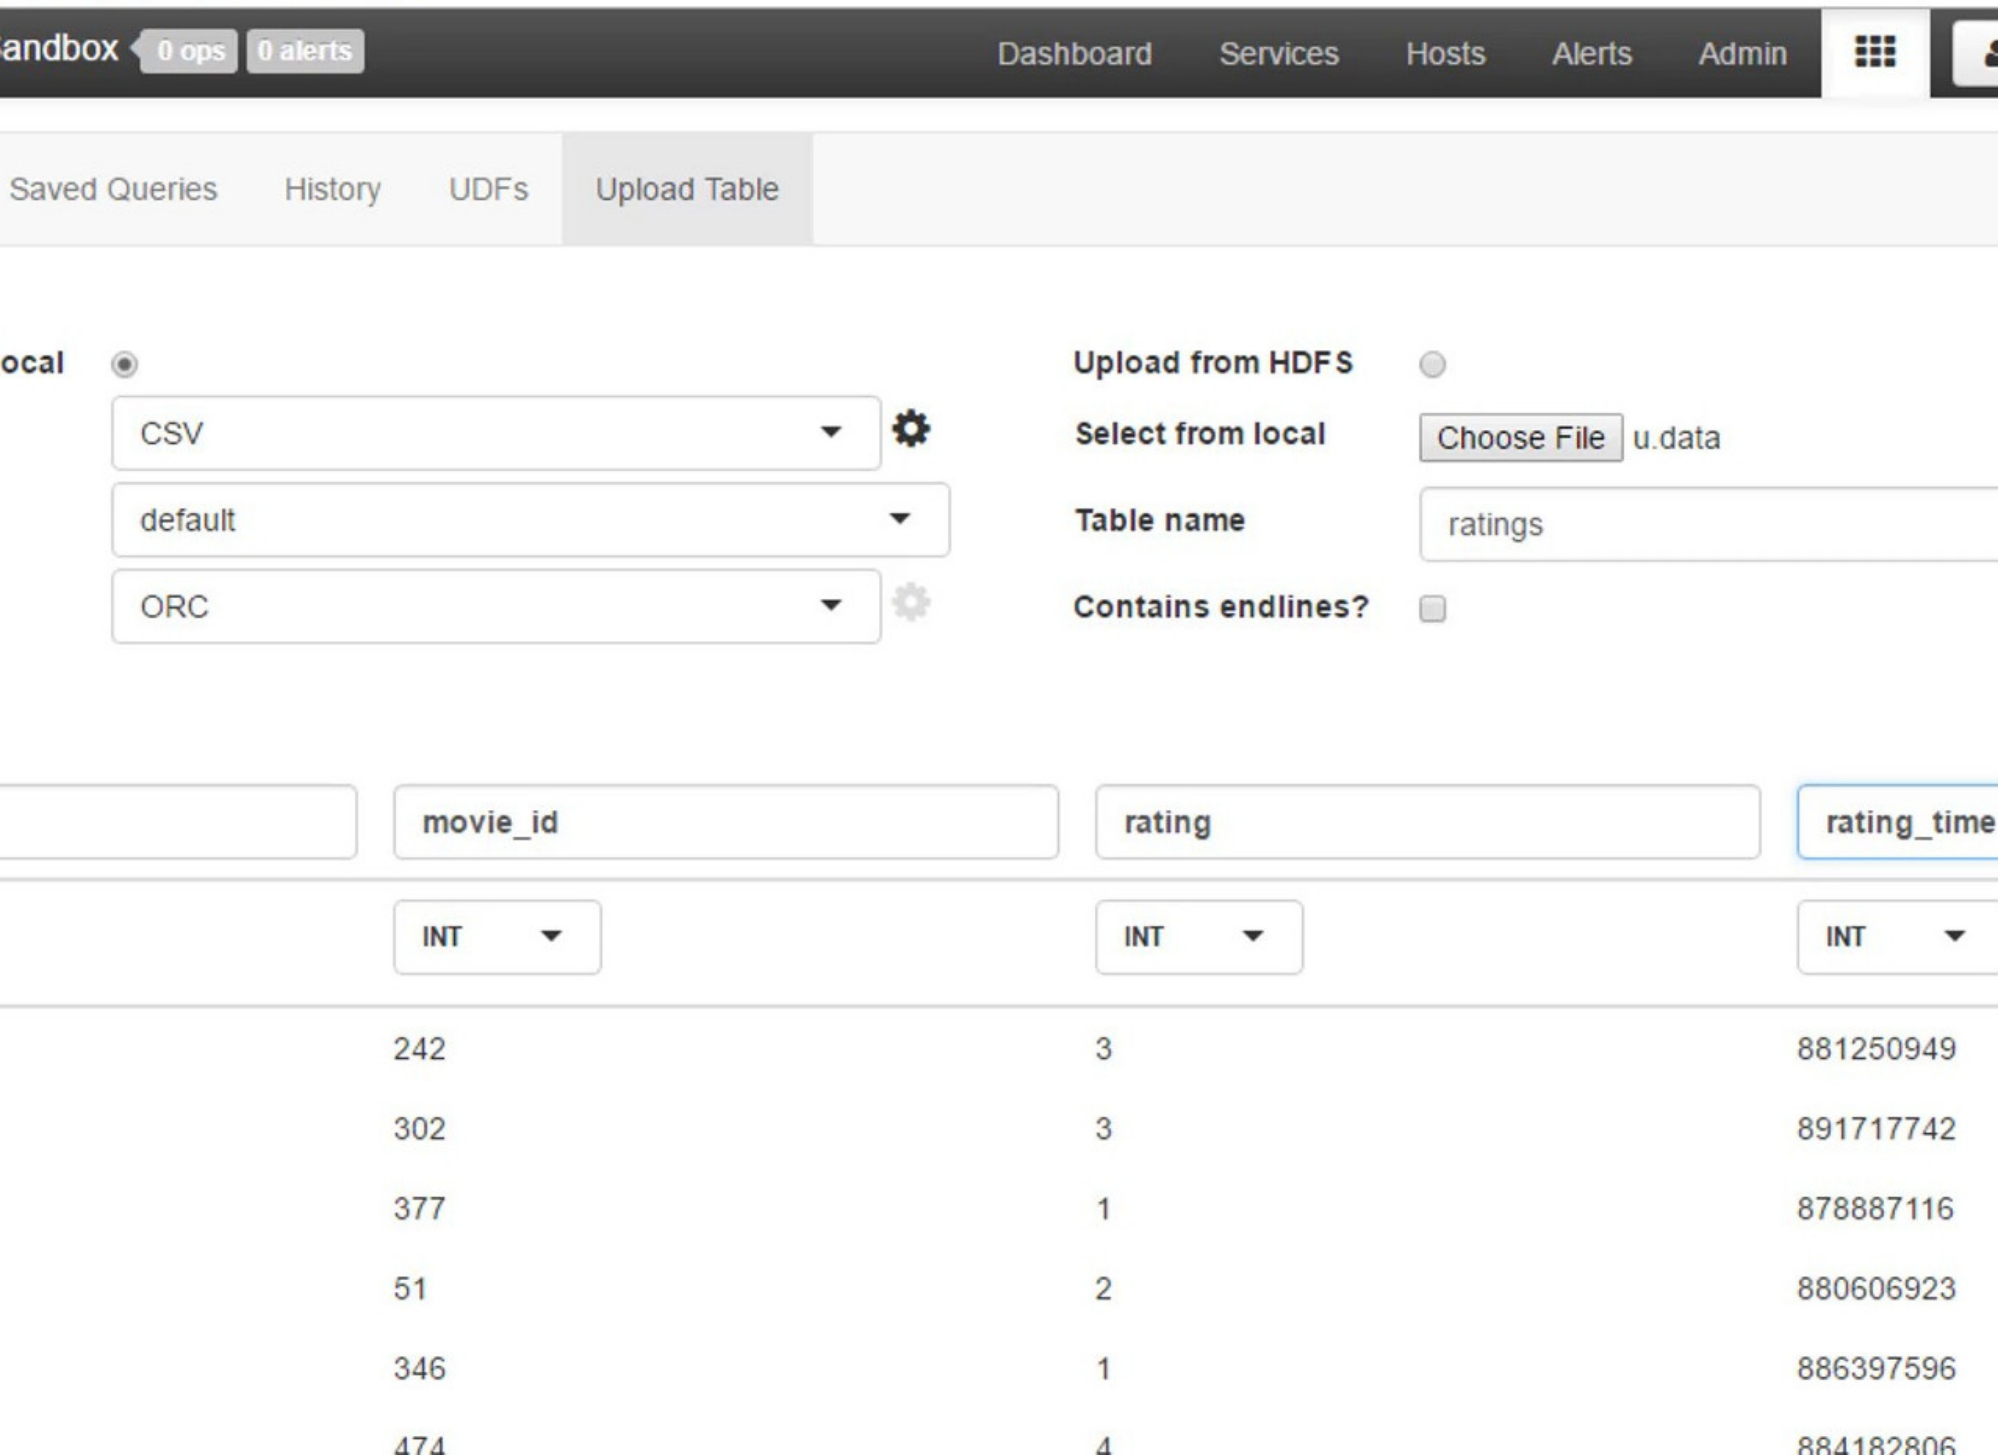
andbox
0 ops
0 alerts
Dashboard
Services
Hosts
Alerts
Admin
Saved Queries
History
UDFs
Upload Table
ocal
CSV
default
ORC
Upload from HDFS
Select from local
Choose File
u.data
Table name
ratings
Contains endlines?
movie_id
rating
rating_time
INT
INT
INT
242
3
881250949
302
3
891717742
377
1
878887116
51
2
880606923
346
1
886397596
474
4
884182806

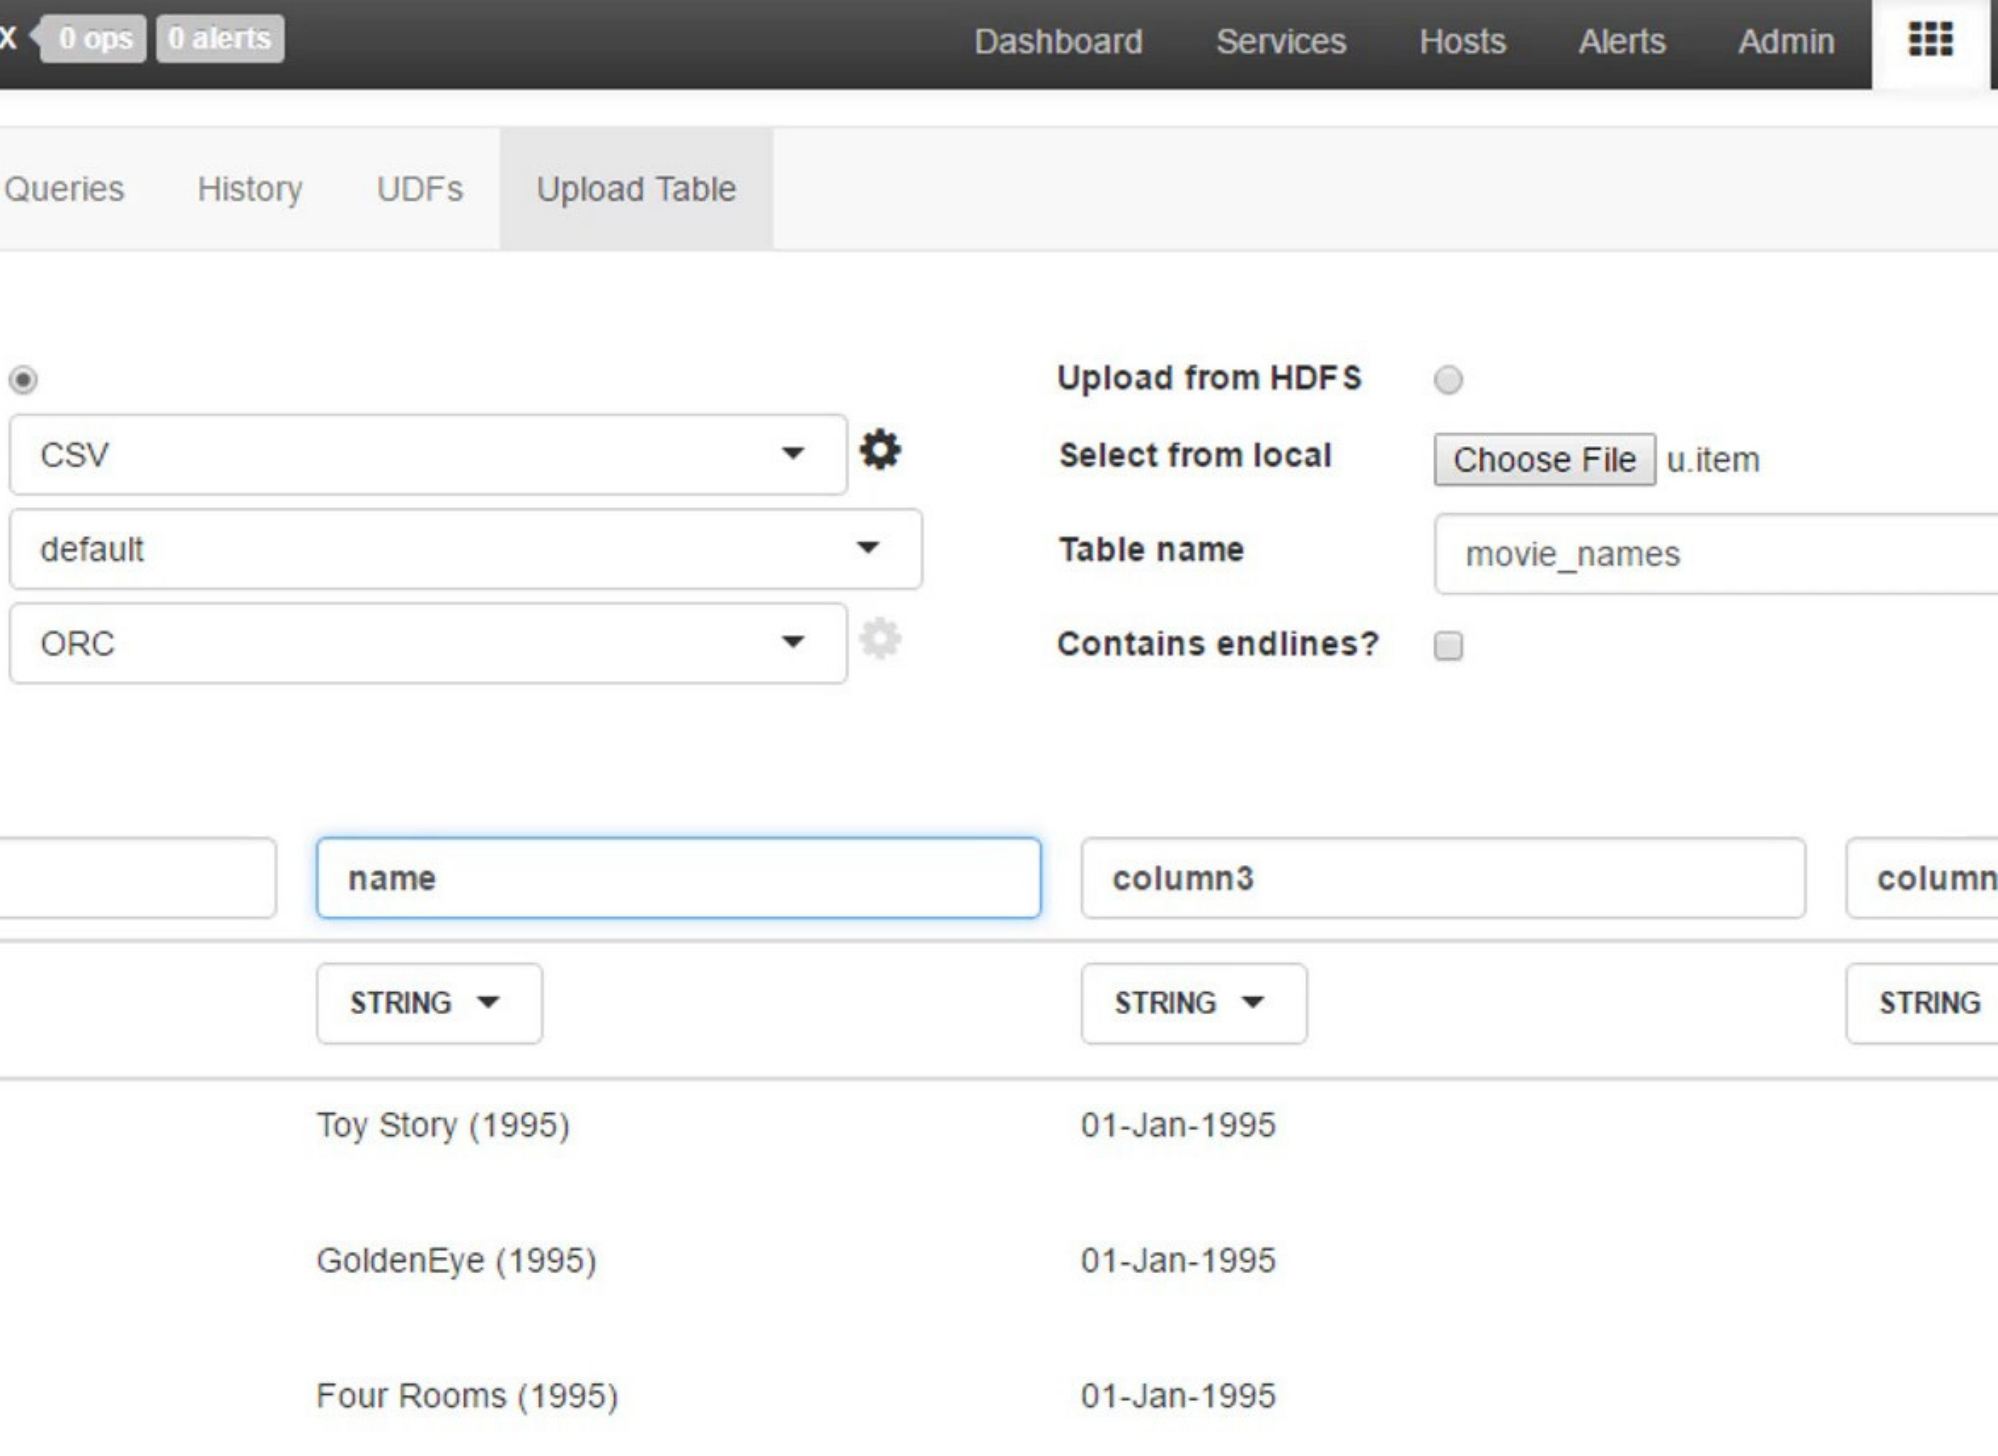
0 ops
0 alerts
Dashboard
Services
Hosts
Alerts
Admin
Queries
History
UDFs
Upload Table
Upload from HDFS
CSV
Select from local
Choose File
u.item
default
Table name
movie_names
ORC
Contains endlines?
name
column3
column
STRING
STRING
STRING
Toy Story (1995)
01-Jan-1995
GoldenEye (1995)
01-Jan-1995
Four Rooms (1995)
01-Jan-1995

Worksheet

```
1 SELECT movie_id, count(movie_id) as ratingCount
2 FROM ratings
3 GROUP BY movie_id
4 ORDER BY ratingCount
5 DESC;
```

Execute | Explain | Save as... | New Worksheet

Query Process Results (Status: SUCCEEDED) Save results...

Logs | Results

Filter columns... previous next

movie_id	ratingcount
50	583
258	509
100	508
181	507

```
movie_id, count(movie_id) as ratingCount
tings
Y movie_id
Y ratingCount
```

xplain | Save as... | New Worksh

DATA VISUALIZER

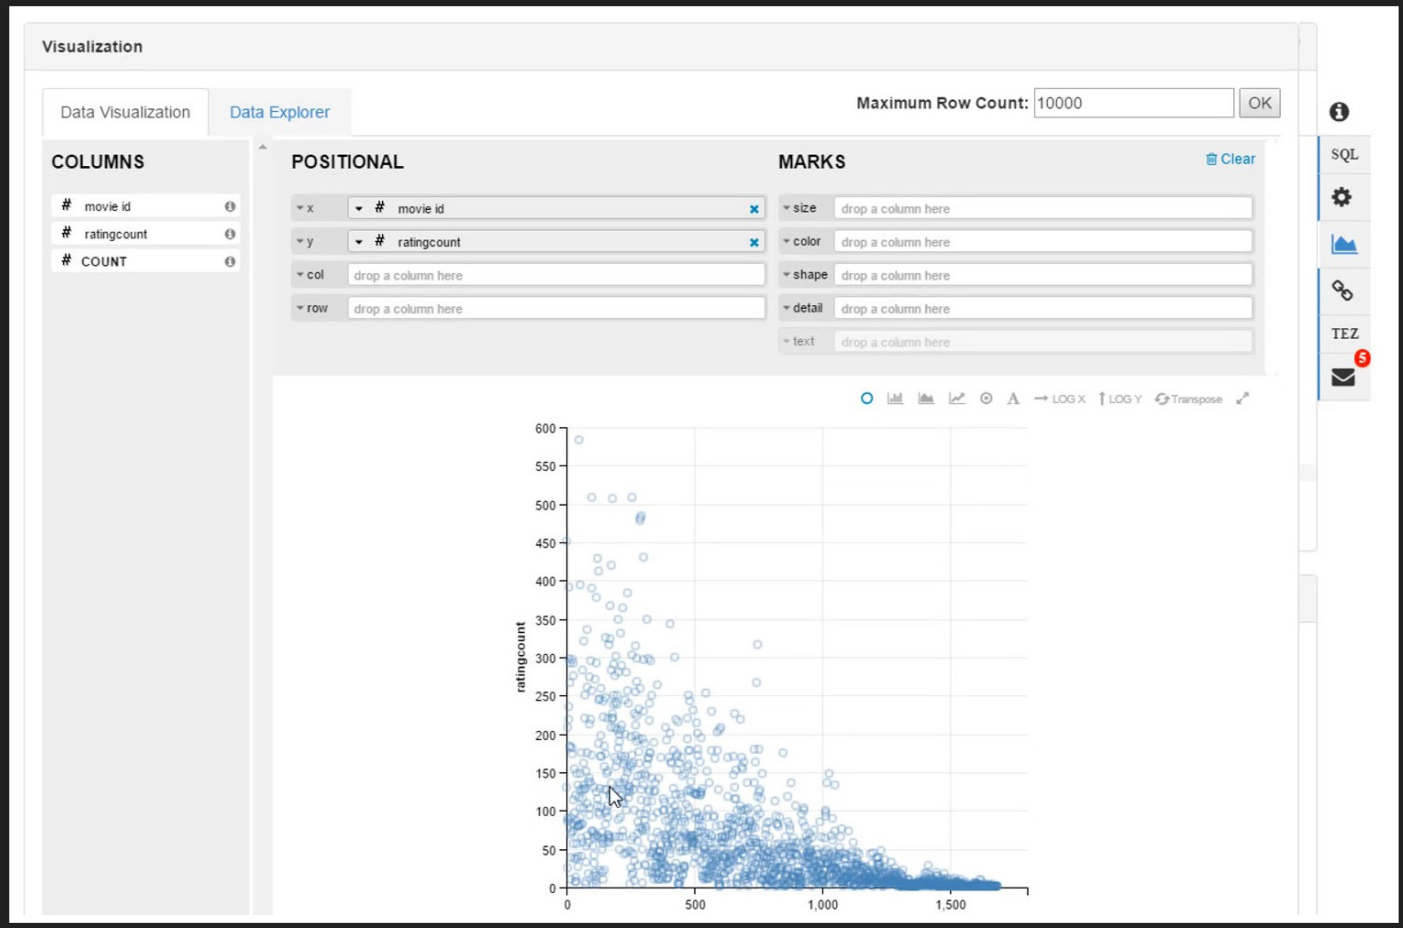

STAR WARS

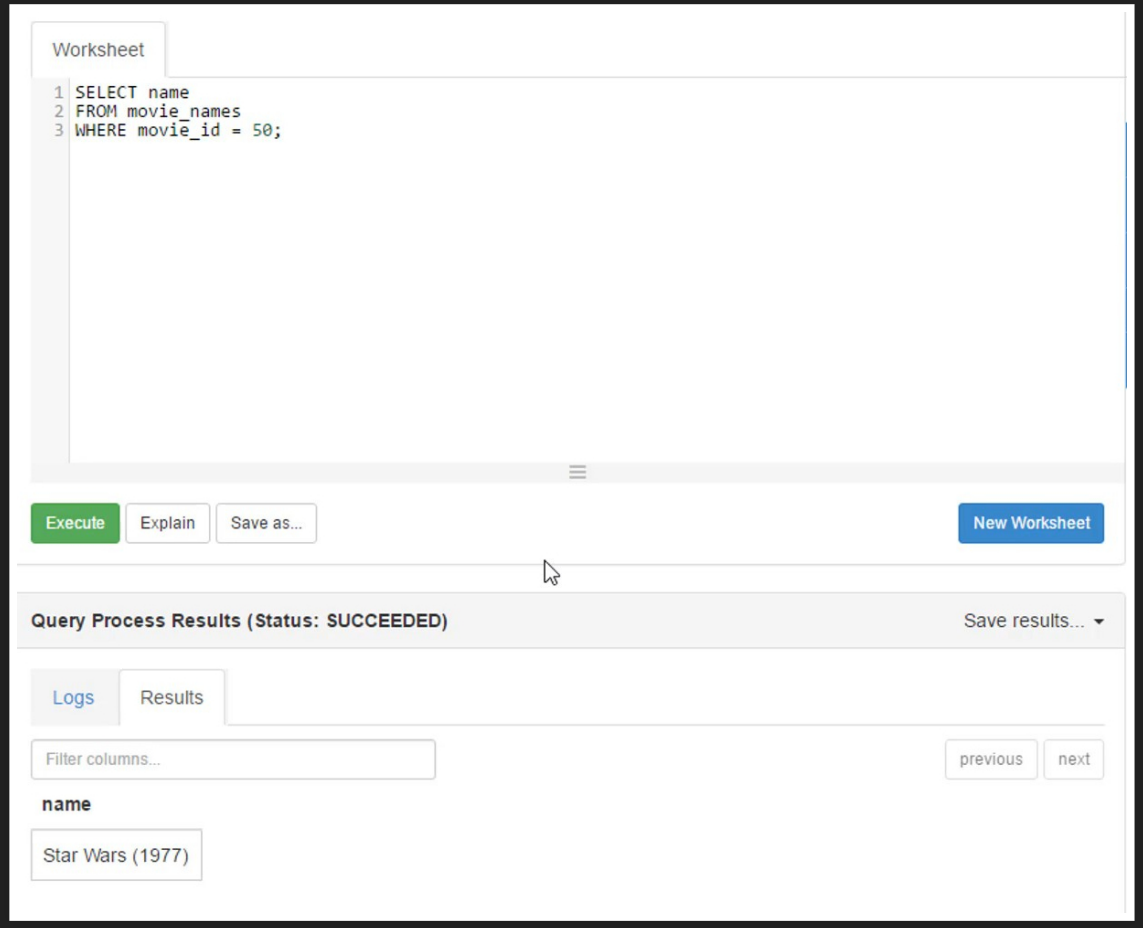
Worksheet
1 SELECT name
2 FROM movie_names
3 WHERE movie_id = 50;
Execute
Explain
Save as...
New Worksheet
Query Process Results (Status: SUCCEEDED)
Save results...
Logs
Results
Filter columns...
previous
next
name
Star Wars (1977)


hadoop
YARN

Complicated stuff, we don't have to worry about the details

Complicated stuff, we don't have to worry about the details

Hadoop V2 separated from MapReduce

Complicated stuff, we don't have to worry about the details

Hadoop V2 separated from MapReduce

Run MapReduce alternatives on it (TEZ)

Complicated stuff, we don't have to worry about the details

Hadoop V2 separated from MapReduce

Run MapReduce alternatives on it (TEZ)

Idea? split the computation across the cluster

Complicated stuff, we don't have to worry about the details

Hadoop V2 separated from MapReduce

Run MapReduce alternatives on it (TEZ)

Idea? split the computation across the cluster

Maintain data locality (Integrated with HDFS, where data lives)

YARN

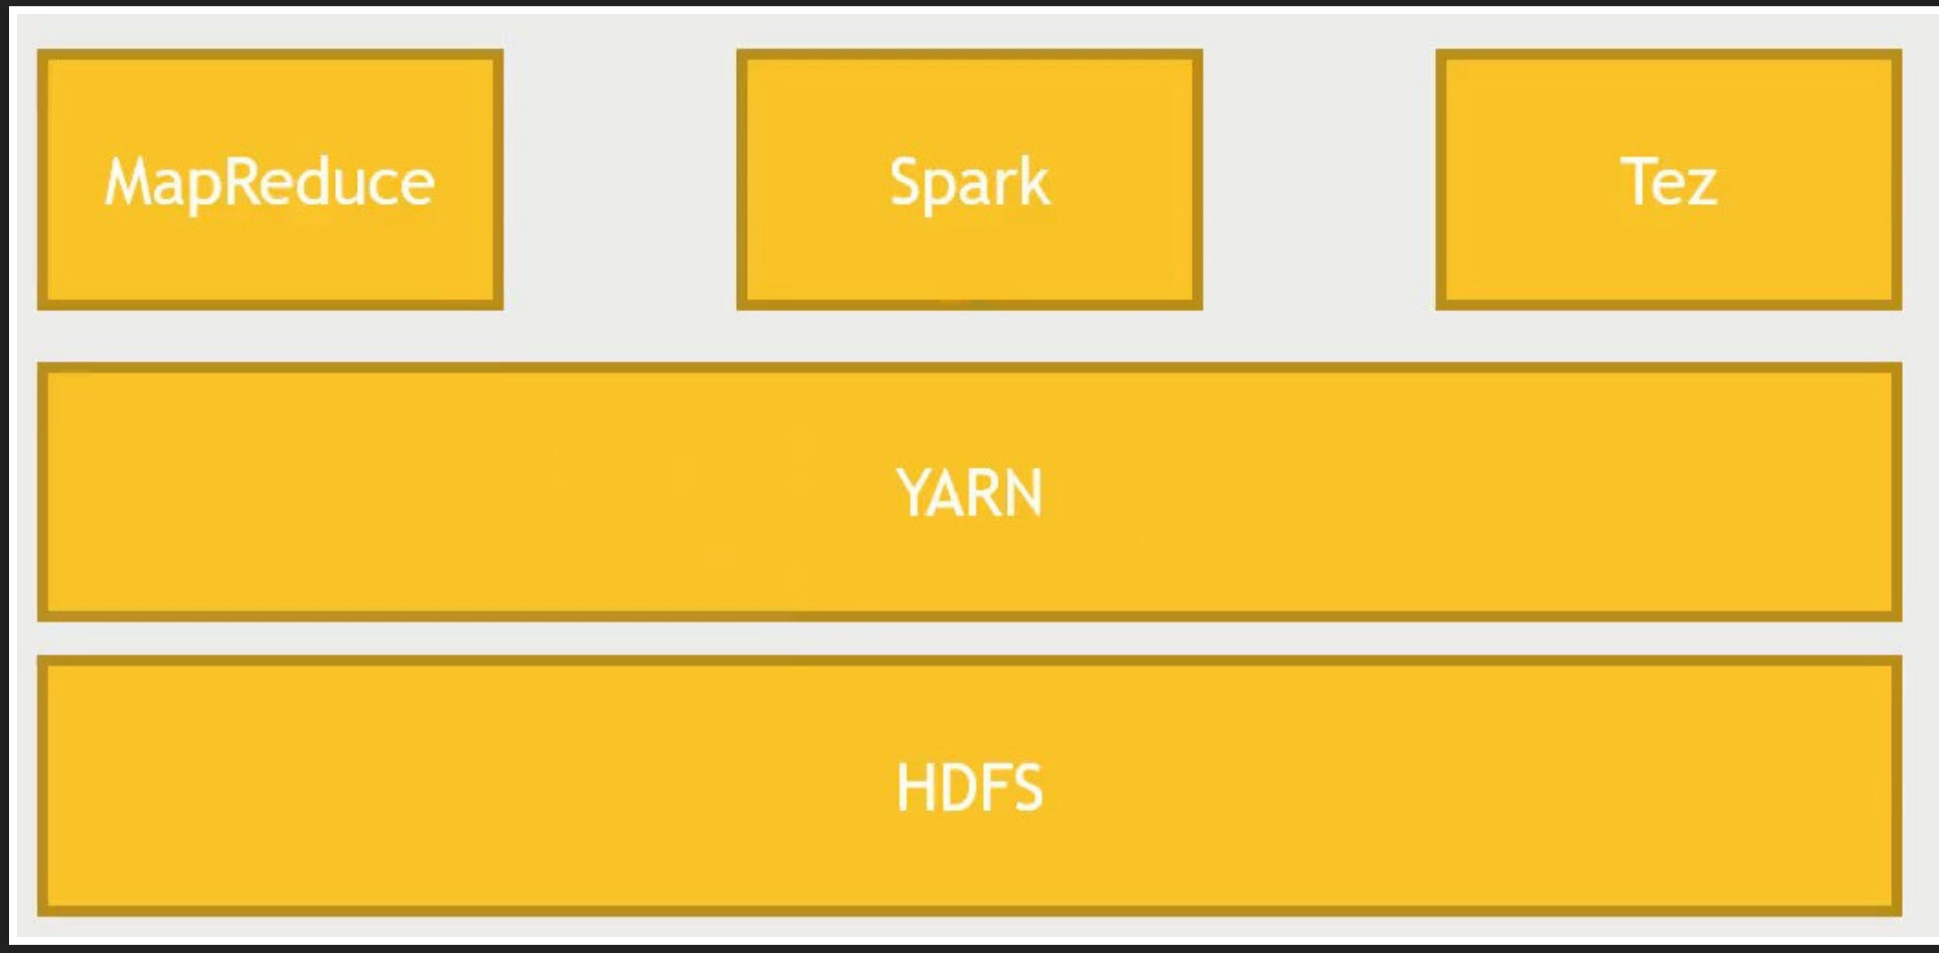

YARN

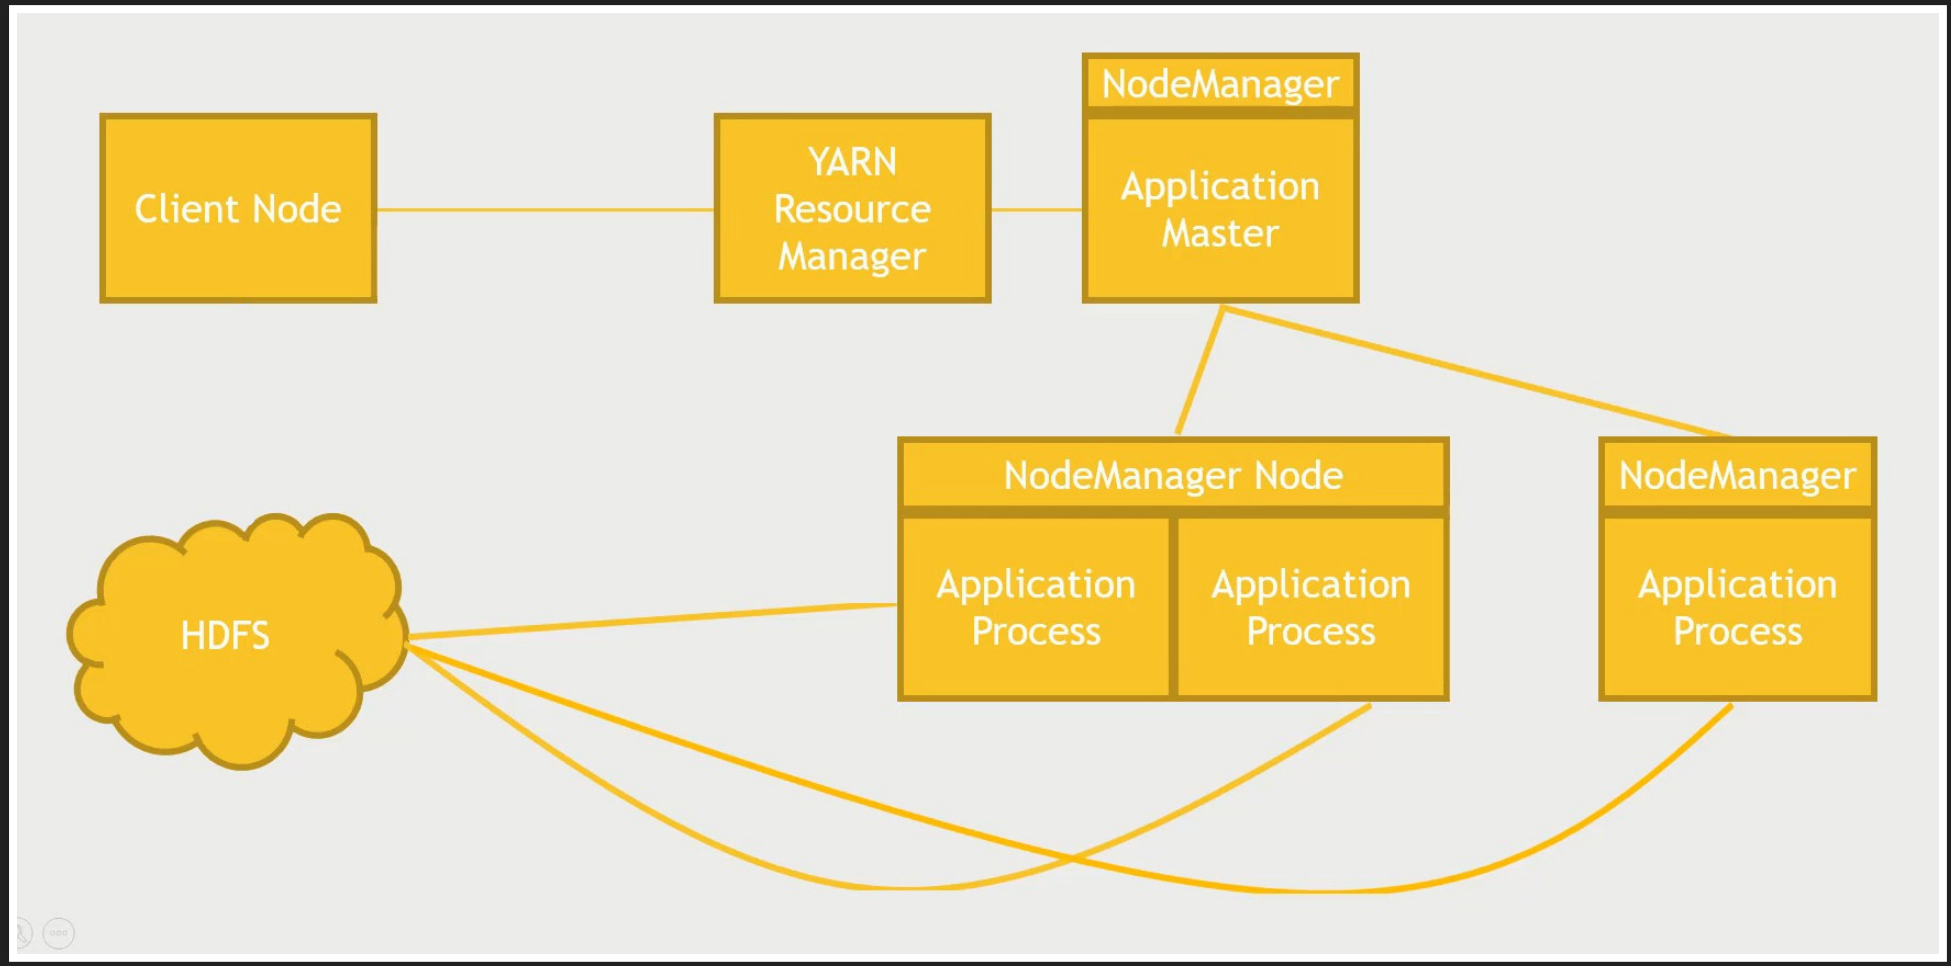
Client Node
YARN Resource Manager
NodeManager
Application Master
NodeManager Node
Application Process
Application Process
NodeManager
Application Process
HDFS

YARN

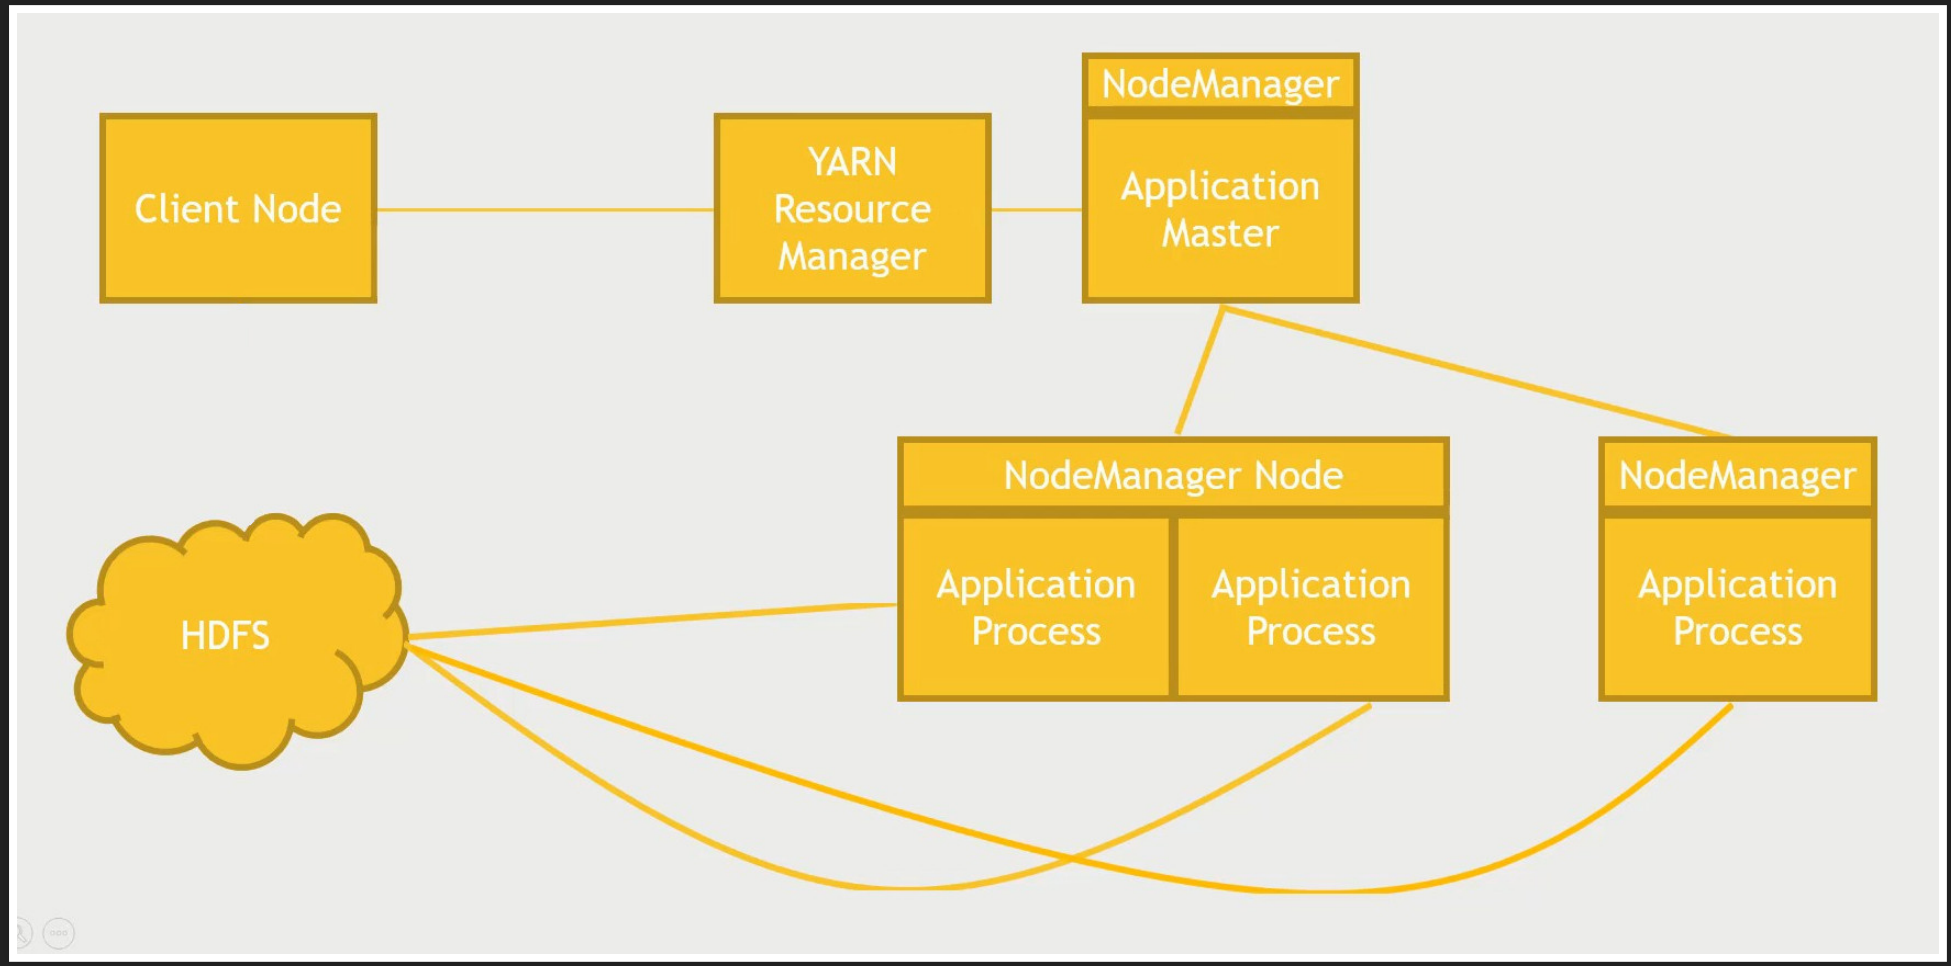

We can have multiple resource manager

TEZ (STILL YARN)

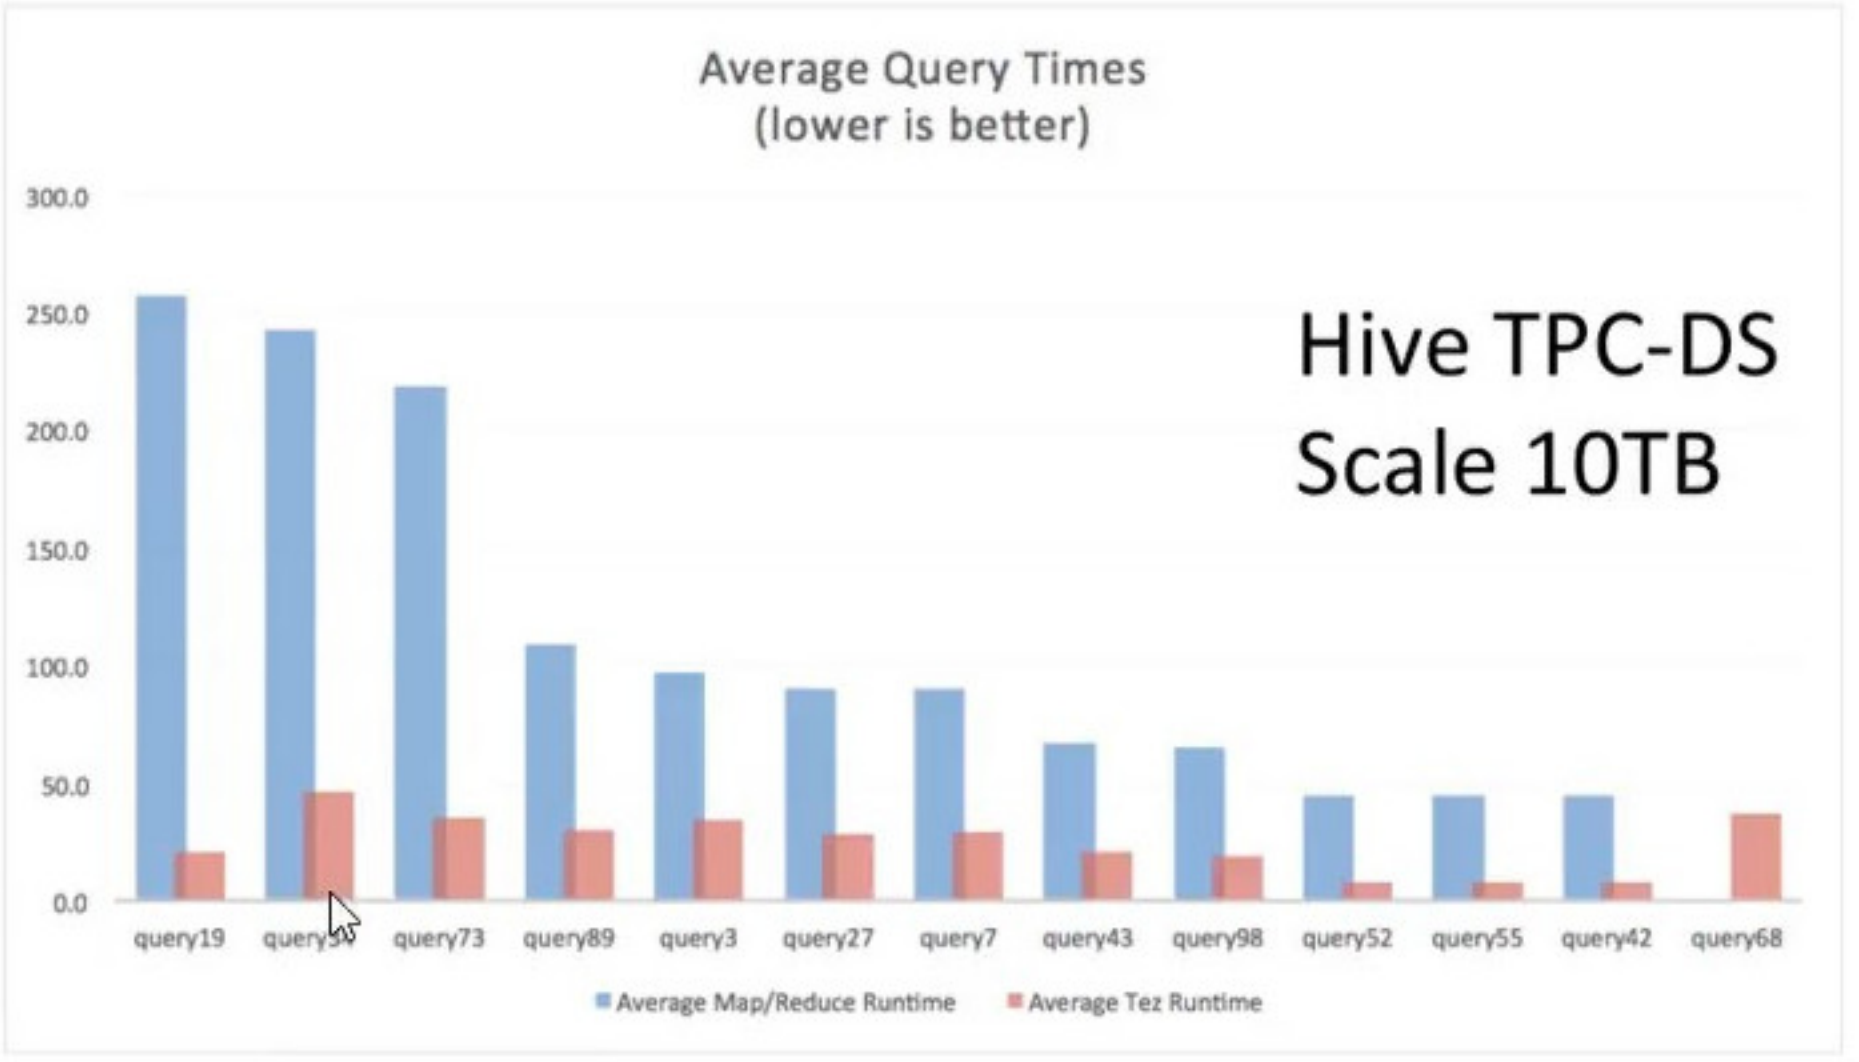


hadoop
HDFS

Allow our big data to be stored across entire cluster in distributed and reliable manner.

Allow our big data to be stored across entire cluster in distributed and reliable manner.

Handling large files

Allow our big data to be stored across entire cluster in distributed and reliable manner.

Handling large files

Breaking data into blocks - 128 MB

Allow our big data to be stored across entire cluster in distributed and reliable manner.

Handling large files

Breaking data into blocks - 128 MB

Keeps multiple copies of these blocks (Clever way)

Allow our big data to be stored across entire cluster in distributed and reliable manner.

Handling large files

Breaking data into blocks - 128 MB

Keeps multiple copies of these blocks (Clever way)

Allow us to use regular computers (No special hardware needed)

HDFS

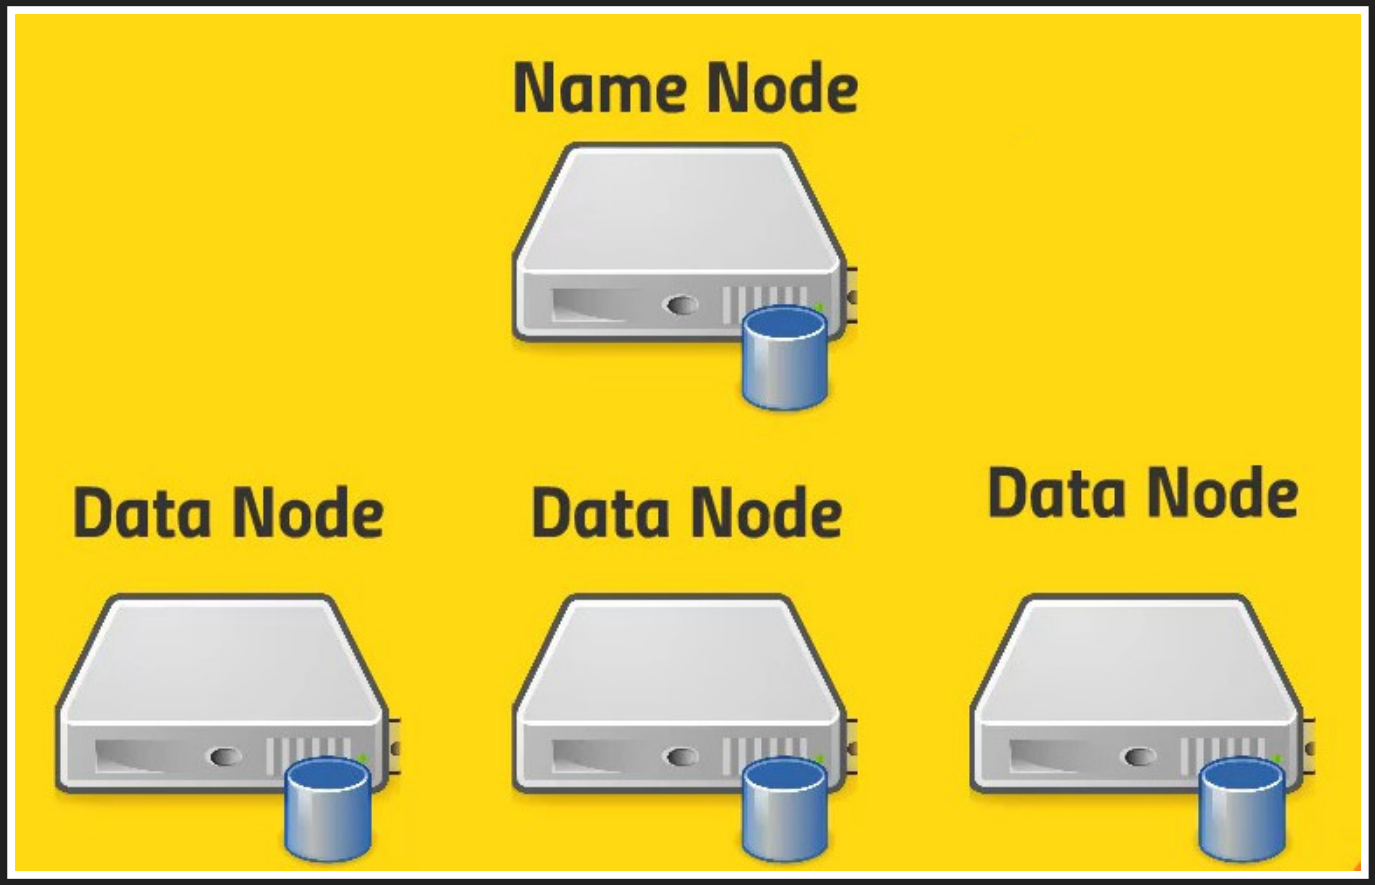

HA HDFS

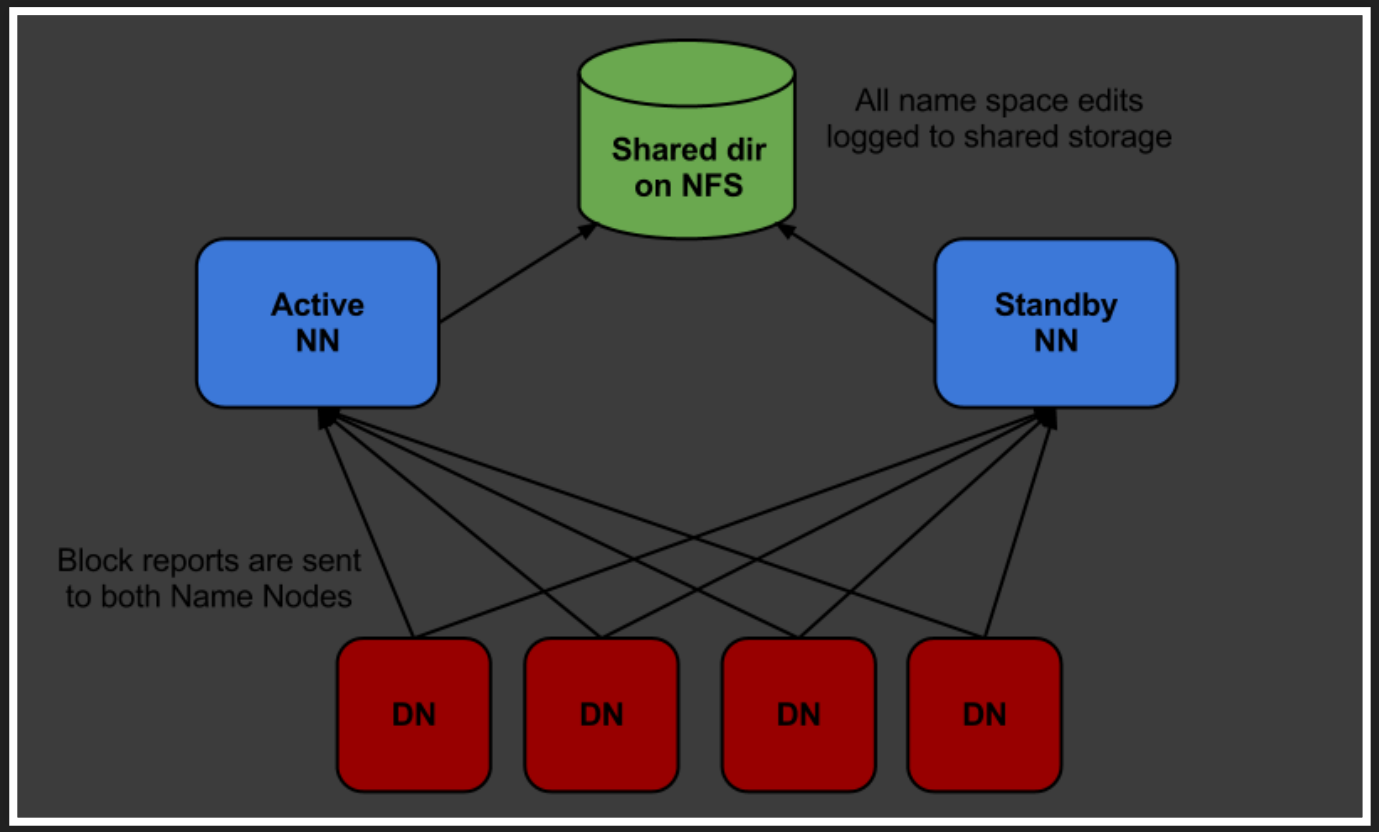

Client <--- > Zookeeper

One namenode is active at a time

HDFS FEDERATION

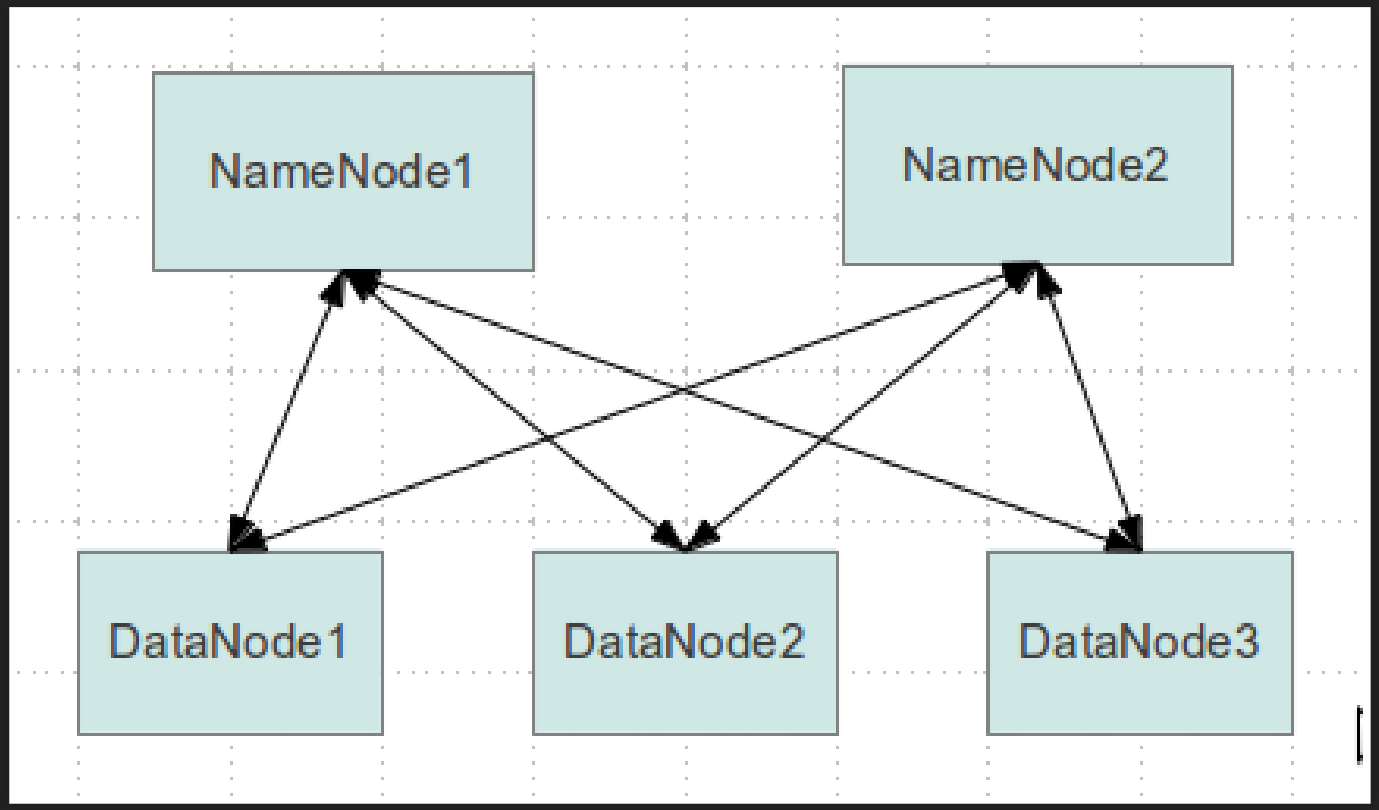

Sub Directories -> namespace Volume -> each namenode manage one namespace volume

HDFS FEDERATION

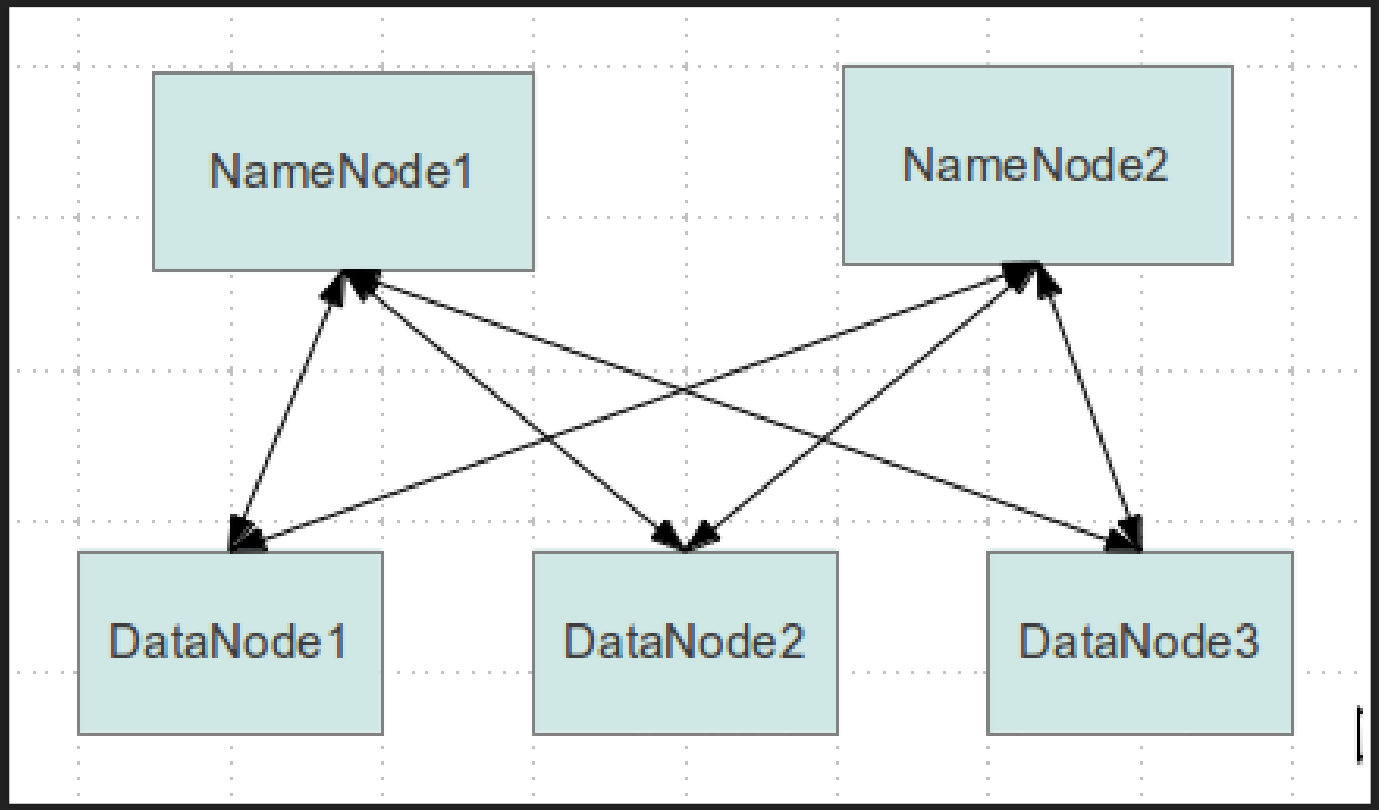

Sub Directories -> namespace Volume -> each namenode manage one namespace volume

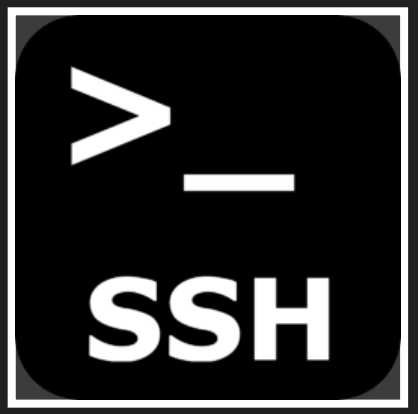
>_
SSH


hadoop
Map Reduce

Map: Transfer Data that we care about

Map: Transfer Data that we care about

Reduce: Aggregate Data

Map: Transfer Data that we care about

Shuffle and sort

Reduce: Aggregate Data

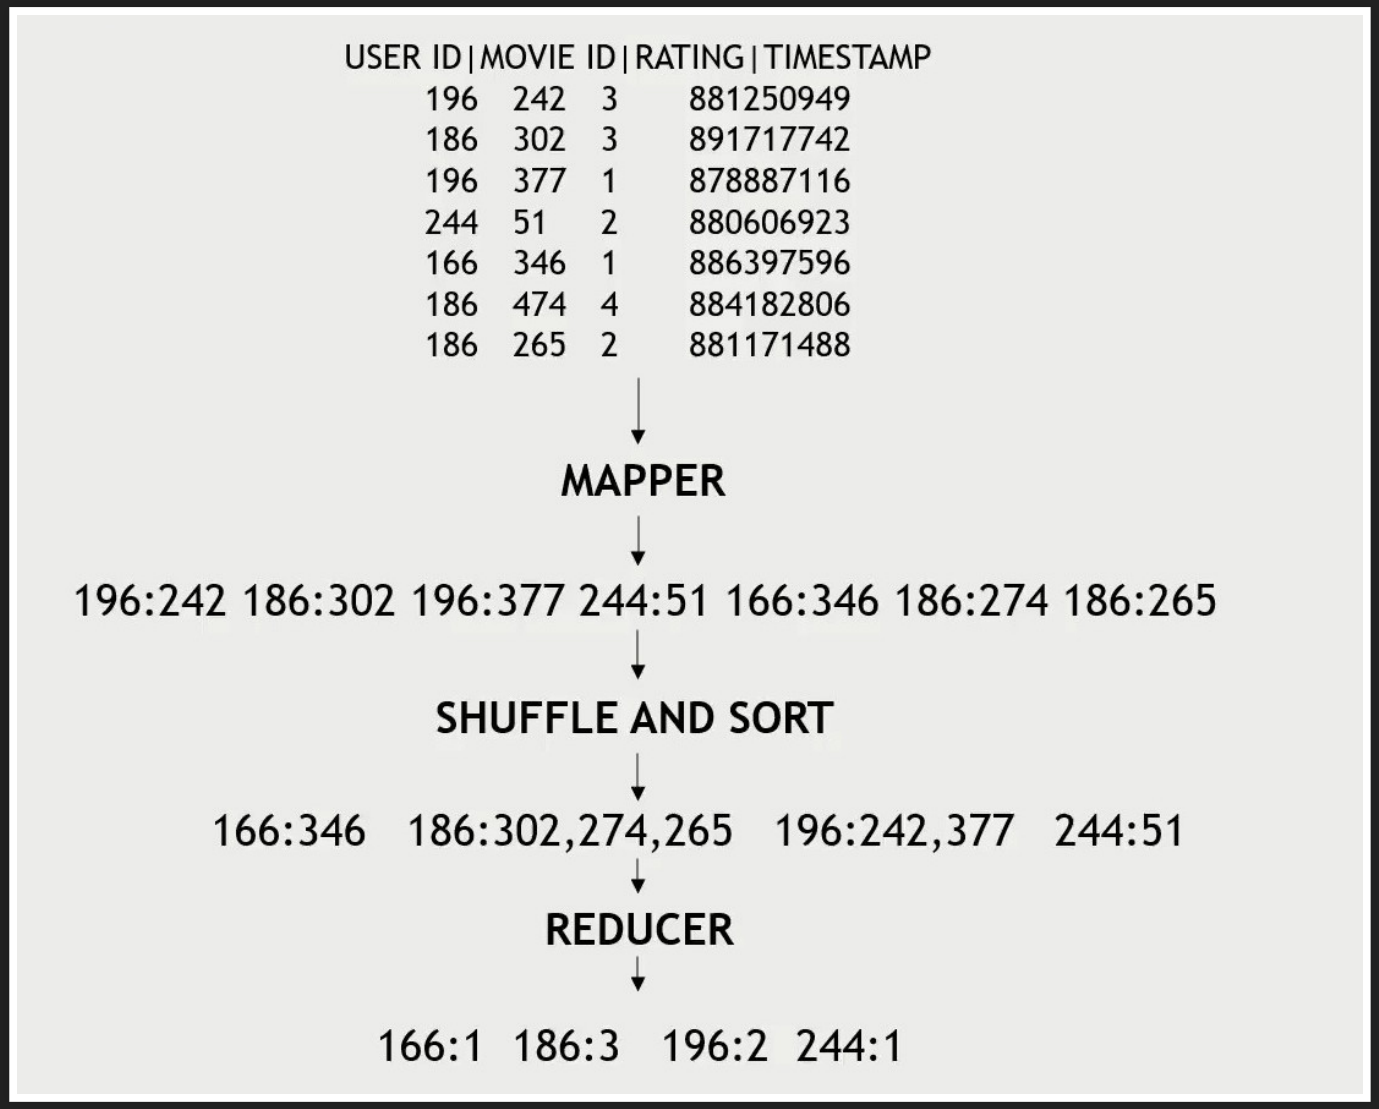
USER ID | MOVIE ID | RATING | TIMESTAMP
196 242 3 881250949
186 302 3 891717742
196 377 1 878887116
244 51 2 880606923
166 346 1 886397596
186 474 4 884182806
186 265 2 881171488
MAPPER
196:242 186:302 196:377 244:51 166:346 186:274 186:265
SHUFFLE AND SORT
166:346 186:302,274,265 196:242,377 244:51
REDUCER
166:1 186:3 196:2 244:1

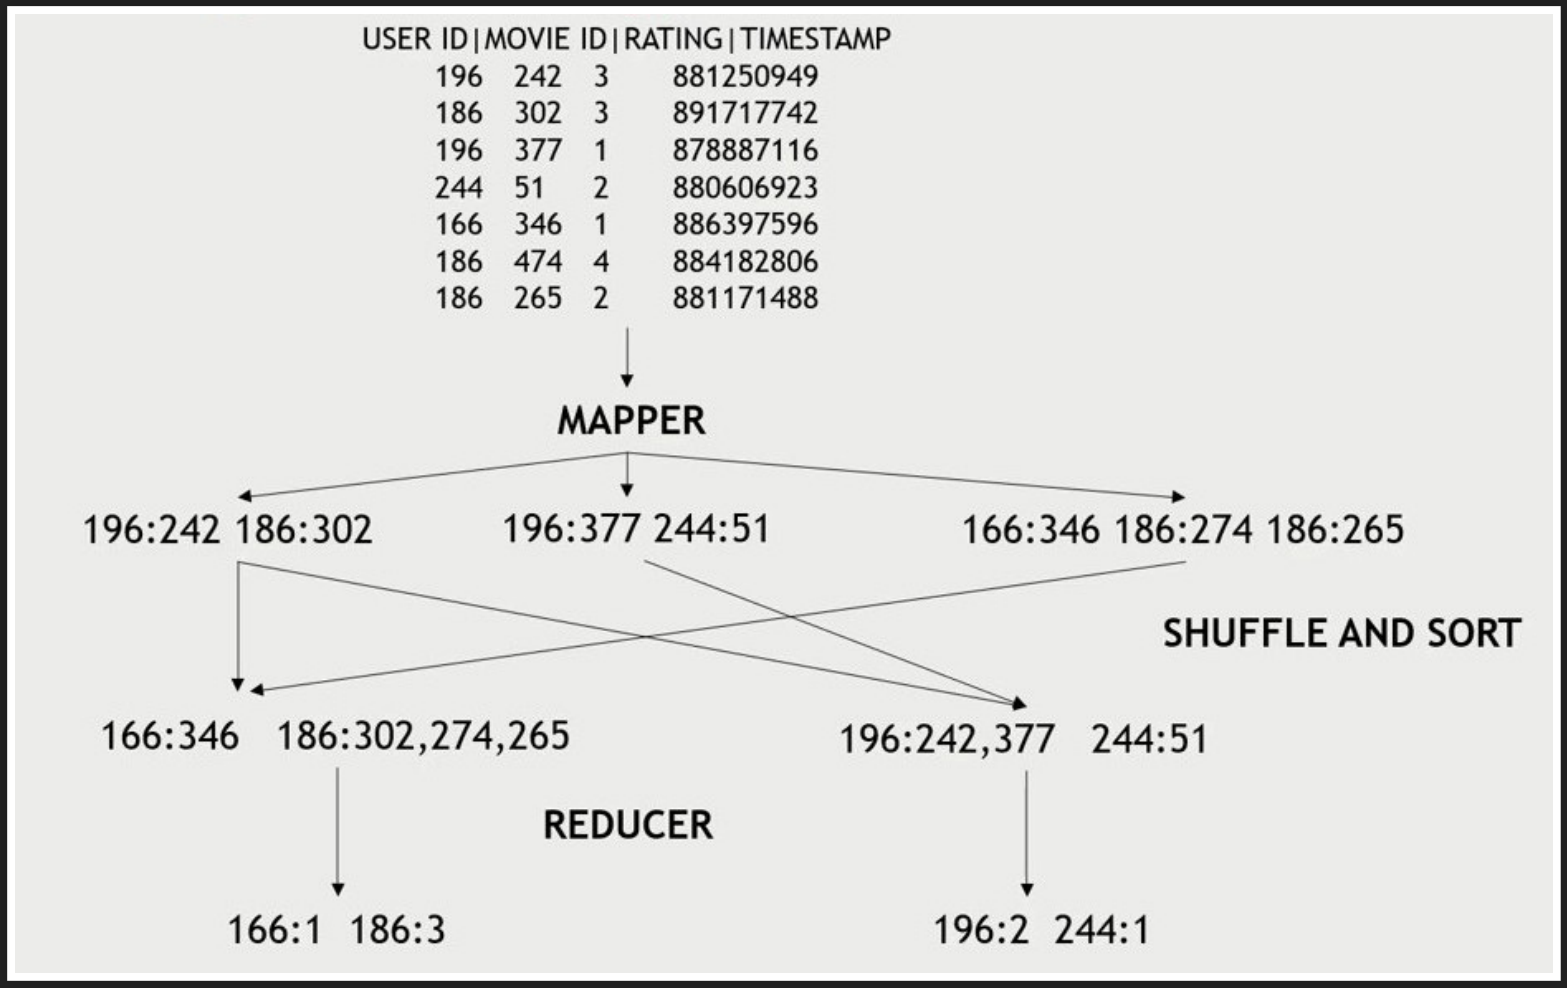
USER ID | MOVIE ID | RATING | TIMESTAMP
196 242 3 881250949
186 302 3 891717742
196 377 1 878887116
244 51 2 880606923
166 346 1 886397596
186 474 4 884182806
186 265 2 881171488
MAPPER
196:242 186:302
196:377 244:51
166:346 186:274 186:265
SHUFFLE AND SORT
166:346
186:302,274,265
196:242,377
244:51
REDUCER
166:1
186:3
196:2
244:1

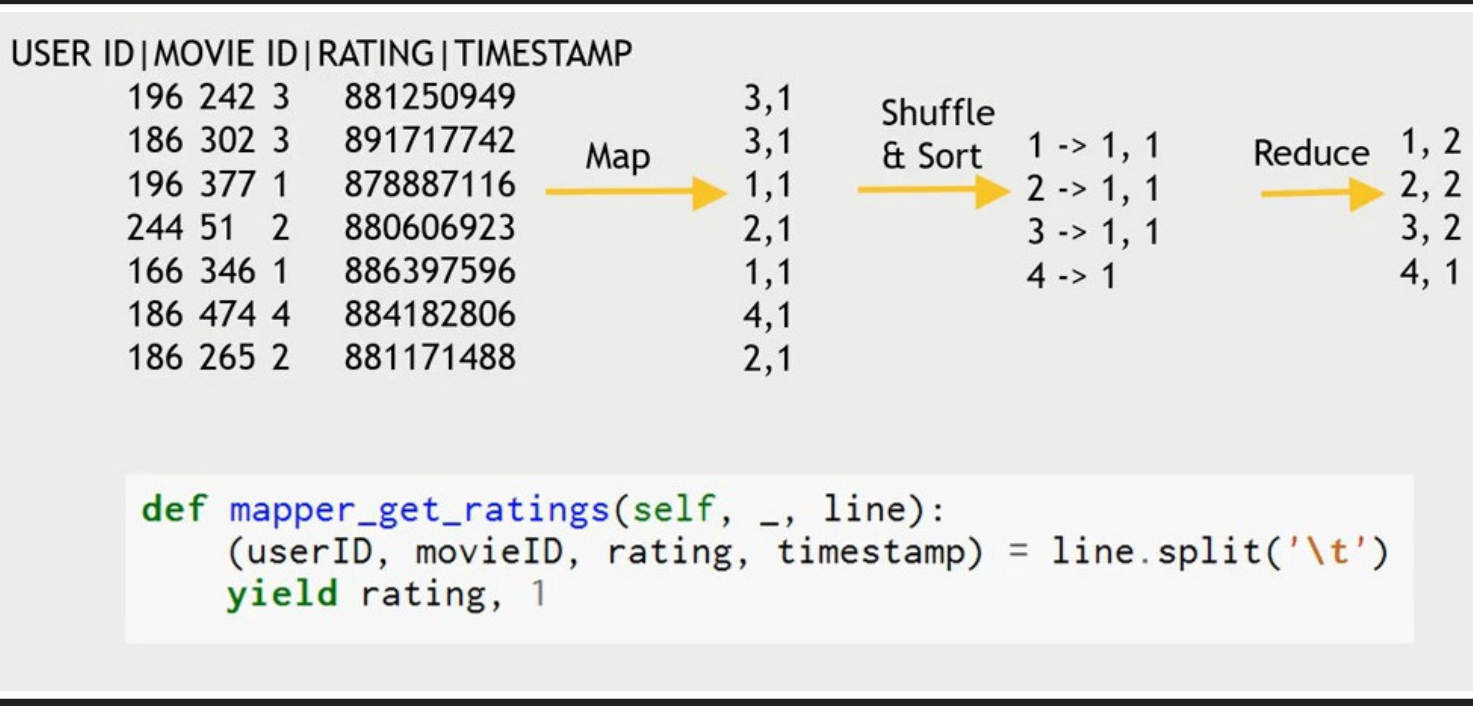


```
def mapper_get_ratings(self, _, line):
    (userID, movieID, rating, timestamp) = line.split('\t')
    yield rating, 1
```

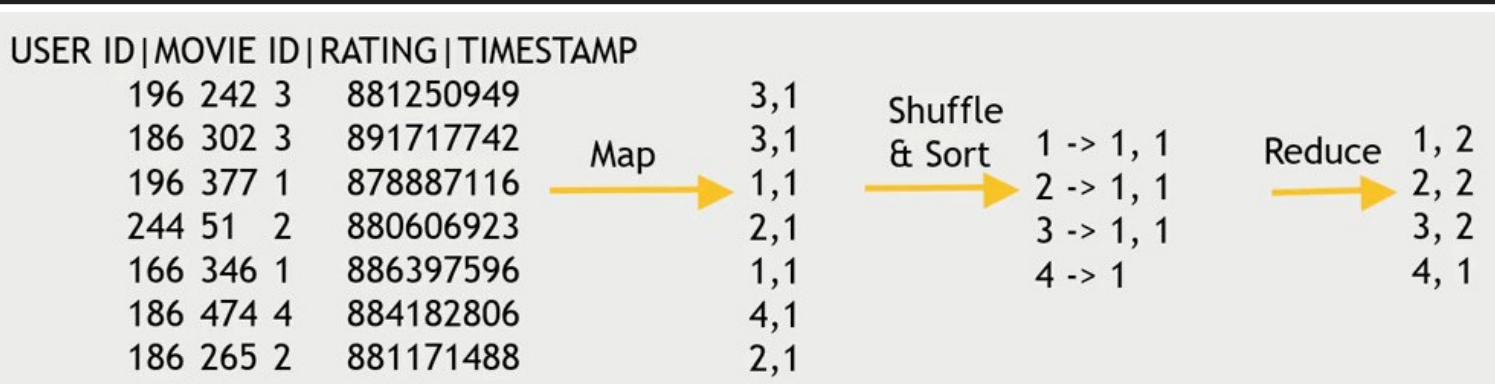


```
def reducer_count_ratings(self, key, values):
    yield key, sum(values)
```

```python
from mrjob.job import MRJob
from mrjob.step import MRStep

class RatingsBreakdown(MRJob):
    def steps(self):
        return [
            MRStep(mapper=self.mapper_get_ratings,
                   reducer=self.reducer_count_ratings)
        ]

    def mapper_get_ratings(self, _, line):
        (userID, movieID, rating, timestamp) = line.split('\t')
        yield rating, 1

    def reducer_count_ratings(self, key, values):
        yield key, sum(values)

if __name__ == '__main__':
    RatingsBreakdown.run()
```

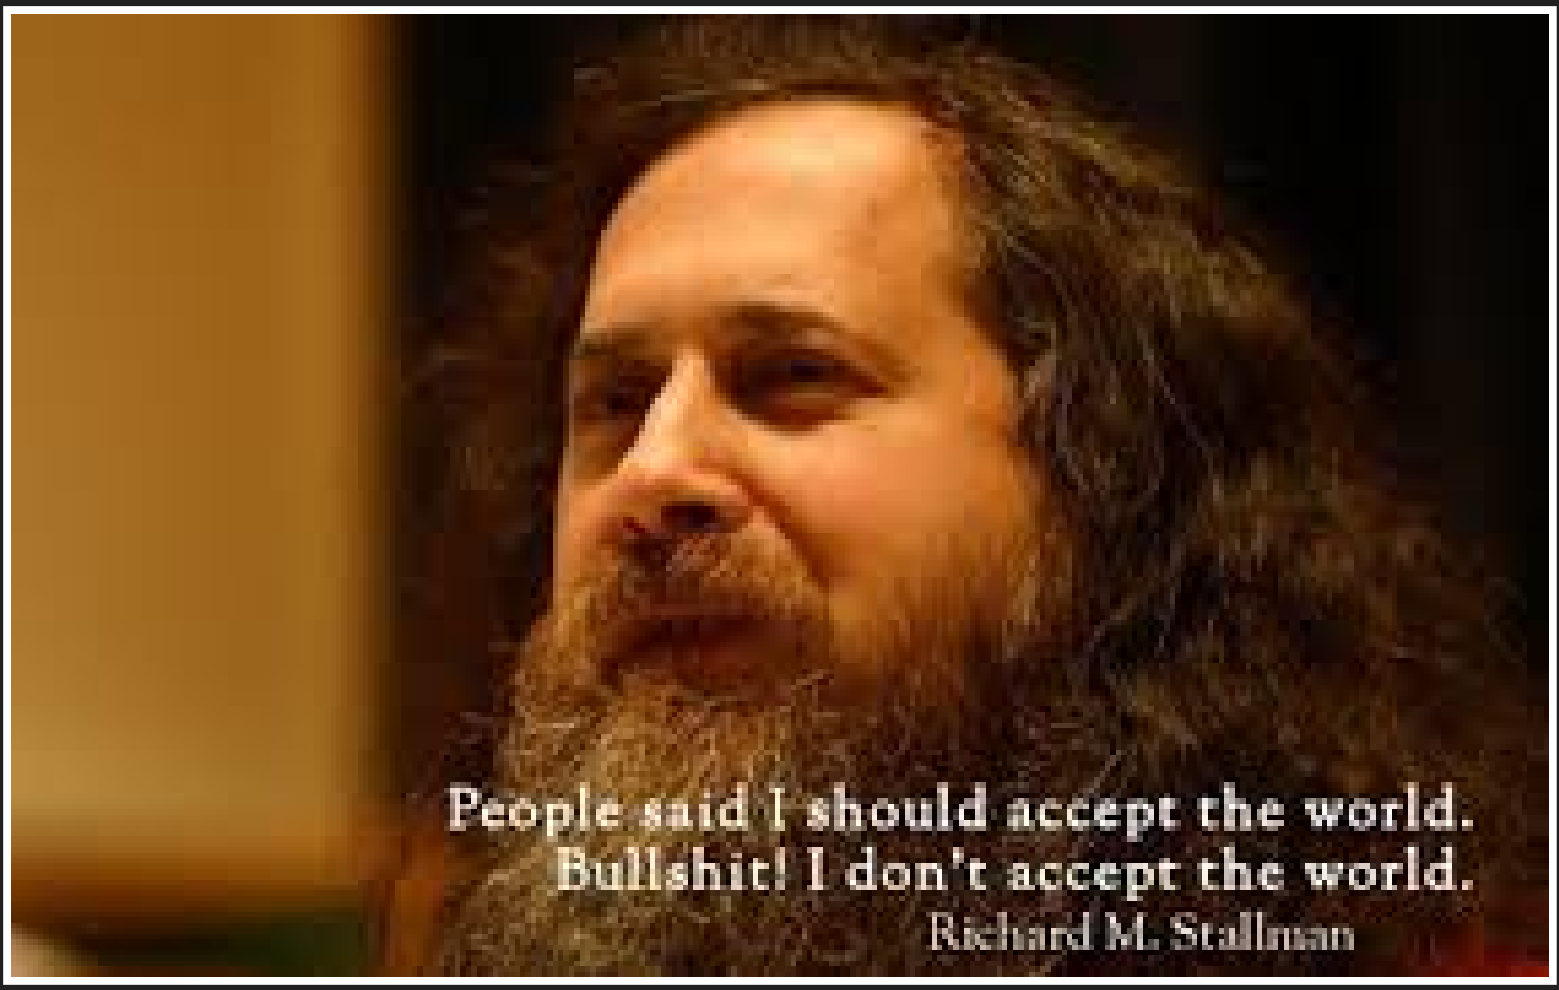
People said I should accept the world.
Bullshit! I don't accept the world.
Richard M. Stallman