

BIG DATA AND THE HADOOP FRAMEWORK

By
Revanth Reddy

Contents

- DATA
- BIGDATA
- HDFS
- MAPREDUCE
- PIG
- HIVE

History of Big data-

- 1941- “ Information Explosion”

Seventy years ago there was an encounter to quantify the growth rate in the *volume of Data* or what has popularly known as the “**Information Explosion**” (a term first used in 1941, according to the Oxford English Dictionary) .


- 1975- “Ministry of posts and telecommunications in Japan”.

After tracking the volume of information circulating in Japan, they found that the information provided by the mass media(1-way) has become stagnant, and the information provided by other **personal communication media(2-way)** has drastically increased.

- 1997- Michael Cox and David Ellsworth

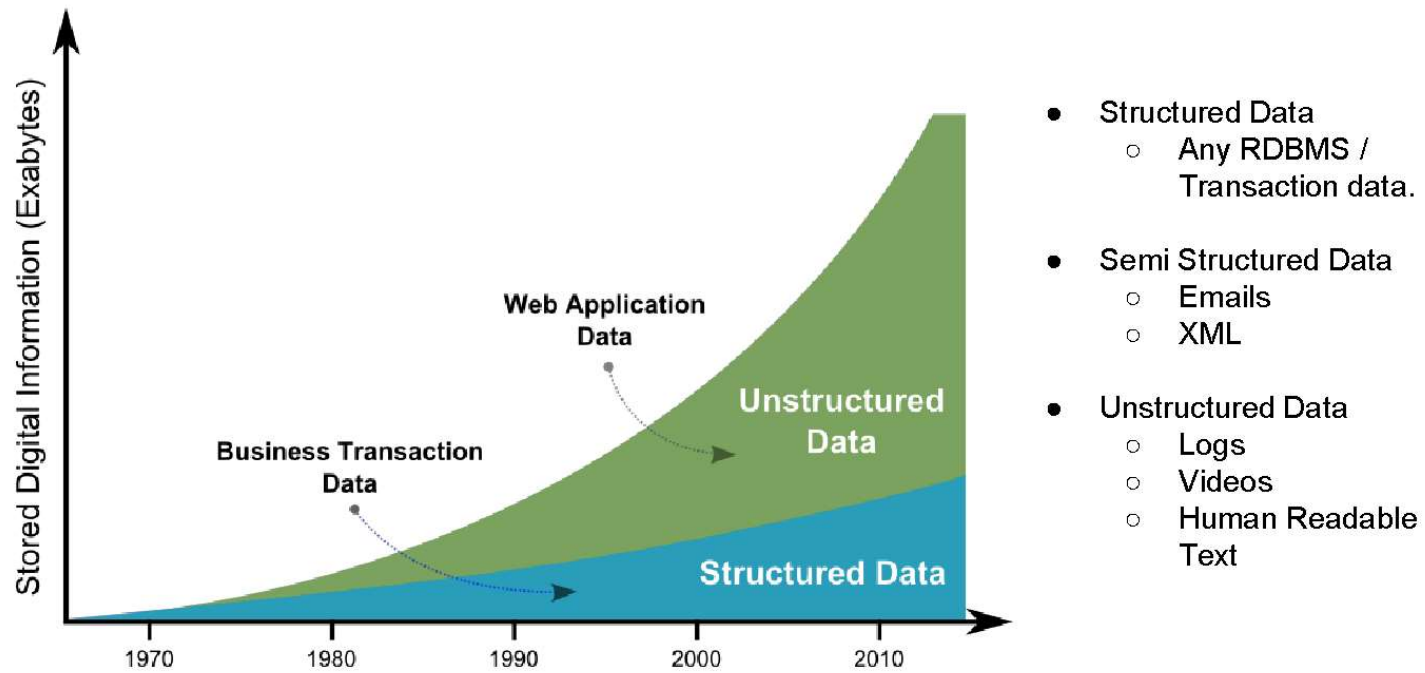
They published a paper which says that “when data sets do not fit in main memory or local disk most common solution is to acquire more resources. We call this the problem of **big data**”.(first article in ACM digital library to use the term big data.)

Twitter Example



- Twitter has over 500 million registered users.
- The USA, whose 141.8 million accounts represents 27.4 percent of all Twitter users, good enough to finish well ahead of Brazil, Japan, the UK and Indonesia.
- 79% of US Twitter users are more like to recommend brands they follow .
- 67% of US Twitter users are more likely to buy from brands they follow .
- 57% of all companies that use social media for business use Twitter

Structured vs Unstructured



What is Big Data ?

Big data is a phrase, used to describe a massive volume of both structured and unstructured data. In most enterprise scenarios the data is too big or it moves too fast or it exceeds current processing capacity.

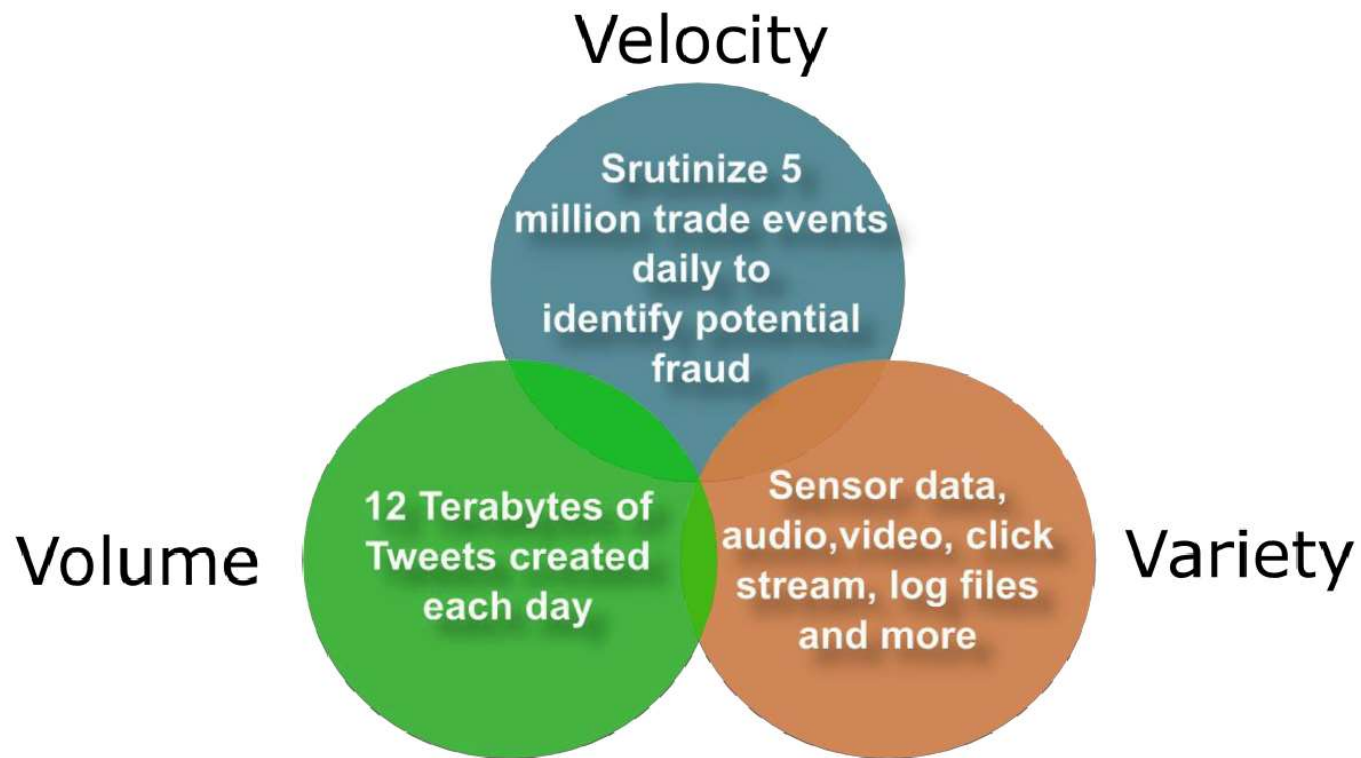
e.g Rolling web log data
Network & System Logs
Click information

What is considered "big data" varies depending on the capabilities of the organization managing the set, and on the capabilities of the applications that are traditionally used to process and analyze the data set in its domain




photo credit: [CarbonNYC](#) via [photopin cc](#)

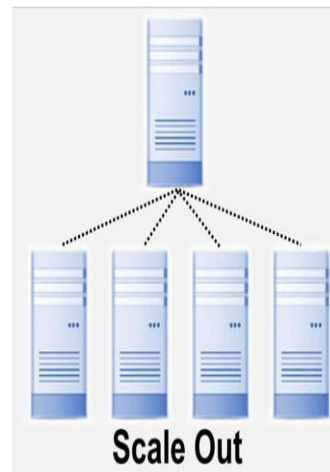
Big Data Dimensions



Scale Vertically vs Horizontally



VS



How to solve the
problem of Big data?

Why DFS ?

While Reading 1TB data ...



1 Machine
4 I/O Channel
Each Channel - 100MB/S

Why DFS ?

While Reading 1TB data ...



45 Minutes

1 Machine
4 I/O Channel
Each Channel - 100MB/S

Why DFS ?

While Reading 1TB data ...



45 Minutes



1 Machine
4 I/O Channel
Each Channel - 100MB/S



10 Machines
4 I/O Channel
Each Channel - 100MB/S

Why DFS ?

While Reading 1TB data ...



45 Minutes



4.5 Minutes

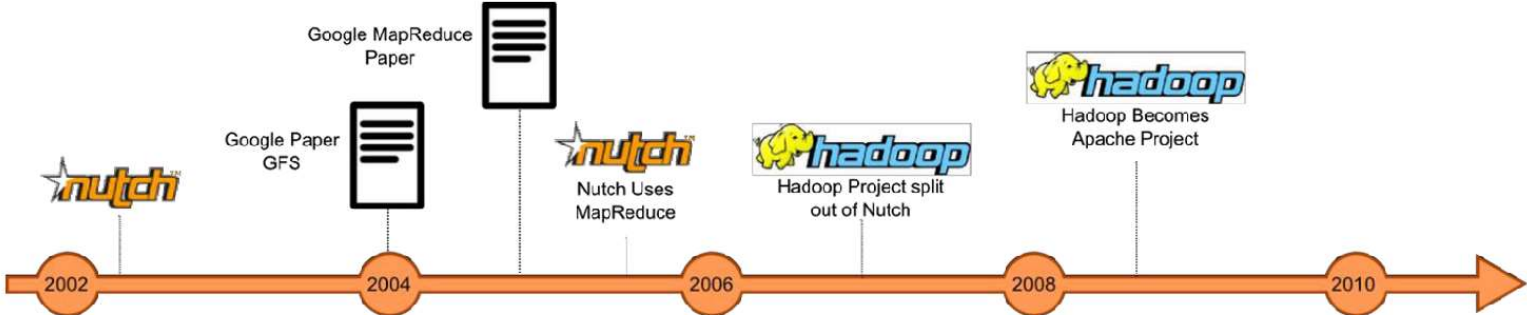


1 Machine
4 I/O Channel
Each Channel - 100MB/S

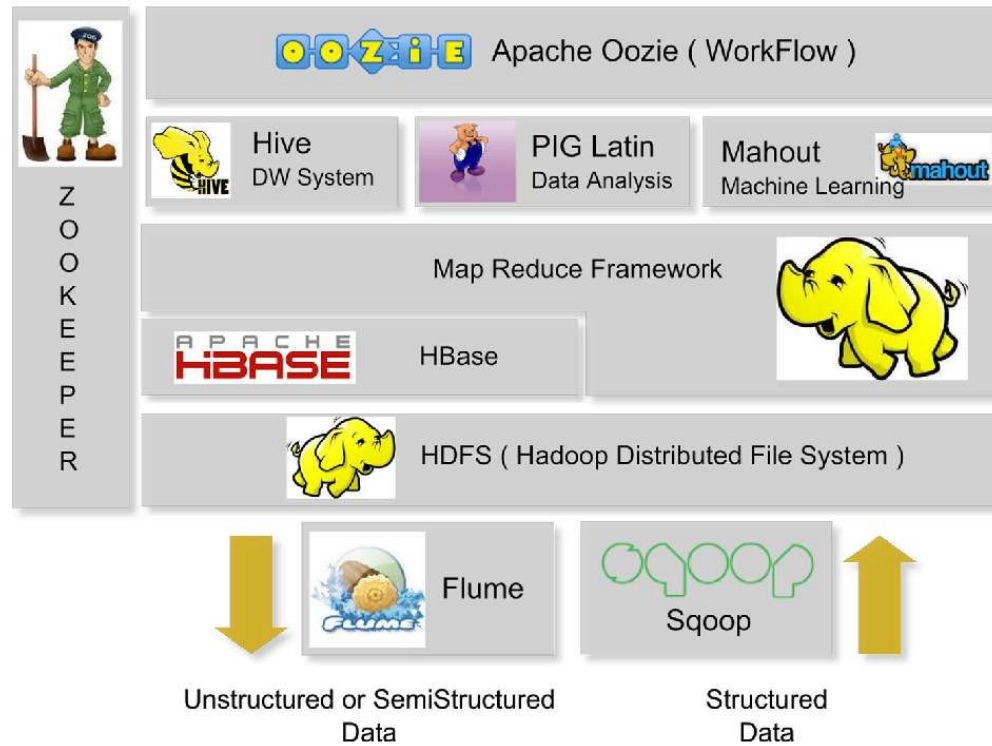


10 Machines
4 I/O Channel
Each Channel - 100MB/S

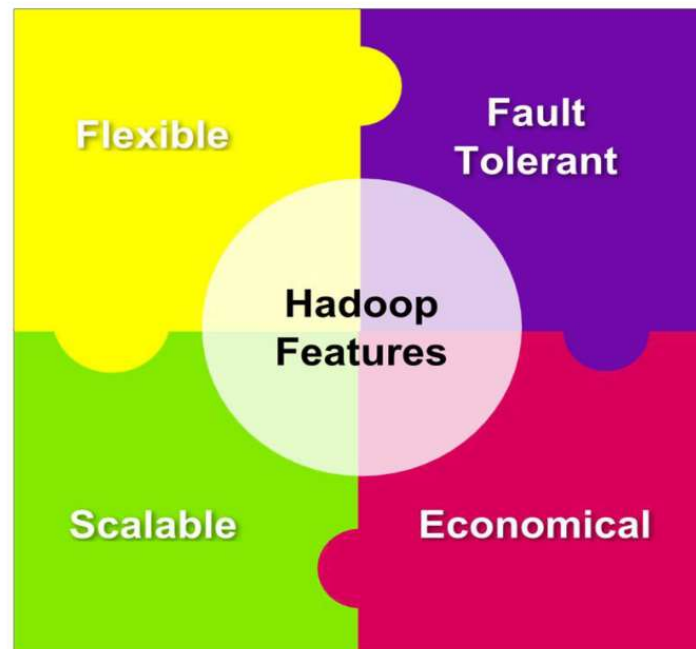
History of Hadoop



Hadoop EcoSystem

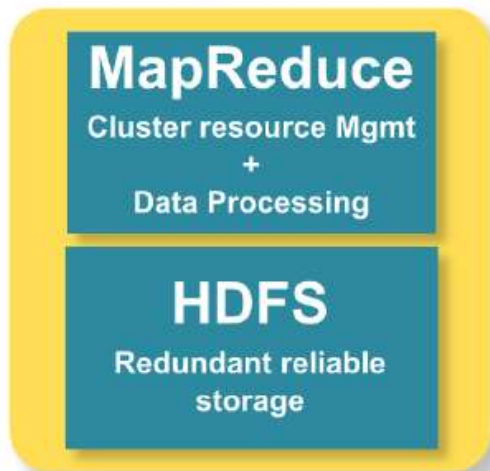


Hadoop Features





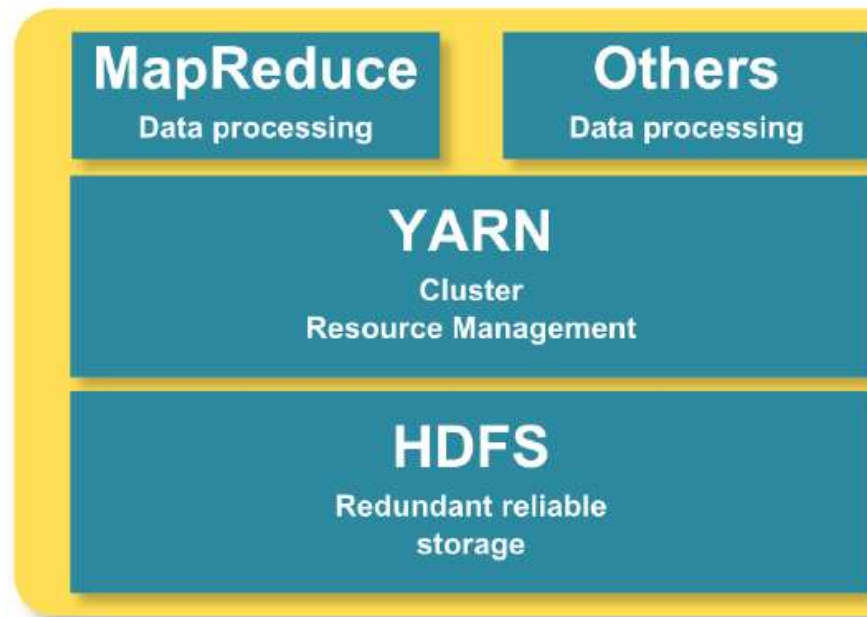
Hadoop1



Single Use System
Batch Apps



Hadoop2



Multi Purpose Platform
Batch, interactive, Streaming

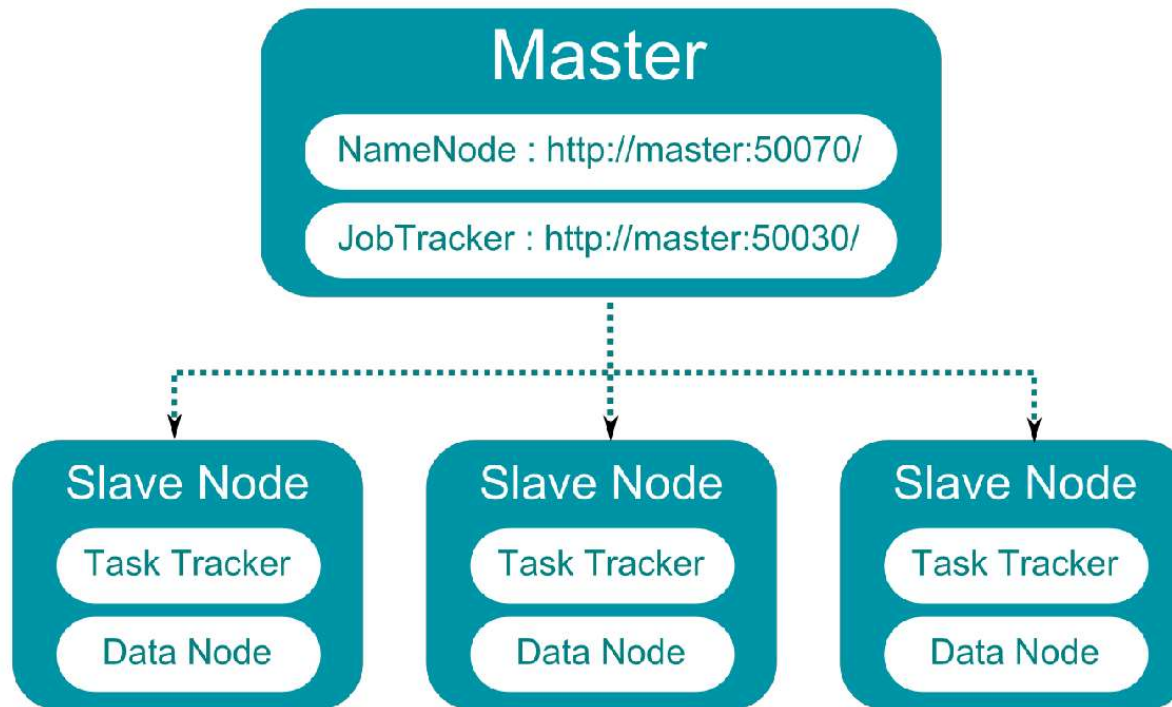
RDBMS vs Big Data

RDBMS

Big Data

RDBMS		Big Data
Structured	Data Types	Multi and Unstructured
Limited, No Data Processing	Processing	Processing coupled with Data
Standards & Structured	Governance	Loosely Structured
Required on write	Schema	Required on Read
Reads are Fast	Speed	Writes are Fast
Software License	Cost	Support only
Known Entity	Resources	Growing, complexities, wide
Interactive OLAP Analytics		Data discovery
Complex ACID Transactions		Processing unstructured data
Operational Data Store	Best Fit Use	Massive storage/processing

Sample Cluster Configuration



What is MapReduce

Provides Parallel Programming Model

Process data wherever it is stored

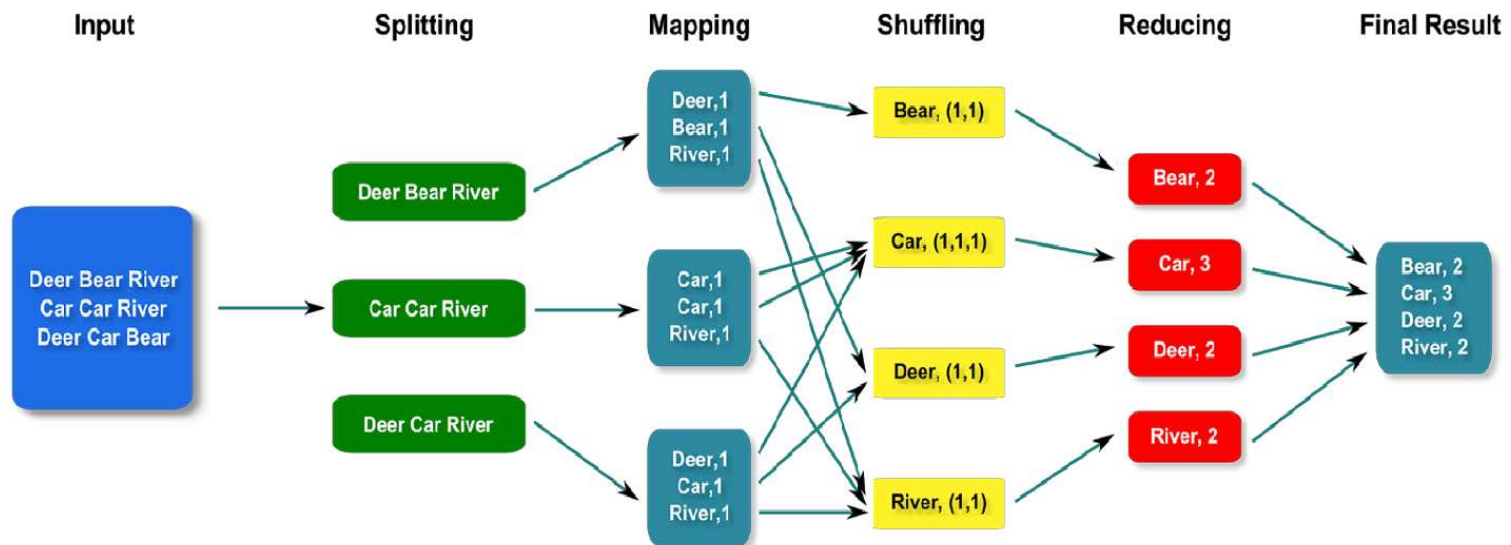
Fault Tolerant

Handles Scheduling


Status Reporting

MapReduce Paradigm

The overall MapReduce wordcount process



Need of Pig ?



Do you Know !!
10 Lines of PIG = 200 Lines of Java
+ Built in Operations like
Join, Group, Filter, Sort ...

Why Was Pig Created?


- An ad-hoc way of creating and executing MapReduce jobs on very large data sets
- Rapid Development
- No Java is required
- Developed by Yahoo!

What is PIG?



- Pig is an open-source high-level dataflow system
- It provides a simple language for queries and data manipulation Pig Latin, that is compiled into mapreduce jobs that are run on Hadoop
- Why is it important?
 - Companies like Yahoo, Google and Microsoft are collecting enormous data sets in the form of click streams, search logs, and web crawls.
 - Some form of ad-hoc processing and analysis of all of this information is required.

PIG vs MapReduce



MapReduce

- Powerful model for parallelism.
- Based on a rigid procedural structure.
- Provides a good opportunity to parallelize algorithm.
- Must think in terms of map and reduce functions
- More than likely will require Java programmers

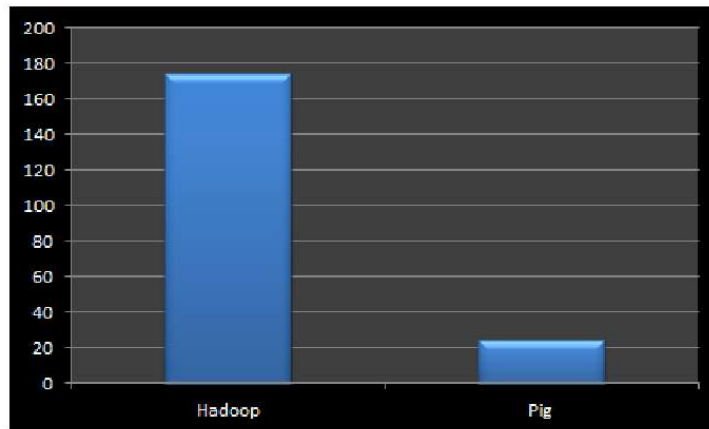


PIG

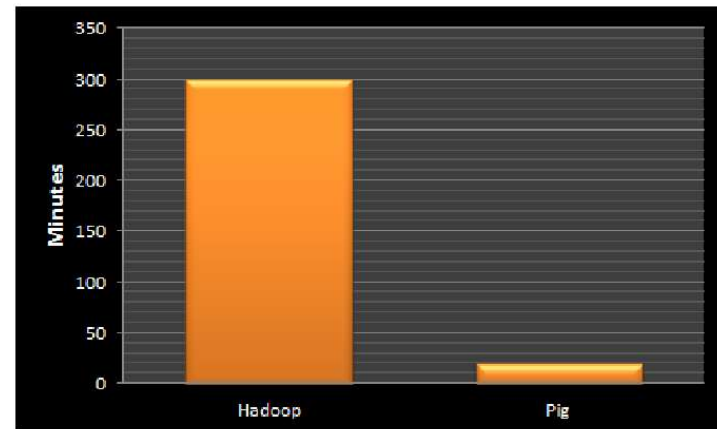
- Higher level procedural data flow language
- Similar to SQL query where the user specifies the what and leaves the “how” to the underlying processing engine.

PIG vs MapReduce



Lines of Code (1/20)



Development Time (1/16)

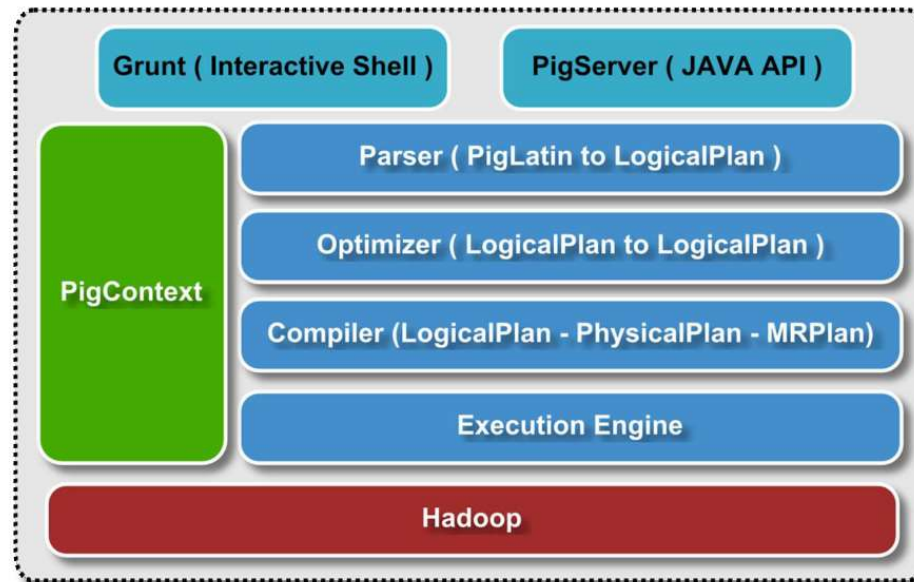


Where to use PIG ?



- Pig is a data flow language. It is at the top of Hadoop and makes it possible to create complex jobs to process large volumes of data quickly and efficiently.
- It will consume any data that you feed it: Structured, semi-structured, or unstructured.
- Pig provides the common data operations (filters, joins, ordering) and nested data types (tuple, bags, and maps) which are missing in mapreduce.
- Pig's multi-query approach combines certain types of operations together in a single pipeline, reducing the number of times data is scanned.
- PIG scripts are easier and faster to write than standard Java Hadoop jobs and PIG has lot of clever optimizations like multi query execution, which can make your complex queries execute quicker.

Pig Architecture



Pig Latin Relational Operators

Category	Operator	Description
Loading and Storing	LOAD STORE DUMP	Loads data from the file system. Saves a relation to the file system or other storage. Prints a relation to the console
Filtering	FILTER DISTINCT FOREACH...GENERATE STREAM	Removes unwanted rows from a relation. Removes duplicate rows from a relation. Adds or removes fields from a relation. Transforms a relation using an external program.
Grouping and Joining	JOIN COGROUP GROUP CROSS	Joins two or more relations. Groups the data in two or more relations. Groups the data in a single relation. Creates the cross product of two or more relations.
Sorting	ORDER LIMIT	Sorts a relation by one or more fields. Limits the size of a relation to a maximum number of tuples.
Combining and Splitting	UNION SPLIT	Combines two or more relations into one. Splits a relation into two or more relations.

Facebook problem with Data

Problem

- Data, data and more data
- 200GB per day in March 2008
- 2+TB(compressed) raw data per day today

The Hadoop Experiment

- Superior to availability/scalability/manageability of commercial DBs
- Efficiency not that great, but throw more hardware
- Partial Availability/resilience/scale more important than ACID

Problem

- Programmability and Metadata
- Mapreduce hard to program (users know sql/bash/python)
- Need to publish data in well known schemas

Solution

- HIVE



What is Hive

- A system for managing and querying structured data built on top of Hadoop
- Uses MapReduce for execution
- HDFS for storage
- Extensible to other Data Repositories


Key Building Principles

- SQL on structured data as a familiar data warehousing tool
- Extensibility (Pluggable map/reduce scripts in the language of your choice, Rich and User Defined data types, User Defined Functions)
- Interoperability (Extensible framework to support different file and data formats)
- Performance

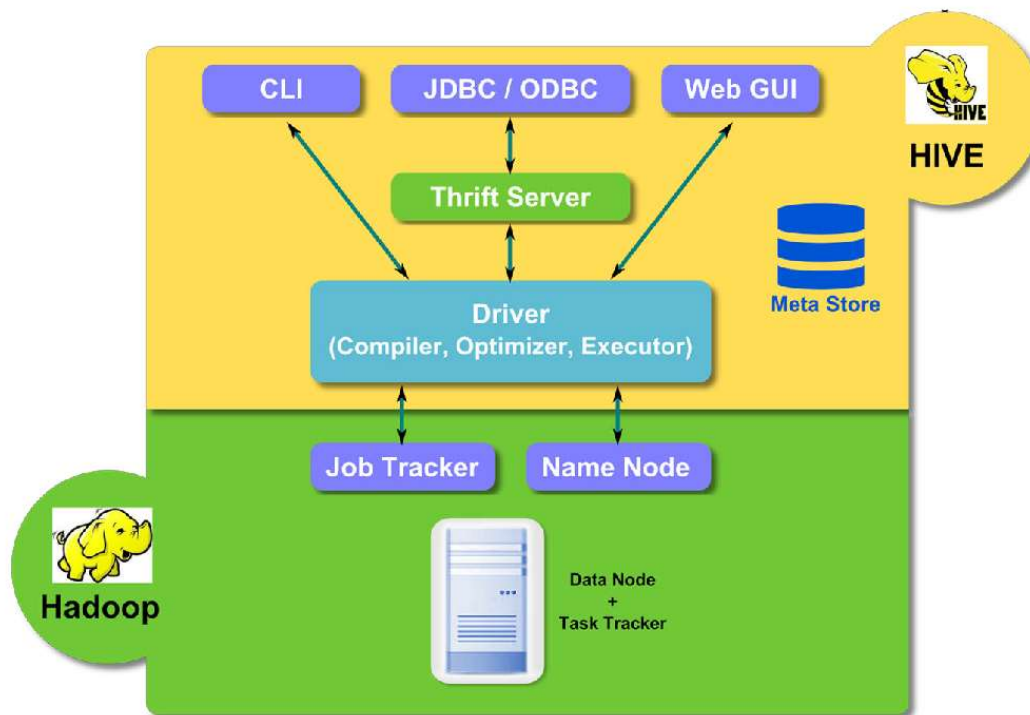


What Hive is NOT

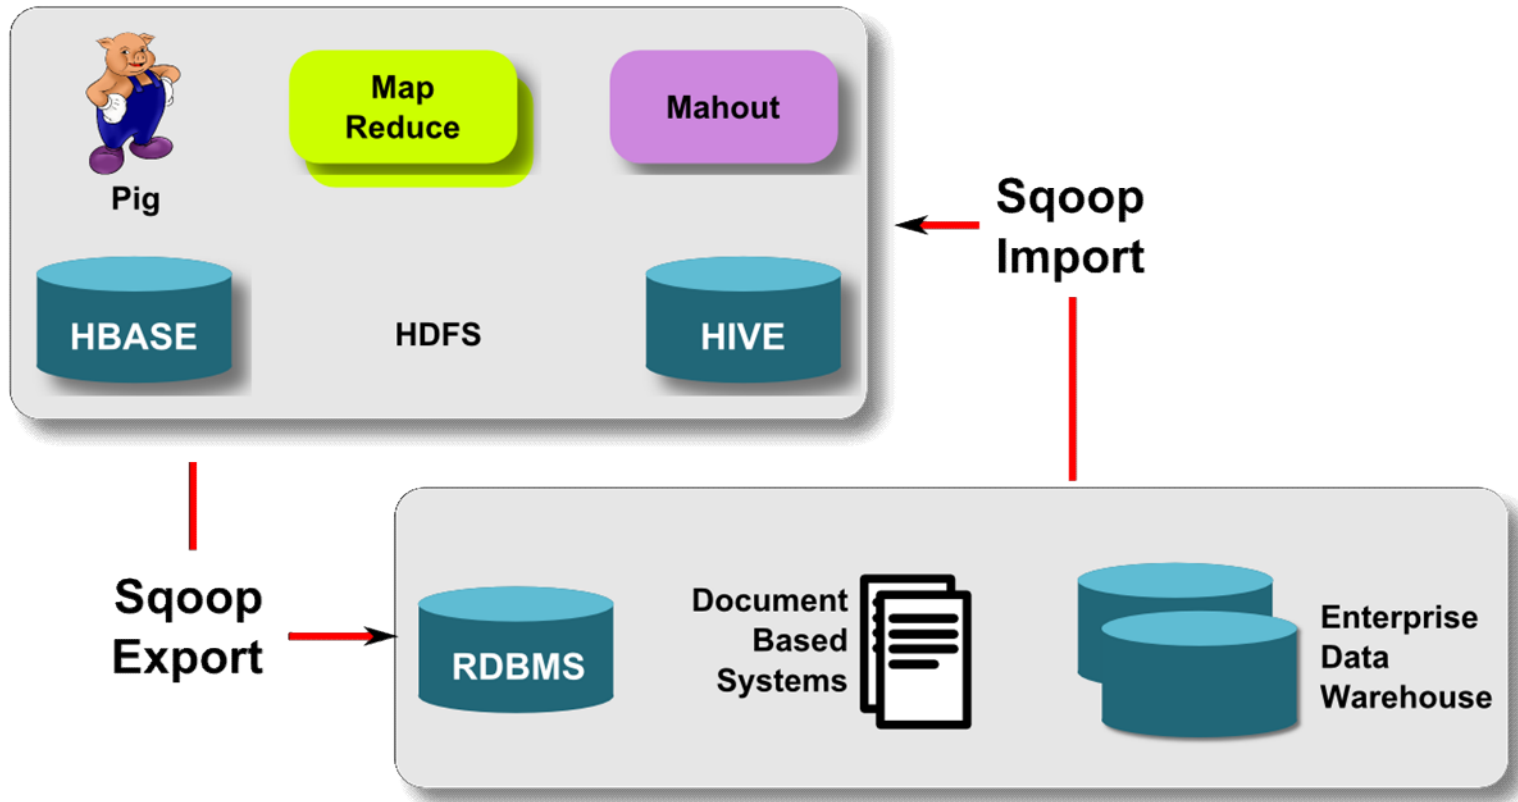
- Not OLTP
- Not suitable for OLAP* (not “online”)
- No extensive writes*
- Not a real RDBMS!
- Not real-time Analytics!
- Hadoop is Batchy => Hive too



Hive Components



Apache Sqoop



Why Sqoop?

- Most reporting and data visualization built on top of RDBMS
- A way needed for transferring processed data from HDFS to RDBMS
- Need parallelism for loading data into RDBMS
- For application which require to move data from RDBMS to Hadoop
- Using scripts for transferring data is inefficient and time consuming



photo credit: [CarbonNYC](#) via [photopin](#) cc

Oozie

- Oozie is a server based Workflow Engine specialized in running workflow jobs with actions that run Hadoop MapReduce, Pig, Hive, Sqoop, FileSystem and shell jobs.
- Oozie is a Java Web-Application that runs in a Java servlet-container.
- For the purposes of Oozie, a workflow is a collection of actions (i.e. Hadoop Map/Reduce jobs, Pig jobs) arranged in a control dependency DAG (Direct Acyclic Graph). "control dependency" from one action to another means that the second action can't run until the first action has completed.

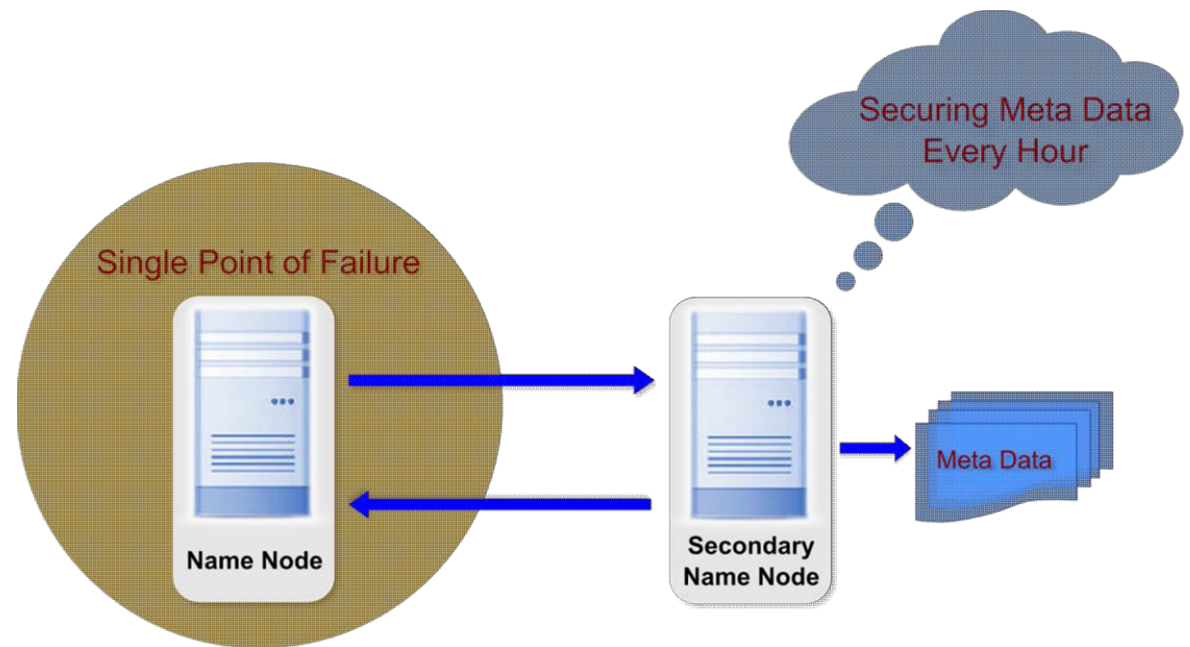


Overview of HBase

- HBase is a part of Hadoop
- Apache Hadoop is an open source system to reliably store and process data across many commodity computers
- Uses HDFS for data storage
- A **NoSQL** store for big data
- Open source written in java
- **Distributed** database
- **Automatic Sharding** , table data spread over cluster
- Automatic region server failover

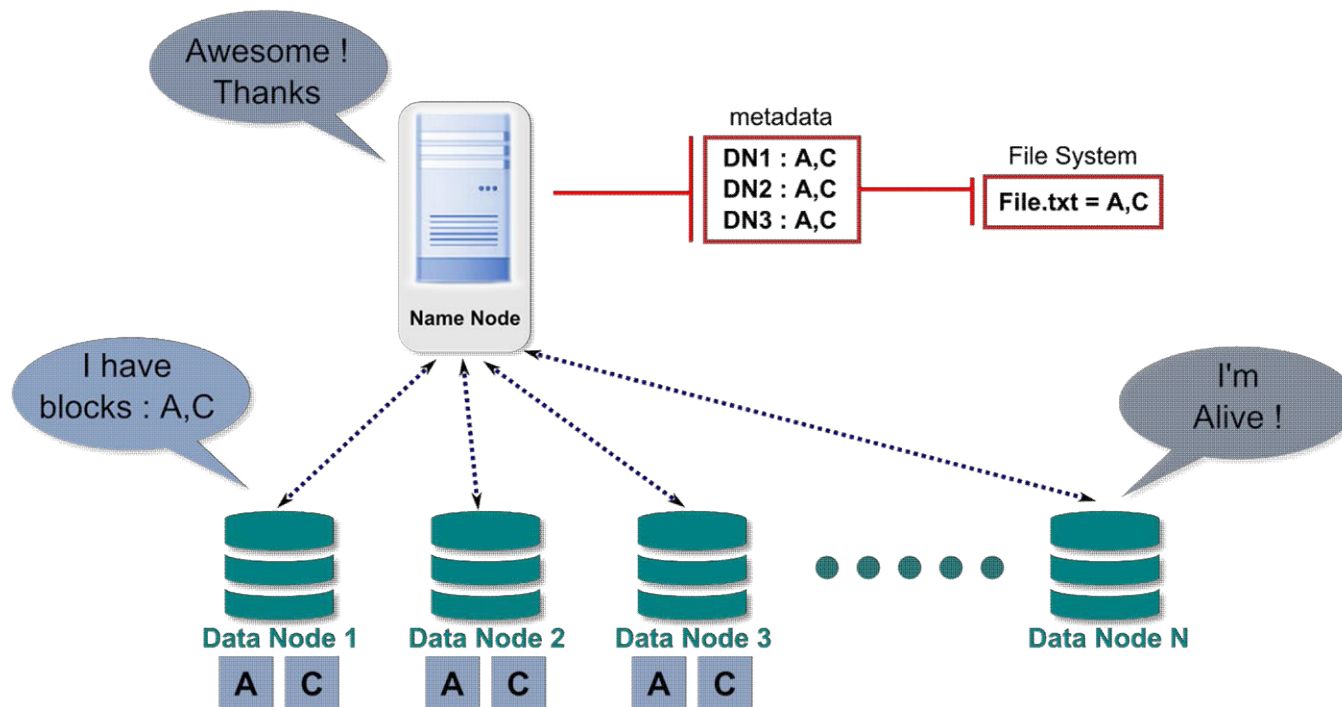
Hadoop1 Limitations

Namenode is **Single Point of Failure** in hadoop1 Architecture. If Namenode goes down, Entire Hadoop cluster is down.



Hadoop1 Limitations

HDFS is built upon the **single-node namespace server** architecture. Since the name-node is a single container of the file system metadata, it naturally becomes a limiting factor for file system growth.



Final architecture

