

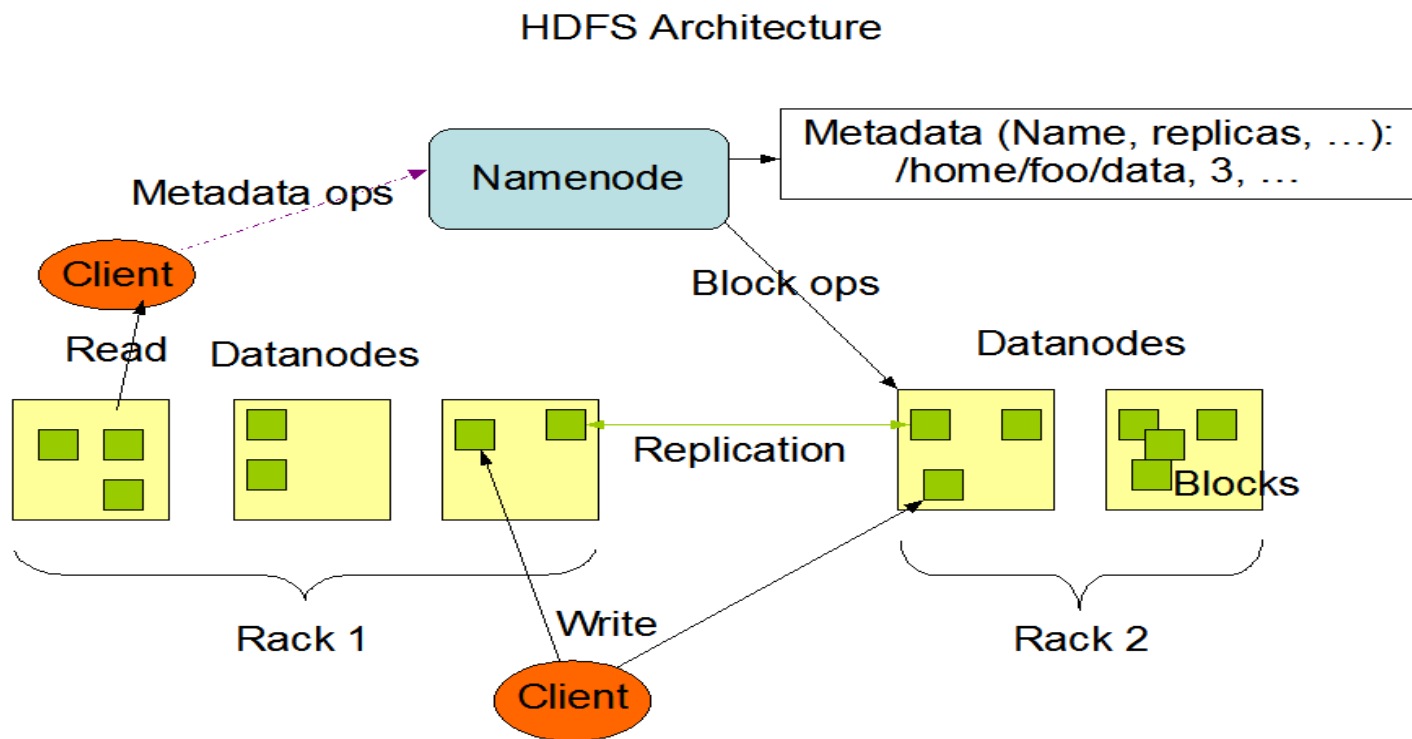
CS698F Advanced Data Management

Instructor: Medha Atre

Announcements

- Assignment-2 paper selection due: 20-Oct-2017, 23:59
- Assignment-2 presentations:
 - 25, 27 Oct 2017 in class
- Course project report and code due: 14-Nov-2017, 23:59
 - In class presentations and demo on 15, 17 Nov, 2017
 - **No extension to this allowed!**
- Endsem written exam: 18-Nov-2017 16:00-19:00

Hadoop Architecture



Hadoop basics

- Two fundamental functions
 - Map (key1, value1) => list-of((key2, value2))
 - Reduce (key2, list(values2)) => list(key3, values3)
- HDFS (Hadoop Distributed File System) makes underlying cluster opaque to the end user.
- NameNode = master node that manages the cluster
- DataNode/s = slaves that store data and run map-reduce jobs
- The data files can be in any format, and HDFS splits them according to its default logic => 76 MB blocks on each data node.

Table distribution

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Avg
1979	23	23	2	43	24	25	26	26	26	26	25	26	25
1980	26	27	28	28	28	30	31	31	31	30	30	30	29
1981	31	32	32	32	33	34	35	36	36	34	34	34	34
1984	39	38	39	39	39	41	42	43	40	39	38	38	40
1985	38	39	39	39	39	41	41	41	00	40	39	39	45

Table as simple txt file

1979	23	23	2	43	24	25	26	26	26	26	25	26	25
1980	26	27	28	28	28	30	31	31	31	30	30	30	29
1981	31	32	32	32	33	34	35	36	36	34	34	34	34
1984	39	38	39	39	39	41	42	43	40	39	38	38	40
1985	38	39	39	39	39	41	41	41	00	40	39	39	45

Data Distribution

- `$HADOOP_HOME/bin/hadoop fs -mkdir input_dir`
- `$HADOOP_HOME/bin/hadoop fs -put
/home/hadoop/sample.txt input_dir`
- `$HADOOP_HOME/bin/hadoop jar units.jar
hadoop.ProcessUnits input_dir output_dir`

ProcessUnits example

Let us go to and see the example code

https://www.tutorialspoint.com/hadoop/hadoop_mapreduce.htm

Key takeaways

- You can put any file in HDFS in any format
- The interpretation of the file depends on `InputFormat`.
- Default format is `TextInputFormat` and the corresponding mapper key-value pair is `<LongWritable, Text>`
- `InputFormat` defines how to read the input data
 - `InputSplit` (abstract class)
 - `RecordReader` (abstract class)
- `InputFormat` custom class can read the data in *your* expected format and generate key-value pairs as expected by the `map` function.

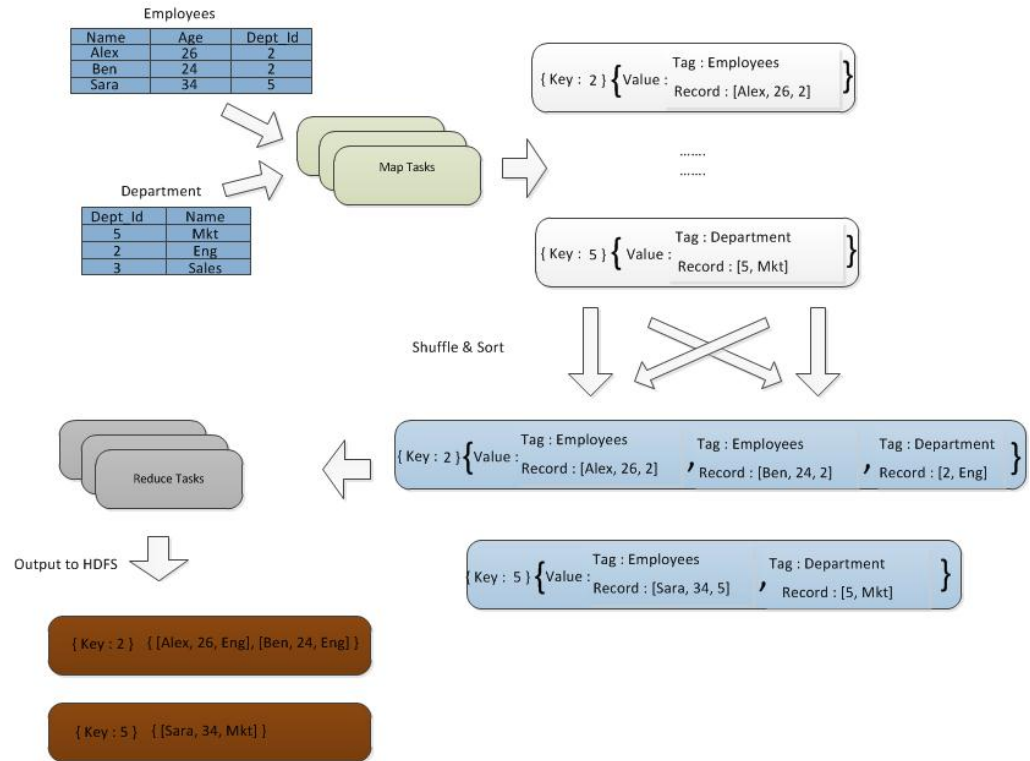
Joins with MapReduce

- Two ways
 - Reduce side join
 - Map side join

Reduce side join

Employees			Department	
Name	Age	Dept Id	Dept Id	Name
Alex	26	2	5	Mkt
Ben	24	2	2	Eng
Sara	34	5	3	Sales

```
SELECT Employees.Name,
Employees.Age, Department.Name
FROM Employees, Department
WHERE
Employees.Dept_Id=Department.Dept_Id
```



Map function

```
map (K table, V rec) {  
    dept_id = rec.Dept_Id  
    tagged_rec.tag = table  
    tagged_rec.rec = rec  
    emit(dept_id, tagged_rec)  
}
```

Reduce function (join)

```
reduce (K dept_id, list<tagged_rec> tagged_recs) {  
    for (tagged_rec : tagged_recs) {  
        for (tagged_rec1 : taagged_recs) {  
            if (tagged_rec.tag != tagged_rec1.tag) {  
                joined_rec = join(tagged_rec, tagged_rec1)  
            }  
            emit (tagged_rec.rec.Dept_Id, joined_rec)  
        }  
    }  
}
```

Map side join

- One relation has to fit in to memory.
- The smaller table is replicated to each node and loaded to the memory.
- The join happens at map side without reducer involvement.
- Significantly speed up in the process
 - Avoids shuffling all data across the network.

Map side join mapper

```
map (K table, V rec) {  
  
  list recs = lookup(rec.Dept_Id) // Get smaller table records having this Dept_Id  
  
  for (small_table_rec : recs) {  
  
    joined_rec = join (small_table_rec, rec)  
  
  }  
  
  emit (rec.Dept_id, joined_rec)  
  
}
```

Reducer is an identity function

Summary (distributed proc)

- P2P – Distributed Hash Table
 - User has the onus of data distribution and format
 - Data reshuffling (shipping/redistribution) depends on type of queries and joins – sometimes reshuffling required, sometimes not
- MapReduce – hierarchical – name-node (master), data-nodes (slaves)
 - HDFS keeps data distribution opaque to the user
 - User writes custom InputFormat, which decides mapper key-value pairs
 - Joins happen in an iterative manner.