

Whirlwind tour of Hadoop

- Inspired by Google's GFS
- Clusters from 1-10000 systems
- Batch Processing
- High Throughput
- Partition-able problems
- Fault Tolerance in storage
- Fault Tolerance in processing

HDFS

- Distributed File System (HDFS)
- Centralized “INODE” store NameNode
- DataNodes provide storage 1-10000
- Replication set Per File
- Block Size normally large 128 MB
- Optimized for throughput

Map Reduce

- Jobs are typically long running
- Move processing not data
- Algorithm uses map() and optionally reduce()
- Source Input is split
- Splits processed simultaneously
- Failed Tasks are retried

Things you do not have

- Locking
- Blocking
- Mutexes
- wait/notify
- POSIX file semantics

Problem: Large Scale Log Processing

- Challenges
- Large volumes of Data
- Store it indefinitely
- Process large data sets on multiple systems
- Data and Indexes for short periods (one week) do NOT fit in main memory

Log Format

```
File Edit View Terminal Tabs Help
# The following directives define some format nicknames for use with
# a CustomLog directive (see below).
#
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" combined
LogFormat "%h %l %u %t \"%r\" %>s %b" common
LogFormat "%{Referer}i -> %U" referer
LogFormat "%{User-agent}i" agent

LogFormat "%Y-%m-%d^A%H:%M:%S)t^Aa^Al^Av^AA^AA^Am^AU^Aq^As^A0^Ab^A%{Content-
Length}i^AT^A%{User-Agent}i^A%{Cookie}i^A%{Referer}i" w4c_ext
CustomLog "| /usr/sbin/rotatelog /var/log/httpd/phpbb.extended.log 3600" w4c_ext

# "combinedio" includes actual counts of actual bytes received (%I) and sent (%O); this
```

Getting Files into HDFS

File Edit View Terminal Tabs Help

The general command line syntax is

```
bin/hadoop command [genericOptions] [commandOptions]
```

```
bash-3.2# hadoop dfs -mkdir /user/root/phpbb
```

```
bash-3.2# hadoop dfs -copyFromLocal /var/log/httpd/phpbb.extended.log.2009-12-16-22 /usr/root/php/
```

```
bash-3.2# top
```

Hadoop Shell

- Hadoop has a shell ls, put, get, cat..

```
File Edit View Terminal Tabs Help
bash-3.2# hadoop dfs -ls /user
Found 2 items
drwxr-xr-x - root supergroup          0 2009-12-13 13:31 /user/hive
drwxr-xr-x - root supergroup          0 2009-12-16 17:36 /user/root
bash-3.2# hadoop dfs -ls /user/root
Found 2 items
drwxr-xr-x - root supergroup          0 2009-12-16 17:35 /user/root/php
drwxr-xr-x - root supergroup          0 2009-12-22 17:58 /user/root/phpbb
bash-3.2# hadoop dfs -ls /user/root/phpbb
Found 1 items
-rw-r--r--  1 root supergroup          8594 2009-12-22 17:58 /user/root/phpbb/phpbb.extended.log.2009-12-19-19
bash-3.2# hadoop dfs -cat '/user/root/phpbb/*' | head -3

2009-12-19*14:46:22*38.102.48.113*->phpbb.joingthegrid.com*10.0.21.1*10.0.21.1*GET*/phpbb-jtg/index.php*200*0*1564
4*->0*Mozilla/4.0 (compatible; MSIE 6.0; Windows 98; Win 9x 4.90)*JSESSIONID=C221101EE798A3DF3134F1A4309A62C9*http
://www.joingthegrid.com/phpbb-jtg/
2009-12-19*14:46:25*38.102.48.113*->phpbb.joingthegrid.com*10.0.21.1*10.0.21.1*GET*/phpbb-jtg/viewforum.php*f=3&sid
=448e74e726037cc5dblad4e2dd5681ff*200*0*33714*->0*Mozilla/4.0 (compatible; MSIE 6.0; Windows 98; Win 9x 4.90)*JSE
SSIONID=C221101EE798A3DF3134F1A4309A62C9; phpbb3_lkduj_u=1; phpbb3_lkduj_k=; phpbb3_lkduj_sid=448e74e726037cc5dbla
d4e2dd5681ff*http://www.joingthegrid.com/phpbb-jtg/viewforum.php?f=3&sid=448e74e726037cc5dblad4e2dd5681ff
2009-12-19*14:46:32*38.102.48.113*->phpbb.joingthegrid.com*10.0.21.1*10.0.21.1*GET*/phpbb-jtg/ucp.php*?mode=registe
r*200*0*9321*->0*Mozilla/4.0 (compatible; MSIE 6.0; Windows 98; Win 9x 4.90)*JSESSIONID=C221101EE798A3DF3134F1A430
9A62C9; phpbb3_lkduj_u=1; phpbb3_lkduj_k=; phpbb3_lkduj_sid=448e74e726037cc5dblad4e2dd5681ff*http://www.joingthegri
d.com/phpbb-jtg/ucp.php?mode=register
```

Sample Program Group and Count

- Source data looks like
- `jan 10 2009:.....:200:/index.htm`
`jan 10 2009:.....:200:/index.htm`
`jan 10 2009:.....:200:/igloo.htm`
`jan 10 2009:.....:200:/ed.htm`

In case your a Math 'type'

(input) $\langle k1, v1 \rangle \rightarrow$
map $\rightarrow \langle k2, v2 \rangle \rightarrow$ combine $\rightarrow \langle k2, v2 \rangle \rightarrow$
reduce $\rightarrow \langle k3, v3 \rangle$ (output)

Map($k1, v1$) \rightarrow list($k2, v2$)
Reduce($k2, \text{list}(v2)$) \rightarrow list($v3$)

Calculate Splits

- Input does not have to be a file to be “splittable”

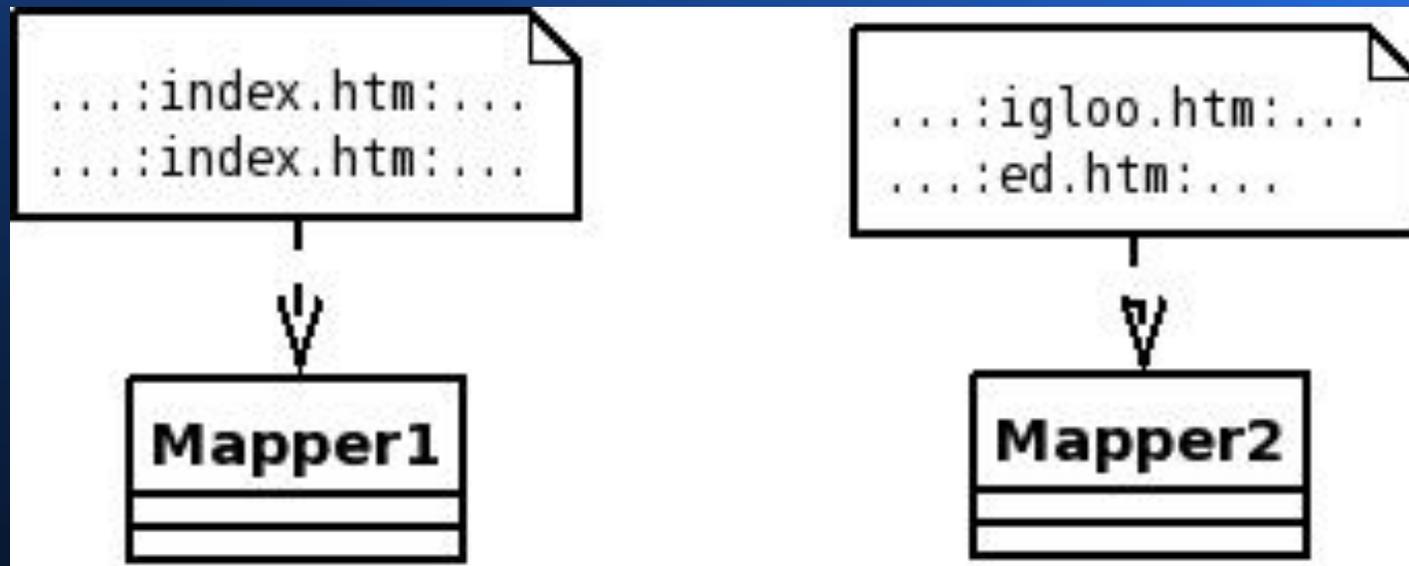
```
...:index.htm:...  
...:index.htm:...  
...:igloo.htm:...  
...:ed.html:...
```

```
...:index.htm:...  
...:index.htm:...
```

```
...:igloo.htm:...  
...:ed.htm:...
```

Splits → Mappers

- Mappers run the map() process on each split



Mapper runs map function

- For each row/line (in this case)

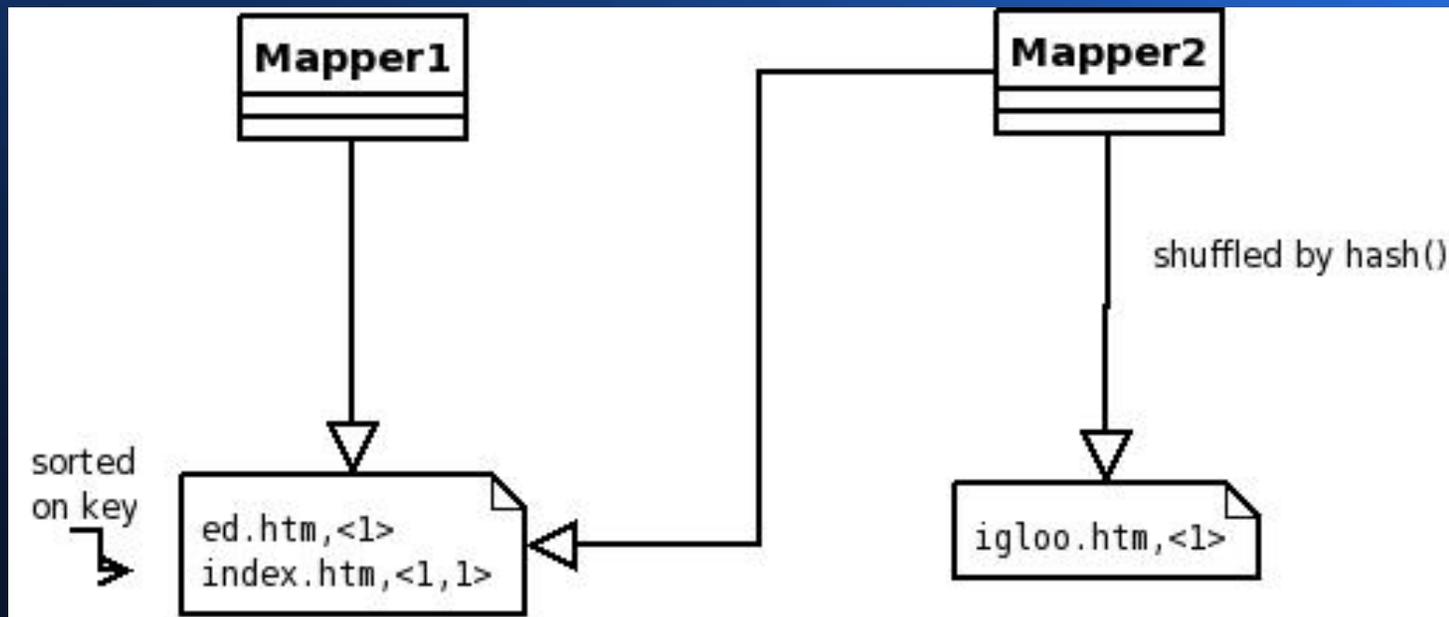
```
File Edit View Terminal Tabs Help
public static final char SOH = 0x01;
private final static IntWritable one = new IntWritable(1);

public static class Map extends MapReduceBase implements Mapper<LongWritable, Text, Text, IntWritable> {
    private Text word = new Text();

    public void map(LongWritable key, Text value, OutputCollector<Text, IntWritable> output, Reporter reporter) throws IOException {
        String line = value.toString();
        String [] parts = line.split("[ "+SOH+"");
        if (parts.length == 18 ){
            word.set(parts[17]);
            output.collect(word,one);
        }
    }
}
```

Shuffle sort, NOT the hottest new dance move

- Data Shuffled to reducer on `key.hash()`
- Data sorted by key per reducer.



Reduce

- One Reducer gets data that looks like
- "index.htm", <1, 1>
- "ed,htm", <1>
- //igloo.htm went to another reducer

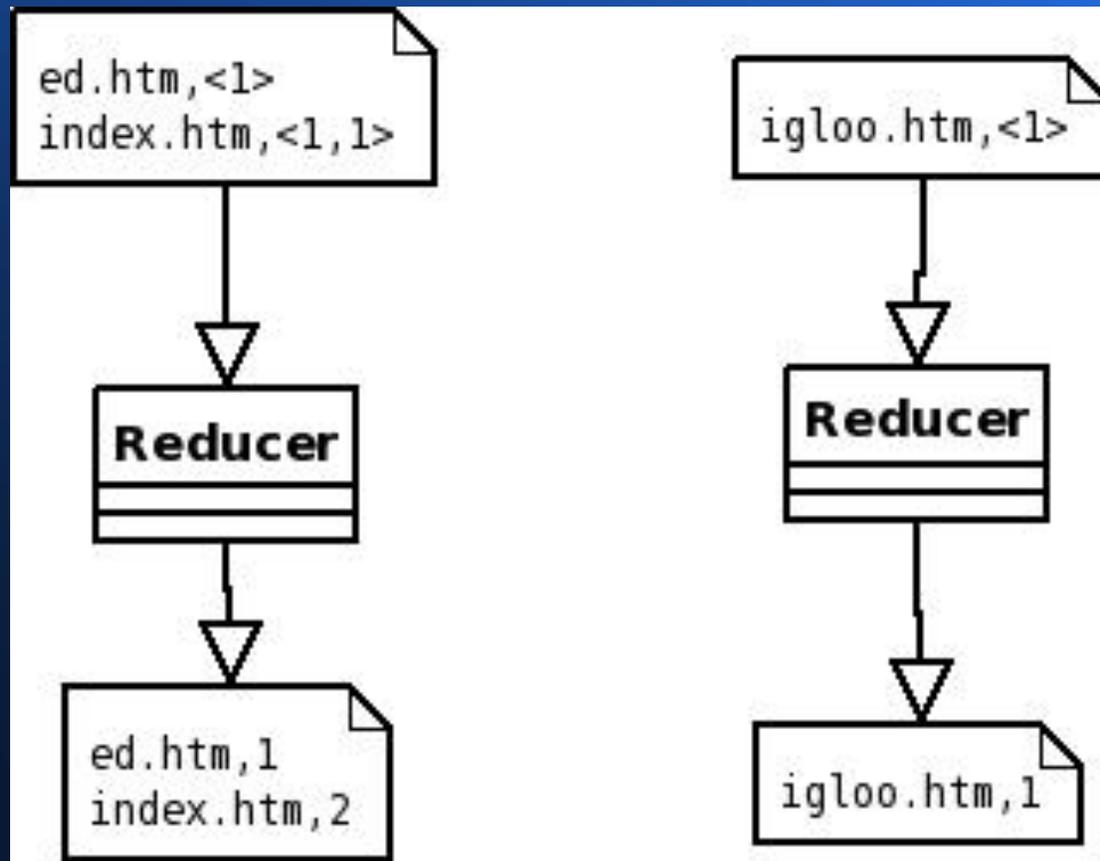
File Edit View Terminal Tabs Help

```
public void reduce
(Text key, Iterator<IntWritable> values, OutputCollector<Text, IntWritable> output, Reporter reporter)
throws IOException {
    int sum=0;
    while (values.hasNext()) {
        sum+=values.next().get();
    }
    output.collect(key, new IntWritable(sum));
}
}
```

38,3

52%

Reduce Phase



Scalable

- Tune mappers and reducers per job
- More Data...get...More Systems
- Lose DataNode files replicate elsewhere
- Lose TaskTracker running tasks will restart
- Lose NameNode, Jobtracker- Call bigsaad

Hive

- Translates queries into map/reduce jobs
- SQL like language atop Hadoop
- like MySQL CSV engine, on Hadoop, on roids
- Much easier and reusable then coding your own joins, max(), where
- CLI, Web Interface

Hive has no “Hive Format”

- Works with most input
- Configurable delimiters
- Types like string, long also map, struct, array
- Configurable SerDe

Create & Load

- `hive>create table web_data
(sdate STRING, stime STRING,
envvar STRING, remotelogname STRING ,servername
STRING, localip STRING, literaldash STRING, method
STRING, url STRING, querystring STRING, status STRING,
litteralzero STRING ,bytessent INT,header STRING,
timetoserver INT, useragent STRING ,cookie STRING, referer
STRING);`

```
LOAD DATA INPATH '/user/root/phpbb/phpbb.extended.log.  
2009-12-19-19' INTO TABLE web_data
```

Query

- `SELECT url,count(1) FROM web_data GROUP BY url`
- Hive translates query into map reduce program (s)

File Edit View Terminal Tabs Help

```
hive> select url,count(1) FROM web_data group by url;
```

```
Total MapReduce jobs = 1
```

```
Number of reduce tasks not specified. Estimated from input data size: 1
```

```
In order to change the average load for a reducer (in bytes):
```

```
  set hive.exec.reducers.bytes.per.reducer=<number>
```

```
In order to limit the maximum number of reducers:
```

```
  set hive.exec.reducers.max=<number>
```

```
In order to set a constant number of reducers:
```

```
  set mapred.reduce.tasks=<number>
```

```
Starting Job = job_200912211519_0004, Tracking URL = http://hbase.jointhegrid.com:50030/jobdetails.jsp?jobid=job_200912211519_0004
```

```
Kill Command = /usr/lib/hadoop-0.20/bin/hadoop job -Dmapred.job.tracker=hbase.jointhegrid.com:8021 -kill job_200912211519_0004
```

```
2010-01-03 01:41:38,117 map = 0%, reduce =0%
```

```
2010-01-03 01:41:52,237 map = 50%, reduce =0%
```

```
2010-01-03 01:41:55,310 map = 100%, reduce =0%
```

```
2010-01-03 01:42:09,761 map = 100%, reduce =100%
```

```
Ended Job = job_200912211519_0004
```

```
OK
```

```
/phpbb-jtg/index.php 3
```

```
/phpbb-jtg/memberlist.php 1
```

```
/phpbb-jtg/posting.php 4
```

```
/phpbb-jtg/ucp.php 12
```

```
/phpbb-jtg/viewforum.php 1
```

```
/phpbb-jtg/viewtopic.php 1
```

```
Time taken: 39.268 seconds
```

```
hive> █
```

Manage Session mysession - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://hadoop1:9999/hwi/session_manage.js

Most Visited CentOS Support 2008/2009 Galat...

User

[Authorize](#)
[Browse Schema](#)
[Create Session](#)
[Diagnostics](#)
[List Sessions](#)

Administration

[List Running Jobs](#)

Manage Session mysession

Session is set to start.

Session Diagnostics: [mysession](#)
Set Processor: [mysession](#)
Session Remove: [mysession](#)

Result File:
[View File](#)

Error File:

Query:

```
select * from bigdata limit 3
```

Silent Mode ?

Start Query?

Query Return Code -40

Done

Hive Web Interface - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://hadoop1:9999/hwi/view_file.jsp?sessic

Most Visited CentOS Support 2008/2009 Galat...

User

- [Authorize](#)
- [Browse Schema](#)
- [Create Session](#)
- [Diagnostics](#)
- [List Sessions](#)

Administration

- [List Running Jobs](#)

Hive Web Interface

/tmp/mysession.out

```
266.0 909.0 159625.0
159627.0 2635.0
2.0000827011759E7 1239029.0
1341.0 -26.399744 1228.0
0.56000537 48.0
317.0 1243.0 159625.0
159627.0 2636.0
2.0000827011759E7 1239029.0
```

This file contains 0 of 1024 blocks. [Next Block](#)

Done

The End...

Questions???