

## Chapter 1: Hadoop Architecture and Deployment

```
[hadoop@rt1 ~]$ cat /etc/yum.repos.d/bigtop.repo
[bigtop]
name=Bigtop
enabled=1
gpgcheck=1
type=NONE
baseurl=http://bigtop-repos.s3.amazonaws.com/releases/1.1.0/centos/6/x86_64
gpgkey=https://dist.apache.org/repos/dist/release/bigtop/KEYS
```

```
[root@master1 ~]# cat /etc/sysconfig/network
NETWORKING=yes
HOSTNAME=master1.cyrus.com
```

```
[root@master1 ~]# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain
10.0.0.104  master1.cyrus.com
```

```
[root@nn1 ~]# ls -l /usr/java/
total 4
lrwxrwxrwx 1 root root 16 Mar 11 2016 default -> /usr/java/latest
drwxr-xr-x 8 root root 4096 Mar 11 2016 jdk1.8.0_11
lrwxrwxrwx 1 root root 21 Mar 11 2016 latest -> /usr/java/jdk1.8.0_11
```

```
[root@nn1 ~]# ls -l /opt/cluster/
total 0
drwxr-xr-x 10 hadoop hadoop 150 Oct 16 23:33 hadoop-2.7.3
```

```
[root@nn1 ~]# ls -l /opt/cluster/hadoop-2.7.3/
total 116
drwxr-xr-x 2 hadoop hadoop 4096 Oct 16 13:55 bin
drwxr-xr-x 3 hadoop hadoop 19 Oct 16 13:55 etc
drwxr-xr-x 2 hadoop hadoop 101 Oct 16 13:55 include
drwxr-xr-x 3 hadoop hadoop 19 Oct 16 13:55 lib
drwxr-xr-x 2 hadoop hadoop 4096 Oct 16 13:55 libexec
-rw-r--r-- 1 hadoop hadoop 84854 Oct 16 13:55 LICENSE.txt
drwxrwxr-x 2 hadoop hadoop 138 Oct 16 23:35 logs
-rw-r--r-- 1 hadoop hadoop 14978 Oct 16 13:55 NOTICE.txt
-rw-r--r-- 1 hadoop hadoop 1366 Oct 16 13:55 README.txt
drwxr-xr-x 2 hadoop hadoop 4096 Oct 16 13:55 sbin
drwxr-xr-x 4 hadoop hadoop 29 Oct 16 13:55 share
```

```
drwxr-xr-x 2 hadoop hadoop 4096 Oct 16 23:32 hadoop
[root@nn1 ~]# ls -l /opt/cluster/hadoop-2.7.3/etc/hadoop/
total 152
-rw-r--r-- 1 hadoop hadoop 4436 Oct 16 13:55 capacity-scheduler.xml
-rw-r--r-- 1 hadoop hadoop 1335 Oct 16 13:55 configuration.xml
-rw-r--r-- 1 hadoop hadoop 318 Oct 16 13:55 container-executor.cfg
-rw-r--r-- 1 hadoop hadoop 348 Oct 16 23:31 core-site.xml
-rw-r--r-- 1 hadoop hadoop 3589 Oct 16 13:55 hadoop-env.cmd
-rw-r--r-- 1 hadoop hadoop 4224 Oct 16 13:55 hadoop-env.sh
-rw-r--r-- 1 hadoop hadoop 2598 Oct 16 13:55 hadoop-metrics2.properties
-rw-r--r-- 1 hadoop hadoop 2490 Oct 16 13:55 hadoop-metrics.properties
-rw-r--r-- 1 hadoop hadoop 9683 Oct 16 13:55 hadoop-policy.xml
-rw-r--r-- 1 hadoop hadoop 590 Oct 16 23:32 hdfs-site.xml
```

```
[root@nn1 ~]# ls -l /opt/cluster/hadoop-2.7.3/bin/
total 376
-rwxr-xr-x 1 hadoop hadoop 131122 Oct 16 13:55 container-executor
-rwxr-xr-x 1 hadoop hadoop 6488 Oct 16 13:55 hadoop
-rwxr-xr-x 1 hadoop hadoop 8514 Oct 16 13:55 hadoop.cmd
-rwxr-xr-x 1 hadoop hadoop 12223 Oct 16 13:55 hdfs
-rwxr-xr-x 1 hadoop hadoop 7238 Oct 16 13:55 hdfs.cmd
-rwxr-xr-x 1 hadoop hadoop 5953 Oct 16 13:55 mapred
-rwxr-xr-x 1 hadoop hadoop 6094 Oct 16 13:55 mapred.cmd
-rwxr-xr-x 1 hadoop hadoop 1776 Oct 16 13:55 rcc
-rwxr-xr-x 1 hadoop hadoop 157730 Oct 16 13:55 test-container-executor
-rwxr-xr-x 1 hadoop hadoop 13352 Oct 16 13:55 yarn
-rwxr-xr-x 1 hadoop hadoop 11054 Oct 16 13:55 yarn.cmd
```

```
[root@nn1 ~]# ls -l /opt/cluster/hadoop-2.7.3/sbin/
total 120
-rwxr-xr-x 1 hadoop hadoop 2752 Oct 16 13:55 distribute-exclude.sh
-rwxr-xr-x 1 hadoop hadoop 6452 Oct 16 13:55 hadoop-daemon.sh
-rwxr-xr-x 1 hadoop hadoop 1360 Oct 16 13:55 hadoop-daemons.sh
-rwxr-xr-x 1 hadoop hadoop 1597 Oct 16 13:55 hdfs-config.cmd
-rwxr-xr-x 1 hadoop hadoop 1427 Oct 16 13:55 hdfs-config.sh
-rwxr-xr-x 1 hadoop hadoop 2291 Oct 16 13:55 httpfs.sh
-rwxr-xr-x 1 hadoop hadoop 3128 Oct 16 13:55 kms.sh
```

```
[root@nn1 ~]# cat /etc/profile.d/hadoopenv.sh
export JAVA_HOME=/usr/java/latest
export HADOOP_HOME=/opt/cluster/hadoop

export HADOOP_MAPRED_HOME=$HADOOP_HOME
export HADOOP_COMMON_HOME=$HADOOP_HOME
export HADOOP_HDFS_HOME=$HADOOP_HOME
export YARN_HOME=$HADOOP_HOME
export HADOOP_CONF_DIR=$HADOOP_HOME/etc/hadoop
export YARN_CONF_DIR=$HADOOP_HOME/etc/hadoop

export HADOOP_HOME_WARN_SUPPRESS=True
PATH=$JAVA_HOME/bin:$HADOOP_HOME/bin:$HADOOP_HOME/sbin:$PATH
export PATH
```

```
[root@nn1 ~]# . /etc/profile.d/hadoopenv.sh
```

```
[root@nn1 ~]# su - hadoop
Last login: Sat Oct 22 13:34:57 IST 2016 on pts/0
[hadoop@nn1 ~]$ id
uid=1003(hadoop) gid=1003(hadoop) groups=1003(hadoop)
```

```
[hadoop@nn1 cluster]$ pwd
/opt/cluster
[hadoop@nn1 cluster]$ ls -l
total 0
drwxr-xr-x 10 hadoop hadoop 150 Oct 16 23:33 hadoop-2.7.3
[hadoop@nn1 cluster]$ ln -s hadoop-2.7.3 hadoop
[hadoop@nn1 cluster]$ ls -l
total 0
lrwxrwxrwx 1 hadoop hadoop 12 Oct 22 14:08 hadoop -> hadoop-2.7.3
drwxr-xr-x 10 hadoop hadoop 150 Oct 16 23:33 hadoop-2.7.3
```

```
[hadoop@nn1 cluster]$ cat hadoop/etc/hadoop/core-site.xml
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl"?>

<!-- Put site-specific property overrides in this file. -->

<configuration>

<property>
<name>fs.defaultFS</name>
<value>hdfs://nn1.cluster1.com:9000</value>
</property>

</configuration>
```



```
[hadoop@nn1 cluster]$ cat hadoop/etc/hadoop/hdfs-site.xml
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl"?>

<!-- Put site-specific property overrides in this file. -->

<configuration>

<property>
<name>dfs.name.dir</name>
<value>/data/namenode</value>
</property>
```

```
[hadoop@nn1 hadoop]$ cat hdfs-site.xml
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl"?>

<!-- Put site-specific property overrides in this file. -->

<configuration>

<property>
<name>dfs.name.dir</name>
<value>/data/namenode</value>
</property>

<property>
<name>dfs.datanode.data.dir</name>
<value>/data/datanode</value>
</property>

</configuration>
```

```
[hadoop@nn1 cluster]$ jps
2161 DataNode
2107 NameNode
2254 Jps
```

```
[hadoop@nn1 hadoop]$ cat mapred-site.xml
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl"?>

<!-- Put site-specific property overrides in this file. -->

<configuration>

<property>
<name>mapreduce.framework.name</name>
<value>yarn</value>
</property>

</configuration>
```

```
[hadoop@nn1 hadoop]$ cat yarn-site.xml
<?xml version="1.0"?>
<configuration>

<property>
<name>yarn.nodemanager.aux-services</name>
<value>mapreduce_shuffle</value>
</property>

<property>
<name>yarn.nodemanager.aux-services.mapreduce.shuffle.class</name>
<value>org.apache.hadoop.mapred.ShuffleHandler</value>
</property>

<property>
<name>yarn.resourcemanager.resource-tracker.address</name>
<value>nn1.cluster1.com:9001</value>
</property>

<property>
<name>yarn.resourcemanager.scheduler.address</name>
<value>nn1.cluster1.com:9002</value>
</property>

<property>
<name>yarn.resourcemanager.address</name>
<value>nn1.cluster1.com:9003</value>
</property>

</configuration>
```

```
[hadoop@nn1 hadoop]$ yarn-daemon.sh start resourcemanager
starting resourcemanager, logging to /opt/cluster/hadoop-2
t
[hadoop@nn1 hadoop]$ yarn-daemon.sh start nodemanager
starting nodemanager, logging to /opt/cluster/hadoop-2.7.3
[hadoop@nn1 hadoop]$ jps
2673 Jps
2582 NodeManager
2344 ResourceManager
2107 NameNode
```

```
<property>
<name>dfs.name.dir</name>
<value>/data/namenode</value>
</property>
```

```
<property>
<name>dfs.datanode.data.dir</name>
<value>/data/datanode</value>
</property>
```

```
[hadoop@nn1 hadoop]$ cat yarn-site.xml
<?xml version="1.0"?>
<configuration>

<property>
<name>yarn.nodemanager.aux-services</name>
<value>mapreduce_shuffle</value>
</property>

<property>
<name>yarn.nodemanager.aux-services.mapreduce.shuffle.class</name>
<value>org.apache.hadoop.mapred.ShuffleHandler</value>
</property>

<property>
<name>yarn.resourcemanager.resource-tracker.address</name>
<value>jt1.cluster1.com:9001</value>
</property>

<property>
<name>yarn.resourcemanager.scheduler.address</name>
<value>jt1.cluster1.com:9002</value>
</property>

<property>
<name>yarn.resourcemanager.address</name>
<value>jt1.cluster1.com:9003</value>
</property>

</configuration>
```

```

[hadoop@nn1 hadoop]$ cat yarn-site.xml
<?xml version="1.0"?>
<configuration>

<property>
<name>yarn.nodemanager.aux-services</name>
<value>mapreduce_shuffle</value>
</property>

<property>
<name>yarn.nodemanager.aux-services.mapreduce.shuffle.class</name>
<value>org.apache.hadoop.mapred.ShuffleHandler</value>
</property>

<property>
<name>yarn.resourcemanager.resource-tracker.address</name>
<value>jt1.cluster1.com:9001</value>
</property>

<property>
<name>yarn.resourcemanager.scheduler.address</name>
<value>jt1.cluster1.com:9002</value>
</property>

<property>
<name>yarn.resourcemanager.address</name>
<value>jt1.cluster1.com:9003</value>
</property>

</configuration>

```

<b>Hadoop</b>	Overview	Datanodes	Datanode Volume Failures	Snapshot	Startup Progress	Utilities
---------------	----------	-----------	--------------------------	----------	------------------	-----------

Upgrade in progress. Not yet finalized.

## Overview 'nn1.cluster1.com:9000' (active)

<b>Started:</b>	Sat Oct 22 16:22:17 IST 2016
<b>Version:</b>	2.7.3, rUnknown
<b>Compiled:</b>	2016-10-16T08:20Z by root from Unknown
<b>Cluster ID:</b>	CID-ff2466fc-0ae9-4c5f-bce7-281c0238a776
<b>Block Pool ID:</b>	BP-32539073-192.168.1.70-1476641130239

## Chapter 2: Maintaining Hadoop Cluster HDFS

```
<property>  
<name>dfs.blocksize</name>  
<value>67108864</value>  
</property>
```

```
[hadoop@nn1 ~]$ hadoop fs -Ddfs.blocksize=67108864 -put file1 /
```

```
<property>  
<name>dfs.namenode.name.dir</name>  
<value>/data/namenode</value>  
</property>
```

```
<property>  
<name>dfs.namenode.name.dir</name>  
<value>/data/namenode,/data/namenode1</value>  
</property>
```

```
[hadoop@nn1 ~]$ hadoop fs -put file1.txt /  
[hadoop@nn1 ~]$ hadoop fs -ls /file1.txt  
-rw-r--r--  2 hadoop supergroup      24 2016-11-09 22:38 /file1.txt
```

```
[hadoop@nn1 ~]$ hadoop fs -test -f /file1.txt  
[hadoop@nn1 ~]$ echo $?  
0
```

```
<property>
<name>dfs.replication</name>
<value>2</value>
</property>
```

```
[hadoop@nn1 hadoop]$ start-dfs.sh
Starting namenodes on [nn1.cluster1.com]
nn1.cluster1.com: starting namenode, logging
.com.out
dn3.cluster1.com: starting datanode, logging
.com.out
dn4.cluster1.com: starting datanode, logging
.com.out
dn1.cluster1.com: starting datanode, logging
```

```
[hadoop@master1 ~]$ hdfs balancer
16/11/17 23:48:40 INFO balancer.Balancer: namenodes = [hdfs://master1.cyrus.com:9000]
16/11/17 23:48:40 INFO balancer.Balancer: parameters = Balancer.Parameters[BalancingPolicy.Node,
of nodes to be excluded = 0, number of nodes to be included = 0]
Time Stamp      Iteration#  Bytes Already Moved  Bytes Left To Move  Bytes Being Moved
16/11/17 23:48:41 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.37:50010
16/11/17 23:48:41 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.75:50010
16/11/17 23:48:41 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.39:50010
16/11/17 23:48:41 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.40:50010
16/11/17 23:48:41 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.74:50010
16/11/17 23:48:41 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.38:50010
16/11/17 23:48:41 INFO balancer.Balancer: 0 over-utilized: []
16/11/17 23:48:41 INFO balancer.Balancer: 0 underutilized: []
The cluster is balanced. Exiting...
Nov 17, 2016 11:48:41 PM          0          0 B          0 B          -1 B
Nov 17, 2016 11:48:41 PM Balancing took 1.703 seconds
```

```
[hadoop@master1 ~]$ hdfs balancer -threshold 40
16/11/18 01:14:53 INFO balancer.Balancer: Using a threshold of 40.0
16/11/18 01:14:53 INFO balancer.Balancer: namenodes = [hdfs://master1.cyrus.com:9000]
16/11/18 01:14:53 INFO balancer.Balancer: parameters = Balancer.Parameters[BalancingPolicy.Node,
of nodes to be excluded = 0, number of nodes to be included = 0]
Time Stamp      Iteration#  Bytes Already Moved  Bytes Left To Move  Bytes Being Moved
16/11/18 01:14:54 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.74:50010
16/11/18 01:14:54 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.40:50010
16/11/18 01:14:54 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.38:50010
16/11/18 01:14:54 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.39:50010
16/11/18 01:14:54 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.75:50010
16/11/18 01:14:54 INFO net.NetworkTopology: Adding a new node: /default-rack/10.0.0.37:50010
16/11/18 01:14:54 INFO balancer.Balancer: 0 over-utilized: []
16/11/18 01:14:54 INFO balancer.Balancer: 0 underutilized: []
The cluster is balanced. Exiting...
```

```
$ hadoop fs -mkdir /projects
$ hadoop fs -ls /projects
```



```
[hadoop@master1 ~]$ hadoop fs -count -q /projects
      none          inf          none          inf
```

```
[hadoop@master1 ~]$ hdfs dfsadmin -setQuota 100 /projects
[hadoop@master1 ~]$ hadoop fs -count -q /projects
      100           99           none          inf
[hadoop@master1 ~]$ hdfs dfsadmin -setSpaceQuota 4G /projects
[hadoop@master1 ~]$ hadoop fs -count -q /projects
      100           99          4294967296    4294967296
```

```
[hadoop@master1 ~]$ hdfs dfsadmin -clrSpaceQuota /projects
[hadoop@master1 ~]$ hadoop fs -count -q /projects
      100           99           none          inf
[hadoop@master1 ~]$ hdfs dfsadmin -clrQuota /projects
[hadoop@master1 ~]$ hadoop fs -count -q /projects
      none          inf          none          inf
```



```

.....Status: CORRUPT
Total size:      850199694 B
Total dirs:     179
Total files:    1213
Total symlinks: 0
Total blocks (validated): 1209 (avg. block size 703225 B)
*****
UNDER MIN REPL'D BLOCKS: 2 (0.16542597 %)
dfs.namenode.replication.min: 1
CORRUPT FILES: 2
MISSING BLOCKS: 2
MISSING SIZE: 2471 B
CORRUPT BLOCKS: 2
*****
Minimally replicated blocks: 1207 (99.83457 %)
Over-replicated blocks: 0 (0.0 %)
Under-replicated blocks: 0 (0.0 %)
Mis-replicated blocks: 0 (0.0 %)
Default replication factor: 3
Average block replication: 2.8908188
Corrupt blocks: 2
Missing replicas: 0 (0.0 %)
Number of data-nodes: 6
Number of racks: 1
FSCK ended at Fri Nov 18 21:38:11 UTC 2016 in 148 milliseconds

The filesystem under path '/' is CORRUPT

```

```

[hadoop@master1 ~]$ hdfs fsck /input/new.txt -files -blocks -racks -locations
Connecting to namenode via http://master1.cyrus.com:50070/fsck?ugi=hadoop&files=
h=%2Finput%2Fnew.txt
FSCK started by hadoop (auth:SIMPLE) from /10.0.0.104 for path /input/new.txt at
/input/new.txt 3707474 bytes, 1 block(s): OK
0. BP-1100651502-10.0.0.104-1463464153571:blk_1073749762_8939 len=3707474 repl=2
default-rack/10.0.0.75:50010]

```

```
.....  
/user/hadoop/examples/apps/streaming/job.properties: CORRUPT blockpool BP-1100651502-  
k blk_1073749593  
  
/user/hadoop/examples/apps/streaming/job.properties: MISSING 1 blocks of total size 1  
/user/hadoop/examples/input-data/text/data.txt: CORRUPT blockpool BP-1100651502-10.0.  
_1073749610  
  
/user/hadoop/examples/input-data/text/data.txt: MISSING 1 blocks of total size 1409 B
```

```
[hadoop@master1 cluster]$ cat topology.sh  
while [ $# -gt 0 ] ; do  
  nodeArg=$1  
  exec< /opt/cluster/topology.data  
  result=""  
  while read line ; do  
    ar=( $line )  
    if [ "${ar[0]}" = "$nodeArg" ] ; then  
      result="${ar[1]}"  
    fi  
  done  
  shift  
  if [ -z "$result" ] ; then  
    echo -n "/default"  
  else  
    echo -n "$result "  
  fi  
done  
[hadoop@master1 cluster]$ cat topology.data  
10.0.0.37      /rack1  
10.0.0.38      /rack1  
10.0.0.39      /rack2  
10.0.0.40      /rack2  
10.0.0.74      /rack3  
10.0.0.75      /rack3
```

```
Name: 10.0.0.74:50010 (dn5.cyrus.com)  
Hostname: dn5.cyrus.com  
Rack: /rack3
```

```
[hadoop@master1 cluster]$ hadoop fs -rm /input/new.txt  
16/11/19 00:00:00 INFO fs.TrashPolicyDefault: Namenode trash configuration  
tier interval = 45 minutes.  
Moved: 'hdfs://master1.cyrus.com:9000/input/new.txt' to trash at: hdfs://m  
/Current
```

```
[hadoop@master1 ~]$ hadoop distcp -update /projects /new
16/11/19 23:06:01 INFO tools.DistCp: Input Options: Dist
ng=false, ignoreFailures=false, maxMaps=20, sslConfigura
ting=null, sourcePaths=[/projects], targetPath=/new, tar
16/11/19 23:06:01 INFO client.RMProxy: Connecting to Res
```

## Chapter 3: Maintaining Hadoop Cluster – YARN and MapReduce

```
[hadoop@master1 ~]$ yarn jar /opt/cluster/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples-2.7.2.jar pi 10
10
Number of Maps = 10
Samples per Map = 10
Wrote input for Map #0
Wrote input for Map #1
Wrote input for Map #2
```

```
[hadoop@master1 ~]$ cat mapper.py
#!/usr/bin/env python

import sys

# input comes from STDIN (standard input)
for line in sys.stdin:
    # remove leading and trailing whitespace
    line = line.strip()
    # split the line into words
    words = line.split()
    # increase counters
    for word in words:
        # write the results to STDOUT (standard output);
        # what we output here will be the input for the
        # Reduce step, i.e. the input for reducer.py
        #
        # tab-delimited; the trivial word count is 1
    print '%s\t%s' % (word, 1)
```

```
[hadoop@master1 ~]$ hadoop jar hadoop/tools/lib/hadoop-streaming-2.7.2.jar -file
/home/hadoop/mapper.py -mapper /home/hadoop/mapper.py -file /home/hadoop/reduce
r.py -reducer /home/hadoop/reducer.py -input /input -output /output
```

```
<property>
  <name>mapreduce.jobhistory.address</name>
  <value>master1.cyrus.com:10020</value>
</property>
```

```
<property>
  <name>mapreduce.jobhistory.webapp.address</name>
  <value>master1.cyrus.com:19888</value>
</property>
```

```
<property>
  <description>Where to aggregate logs to.</description>
  <name>yarn.nodemanager.remote-app-log-dir</name>
  <value>/tmp/logs</value>
</property>
```

```
[hadoop@master1 hadoop]$ jps
5155 ApplicationHistoryServer
3462 SecondaryNameNode
5288 JobHistoryServer
6331 Jps
3228 NameNode
4781 ResourceManager
```



Logged in as: dr.who

## JobHistory

- Application
- About Jobs
- Tools

### Retired Jobs

Show 20 entries Search:

Submit Time	Start Time	Finish Time	Job ID	Name	User	Queue	State	Maps Total	Maps Completed	Reduces Total	Reduces Completed
2016.12.04 20:53:33 UTC	2016.12.04 20:53:38 UTC	2016.12.04 20:53:50 UTC	job_1480884784329_0001	QuasiMonteCarlo	hadoop	default	SUCCEEDED	10	10	1	1

Showing 1 to 1 of 1 entries First Previous 1 Next Last



```
▼<configuration>
  ▼<property>
    <name>dfs.journalnode.rpc-address</name>
    <value>0.0.0.0:8485</value>
    <source>hdfs-default.xml</source>
  </property>
  ▼<property>
    <name>yarn.ipc.rpc.class</name>
    <value>org.apache.hadoop.yarn.ipc.HadoopYarnProtoRPC</value>
    <source>yarn-default.xml</source>
  </property>
  ▼<property>
    <name>mapreduce.job.maxtaskfailures.per.tracker</name>
    <value>3</value>
    <source>mapred-default.xml</source>
  </property>

```

```
{
  "beans": [
    {
      "name": "Hadoop:service=JobHistoryServer,name=UgiMetrics",
      "modelerType": "UgiMetrics",
      "tag.Context": "ugi",
      "tag.Hostname": "master1.cyrus.com",
      "LoginSuccessNumOps": 0,
      "LoginSuccessAvgTime": 0.0,
      "LoginFailureNumOps": 0,
      "LoginFailureAvgTime": 0.0,
      "GetGroupsNumOps": 0,
      "GetGroupsAvgTime": 0.0
    }
  ],

```

```
[hadoop@master1 hadoop]$ curl http://master1.cyrus.com:19888/jmx?qry=Hadoop:service=JobHistoryServer,name=JvmMetrics
{
  "beans" : [ {
    "name" : "Hadoop:service=JobHistoryServer,name=JvmMetrics",
    "modelerType" : "JvmMetrics",
    "tag.Context" : "jvm",
    "tag.ProcessName" : "JobHistoryServer",
    "tag.SessionId" : null,
    "tag.Hostname" : "master1.cyrus.com",
    "MemNonHeapUsedM" : 50.889282,
    "MemNonHeapCommittedM" : 52.4375,
    "MemNonHeapMaxM" : -9.536743E-7,
    "MemHeapUsedM" : 97.7354,
    "MemHeapCommittedM" : 315.5,
    "MemHeapMaxM" : 889.0,
    "MemMaxM" : 889.0,
    "GcCount" : 8,
    "GcTimeMillis" : 136,
```

```
<property>
<name>yarn.resourcemanager.resource-tracker.address</name>
<value>master1.cyrus.com:9001</value>
</property>

<property>
<name>yarn.resourcemanager.scheduler.address</name>
<value>master1.cyrus.com:9002</value>
</property>

<property>
<name>yarn.resourcemanager.address</name>
<value>master1.cyrus.com:9003</value>
</property>
```

```
<property>
<name>yarn.nodemanager.resource.memory.mb</name>
<value>11264</value>
</property>
```

```
<property>
<name>yarn.app.mapreduce.am.resource.mb</name>
<value>1536</value>
</property>
```

```
<property>  
<name>yarn.scheduler.minimum-allocation-mb</name>  
<value>64</value>  
</property>
```

```
<property>  
<name>yarn.scheduler.maximum-allocation-mb</name>  
<value>11264</value>  
</property>
```

```
<property>  
<name>yarn.scheduler.minimum-allocation-vcores</name>  
<value>1</value>  
</property>  
  
<property>  
<name>yarn.scheduler.maximum-allocation-vcores</name>  
<value>32</value>  
</property>
```

```
<property>  
<name>yarn.nodemanager.vmem-pmem-ratio</name>  
<value>3.0</value>  
</property>  
  
<property>  
<name>yarn.nodemanager.pmem-check-enabled</name>  
<value>true</value>  
</property>  
  
<property>  
<name>yarn.nodemanager.vmem-check-enabled</name>  
<value>true</value>  
</property>
```



# All Applications

- Cluster
  - About
  - Nodes
  - Node Labels
  - Applications
    - NEW
    - NEW\_SAVING
    - SUBMITTED
    - ACCEPTED
    - RUNNING
    - FINISHED
    - FAILED
    - KILLED
  - Scheduler

## Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Memory Reserved	VCores Used	VCores Total	VCores Reserved	Active Nodes
0	0	0	0	0	0 B	48 GB	0 B	0	48	0	6

## Scheduler Metrics

Scheduler Type	Scheduling Resource Type	Minimum Allocation
Capacity Scheduler	[MEMORY]	<memory:64, vCores:1>

Show 20 entries

ID	User	Name	Application Type	Queue	StartTime	FinishTime	State	FinalStatus	Progress
No data available in table									

Showing 0 to 0 of 0 entries

Logged in as: dr.who



# Application application\_1480889723402\_0002

- Cluster
  - About
  - Nodes
  - Node Labels
  - Applications
    - NEW
    - NEW\_SAVING
    - SUBMITTED
    - ACCEPTED
    - RUNNING
    - FINISHED
    - FAILED
    - KILLED
  - Scheduler
- Tools

## Kill Application

Application Overview	
User:	hadoop
Name:	QuasiMonteCarlo
Application Type:	MAPREDUCE
Application Tags:	
YarnApplicationState:	RUNNING: AM has registered with RM and started running.
FinalStatus Reported by AM:	Application has not completed yet.
Started:	Mon Dec 05 01:49:17 +0000 2016
Elapsed:	12sec
Tracking URL:	ApplicationMaster
Diagnostics:	

Application Metrics	
Total Resource Preempted:	<memory:0, vCores:0>
Total Number of Non-AM Containers Preempted:	0
Total Number of AM Containers Preempted:	0
Resource Preempted from Current Attempt:	<memory:0, vCores:0>
Number of Non-AM Containers Preempted from Current Attempt:	0
Aggregate Resource Allocation:	116853 MB-seconds, 68 vcore-seconds

Logged in as: dr.who



# NEW,NEW\_SAVING,SUBMITTED,ACCEPTED,RUNNING Applications

- Cluster
  - About
  - Nodes
  - Node Labels
  - Applications
    - NEW
    - NEW\_SAVING
    - SUBMITTED
    - ACCEPTED
    - RUNNING
    - FINISHED
    - FAILED
    - KILLED
  - Scheduler
- Tools

## Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Memory Reserved	VCores Used	VCores Total	VCores Reserved
0	0	0	0	0	0 B	80 GB	0 B	0	80	0

## Cluster Nodes Metrics

Active Nodes	Decommissioning Nodes	Decommissioned Nodes	Lost Nodes	Unhealthy Nodes	Rebooted Nodes	Shutdown Nodes
10	0	0	0	0	0	0

## Scheduler Metrics

Scheduler Type	Scheduling Resource Type	Minimum Allocation	Maximum Allocation	Maximum Cluster Application Priority
Capacity Scheduler	[MEMORY]	<memory:32, vCores:1>	<memory:8192, vCores:8>	0

Dump scheduler logs 1 min

## Application Queues

Legend: Capacity Used Used (over capacity) Max Capacity

- Queue: root 0.0% used
- Queue: default 0.0% used

```
[hadoop@master1 ~]$ curl http://master1.cyrus.com:8088/ws/v1/cluster
{"clusterInfo":{"id":"1480889723402","startedOn":"1480889723402","state":"STARTED","haState":"ACTIVE","rmStateStoreName":"org.apache.hadoop.yarn.server.resourcemanager.recovery.NullRMStateStore","resourceManagerVersion":"2.7.2","resourceManagerBuildVersion":"2.7.2 from Unknown by root source checksum c63f7cc71b8f63249e35126f0f7492d","resourceManagerVersionBuiltOn":"2016-05-16T04:01Z","hadoopVersion":"2.7.2","hadoopBuildVersion":"2.7.2 from Unknown by root source checksum d0fda26633fa762bff87ec759ebe689c","hadoopVersionBuiltOn":"2016-05-16T03:56Z","haZooKeeperConnectionState":"ResourceManager HA is not enabled."}}[hadoop@master1 ~]$
```

```
[hadoop@master1 ~]$ curl http://master1.c
{"clusterMetrics":{"appsSubmitted":2,"app
illed":0,"reservedMB":0,"availableMB":491
:48,"allocatedVirtualCores":0,"containers
```

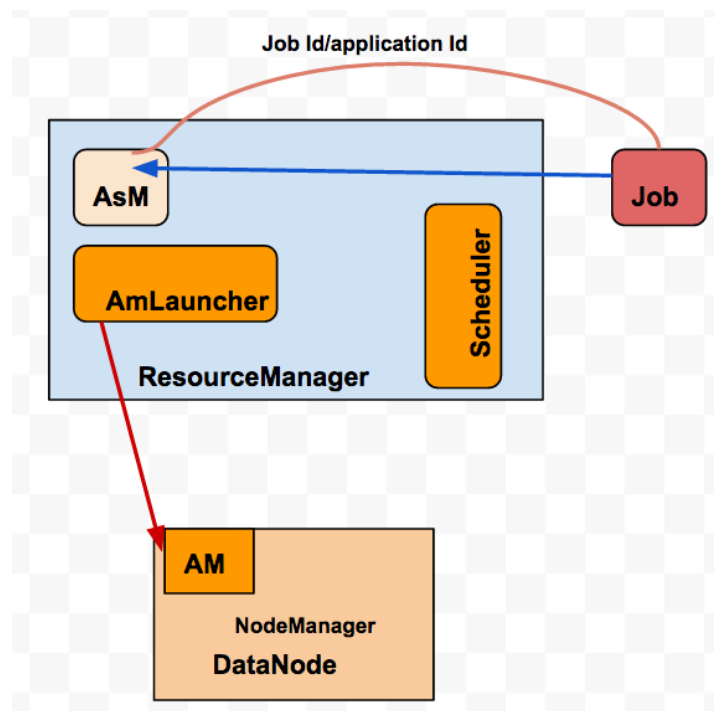
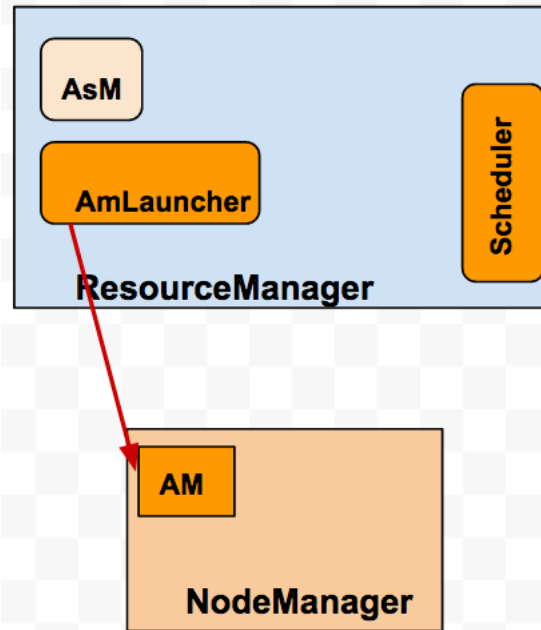
```
[hadoop@master1 ~]$ curl http://master1.cyrus.com:8088/ws/v1/cluster/scheduler
{"scheduler":{"schedulerInfo":{"type":"capacityScheduler","capacity":100.0,"used
100.0,"queueName":"root","queues":{"queue":{"type":"capacitySchedulerLeafQueueI
apacity":0.0,"maxCapacity":100.0,"absoluteCapacity":100.0,"absoluteMaxCapacity":
:0.0,"numApplications":0,"queueName":"default","state":"RUNNING","resourcesUsed"
deReservationQueues":false,"nodeLabels":["*"],"numActiveApplications":0,"numPend
iners":0,"maxApplications":10000,"maxApplicationsPerUser":10000,"userLimit":100,
r":1.0,"AMResourceLimit":{"memory":4928,"vCores":1},"usedAMResource":{"memory":0
```

```
[hadoop@master1 ~]$ curl http://master1.cyrus.com:8088/ws/v1/cluster/apps
{"apps":{"app":{"id":"application_1480889723402_0002","user":"hadoop","name":"QuasiMonteCarlo","queue":"def
ault","state":"FINISHED","finalStatus":"SUCCEEDED","progress":100.0,"trackingUI":"History","trackingUrl":"ht
tp://master1.cyrus.com:8088/proxy/application_1480889723402_0002/","diagnostics":"","clusterId":"148088972340
2","applicationType":"MAPREDUCE","applicationTags":"","startedTime":"1480902557125","finishedTime":"148090257078
6","elapsedTime":13661,"amContainerLogs":"http://dn6.cyrus.com:8042/node/containerlogs/container_148088972340
2_0002_01_000001/hadoop","amHostHttpAddress":"dn6.cyrus.com:8042","allocatedMB":-1,"allocatedVCores":-1,"run
ningContainers":-1,"memorySeconds":130085,"vcoreSeconds":78,"preemptedResourceMB":0,"preemptedResourceVCores
":0,"numNonAMContainerPreempted":0,"numAMContainerPreempted":0},"id":"application_1480889723402_0001","user
":"hadoop","name":"QuasiMonteCarlo","queue":"default","state":"FINISHED","finalStatus":"SUCCEEDED","progress
":100.0,"trackingUI":"History","trackingUrl":"http://master1.cyrus.com:8088/proxy/application_1480889723402_
0001/","diagnostics":"","clusterId":"1480889723402","applicationType":"MAPREDUCE","applicationTags":"","starte
dTime":"1480902526464","finishedTime":"1480902542342","elapsedTime":15878,"amContainerLogs":"http://dn4.cyrus.co
m:8042/node/containerlogs/container_1480889723402_0001_01_000001/hadoop","amHostHttpAddress":"dn4.cyrus.com:
8042","allocatedMB":-1,"allocatedVCores":-1,"runningContainers":-1,"memorySeconds":150294,"vcoreSeconds":90,
"preemptedResourceMB":0,"preemptedResourceVCores":0,"numNonAMContainerPreempted":0,"numAMContainerPreempted"
```

```
<property>
  <name>yarn.resourcemanager.recovery.enabled</name>
  <value>>true</value>
</property>
```

```
<property>
  <name>yarn.resourcemanager.store.class</name>
  <value>org.apache.hadoop.yarn.server.resourcemanager.recovery.FileSystemRMStateStore</value>
</property>
```

```
<property>  
  <name>yarn.resourcemanager.fs.state-store.uri</name>  
  <value>hdfs://localhost:9000/rmstore</value>  
</property>
```



## Chapter 4: High Availability

```
[ec2-user@master1 ~]$ sudo /etc/init.d/nfs start
Initializing kernel nfsd: [ OK ]
Starting NFS services: [ OK ]
Starting NFS quotas: [ OK ]
Starting NFS mountd: [ OK ]
Starting NFS daemon: [ OK ]
Starting RPC idmapd: [ OK ]
[ec2-user@master1 ~]$ sudo /etc/init.d/rpc
rpcbind      rpcgssd      rpcidmapd    rpcsvcgssd
[ec2-user@master1 ~]$ sudo /etc/init.d/rpcbind status
rpcbind (pid 2402) is running...
```

```
[ec2-user@master1 ~]$ cat /etc/exports
/nfs      *(rw,sync)
[ec2-user@master1 ~]$ sudo exportfs -a
[ec2-user@master1 ~]$ showmount -e master1.cyrus.com
Export list for master1.cyrus.com:
/nfs *
```

```
[ec2-user@master1 ~]$ sudo mkdir -p /shared
[ec2-user@master1 ~]$ sudo mount -t nfs master1.cyrus.com:/nfs /shared
```

```
[ec2-user@master1 ~]$ ls -ld /shared/
drwxr-xr-x 2 hadoop hadoop 4096 Dec  5 09:31 /shared/
```

```
<property>
  <name>fs.defaultFS</name>
  <value>hdfs://mycluster</value>
</property>
```



```
[hadoop@master1 zoo]$ id
uid=501(hadoop) gid=501(hadoop) groups=501(hadoop)
[hadoop@master1 zoo]$ pwd
/opt/cluster/zoo
[hadoop@master1 zoo]$ bin/zkServer.sh start
ZooKeeper JMX enabled by default
Using config: /opt/cluster/zoo/bin/../../conf/zoo.cfg
Starting zookeeper ... STARTED
```

```
[hadoop@master1 zoo]$ ls -l /data/zookeeper/
total 12
-rw-rw-r-- 1 hadoop hadoop 2 Aug 23 02:12 myid
drwxrwxr-x 2 hadoop hadoop 4096 Aug 29 07:53 version-2
-rw-rw-r-- 1 hadoop hadoop 4 Jan 3 22:48 zookeeper_server.pid
[hadoop@master1 zoo]$ ls -l /data/zookeeper/myid
-rw-rw-r-- 1 hadoop hadoop 2 Aug 23 02:12 /data/zookeeper/myid
[hadoop@master1 zoo]$ cat /data/zookeeper/myid
1
```

```
[hadoop@master1 zoo]$ bin/zkServer.sh status
ZooKeeper JMX enabled by default
Using config: /opt/cluster/zoo/bin/../../conf/zoo.cfg
Mode: follower
```

```
[hadoop@master1 ~]$ jps
17027 Jps
16981 SharedCacheManager
9430 QuorumPeerMain
16808 ResourceManager
16217 SecondaryNameNode
15833 NameNode
```

```
[hadoop@master1 ~]$ hdfs cacheadmin -addPool sales
Successfully added cache pool sales.
[hadoop@master1 ~]$ hdfs cacheadmin -listPools -stats
Found 1 result.


| NAME  | OWNER  | GROUP  | MODE      | LIMIT     | MAXTTL | BYTES_NEEDED | BYTES_CACHED |
|-------|--------|--------|-----------|-----------|--------|--------------|--------------|
| sales | hadoop | hadoop | rw-r-xr-x | unlimited | never  | 0            | 0            |


```

```
[hadoop@master1 ~]$ hdfs cacheadmin -addDirective -path /projects -pool sales -replication 2
Added cache directive 1
[hadoop@master1 ~]$ hdfs cacheadmin -listDirectives
Found 1 entry
ID POOL      REPL EXPIRY  PATH
1 sales      2 never  /projects
```

```
[hadoop@master1 ~]$ hadoop fs -put test /projects
[hadoop@master1 ~]$ hdfs cacheadmin -listPools -stats
Found 1 result.
NAME OWNER  GROUP  MODE          LIMIT  MAXTTL  BYTES_NEEDED  BYTES_CACHED  BYTES_OVERLIMIT  FILES_NEEDED
sales hadoop hadoop  rwxr-xr-x  unlimited  never    84            0              0                1
```

```
[hadoop@master1 ~]$ hdfs cacheadmin -listDirectives -stats
Found 1 entry
ID POOL      REPL EXPIRY  PATH          BYTES_NEEDED  BYTES_CACHED  F
1 sales      2 never  /projects     84            0
[hadoop@master1 ~]$ hdfs cacheadmin -listPools -stats
Found 1 result.
NAME  OWNER  GROUP  MODE          LIMIT  MAXTTL  BYTES_NEEDED
sales hadoop hadoop  rwxr-xr-x  unlimited  never    84
```

```
[hadoop@master1 ~]$ hdfs dfsadmin -allowSnapshot /projects
Allowing snapshot on /projects succeeded
[hadoop@master1 ~]$ hdfs dfs -ls /projects
Found 1 items
-rw-r--r--  1 hadoop supergroup          42 2017-01-05 02:10 /projects/test
[hadoop@master1 ~]$ hdfs dfs -ls /projects/.snapshot
[hadoop@master1 ~]$ █
```

```
[hadoop@master1 ~]$ hdfs dfs -ls /projects/.snapshot
Found 1 items
drwxr-xr-x  - hadoop supergroup          0 2017-01-05 03:39 /projects/.snapshot/jan52016
```

```
$ hdfs dfs -rm -r /projects/
INFO fs.TrashPolicyDefault: Namenode trash configuration: Deletion interval = 10080 minutes, Emptier interval = 45 minutes.
to trash: hdfs://master1.cyrus.com:9000/projects: The directory /projects cannot be deleted since /projects is snapshottable
```

```
<property>
<name>dfs.datanode.data.dir</name>
<value>[DISK]file:///grid/dn/disk0,
[SSD]file:///grid/dn/ssd0,
[ARCHIVE]file:///grid/dn/archive0,
[RAM_DISK]file:///grid/dn/ram0</value>
<description>A datanode storage location /grid/dn/disk0 on DISK</description>
<description>A datanode storage location /grid/dn/ssd0 on SSD</description>
<description>A datanode storage location /grid/dn/archive0 on ARCHIVE</description>
<description>A datanode storage location /grid/dn/ram0 on RAM_DISK</description>
</property>
```

```
[hadoop@master1 ~]$ yarn radmin -getServiceState rm1
active
[hadoop@master1 ~]$ yarn radmin -getServiceState rm2
standby
```

```
[hadoop@master1 ~]$ yarn radmin -transitionToActive rm2
Automatic failover is enabled for org.apache.hadoop.yarn.client.RMHAServiceTarget@1e67a849
Refusing to manually manage HA state, since it may cause
a split-brain scenario or other incorrect state.
If you are very sure you know what you are doing, please
specify the --forcemanual flag.
```

```
[hadoop@master1 ~]$ yarn radmin -transitionToStandby rm1 --forcemanual
You have specified the --forcemanual flag. This flag is dangerous, as it can induce a split-brain scenario that WILL CORRUPT your HDFS namespace, possibly irrecoverably.

It is recommended not to use this flag, but instead to shut down the cluster and disable automatic failover if you prefer to manually manage your HA state.

You may abort safely by answering 'n' or hitting ^C now.

Are you sure you want to continue? (Y or N) Y
17/05/10 10:02:58 WARN ha.HAAdmin: Proceeding with manual HA state management even though
automatic failover is enabled for org.apache.hadoop.yarn.client.RMHAServiceTarget@1e67a849
```

## Chapter 5: Schedulers

```
[user1@master1 ~]$ id
uid=503(user1) gid=503(user1) groups=503(user1),501(hadoop)
[user1@master1 ~]$ hadoop fs -put test.txt /user/user1/
[user1@master1 ~]$ hadoop fs -ls /user/user1/
Found 1 items
-rw-r--r--  1 user1 hadoop          20 2017-02-02 10:36 /user/user1/test.txt
```

```
[root@master1 ~]# groupadd prod
[root@master1 ~]# groupadd dev
[root@master1 ~]# useradd -G prod p1
[root@master1 ~]# useradd -G prod p2
[root@master1 ~]# useradd -G dev d1
[root@master1 ~]# useradd -G dev d2
```

```
[hadoop@master1 ~]$ cat /opt/cluster/hadoop/etc/hadoop/fair-scheduler.xml
<allocations>
  <queue name="prod_queue">
    <minResources>10000 mb,0vcores</minResources>
    <maxResources>90000 mb,0vcores</maxResources>
    <maxRunningApps>50</maxRunningApps>
    <maxAMShare>0.3</maxAMShare>
    <weight>2.0</weight>
    <schedulingPolicy>fair</schedulingPolicy>
  </queue>

  <queueMaxAMShareDefault>0.5</queueMaxAMShareDefault>

  <queue name="dev_queue">
    <minResources>5000 mb,0vcores</minResources>
    <maxResources>10000 mb,0vcores</maxResources>
    <maxRunningApps>30</maxRunningApps>
    <maxAMShare>0.1</maxAMShare>
    <weight>2.0</weight>
    <schedulingPolicy>fair</schedulingPolicy>
  </queue>
</allocations>
```

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share Used

- root
  - + root.default
  - + root.dev\_queue
  - + root.prod\_queue

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- root
  - + root.default
  - + root.dev\_queue
  - + root.prod\_queue
  - + root.hadoop

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- root
  - + root.default
  - + root.dev\_queue
  - + root.hadoop
  - + root.prod\_queue
  - + root.d1

```
[d1@master1 ~]$ hadoop fs -ls -d /user/d1
drwxr-xr-x  - d1 d1          0 2017-02-03 02:06 /user/d1
[d1@master1 ~]$ hadoop fs -ls -d /tmp
drwxrwxrwt  - hadoop supergroup 0 2017-02-03 01:59 /tmp
```

```
[hadoop@master1 hadoop]$ cat fair-scheduler.xml
<allocations>

  <pool name="prod">
    <minMaps>10</minMaps>
    <minReduces>5</minReduces>
  </pool>

  <pool name="dev">
    <minMaps>10</minMaps>
    <minReduces>5</minReduces>
  </pool>

</allocations>
```

### Application Queues

**Legend:**

Steady Fair Share

Instantaneous Fair Share

└─ root

└─ + root.default

└─ + root.dev

└─ + root.prod

### Application Queues

**Legend:**

Steady Fair Share

Instantaneous Fair Share

Used

└─ root

└─ + root.d1

└─ + root.default

└─ + root.dev

└─ + root.prod

└─ + root.hadoop

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- root
  - + root.default
  - + root.dev
  - + root.prod
  - + root.d1

```
[hadoop@jt1 hadoop]$ yarn jar /opt/cluster/hadoop/share/hadoop/mapreduce/hadoop-mapreduce-examples-2.7.3.jar wordcount /file1.txt /out
```

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- root
  - + root.dev
  - + root.prod
  - + root.default

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- root
  - + root.default
  - + root.dev
  - + root.prod
  - + root.hadoop

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share Used

- root
  - + root.dev
  - + root.prod
  - + root.default



**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- └─ root
  - └─ + root.default
  - └─ + root.dev
  - └─ + root.prod
  - └─ + root.hadoop

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- └─ root
  - └─ + root.default
  - └─ + root.dev
  - └─ + root.prod

**Application Queues**

**Legend:** Steady Fair Share Instantaneous Fair Share

- └─ root
  - └─ + root.default
  - └─ + root.prod
  - └─ + root.dev

```
[d1@jt1 ~]$ mapred queue -showacIs
17/02/04 09:03:25 INFO client.RMProxy: Connecting to Queue acIs for user : d1

Queue Operations
=====
root ADMINISTER_QUEUE,SUBMIT_APPLICATIONS
root.default ADMINISTER_QUEUE,SUBMIT_APPLICATIONS
root.dev ADMINISTER_QUEUE,SUBMIT_APPLICATIONS
root.prod ADMINISTER_QUEUE,SUBMIT_APPLICATIONS
```

```
[d1@jt1 ~]$ mapred queue -showacls
17/02/04 09:10:12 INFO client.RMProxy: Connecting
Queue acls for user : d1

Queue Operations
=====
root
root.default
root.dev ADMINISTER_QUEUE,SUBMIT_APPLICATIONS
root.prod
```

```
17/02/04 09:12:31 INFO mapreduce.JobSubmitter: Cleaning up the staging
job_1486179599793_0001
java.io.IOException: org.apache.hadoop.yarn.exceptions.YarnException:
0001 to YARN : User d1 cannot submit applications to queue root.prod
```

Scheduler Type	Scheduling Resource Type
Capacity Scheduler	[MEMORY]

### Application Queues

**Legend:** Capacity Used Used (over capacity)

- └─ root
  - └─ + Queue: default

Scheduler Type	Scheduling Resource Type
Capacity Scheduler	[MEMORY]

### Application Queues

**Legend:** Capacity Used Used (over capacity)

- └─ root
  - └─ + Queue: default
  - └─ + Queue: dev
  - └─ + Queue: prod

Scheduler Type	Scheduling Resource Type
Capacity Scheduler	[MEMORY]

Application Queues	
<b>Legend:</b>	Capacity    Used    Used (over capacity)
└─ root	
+ Queue: default	
+ Queue: dev	
+ Queue: prod	

```
[hadoop@jt1 ~]$ mapred queue -list
17/02/04 19:56:26 INFO client.RMProxy:
=====
Queue Name : default
Queue State : running
Scheduling Info : Capacity: 30.000002,
=====
Queue Name : dev
Queue State : running
Scheduling Info : Capacity: 20.0, Maxim
=====
Queue Name : prod
Queue State : running
Scheduling Info : Capacity: 50.0, Maxim
```

```
[hadoop@jt1 hadoop]$ yarn node -status dn1.cluster1.com:57071
17/02/05 09:34:04 INFO client.RMProxy: Connecting to Resource
Node Report :
Node-Id : dn1.cluster1.com:57071
Rack : /default-rack
Node-State : RUNNING
Node-Http-Address : dn1.cluster1.com:8042
Last-Health-Update : Sun 05/Feb/17 09:32:18:56IST
Health-Report :
Containers : 0
Memory-Used : 0MB
Memory-Capacity : 8192MB
CPU-Used : 0 vcores
CPU-Capacity : 8 vcores
Node-Labels : memory
```

```
[hadoop@jtl hadoop]$ ./share/hadoop/tools/sls/bin/slsrun.sh --input-rumen=`pwd`/share/hadoop/tools/sls/sample-data/2jobs2min-rumen-jh.json --output-dir=out_sim
17/02/05 10:08:03 INFO conf.Configuration: found resource core-site.xml at file:/opt/cluster/hadoop-2.7.3/etc/hadoop/core-site.xml
17/02/05 10:08:03 INFO security.Groups: clearing userToGroupsMap cache
17/02/05 10:08:03 INFO conf.Configuration: found resource yarn-site.xml at file:/opt/cluster/hadoop-2.7.3/etc/hadoop/yarn-site.xml
17/02/05 10:08:04 INFO event.AsyncDispatcher: Registering class org.apache.hadoop.yarn.server.resourcemanager.RMFatalEventType for class org.apache.hadoop.yarn.server.resourcemanager.ResourceManager$RMFatalEventDispatcher
17/02/05 10:08:04 INFO security.NMTokenSecretManagerInRM: NMTokenKeyRollingInterval: 86400000ms and NMTokenKeyActivationDelay: 900000ms
17/02/05 10:08:04 INFO security.RMContainerTokenSecretManager: ContainerTokenKeyRollingInterval: 86400000ms and ContainerTokenKeyActivationDelay: 900000ms
17/02/05 10:08:04 INFO security.AMRMTokenSecretManager: AMRMTokenKeyRollingInterval: 86400000ms and AMRMTokenKeyActivationDelay: 900000ms
```

```
[hadoop@jtl hadoop]$ ls -l out_sim/metrics/
total 356
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.allocate.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.APP_ADDED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.APP_ATTEMPT_ADDED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.APP_ATTEMPT_REMOVED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.APP_REMOVED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.CONTAINER_EXPIRED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.CONTAINER_RESCHEDULED.csv
-rw-rw-r-- 1 hadoop hadoop 2356 Feb  5 10:08 counter.scheduler.operation.handle.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.DROP_RESERVATION.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.KILL_CONTAINER.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.NODE_ADDED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.NODE_LABELS_UPDATE.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.NODE_REMOVED.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.NODE_RESOURCE_UPDATE.csv
-rw-rw-r-- 1 hadoop hadoop 2347 Feb  5 10:08 counter.scheduler.operation.handle.NODE_UPDATE.csv
-rw-rw-r-- 1 hadoop hadoop 2153 Feb  5 10:08 counter.scheduler.operation.handle.PREEMPT_CONTAINER.csv
-rw-rw-r-- 1 hadoop hadoop 14738 Feb  5 10:08 sampler.scheduler.operation.allocate.timecost.csv
-rw-rw-r-- 1 hadoop hadoop 14738 Feb  5 10:08 sampler.scheduler.operation.handle.APP_ADDED.timecost.csv
-rw-rw-r-- 1 hadoop hadoop 14738 Feb  5 10:08 sampler.scheduler.operation.handle.APP_ATTEMPT_ADDED.timecost.csv
-rw-rw-r-- 1 hadoop hadoop 14738 Feb  5 10:08 sampler.scheduler.operation.handle.APP_ATTEMPT_REMOVED.timecost.csv
-rw-rw-r-- 1 hadoop hadoop 14738 Feb  5 10:08 sampler.scheduler.operation.handle.APP_REMOVED.timecost.csv
```

## Chapter 6: Backup and Recovery

```
[hadoop@master1 ~]$ hdfs dfsadmin -safemode enter
Safe mode is ON
[hadoop@master1 ~]$ hdfs dfsadmin -saveNamespace
Save namespace successful
```

```
[hadoop@master1 ~]$ hdfs dfsadmin -rollEdits
Successfully rolled edit logs.
New segment starts at txid 132200
```

```
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 1200
-rw-rw-r-- 1 hadoop hadoop 1048576 Mar 12 21:15 edits_inprogress_0000000000000132203
-rw-rw-r-- 1 hadoop hadoop 77888 Mar 12 21:01 fsimage_0000000000000132197
-rw-rw-r-- 1 hadoop hadoop 62 Mar 12 21:01 fsimage_0000000000000132197.md5
-rw-rw-r-- 1 hadoop hadoop 77888 Mar 12 21:03 fsimage_0000000000000132201
-rw-rw-r-- 1 hadoop hadoop 62 Mar 12 21:03 fsimage_0000000000000132201.md5
-rw-rw-r-- 1 hadoop hadoop 7 Mar 12 21:15 seen_txid
-rw-rw-r-- 1 hadoop hadoop 202 Mar 12 21:03 VERSION
```

```
[hadoop@master1 current]$ hdfs oev -i edits_00000000000000009-00000000000000208 -o edits.txt -p stats
[hadoop@master1 current]$ head edits.txt
VERSION : -63
OP_ADD ( 0): 33
OP_RENAME_OLD ( 1): 33
OP_DELETE ( 2): null
OP_MKDIR ( 3): null
OP_SET_REPLICATION ( 4): null
OP_DATANODE_ADD ( 5): null
OP_DATANODE_REMOVE ( 6): null
OP_SET_PERMISSIONS ( 7): null
OP_SET_OWNER ( 8): null
```

```
[hadoop@master1 ~]$ jar -ft $HADOOP_HOME/share/hadoop/hdfs/hadoop-hdfs-2.7.2.jar | grep default
hdfs-default.xml
[hadoop@master1 ~]$ jar -ft $HADOOP_HOME/share/hadoop/hdfs/hadoop-hdfs-2.7.2.jar | grep default
hdfs-default.xml
[hadoop@master1 ~]$ clear
[hadoop@master1 ~]$ jar -ft $HADOOP_HOME/share/hadoop/hdfs/hadoop-hdfs-2.7.2.jar | grep default
hdfs-default.xml
[hadoop@master1 ~]$ jar -fx $HADOOP_HOME/share/hadoop/hdfs/hadoop-hdfs-2.7.2.jar hdfs-default.xml
[hadoop@master1 ~]$ ls -l hdfs-default.xml
-rw-rw-r-- 1 hadoop hadoop 81867 May 16 2016 hdfs-default.xml
```

```
# Rolling File Appender - cap space usage at 5gb.
#
hadoop.log.maxfilesize=256MB
hadoop.log.maxbackupindex=20
log4j.appender.RFA=org.apache.log4j.RollingFileAppender
log4j.appender.RFA.File=${hadoop.log.dir}/${hadoop.log.file}

log4j.appender.RFA.MaxFileSize=${hadoop.log.maxfilesize}
```

```
# Pattern format: Date LogLevel LoggerName LogMessage
log4j.appender.RFA.layout.ConversionPattern=%d{ISO8601} %p %c: %m%n
# Debugging Pattern format
#log4j.appender.RFA.layout.ConversionPattern=%d{ISO8601} %-5p %c{2} (%F:%M(%L)) - %m%n
```

```
# Rollover at midnight
log4j.appender.DRFA.DatePattern=.yyyy-MM-dd

log4j.appender.DRFA.layout=org.apache.log4j.PatternLayout
```



```
#
# HDFS block state change log from block manager
#
# Uncomment the following to suppress normal block state change
# messages from BlockManager in NameNode.
#log4j.logger.BlockStateChange=WARN
```

```
# hdfs audit logging
#
hdfs.audit.logger=INFO,NullAppender
hdfs.audit.log.maxfilesize=256MB
hdfs.audit.log.maxbackupindex=20
```

```
*****/
17/03/14 06:55:57 INFO namenode.NameNode: registered UNIX signal handlers for [TERM, HUP, INT]
17/03/14 06:55:57 INFO namenode.NameNode: createNameNode [-importCheckpoint]
17/03/14 06:55:57 INFO impl.MetricsConfig: loaded properties from hadoop-metrics2.properties
17/03/14 06:55:57 INFO impl.MetricsSystemImpl: Scheduled snapshot period at 10 second(s).
17/03/14 06:55:57 INFO impl.MetricsSystemImpl: NameNode metrics system started
17/03/14 06:55:57 INFO namenode.NameNode: fs.defaultFS is hdfs://master1.cyrus.com:9000
```

```
6411 Jps
6365 SecondaryNameNode
```

```
[ec2-user@master1 ~]$ hostname -f
master1.cyrus.com
[ec2-user@master1 ~]$ ifconfig
eth0      Link encap:Ethernet  HWaddr 0A:D3:E1:29:9D:97
          inet addr:10.0.0.104  Bcast:10.0.0.255  Mask:255.255.255.0
          inet6 addr: fe80::8d3:elff:fe29:9d97/64  Scope:Link
```

```
[hadoop@master1 hadoop]$ sudo service nfs stop
Shutting down NFS daemon: [ FAILED ]
Shutting down NFS mountd: [ FAILED ]
Shutting down NFS quotas: [ FAILED ]
[hadoop@master1 hadoop]$ sudo service rpcbind stop
Stopping rpcbind: [ OK ]
```

```
[root@master1 ~]# jps
7425 SecondaryNameNode
7889 Jps
7207 NameNode
7819 Portmap
```

```
[hadoop@master1 hadoop]$ jps
7425 SecondaryNameNode
7921 Nfs3
7207 NameNode
8009 Jps
```

```
[hadoop@master1 hadoop]$ rpcinfo -p `hostname`
  program vers  proto  port  service
  100005   3   udp   4242  mountd
  100005   1   tcp   4242  mountd
  100000   2   udp    111  portmapper
  100000   2   tcp    111  portmapper
  100005   3   tcp   4242  mountd
  100005   2   tcp   4242  mountd
  100003   3   tcp   2049  nfs
  100005   2   udp   4242  mountd
  100005   1   udp   4242  mountd
```

```
[hadoop@master1 nfs]$ showmount -e `hostname`
Export list for master1.cyrus.com:
/ *
```

```
[hadoop@master1 hadoop]$ cd /data/nfs/
[hadoop@master1 nfs]$ ls -l
total 4
drwxr-xr-x 4 hadoop 2584148964 128 Mar  2 18:47 out
drwxr-xr-x 4 hadoop 2584148964 128 Mar  2 19:03 out1
drwxr-xr-x 2 hadoop 2584148964   64 Jan  5 04:11 projects
drwxr-xr-x 2 hadoop 2584148964   64 Jan  5 01:11 sharedcache
-rw-r--r-- 1 hadoop 2584148964   42 Mar  2 18:46 test
drwxrwxrwt 4 hadoop 2584148964 128 Feb  3 01:59 tmp
drwxrwxrwt 7 hadoop 2584148964 224 Feb  3 02:01 user
```



```

[hadoop@master1 ~]$ hadoop fs -ls /
Found 4 items
drwxr-xr-x   - hadoop supergroup          0 2017-01-05 04:11 /projects
-rw-r--r--   1 hadoop supergroup          42 2017-03-02 18:46 /test
drwxrwxrwt   - hadoop supergroup          0 2017-02-03 01:59 /tmp
drwxrwxrwt   - hadoop supergroup          0 2017-02-03 02:01 /user
[hadoop@master1 ~]$ cp fsimage.txt /data/nfs/
[hadoop@master1 ~]$ hadoop fs -ls /
Found 5 items
-rw-rw-r--   1 hadoop supergroup          0 2017-03-14 22:01 /fsimage.txt
drwxr-xr-x   - hadoop supergroup          0 2017-01-05 04:11 /projects
-rw-r--r--   1 hadoop supergroup          42 2017-03-02 18:46 /test
drwxrwxrwt   - hadoop supergroup          0 2017-02-03 01:59 /tmp
drwxrwxrwt   - hadoop supergroup          0 2017-02-03 02:01 /user

```

```

Planning to load image: FSImageFile(file=/data/namenode1/current/fsimage_00
ormatPBINode: Loading 3 INodes.
ormatProtobuf: Loaded FSImage in 0 seconds.
Loaded image for txid 244 from /data/namenode1/current/fsimage_000000000000
Reading org.apache.hadoop.hdfs.server.namenode.RedundantEditLogInputStream@
Start loading edits file /data/namenode1/current/edits_000000000000000245-
nputStream: Fast-forwarding stream '/data/namenode1/current/edits_000000000
Edits file /data/namenode1/current/edits_000000000000000245-0000000000000
Reading org.apache.hadoop.hdfs.server.namenode.RedundantEditLogInputStream@
Start loading edits file /data/namenode1/current/edits_000000000000000252-
nputStream: Fast-forwarding stream '/data/namenode1/current/edits_0000000000
Edits file /data/namenode1/current/edits_000000000000000252-0000000000000
stem: Need to save fs image? false (staleImage=false, haEnabled=false, isRol
g: Starting log segment at 253
e: initialized with 0 entries 0 lookups
stem: Finished loading FSImage in 322 msec
Save namespace ...
g: Ending log segment 253, 253
g: logSyncAll toSyncToTxId=254 lastSyncedTxid=253 mostRecentTxid=254
g: Done logSyncAll lastWrittenTxId=254 lastSyncedTxid=254 mostRecentTxid=254

```

```

[hadoop@master1 ~]$ hdfs getconf -confkey dfs.namenode.name.dir
file:/data/namenode1,file:/data/namenode2
[hadoop@master1 ~]$ df -h | grep namenode
/dev/xvdb1          20G   47M   19G    1% /data/namenode1
filer.cyrus.com:/nfs 7.8G  4.5G  3.3G  58% /data/namenode2

```

```
[hadoop@master1 ~]$ hdfs getconf -confkey dfs.namenode.name.dir
file:/data/namenode1
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 4120
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 14 01:15 edits_000000000000000001-000000000000000001
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:05 edits_000000000000000002-000000000000000002
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:10 edits_000000000000000003-0000000000000000217
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:31 edits_inprogress_0000000000000000218
-rw-rw-r-- 1 hadoop hadoop 317 Apr 14 01:15 fsimage_00000000000000000000
-rw-rw-r-- 1 hadoop hadoop 62 Apr 14 01:15 fsimage_00000000000000000000.md5
-rw-rw-r-- 1 hadoop hadoop 317 Apr 18 01:05 fsimage_00000000000000000001
-rw-rw-r-- 1 hadoop hadoop 62 Apr 18 01:05 fsimage_00000000000000000001.md5
-rw-rw-r-- 1 hadoop hadoop 4 Apr 18 01:31 seen_txid
-rw-rw-r-- 1 hadoop hadoop 215 Apr 18 01:05 VERSION
```

```
[hadoop@master1 ~]$ hdfs dfsadmin -rollEdits
Successfully rolled edit logs.
New segment starts at txid 221
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 5148
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 14 01:15 edits_000000000000000001-000000000000000001
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:05 edits_000000000000000002-000000000000000002
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:10 edits_000000000000000003-0000000000000000217
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:31 edits_0000000000000000218-0000000000000000218
-rw-rw-r-- 1 hadoop hadoop 42 Apr 18 01:36 edits_0000000000000000219-0000000000000000220
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:36 edits_inprogress_0000000000000000221
```

```
[hadoop@master1 ~]$ mv /data/namenode1/current/edits_0000000000000000* ~/image_backup/
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 1048
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:36 edits_inprogress_0000000000000000221
-rw-rw-r-- 1 hadoop hadoop 317 Apr 14 01:15 fsimage_00000000000000000000
-rw-rw-r-- 1 hadoop hadoop 62 Apr 14 01:15 fsimage_00000000000000000000.md5
-rw-rw-r-- 1 hadoop hadoop 317 Apr 18 01:05 fsimage_00000000000000000001
-rw-rw-r-- 1 hadoop hadoop 62 Apr 18 01:05 fsimage_00000000000000000001.md5
-rw-rw-r-- 1 hadoop hadoop 4 Apr 18 01:36 seen_txid
-rw-rw-r-- 1 hadoop hadoop 215 Apr 18 01:05 VERSION
```

```
2017-04-18 01:41:45,012 INFO org.apache.hadoop.metrics2.impl.MetricsSystemImpl: NameNode metrics system sto
2017-04-18 01:41:45,012 INFO org.apache.hadoop.metrics2.impl.MetricsSystemImpl: NameNode metrics system shu
2017-04-18 01:41:45,012 ERROR org.apache.hadoop.hdfs.server.namenode.NameNode: Failed to start namenode.
java.io.IOException: Gap in transactions. Expected to be able to read up until at least txid 221 but unable
txid 2
```

```
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 6184
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:58 edits_000000000000000001-000000000000000001
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:58 edits_000000000000000002-000000000000000002
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:58 edits_000000000000000003-000000000000000003
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:58 edits_0000000000000000218-0000000000000000218
-rw-rw-r-- 1 hadoop hadoop 42 Apr 18 01:58 edits_0000000000000000219-0000000000000000220
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 01:36 edits_0000000000000000221-0000000000000000221
-rw-rw-r-- 1 hadoop hadoop 42 Apr 18 02:00 edits_0000000000000000222-0000000000000000223
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 02:00 edits_inprogress_0000000000000000224
-rw-rw-r-- 1 hadoop hadoop 317 Apr 14 01:15 fsimage_0000000000000000000
-rw-rw-r-- 1 hadoop hadoop 62 Apr 14 01:15 fsimage_0000000000000000000.md5
-rw-rw-r-- 1 hadoop hadoop 317 Apr 18 01:05 fsimage_0000000000000000001
-rw-rw-r-- 1 hadoop hadoop 62 Apr 18 01:05 fsimage_0000000000000000001.md5
-rw-rw-r-- 1 hadoop hadoop 2389 Apr 18 02:00 fsimage_00000000000000000223
-rw-rw-r-- 1 hadoop hadoop 62 Apr 18 02:00 fsimage_00000000000000000223.md5
```

```
[hadoop@master1 ~]$ mv /data/namenode1/current/edits_0000000000000000* ~/image_backup/
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 1056
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 02:00 edits_inprogress_0000000000000000224
-rw-rw-r-- 1 hadoop hadoop 317 Apr 14 01:15 fsimage_0000000000000000000
-rw-rw-r-- 1 hadoop hadoop 62 Apr 14 01:15 fsimage_0000000000000000000.md5
-rw-rw-r-- 1 hadoop hadoop 317 Apr 18 01:05 fsimage_0000000000000000001
-rw-rw-r-- 1 hadoop hadoop 62 Apr 18 01:05 fsimage_0000000000000000001.md5
-rw-rw-r-- 1 hadoop hadoop 2389 Apr 18 02:00 fsimage_00000000000000000223
-rw-rw-r-- 1 hadoop hadoop 62 Apr 18 02:00 fsimage_00000000000000000223.md5
```

```
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 6172
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 02:12 edits_0000000000000000224-0000000000000000224
-rw-rw-r-- 1 hadoop hadoop 42 Apr 18 02:12 edits_0000000000000000225-0000000000000000226
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 02:06 edits_0000000000000000227-0000000000000000227
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 02:15 edits_0000000000000000228-0000000000000000228
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 02:43 edits_0000000000000000229-0000000000000000229
-rw-rw-r-- 1 hadoop hadoop 42 Apr 18 02:50 edits_0000000000000000230-0000000000000000231
-rw-rw-r-- 1 hadoop hadoop 788 Apr 18 04:49 edits_0000000000000000232-0000000000000000241
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 04:49 edits_0000000000000000242-0000000000000000242
```

```
[hadoop@master1 ~]$ ls -l /data/namenode1/current/
total 2080
-rw-rw-r-- 1 hadoop hadoop 788 Apr 18 04:49 edits_0000000000000000232-0000000000000000241
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 04:49 edits_0000000000000000242-0000000000000000242
-rw-rw-r-- 1 hadoop hadoop 42 Apr 18 04:52 edits_0000000000000000243-0000000000000000244
-rw-rw-r-- 1 hadoop hadoop 1048576 Apr 18 04:52 edits_inprogress_0000000000000000245
```

```
[hadoop@master1 ~]$ hadoop fs -cat /test
This is test
[hadoop@master1 ~]$ hadoop fs -rm -skipTrash /test
Deleted /test
```

```
<RECORD>
  <OPCODE>OP_DELETE</OPCODE>
  <DATA>
    <TXID>252</TXID>
    <LENGTH>0</LENGTH>
    <PATH>/test</PATH>
    <TIMESTAMP>1492491267769</TIMESTAMP>
    <RPC_CLIENTID>34809cac-7e94-4881-a4d7-10c54d7aac1a</RPC_CLIENTID>
    <RPC_CALLID>1</RPC_CALLID>
  </DATA>
</RECORD>
```

```
[hadoop@master1 ~]$ hadoop-daemon.sh start namenode
starting namenode, logging to /opt/cluster/hadoop-2.8.0/logs/h
[hadoop@master1 ~]$ hadoop fs -ls /
Found 2 items
-rw-r--r--   3 hadoop hadoop          13 2017-04-18 04:53 /test
drwxr-xr-x   - hadoop hadoop          0 2017-04-18 01:10 /user
```

## Chapter 7: Data Ingestion and Workflow

```
[root@edge1 ~]# cat /etc/profile.d/hadoopenv.sh
export HADOOP_HOME=/opt/cluster/hadoop
export HADOOP_MAPRED_HOME=$HADOOP_HOME
export HADOOP_COMMON_HOME=$HADOOP_HOME
export HADOOP_HDFS_HOME=$HADOOP_HOME
export YARN_HOME=$HADOOP_HOME
export HADOOP_CONF_DIR=$HADOOP_HOME/etc/hadoop
export YARN_CONF_DIR=$HADOOP_HOME/etc/hadoop
export HBASE_HOME=/opt/cluster/hbase
export JAVA_HOME=/usr/java/latest
export PATH=$PATH:/usr/local/scala/bin

export HIVE_HOME=/opt/cluster/hive
PATH=$HBASE_HOME/bin:$HIVE_HOME/bin:$HADOOP_HOME/bin
export PATH
```

```
[hadoop@edge1 ~]$ which hive
/opt/cluster/hive/bin/hive
[hadoop@edge1 ~]$ echo $HIVE_HOME
/opt/cluster/hive
```

```
[hadoop@edge1 hive]$ pwd
/opt/cluster/hive
[hadoop@edge1 hive]$ ls -l
total 476
drwxr-xr-x 3 hadoop hadoop 4096 Oct 23 21:07 bin
drwxr-xr-x 2 hadoop hadoop 4096 Oct 25 11:41 conf
drwxr-xr-x 4 hadoop hadoop 4096 Jun 11 2016 examples
drwxr-xr-x 7 hadoop hadoop 4096 Oct 12 23:11 hcatalog
drwxr-xr-x 4 hadoop hadoop 4096 Aug 15 2016 lib
-rw-rw-r-- 1 hadoop hadoop 24754 Aug 26 2016 LICENSE
-rw-rw-r-- 1 hadoop hadoop 397 Aug 26 2016 NOTICE
-rw-rw-r-- 1 hadoop hadoop 4366 Aug 26 2016 README.txt
-rw-rw-r-- 1 hadoop hadoop 421129 Aug 26 2016 RELEASE_NOTES.txt
drwxr-xr-x 3 hadoop hadoop 4096 Jun 11 2016 scripts
```

```
[hadoop@edge1 conf]$ grep warehouse hive-default.xml
<name>hive.metastore.warehouse.dir</name>
<value>/user/hive/warehouse</value>
```



```
hive> create database test_hive_db;
OK
Time taken: 0.141 seconds
hive> quit;
[hadoop@edge1 ~]$ hadoop fs -ls /user/hive/warehouse
Found 1 items
drwxr-xr-x - hadoop hadoop 0 2017-04-19 01:12 /user/hive/warehouse/test_hive_db.db
```

```
[hadoop@edge1 conf]$ ls -l metastore_db/
total 28
-rw-rw-r-- 1 hadoop hadoop 4 Mar 21 21:13 dbex.lck
-rw-rw-r-- 1 hadoop hadoop 38 Mar 21 21:13 db.lck
drwxrwxr-x 2 hadoop hadoop 4096 Mar 21 21:09 log
-rw-rw-r-- 1 hadoop hadoop 608 Mar 21 21:09 README_DO_NOT_TOUCH_FILES.txt
drwxrwxr-x 2 hadoop hadoop 4096 Mar 21 21:09 seg0
-rw-rw-r-- 1 hadoop hadoop 931 Mar 21 21:09 service.properties
drwxrwxr-x 2 hadoop hadoop 4096 Mar 21 21:13 tmp
```

```
[hadoop@edge1 conf]$ hive --service hiveserver2&
[1] 24549
[hadoop@edge1 conf]$ jps
24692 Jps
24549 RunJar
```

```
Caused by: java.sql.SQLException: Failed to start database 'metastore_db' with class loader sun.misc.Launcher$AppClassLoader@214c265e, see the next exception for details.
    at org.apache.derby.impl.jdbc.SQLExceptionFactory.getSQLException(Unknown Source)
    at org.apache.derby.impl.jdbc.SQLExceptionFactory40.wrapArgsForTransportAcrossDRDA(Unknown Source)
    ... 74 more
Caused by: java.sql.SQLException: Another instance of Derby may have already booted the database /opt/clus
ache-hive-1.2.1-bin/conf/metastore_db.
    at org.apache.derby.impl.jdbc.SQLExceptionFactory.getSQLException(Unknown Source)
    at org.apache.derby.impl.jdbc.SQLExceptionFactory40.wrapArgsForTransportAcrossDRDA(Unknown Source)
    at org.apache.derby.impl.jdbc.SQLExceptionFactory40.getSQLException(Unknown Source)
    at org.apache.derby.impl.jdbc.Util.generateCsSQLException(Unknown Source)
```

```
[hadoop@edge1 conf]$ rm -rf metastore_db
[hadoop@edge1 conf]$ rm derby.log
[hadoop@edge1 conf]$ clear
[hadoop@edge1 conf]$ hadoop fs -ls /user/hive/warehouse/
Found 1 items
drwxr-xr-x - hadoop supergroup 0 2017-03-21 21:13 /user/hive/warehouse
/test_hive_db.db
[hadoop@edge1 conf]$ hadoop fs -rm -r /user/hive/warehouse/test_hive_db.db
```

```
[hadoop@edge1 cluster]$ hive --service hiveserver2&
[1] 30977
[hadoop@edge1 cluster]$ hive

Logging initialized using configuration in jar:file
.1.jar!/hive-log4j.properties
hive> CREATE DATABASE test;
OK
Time taken: 0.836 seconds
hive> █
```

```
[hadoop@edge1 ~]$ hadoop fs -ls /user/hive/warehouse/
Found 1 items
drwxr-xr-x - hadoop supergroup 0 2017-03-21
22:20 /user/hive/warehouse/test.db
```

```
[hadoop@edge1 ~]$ mysql -u hadoop -h master1.cyrus.com -p
Enter password:
Welcome to the MySQL monitor. Commands end with ; or \g.
Your MySQL connection id is 18
Server version: 5.5.54 MySQL Community Server (GPL)

Copyright (c) 2000, 2016, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| hive_db |
| mysql |
| performance_schema |
| test |
+-----+
5 rows in set (0.00 sec)
```

```
[hadoop@edge1 ~]$ beeline
Beeline version 1.2.1 by Apache Hive
beeline> !connect jdbc:hive2://localhost:10000 hadoop hadoop
Connecting to jdbc:hive2://localhost:10000
Connected to: Apache Hive (version 1.2.1)
Driver: Hive JDBC (version 1.2.1)
Transaction isolation: TRANSACTION_REPEATABLE_READ
0: jdbc:hive2://localhost:10000> show databases;
+-----+--+
| database_name |
+-----+--+
| default      |
| test         |
| testdb       |
+-----+--+
3 rows selected (0.158 seconds)
```

```
hive> use test;
OK
Time taken: 0.015 seconds
hive> CREATE TABLE pokes (foo INT, bar STRING);
OK
Time taken: 0.199 seconds
hive> LOAD DATA LOCAL INPATH './hive/examples/files/kv1.txt' INTO TABLE pokes;
Loading data to table test.pokes
Table test.pokes stats: [numFiles=1, totalSize=5812]
OK
Time taken: 0.473 seconds
```

```
[hadoop@master1 ~]$ hadoop fs -ls /user/hive/warehouse/
Found 1 items
drwxr-xr-x  - hadoop supergroup          0 2017-03-22
03:53 /user/hive/warehouse/test.db
[hadoop@master1 ~]$ hadoop fs -ls /user/hive/warehouse/
test.db/pokes/
Found 1 items
-rwxr-xr-x  2 hadoop supergroup        5812 2017-03-22
03:53 /user/hive/warehouse/test.db/pokes/kv1.txt
```



```
hive> desc pokes;
OK
foo                int
bar                string
Time taken: 0.067 seconds, Fetched: 2 row(s)
hive> select * from pokes limit 5;
OK
238      val_238
86       val_86
311     val_311
27      val_27
165     val_165
Time taken: 0.262 seconds, Fetched: 5 row(s)
```

```
[hadoop@master1 ~]$ hadoop fs -cat /user/hive/warehouse/test.db/pokes/kv1.txt | head -1 | tr -d "\n" | od -An -t dC
6
 50  51  56   1 118  97 108  95  50  51  5
```

```
hive> use test;
OK
Time taken: 0.71 seconds
hive> drop table pokes;
Moved: 'hdfs://master1.cyr.com:9000/user/hadoop/.Trash'
OK
Time taken: 0.467 seconds
```

```
hive> !hadoop fs -ls /user/hive/warehouse/test.db/;
hive> !hadoop fs -put /opt/cluster/hive/examples/files/kv1.txt /user/hive/warehouse/test.db/;
hive> !hadoop fs -ls /user/hive/warehouse/test.db/;
Found 1 items
-rw-r--r--    2 hadoop supergroup      5812 2017-03
-22 04:23 /user/hive/warehouse/test.db/kv1.txt
```

```
hive> !hadoop fs -ls /user/hive/warehouse/test.db/invites;
Found 2 items
drwxr-xr-x  - hadoop supergroup          0 2017-03-22 05:
25 /user/hive/warehouse/test.db/invites/ds=2008-08-08
drwxr-xr-x  - hadoop supergroup          0 2017-03-22 05:
25 /user/hive/warehouse/test.db/invites/ds=2008-08-15
```

```
hive> explain select * from invites where ds='2008-08-08';
OK
STAGE DEPENDENCIES:
  Stage-0 is a root stage

STAGE PLANS:
  Stage: Stage-0
    Fetch Operator
      limit: -1
      Processor Tree:
        TableScan
          alias: invites
          Statistics: Num rows: 2 Data size: 216 Basic stats: COMPLETE Column stats: NONE
        Select Operator
          expressions: foo (type: int), bar (type: string), '2008-08-08' (type: string)
          outputColumnNames: _col0, _col1, _col2
          Statistics: Num rows: 2 Data size: 216 Basic stats: COMPLETE Column stats: NONE
        ListSink

Time taken: 0.074 seconds, Fetched: 17 row(s)
```

```
hive> select * from emp;
OK
1      Amit      Programmer      50000      A
2      Gopal      Programmer      50000      B
3      Ankit      Marketing       60000      A
4      Harry      Programmer      50000      C
5      Kapil      Programmer      50000      C
6      Saket      Lead           12000      A
7      Karan      Lead           12000      B
8      Anju      Engineer       45000      B
9      Nan       Programmer      50000      A
10     Andy      Manager 40000   A
11     Anil      Manager 40000   B
12     Ratish   Marketing       60000      B
13     Marry    Manager 40000   A
14     Lamma    Engineer       45000      A
15     Catty    Ceo            100000     B
16     Arthi    Manager 40000   B
17     Pawan    Engineer       45000      B
18     Jogi     Programmer      50000      B
19     Andrew   Engineer       45000      A
20     Sandeep  Cto            100000     A
Time taken: 0.495 seconds, Fetched: 20 row(s)
```

```
hive> !hadoop fs -ls /user/hive/warehouse/test.db/emp_bucket;
Found 3 items
-rwxr-xr-x  2 hadoop supergroup      224 2017-03-22 06:31 /user/
hive/warehouse/test.db/emp_bucket/000000_0
-rwxr-xr-x  2 hadoop supergroup       54 2017-03-22 06:31 /user/
hive/warehouse/test.db/emp_bucket/000001_0
-rwxr-xr-x  2 hadoop supergroup      224 2017-03-22 06:31 /user/
hive/warehouse/test.db/emp_bucket/000002_0
```

```
mysql> select * from VERSION;
+-----+-----+-----+
| VER_ID | SCHEMA_VERSION | VERSION_COMMENT |
+-----+-----+-----+
|      1 | 1.2.0          | Set by MetaStore hadoop@10.0.0.52 |
+-----+-----+-----+
1 row in set (0.00 sec)
```

```
mysql> select * from TBLS limit 2;
+-----+-----+-----+-----+-----+
| TBL_ID | CREATE_TIME | DB_ID | LAST_A | SD_ID |
| CCESS_TIME | OWNER | RETENTION | VIEW_EXP |
| TBL_NAME | TBL_TYPE | VIEW_ORIGINAL_TEXT |
+-----+-----+-----+-----+-----+
|      11 | 1490157004 |      6 | 0 | 11 |
| pokes | EXTERNAL_TABLE | NULL |
|      12 | 1490157098 |      6 | 0 | 12 |
|      0 | hadoop |
```

```
mysql> select * from DBS;
+-----+-----+-----+-----+
| DB_ID | DESC | DB_LOCATION_URI |
| NAME | OWNER_NAME | OWNER_TYPE |
+-----+-----+-----+-----+
|      1 | Default Hive database | hdfs://master1.cyrus.com:9000/user/hive/warehouse |
| default | public | ROLE |
|      6 | NULL | hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db |
| test | hadoop | USER |
```

```
mysql> select * from PARTITIONS;
```

PART_ID	CREATE_TIME	LAST_ACCESS_TIME	PART_NAME	SD_ID	TBL_ID
1	1490160309	0	ds=2008-08-15	17	16
2	1490160349	0	ds=2008-08-08	18	16

2 rows in set (0.00 sec)

```
mysql> select TBLS.TBL_NAME,SDS.LOCATION
-> from SDS,TBLS
-> where TBLS.SD_ID = SDS.SD_ID;
```

TBL_NAME	LOCATION
pokes	hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db
pokes1	hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db
invites	hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db/invites
emp	hdfs://master1.cyrus.com:9000/user/hive/warehouse/emp
emp_bucket	hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db/emp_bucket

5 rows in set (0.00 sec)

```
mysql> select TBLS.TBL_NAME,PARTITIONS.PART_NAME,SDS.LOCATION
-> from SDS,TBLS,PARTITIONS
-> where PARTITIONS.SD_ID = SDS.SD_ID
-> and TBLS.TBL_ID=PARTITIONS.TBL_ID
-> order by 1,2;
```

TBL_NAME	PART_NAME	LOCATION
invites	ds=2008-08-08	hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db/invites/ds=2008-08-08
invites	ds=2008-08-15	hdfs://master1.cyrus.com:9000/user/hive/warehouse/test.db/invites/ds=2008-08-15

```
0: jdbc:hive2://localhost:10000> !connect jdbc:hive2://localhost:10000
Connecting to jdbc:hive2://localhost:10000
Enter username for jdbc:hive2://localhost:10000: root
Enter password for jdbc:hive2://localhost:10000:
Connected to: Apache Hive (version 2.1.0-amzn-0)
Driver: Hive JDBC (version 2.1.0-amzn-0)
16/09/30 05:54:28 [main]: WARN jdbc.HiveConnection: Request to set autoCommit
autoCommit=false.
Transaction isolation: TRANSACTION_REPEATABLE_READ
```

```

1: jdbc:hive2://localhost:10000> set role ADMIN;
INFO : Compiling command(queryId=hive_20160930055431_53dc8cdd-4d7f-40a2-b7c4-f55c7f508307): set role ADMIN
INFO : Semantic Analysis Completed
INFO : Returning Hive schema: Schema(fieldSchemas:null, properties:null)
INFO : OK

```

```

1: jdbc:hive2://localhost:10000> show roles;
INFO : Compiling command(queryId=hive_20160930055456_3c6578ed-d7c0-4a40-9f36-9ab0fa6c47cf): show roles
INFO : Semantic Analysis Completed
INFO : Returning Hive schema: Schema(fieldSchemas:[FieldSchema(name:role, type:string, comment:from deserializer)], properties:null)
INFO : Completed compiling command(queryId=hive_20160930055456_3c6578ed-d7c0-4a40-9f36-9ab0fa6c47cf); Time taken: 0.005 seconds
INFO : Concurrency mode is disabled, not creating a lock manager
INFO : Executing command(queryId=hive_20160930055456_3c6578ed-d7c0-4a40-9f36-9ab0fa6c47cf): show roles
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20160930055456_3c6578ed-d7c0-4a40-9f36-9ab0fa6c47cf); Time taken: 0.021 seconds
INFO : OK

```

role
admin
public

```

1: jdbc:hive2://localhost:10000> create role src_role_wadmin;
INFO : Compiling command(queryId=hive_20160930055525_f0d362be-cbb5-42c9-9a8b-deel711d848b): create role src_role_wadmin
INFO : Semantic Analysis Completed
INFO : Returning Hive schema: Schema(fieldSchemas:null, properties:null)
INFO : Completed compiling command(queryId=hive_20160930055525_f0d362be-cbb5-42c9-9a8b-deel711d848b); Time taken: 0.003 seconds
INFO : Concurrency mode is disabled, not creating a lock manager
INFO : Executing command(queryId=hive_20160930055525_f0d362be-cbb5-42c9-9a8b-deel711d848b): create role src_role_wadmin
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20160930055525_f0d362be-cbb5-42c9-9a8b-deel711d848b); Time taken: 0.021 seconds
INFO : OK
No rows affected (0.039 seconds)

```

```

1: jdbc:hive2://localhost:10000> grant src_role_wadmin to user hadoop with admin option;
INFO : Compiling command(queryId=hive_20160930055542_fb363fa1-39de-45d2-835f-80ad4610d40): grant src_role_wadmin to user hadoop with admin option
INFO : Semantic Analysis Completed
INFO : Returning Hive schema: Schema(fieldSchemas:null, properties:null)
INFO : Completed compiling command(queryId=hive_20160930055542_fb363fa1-39de-45d2-835f-80ad4610d40); Time taken: 0.005 seconds
INFO : Concurrency mode is disabled, not creating a lock manager
INFO : Executing command(queryId=hive_20160930055542_fb363fa1-39de-45d2-835f-80ad4610d40): grant src_role_wadmin to user hadoop with admin option
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20160930055542_fb363fa1-39de-45d2-835f-80ad4610d40); Time taken: 0.021 seconds
INFO : OK

```

```

1: jdbc:hive2://localhost:10000> show role grant user hadoop;
INFO : Compiling command(queryId=hive_20160930055553_f5d0bfbcb7e-413e-84f3-ed142851a43d): show role grant user hadoop
INFO : Semantic Analysis Completed
INFO : Returning Hive schema: Schema(fieldSchemas:[FieldSchema(name:role, type:string, comment:from deserializer), FieldSchema(name:grant_option, type:boolean, comment:from deserializer), FieldSchema(name:grant_time, type:bigint, comment:from deserializer), FieldSchema(name:grantor, type:string, comment:from deserializer)], properties:null)
INFO : Completed compiling command(queryId=hive_20160930055553_f5d0bfbcb7e-413e-84f3-ed142851a43d); Time taken: 0.005 seconds
INFO : Concurrency mode is disabled, not creating a lock manager
INFO : Executing command(queryId=hive_20160930055553_f5d0bfbcb7e-413e-84f3-ed142851a43d): show role grant user hadoop
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20160930055553_f5d0bfbcb7e-413e-84f3-ed142851a43d); Time taken: 0.008 seconds
INFO : OK

```

role	grant_option	grant_time	grantor
public	false	0	root
src_role_wadmin	true	1475214942000	root

2 rows selected (0.047 seconds)

```

1: jdbc:hive2://localhost:10000> show principals src role wadmin;
INFO : Compiling command(queryId=hive_20160930055606_d03a3339-fcc1-46eb-bec7-7104ab23297d): show principals src_role_wadmin
INFO : Semantic Analysis Completed
INFO : Starting task [Stage-0:DDL] in serial mode
INFO : Completed executing command(queryId=hive_20160930055606_d03a3339-fcc1-46eb-bec7-7104ab23297d); Time taken: 0.021 seconds
INFO : OK
+-----+-----+-----+-----+-----+-----+
| principal_name | principal_type | grant_option | grantor | grantor_type | grant_time |
+-----+-----+-----+-----+-----+-----+
| hadoop         | USER         | true        | root   | USER        | 1475214942000 |
+-----+-----+-----+-----+-----+-----+
1 row selected (0.048 seconds)

```

```

# http://flume.apache.org/FlumeUserGuide.html#static-interceptor
source_agent.sources.apache_server.interceptors.itype.type = static
source_agent.sources.apache_server.interceptors.itype.key = log_type
source_agent.sources.apache_server.interceptors.itype.value = apache_access_combined

# http://flume.apache.org/FlumeUserGuide.html#memory-channel
source_agent.channels = memoryChannel
source_agent.channels.memoryChannel.type = memory
source_agent.channels.memoryChannel.capacity = 100

## Send to Flume Collector on Hadoop Node
# http://flume.apache.org/FlumeUserGuide.html#avro-sink
source_agent.sinks = avro_sink
source_agent.sinks.avro_sink.type = avro
source_agent.sinks.avro_sink.channel = memoryChannel
source_agent.sinks.avro_sink.hostname = 192.168.1.99
source_agent.sinks.avro_sink.port = 4545

```

```

[hadoop@rt1 ~]$ mvn -version
Apache Maven 3.3.9 (bb52d8502b132ec0a5a3f4c09453c07478323dc5; 2015-11-10T16:41:47+00:00)
Maven home: /opt/apache-maven-3.3.9
Java version: 1.8.0_101, vendor: Oracle Corporation
Java home: /usr/java/jdk1.8.0_101/jre
Default locale: en_US, platform encoding: UTF-8
OS name: "linux", version: "4.9.20-10.30.amzn1.x86_64", arch: "amd64", family: "unix"

```

```

[root@rt1 oozie-4.1.0]# pwd
/data/oozie-build/oozie-4.1.0
[root@rt1 oozie-4.1.0]# cd bin/
[root@rt1 bin]# ls -l mkdistro.sh
-rwxr-xr-x 1 hadoop hadoop 2358 Dec  5 2014 mkdistro.sh
[root@rt1 oozie-4.1.0]# ls distro/target/
antrun          maven-archiver  oozie-4.1.0-distro.tar.gz  tomcat
archive-tmp    oozie-4.1.0-distro  oozie-distro-4.1.0.jar

```



```
[hadoop@edge1 libext]$ pwd
/opt/cluster/oozie/libext
[hadoop@edge1 libext]$ ls -l ext-2.2.zip
-rw-rw-r-- 1 hadoop hadoop 6800612 Aug  4
2016 ext-2.2.zip
```

```
[hadoop@edge1 ~]$ which oozie-setup.sh
/opt/cluster/oozie/bin/oozie-setup.sh
[hadoop@edge1 ~]$ which oozie
/opt/cluster/oozie/bin/oozie
```

```
[hadoop@edge1 ~]$ oozie-setup.sh prepare-war
setting CATALINA_OPTS="$CATALINA_OPTS -Xmx1024m"
setting OOZIE_CONFIG_FILE=oozie-site.xml
setting OOZIE_LOG=${OOZIE_HOME}/logs
setting OOZIE_HTTP_PORT=11000
setting OOZIE_HTTP_HOSTNAME=`hostname -f`
setting JAVA_HOME=/usr/java/latest
```

```
INFO: Adding extension: /opt/cluster/oozie/libext/s
INFO: Adding extension: /opt/cluster/oozie/libext/x
INFO: Adding extension: /opt/cluster/oozie/libext/x
INFO: Adding extension: /opt/cluster/oozie/libext/x
INFO: Adding extension: /opt/cluster/oozie/libext/x
INFO: Adding extension: /opt/cluster/oozie/libext/z
```

```
New Oozie WAR file with added 'ExtJS library, JARs'
oozie-server/webapps/oozie.war
```

```
INFO: Oozie is ready to be started
```

```
<property>
  <name>hadoop.proxyuser.oozie.hosts</name>
  <value>*</value>
</property>

<property>
  <name>hadoop.proxyuser.oozie.groups</name>
  <value>*</value>
</property>
```

```
<property>
  <name>oozie.service.HadoopAccessorService.hadoop.configurations</name>
  <value>*/opt/cluster/hadoop/etc/hadoop</value>
```

```
<name>oozie.service.WorkflowAppService.system.libpath</name>
<value>hdfs://master1.cyrus.com:9000/user/${user.name}/share/lib</value>
```

```
mysql> create database OOZIEDB;
Query OK, 1 row affected (0.01 sec)

mysql> quit;
Bye
```



```
[hadoop@edgel conf]$ ooziedb.sh create -sqlfile oozie.sql -run
  setting CATALINA_OPTS="$CATALINA_OPTS -Xmx1024m"
  setting OOZIE_CONFIG_FILE=oozie-site.xml
  setting OOZIE_LOG=${OOZIE_HOME}/logs
  setting OOZIE_HTTP_PORT=11000
  setting OOZIE_HTTP_HOSTNAME=`hostname -f`
  setting JAVA_HOME=/usr/java/latest
```

Validate DB Connection

DONE

Check DB schema does not exist

DONE

Check OOZIE\_SYS table does not exist

DONE

Create SQL schema

DONE

Create OOZIE\_SYS table

DONE

Oozie DB has been created for Oozie version '4.1.0'

The SQL commands have been written to: oozie.sql

```
[hadoop@edgel ~]$ oozie admin -shareliblist -oozie http://edgel.cyrus
.com:11000/oozie
[Available ShareLib]
hive
distcp
mapreduce-streaming
oozie
hcatalog
sqoop
pig
```

```
[hadoop@edgel examples]$ pwd
/opt/cluster/oozie/examples
[hadoop@edgel examples]$ cat apps/map-reduce/job.properties
nameNode=hdfs://master1.cyrus.com:9000
jobTracker=master1.cyrus.com:9003
queueName=default
examplesRoot=examples
```

```
[hadoop@edge1 examples]$ ls -l
total 12
drwxr-xr-x 24 hadoop hadoop 4096 Aug  2  2016 apps
drwxrwxr-x  4 hadoop hadoop 4096 Aug  4  2016 input-data
drwxr-xr-x  3 hadoop hadoop 4096 Dec  5  2014 src
[hadoop@edge1 examples]$ ls -l input-data/
total 8
drwxrwxr-x 3 hadoop hadoop 4096 Aug  4  2016 rawLogs
drwxrwxr-x 2 hadoop hadoop 4096 Aug  4  2016 text
[hadoop@edge1 examples]$ ls -l input-data/text/
total 4
-rw-r--r-- 1 hadoop hadoop 1409 Dec  5  2014 data.txt
```

## Chapter 8: Performance Tuning

```
[root@master1 ~]# egrep 'trans|thp' /proc/vmstat
nr_anon_transparent_hugepages 0
thp_fault_alloc 0
thp_fault_fallback 0
thp_collapse_alloc 0
thp_collapse_alloc_failed 0
thp_split 0
thp_zero_page_alloc 0
thp_zero_page_alloc_failed 0
[root@master1 ~]# grep -i HugePages_Total /proc/meminfo
HugePages_Total: 0
```

```
[root@master1 ~]# ipcs -lm

----- Shared Memory Limits -----
max number of segments = 4096
max seg size (kbytes) = 67108864
max total shared memory (kbytes) = 17179869184
min seg size (bytes) = 1
```

```
[root@master1 ~]# vmstat 2
procs -----memory----- ---swap-- --
 r  b      swpd      free      buff      cache      si      so
 0  0         0 15762084   22284   185072         0     0
 0  0         0 15761944   22292   185092         0     0
 0  0         0 15761944   22292   185132         0     0
 0  0         0 15761944   22292   185164         0     0
 0  0         0 15761944   22300   185192         0     0
 0  0         0 15761820   22300   185224         0     0
```

```
[hadoop@master1 ~]$ jps
3424 NameNode
3666 SecondaryNameNode
3834 Jps
[hadoop@master1 ~]$ lsof -p 3424 | wc -l
595
```

```
/dev/xvdb1 on /data/namenode1 type ext4 (rw,noatime,data=ordered)
/dev/xvdc1 on /data/namenode2 type ext4 (rw,noatime,data=ordered)
```

```
/dev/xvdb1      20G   51M   19G   1% /space/d1
/dev/xvdc1      20G   58M   19G   1% /space/d2
/dev/xvdd1      20G   63M   19G   1% /space/d3
```

```
[root@master1 ~]# hdfs getconf -confkey dfs.blocksize
134217728
[root@master1 ~]# hdfs getconf -confkey dfs.replication
1
```

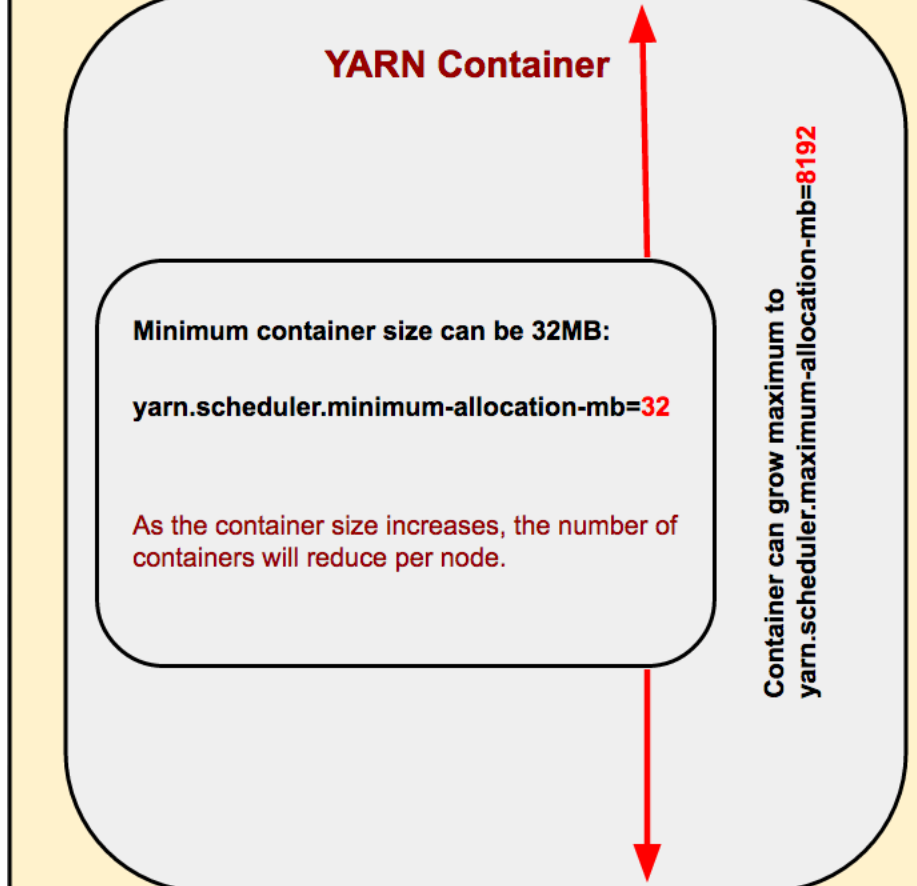
```
[hadoop@master1 ~]$ jmap -heap 3424
Attaching to process ID 3424, please wait...
Debugger attached successfully.
Server compiler detected.
JVM version is 25.92-b14
```

```
using thread-local object allocation.
Parallel GC with 4 thread(s)
```

Heap Configuration:

```
MinHeapFreeRatio          = 0
MaxHeapFreeRatio          = 100
MaxHeapSize                = 2147483648 (2048.0MB)
NewSize                    = 88080384 (84.0MB)
MaxNewSize                 = 715653120 (682.5MB)
OldSize                    = 176160768 (168.0MB)
NewRatio                   = 2
SurvivorRatio              = 8
MetaspaceSize              = 21807104 (20.796875MB)
CompressedClassSpaceSize  = 1073741824 (1024.0MB)
MaxMetaspaceSize          = 17592186044415 MB
G1HeapRegionSize          = 0 (0.0MB)
```

The maximum memory and vCores Nodemanager can allocate:  
yarn.nodemanager.resource.memory-mb=8192  
yarn.scheduler.maximum-allocation-vcores=32



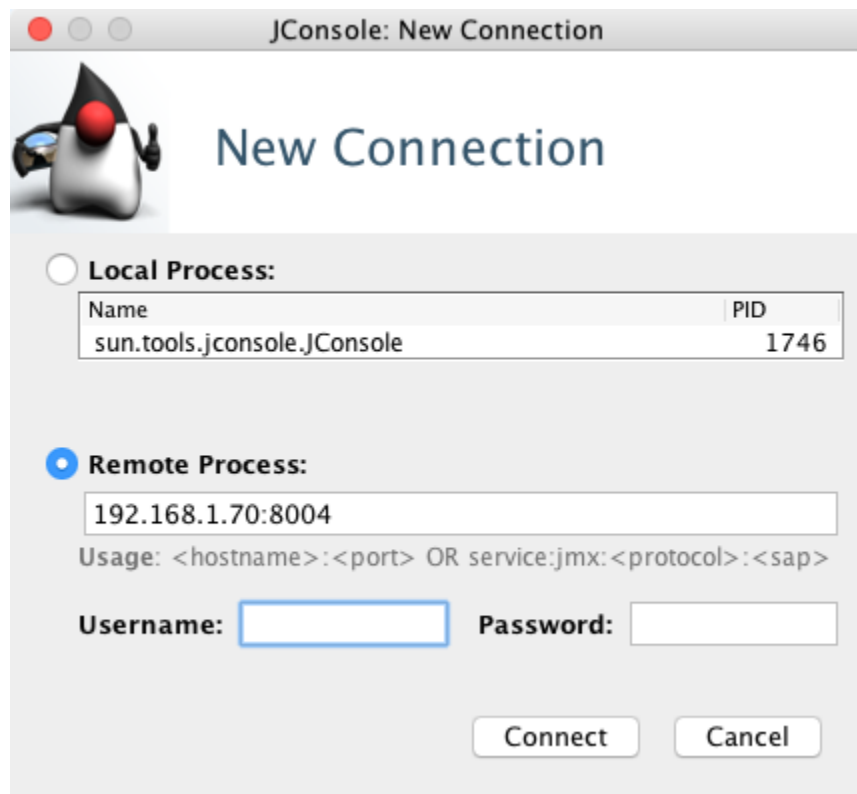
```
# if [ "$SERVICE" = "cli" ]; then  
#   if [ -z "$DEBUG" ]; then  
#     export HADOOP_OPTS="$HADOOP_OPTS -XX:NewRatio=12 -Xms10m -XX:MaxHeapFreeRatio=40 -XX:MinHeapFreeRatio=15 -XX:+UseParNewGC -XX:-UseGCOverheadLimit"  
#   else  
#     export HADOOP_OPTS="$HADOOP_OPTS -XX:NewRatio=12 -Xms10m -XX:MaxHeapFreeRatio=40 -XX:MinHeapFreeRatio=15 -XX:-UseGCOverheadLimit"  
#   fi  
# fi
```


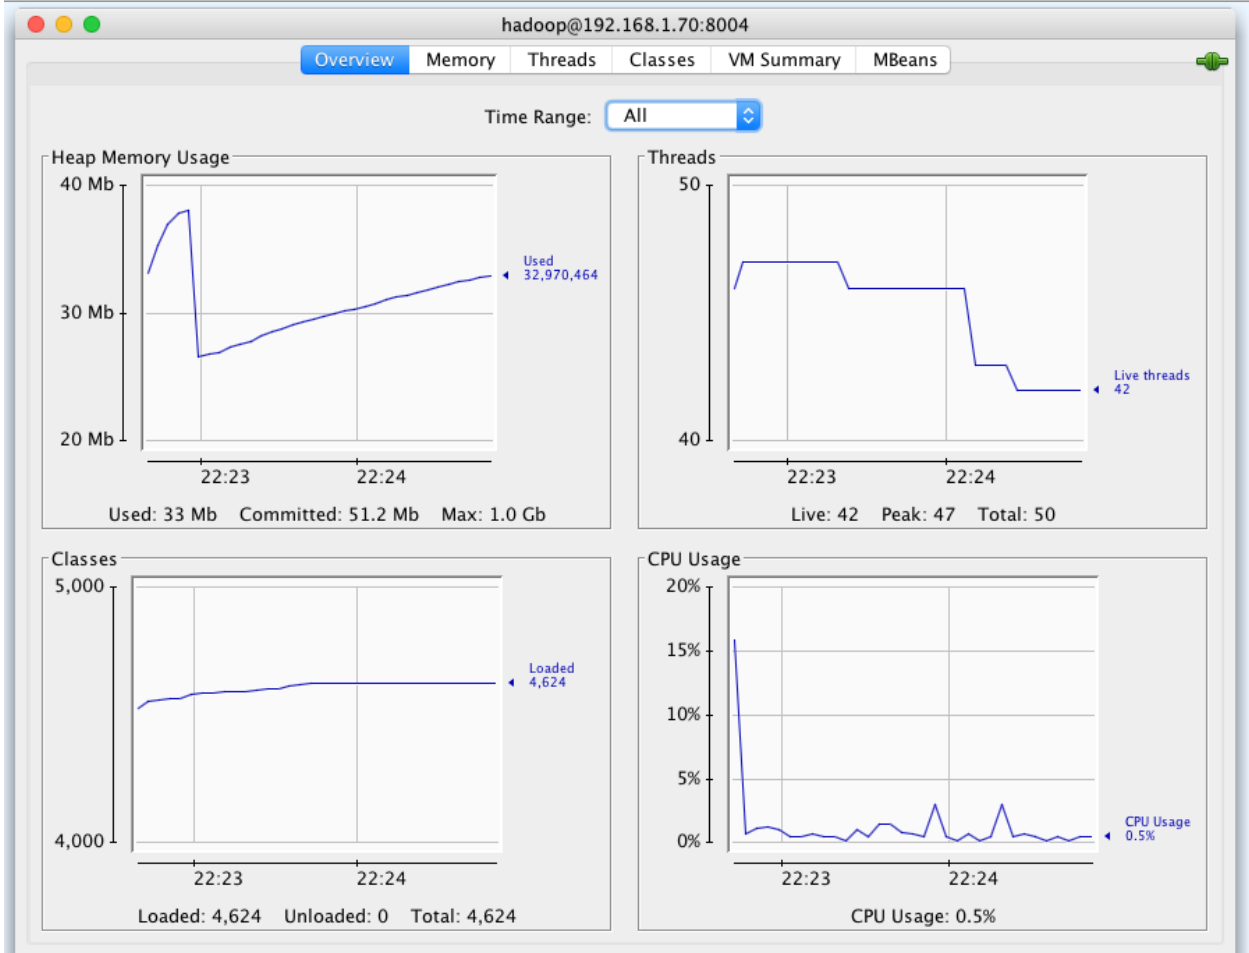
```
hive> set hive.execution.engine=tez;  
hive> set hive.execution.engine;  
hive.execution.engine=tez
```

```
export HADOOP_NAMENODE_OPTS="-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.password.file=$HADOOP_CONF_DIR/jmxremote.password
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.port=8004 $HADOOP_NAMENODE_OPTS"
```

```
[hadoop@nn1 ~]$ sudo netstat -tlnp | grep 8004
tcp        0      0 0.0.0.0:8004          0.0.0.0:*
           LISTEN          2846/java
```

```
[hadoop@nn1 ~]$ hdfs jmxget -server localhost -port 8004 -service NameNode
init: server=localhost;port=8004;service=NameNode;localVMUrl=null
Create RMI connector and connect to the RMI connector serverservice:jmx:rmi
```








```
hive> analyze table emp_hive compute statistics;
Query ID = hadoop_20170407204933_1d837d35-a987-4626-9b83-2bc7b16d4a7d
Total jobs = 1
Launching Job 1 out of 1
```

```
hive> desc formatted emp_hive;
OK
# col_name          data_type

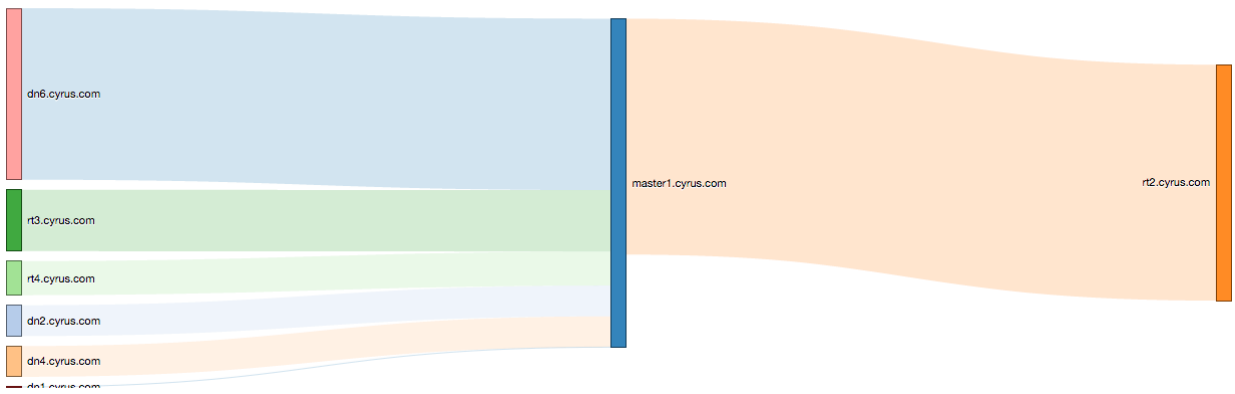
id                  int
name                string
salary              string

# Detailed Table Information
Database:           default
Owner:              hadoop
CreateTime:         Fri Apr 07 09:17:4
LastAccessTime:    UNKNOWN
Retention:          0
Location:           hdfs://ip-172-31-9
Table Type:         EXTERNAL_TABLE
Table Parameters:
    COLUMN_STATS_ACCURATE  {\ "BASIC_S
"true\", \"salary\": \"true\"}
    EXTERNAL                TRUE
    numFiles                 2
    numRows                  60000
    rawDataSize              1117788
    totalSize                1177788
    transient_lastDdlTime   1491598179
```

### Top Flow Talkers



### Top Flow Talkers



## Chapter 9: HBase Administration

```
export HBASE_HOME=/opt/cluster/hbase
PATH=$OOZIE_HOME/bin:$SPARK_HOME/bin:$SPARK_HOME/sbin:
$HBASE_HOME/bin:$HIVE_HOME/bin:$HADOOP_HOME/bin/:$HADO
OP_HOME/sbin/:$JAVA_HOME/bin/:$PATH
```

```
[hadoop@master1 ~]$ which hbase
/opt/cluster/hbase/bin/hbase
[hadoop@master1 ~]$ echo $HBASE_HOME
/opt/cluster/hbase
```

```
[hadoop@master1 cluster]$ ls -l hbase/
total 384
drwxr-xr-x  4 hadoop hadoop  4096 Mar 29 09:00 bin
-rw-r--r--  1 hadoop hadoop 183633 Mar 29 09:00 CHANGES.txt
drwxr-xr-x  2 hadoop hadoop  4096 Mar 29 09:00 conf
drwxr-xr-x 12 hadoop hadoop  4096 Mar 29 09:00 docs
drwxr-xr-x  7 hadoop hadoop  4096 Mar 29 09:00 hbase-webapps
-rw-r--r--  1 hadoop hadoop   261 Mar 29 09:00 LEGAL
drwxr-xr-x  3 hadoop hadoop  4096 Mar 29 09:00 lib
-rw-r--r--  1 hadoop hadoop 140387 Mar 29 09:00 LICENSE.txt
-rw-r--r--  1 hadoop hadoop  34474 Mar 29 09:00 NOTICE.txt
-rw-r--r--  1 hadoop hadoop   1477 Mar 29 09:00 README.txt
```

```
[hadoop@master1 conf]$ jps
3538 SecondaryNameNode
3765 DataNode
3847 Jps
3305 NameNode
```

```
[hadoop@master1 conf]$ jps
3538 SecondaryNameNode
3765 DataNode
3305 NameNode
4731 Jps
4492 HMaster
```



[Home](#)

[Table Details](#)

[Local Logs](#)

[Log Level](#)

[Debug Dump](#)

## Master [master1.cyrus.com](#)

### Region Servers

**Base Stats**

[Memory](#)

[Requests](#)

[Storefiles](#)

[Compactions](#)

**ServerName**

**Start time**

[master1.cyrus.com,11959,1490784557188](#)

Wed Mar 29 10:49:17 UTC 2017

Total:1

```
[hadoop@master1 ~]$ echo "status" | hbase shell
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/opt/cluster/hbase-
ticLoggerBinder.class]
SLF4J: Found binding in [jar:file:/opt/cluster/hadoc
ar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple
SLF4J: Actual binding is of type [org.slf4j.impl.Log
HBase Shell; enter 'help<RETURN>' for list of suppor
Type "exit<RETURN>" to leave the HBase Shell
Version 1.1.5, r239b80456118175b340b2e562a5568b5c744

status
1 servers, 0 dead, 2.0000 average load
```

[Home](#)[Table Details](#)[Procedures](#)[Local Logs](#)[Log Level](#)[Debug Dump](#)[Metrics Dump](#)

## Tables

[User Tables](#)[System Tables](#)[Snapshots](#)

Table Name	Description
<a href="#">hbase:meta</a>	The hbase:meta table holds references to all User Table regions.
<a href="#">hbase:namespace</a>	The hbase:namespace table holds information about namespaces.

```
[hadoop@master1 conf]$ cat regionservers
rt1.cyrus.com
rt2.cyrus.com
rt3.cyrus.com
```

```
[hadoop@master1 conf]$ cat backup-masters
rt4.cyrus.com
[hadoop@master1 conf]$
```

```
-----  
edge1.cyrus.com,master1.cyrus.com (2)  
-----  
Mode: follower  
-----  
rt1.cyrus.com  
-----  
Mode: leader
```

# Master master1.cyrus.com

## Region Servers

Base Stats	Memory	Requests	Storefiles	Compactions
ServerName	Start time			
rt1.cyrus.com,16020,1490789300641	Wed Mar 29 12:08:20 UTC 2017			
rt2.cyrus.com,16020,1490789299973	Wed Mar 29 12:08:19 UTC 2017			
rt3.cyrus.com,16020,1490789300319	Wed Mar 29 12:08:20 UTC 2017			

## Backup Masters

ServerName	Port	Start Time
rt4.cyrus.com	16000	Wed Mar 29 12:08:21 UTC 2017
Total:1		

```
hbase(main):001:0> create 'test', 'cf'
0 row(s) in 2.4290 seconds

=> Hbase::Table - test
```

```
hbase(main):002:0> put 'test', 'row1', 'cf:1', 'value1'
0 row(s) in 0.1140 seconds

hbase(main):003:0> put 'test', 'row2', 'cf:2', 'value2'
0 row(s) in 0.0100 seconds
```

```
hbase(main):004:0> list
TABLE
test
1 row(s) in 0.0100 seconds

=> ["test"]
hbase(main):005:0> scan 'test'
ROW                                COLUMN+CELL
 row1                                column=cf:1, timestamp=1490824460000, value=value1
 row2                                column=cf:2, timestamp=1490824484973, value=value2
2 row(s) in 0.0300 seconds
```

```
hbase(main):008:0> get 'test', 'row1'
COLUMN                                CELL
 cf:1                                timestamp=1490824460000, value=value1
1 row(s) in 0.0250 seconds
```

```
hive> CREATE EXTERNAL TABLE hbasehive_table
> (key string, userid string,bookid string,rating int)
> STORED BY 'org.apache.hadoop.hive.hbase.HBaseStorageHandler'
> WITH SERDEPROPERTIES
> ("hbase.columns.mapping" = ":key,ratings:userid,ratings:bookid,ratings:rating")
> TBLPROPERTIES ("hbase.table.name" = "hivetable");
OK
Time taken: 0.557 seconds
```

```
hive> select * from hbasehive_table;
OK
row1      user1      book1      1
Time taken: 0.479 seconds, Fetched: 1 row(s)
```



```
mysql> select TBL_NAME,SDS.LOCATION from SDS,TBLS where SDS.SD_ID = '31' and TBLS.TBL_ID='31';
+-----+-----+
| TBL_NAME | LOCATION |
+-----+-----+
| hbasehive_table | hdfs://master1.cyrus.com:9000/user/hive/warehouse/hbasehive_table |
+-----+-----+
1 row in set (0.00 sec)

mysql> select * from TBLS where TBLS.TBL_ID='31';
+-----+-----+-----+-----+-----+-----+-----+
| TBL_ID | CREATE_TIME | DB_ID | LAST_ACCESS_TIME | OWNER | RETENTION | SD_ID | TBL_NAME | TBL_TYPE | VIEW_EXPANDED_TEXT | VIEW_ORIGINAL_TEXT |
+-----+-----+-----+-----+-----+-----+-----+
| 31 | 1490827021 | 1 | 0 | hadoop | 0 | 31 | hbasehive_table | EXTERNAL_TABLE | NULL | NULL |
+-----+-----+-----+-----+-----+-----+-----+
```

```
2017-03-29 23:23:16,871 INFO [main] hfile.CacheConfig: CacheConfig:disabled
Block Index:
size=1
key=hivetable,,1490827001115.b635d3534fe0393b753ceaf20e66320a.
/info:regioninfo/1490827000989/Put
offset=0, dataSize=829
```

Summary:

Table hbase:meta is okay.

Number of regions: 1

Deployed on: rt2.cyrus.com,16020,1490789299973

Table hivetable is okay.

Number of regions: 1

Deployed on: rt3.cyrus.com,16020,1490789300319

Table test is okay.

Number of regions: 1

Deployed on: rt1.cyrus.com,16020,1490789300641

Table hbase:namespace is okay.

Number of regions: 1

Deployed on: rt1.cyrus.com,16020,1490789300641

0 inconsistencies detected.

Status: OK

```
hbase(main):002:0* list_snapshots
```

```
SNAPSHOT          TABLE + CREATION TIME
test_30march      test (Thu Mar 30 00:00:55 +0000 2017)
1 row(s) in 0.1720 seconds
```

```
SNAPSHOT          TABLE + CREATION TIME
test_30march      test (Thu Mar 30 00:00:55 +0000 2017)
1 row(s) in 0.0050 seconds
```

```
[hadoop@master1 ~]$ hadoop fs -ls /hbase/.hbase-snapshot
Found 2 items
drwxr-xr-x  - hadoop supergroup          0 2017-03-30 0
0:00 /hbase/.hbase-snapshot/.tmp
drwxr-xr-x  - hadoop supergroup          0 2017-03-30 0
0:00 /hbase/.hbase-snapshot/test_30march
```

```
2017-03-30 01:37:11,713 INFO [main] util.HFileV1Detector: Result:
2017-03-30 01:37:11,713 INFO [main] util.HFileV1Detector: Tables Processed:
2017-03-30 01:37:11,713 INFO [main] util.HFileV1Detector: Count of HFileV1: 0
2017-03-30 01:37:11,713 INFO [main] util.HFileV1Detector: Count of corrupted files: 0
2017-03-30 01:37:11,713 INFO [main] util.HFileV1Detector: Count of Regions with HFileV1: 0
2017-03-30 01:37:11,815 INFO [main] migration.UpgradeTo96: No HFileV1 found.
```

```
[hadoop@master1 hadoop]$ jps
10401 ResourceManager
9749 HMaster
10616 Jps
9163 NameNode
3630 QuorumPeerMain
```

```
mysql> CREATE TABLE emp(emp_id VARCHAR(10), name VARCHAR(20), salary VARCHAR(20));
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> insert into emp(emp_id,name,salary) values('1','Am1','10000');
Query OK, 1 row affected (0.00 sec)
```

```
mysql> insert into emp(emp_id,name,salary) values('2','Am2','10000');
Query OK, 1 row affected (0.00 sec)
```

```
hbase(main):001:0> create 'test','emp_records'
0 row(s) in 8.4690 seconds
```

```
=> Hbase::Table - test
```

```
hbase(main):002:0> scan 'test'
ROW COLUMN+CELL
1 column=emp_records:name, timestamp=1494498684358, value=Am1
1 column=emp_records:salary, timestamp=1494498684358, value=10000
2 column=emp_records:name, timestamp=1494498684358, value=Am2
2 column=emp_records:salary, timestamp=1494498684358, value=10000
3 column=emp_records:name, timestamp=1494498684358, value=Am3
3 column=emp_records:salary, timestamp=1494498684358, value=10000
3 row(s) in 0.1310 seconds
```

## Chapter 10: Cluster Planning

```
Name: 10.0.0.75:50010 (dn6.cyrus.com)
Hostname: dn6.cyrus.com
Rack: /rack6
Decommission Status : Normal
Configured Capacity: 63007739904 (58.68 GB)
DFS Used: 73728 (72 KB)
Non DFS Used: 3409612800 (3.18 GB)
DFS Remaining: 59598053376 (55.51 GB)
DFS Used%: 0.00%
DFS Remaining%: 94.59%
```

```
[hadoop@dn6 ~]$ hdfs getconf -confkey dfs.datanode.data.dir
file:/space/d1,file:/space/d2,file:/space/d3
[hadoop@dn6 ~]$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.8G  2.2G  5.5G  29% /
devtmpfs        3.9G   88K  3.9G   1% /dev
tmpfs           3.9G    0   3.9G   0% /dev/shm
/dev/xvdb1      20G   44M   19G   1% /space/d1
/dev/xvdc1      20G   44M   19G   1% /space/d2
/dev/xvdd1      20G   44M   19G   1% /space/d3
```

```
/dev/xvdb1 on /space/d1 type ext4 (rw,noatime,data=ordered)
/dev/xvdc1 on /space/d2 type ext4 (rw,noatime,data=ordered)
/dev/xvdd1 on /space/d3 type ext4 (rw,noatime,data=ordered)
```

```
[hadoop@master1 hadoop]$ cat hadoop-env.sh | grep HEAPSIZE
export HADOOP_HEAPSIZE=2048
#export HADOOP_NAMENODE_INIT_HEAPSIZE=""
[hadoop@master1 hadoop]$ cat yarn-env.sh | grep HEAPSIZE
# YARN_HEAPSIZE=1000
if [ "$YARN_HEAPSIZE" != "" ]; then
    JAVA_HEAP_MAX="-Xmx"$YARN_HEAPSIZE"m"
#export YARN_RESOURCEMANAGER_HEAPSIZE=1000
#export YARN_TIMELINESERVER_HEAPSIZE=1000
#export YARN_NODEMANAGER_HEAPSIZE=1000
```

## Chapter 11: Troubleshooting, Diagnostics, and Best Practices

```
[hadoop@master1 ~]$ ls -ld /data/namenode*  
drwxr-xr-x 3 hadoop hadoop 4096 Mar 29 08:57 /data/namenode1  
drwxr-xr-x 3 hadoop hadoop 4096 Mar 29 08:57 /data/namenode2
```

```
/dev/xvdb1 on /data/namenode1 type ext4 (rw,noatime,data=ordered)  
/dev/xvdc1 on /data/namenode2 type ext4 (rw,noatime,data=ordered)  
/dev/xvdd1 on /var/log/hadoop type ext4 (rw,noatime,data=ordered)
```

```
<property>  
<name>fs.defaultFS</name>  
<value>hdfs://master1.cyrus.com:9000</value>
```

```
[hadoop@master1 cluster]$ which java  
/usr/java/latest/bin/java  
[hadoop@master1 cluster]$ java -version  
java version "1.8.0_92"
```

```
[hadoop@master1 ~]$ jps  
Java HotSpot(TM) 64-Bit Server VM warning: Insufficient space for shared memory file:  
14730  
Try using the -Djava.io.tmpdir= option to select an alternate temp location.
```

```
[hadoop@master1 ~]$ df -h  
Filesystem      Size  Used Avail Use% Mounted on  
/dev/xvda1      7.8G  7.8G   0 100% /  
devtmpfs        7.9G   84K  7.9G   1% /dev  
tmpfs           7.9G    0  7.9G   0% /dev/shm
```

```
17/04/03 00:00:35 INFO namenode.FSNamesystem: Retry cache will use 0.
iry time is 600000 millis
17/04/03 00:00:35 INFO util.GSet: Computing capacity for map NameNode
17/04/03 00:00:35 INFO util.GSet: VM type = 64-bit
17/04/03 00:00:35 INFO util.GSet: 0.029999999329447746% max memory 1.
17/04/03 00:00:35 INFO util.GSet: capacity = 2^16 = 65536 entries
Re-format filesystem in Storage Directory /data/namenode1 ? (Y or N)
```

```
[hadoop@master1 cluster]$ sudo netstat -tlnp | grep 9000
tcp        0      0 10.0.0.104:9000    0.0.0.0:*          LISTEN
```

```
<property>
  <name>dfs.datanode.data.dir</name>
  <value>file:/space/d1,file:/space/d2,file:/space/d3</value>
</property>
```

```
drwxr-xr-x 5 hadoop hadoop 4096 May 17 2016 /space/
[hadoop@dn1 ~]$ ls -ld /space/d*
drwx----- 3 hadoop hadoop 4096 Apr  3 00:14 /space/d1
drwx----- 3 hadoop hadoop 4096 Apr  3 00:14 /space/d2
drwx----- 3 hadoop hadoop 4096 Apr  3 00:14 /space/d3
```

```
/dev/xvdb1 on /space/d1 type ext4 (rw,noatime,data=ordered)
/dev/xvdc1 on /space/d2 type ext4 (rw,noatime,data=ordered)
/dev/xvdd1 on /space/d3 type ext4 (rw,noatime,data=ordered)
```

```
[hadoop@dn1 ~]$ netstat -tlnp
(Not all processes could be identified, non-owned process info
will not be shown, you would have to be root to see it all.)
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State
tcp    0      0 0.0.0.0:5679           0.0.0.0:*               LISTEN
tcp    0      0 0.0.0.0:111            0.0.0.0:*               LISTEN
tcp    0      0 0.0.0.0:22             0.0.0.0:*               LISTEN
tcp    0      0 0.0.0.0:26905          0.0.0.0:*               LISTEN
tcp    0      0 127.0.0.1:25           0.0.0.0:*               LISTEN
tcp    0      0 0.0.0.0:50010          0.0.0.0:*               LISTEN
tcp    0      0 0.0.0.0:50075          0.0.0.0:*               LISTEN
```



```
[hadoop@master1 cluster]$ telnet master1.cyrus.com 9000
Trying 10.0.0.104...
Connected to master1.cyrus.com.
Escape character is '^]'.

```

```
<property>
<name>yarn.resourcemanager.resource-tracker.address</name>
<value>master1.cyrus.com:9001</value>
</property>

<property>
<name>yarn.resourcemanager.scheduler.address</name>
<value>master1.cyrus.com:9002</value>
</property>

<property>
<name>yarn.resourcemanager.address</name>
<value>master1.cyrus.com:9003</value>
</property>

```

```
[hadoop@master1 cluster]$ netstat -tp
(Not all processes could be identified, non-owned process info
will not be shown, you would have to be root to see it all.)
Active Internet connections (w/o servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State       PID/Program
tcp      0      0 master1.cyrus.co:cslistener dn3.cyrus.com:gprs-data ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener dn2.cyrus.com:31448     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener dn6.cyrus.com:47486     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener dn1.cyrus.com:11774     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener dn4.cyrus.com:52050     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener dn5.cyrus.com:36028     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener rt1.cyrus.com:15594     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener rt2.cyrus.com:41914     ESTABLISHED 17263/java
tcp      0      0 master1.cyrus.co:cslistener rt3.cyrus.com:40890     ESTABLISHED 17263/java

```

```
[hadoop@master1 cluster]$ tail -5f hadoop/logs/hadoop-hadoop-namenode-master1.cyrus.com.log
2017-04-03 01:55:41,769 INFO BlockStateChange: BLOCK* processReport: discarded non-initial block report from DataNodeRegistration(10.0.0.38:50010, datanodeUuid=15c07c2a-0227-4733-b05d-c272bfc280c0, infoPort=50075, infoSecurePort=0, ipcPort=50020, storageInfo=lv=-56;cid=CID-58e5f718-1bb0-4a03-ab11-3a6f368d81d2;nsid=880609571;c=0) because namenode still in startup phase

```

```
java 17263 hadoop 300u IPv4 893785 0t0 TCP master1.cyrus.com:cslistener->dn2.cyrus.com:31448 (ESTABLISHED)
java 17263 hadoop 301u IPv4 894391 0t0 TCP master1.cyrus.com:cslistener->dn5.cyrus.com:36028 (ESTABLISHED)
java 17263 hadoop 302u IPv4 893800 0t0 TCP master1.cyrus.com:cslistener->dn4.cyrus.com:52050 (ESTABLISHED)
java 17263 hadoop 303u IPv4 898212 0t0 TCP master1.cyrus.com:cslistener->dn3.cyrus.com:gprs-data (ESTABLISHED)
java 17263 hadoop 304u IPv4 898221 0t0 TCP master1.cyrus.com:cslistener->dn1.cyrus.com:11774 (ESTABLISHED)

```



```
java 17263 hadoop 279uW REG 202,17 23 15 /data/namenode1/in_use.lock
java 17263 hadoop 280u unix 0xffff8803f85e1400 0t0 891726 socket
java 17263 hadoop 281uW REG 202,33 23 15 /data/namenode2/in_use.lock
java 17263 hadoop 282u REG 202,17 1048576 8751 /data/namenode1/current/edits_inprogres
s_0000000000000000024829
java 17263 hadoop 283u REG 202,33 1048576 8751 /data/namenode2/current/edits_inprogres
s_0000000000000000024829
```

```
[hadoop@master1 ~]$ jmap -heap 17263
Attaching to process ID 17263, please wait...
Debugger attached successfully.
Server compiler detected.
JVM version is 25.92-b14
```

```
using thread-local object allocation.
Parallel GC with 4 thread(s)
```

```
Heap Configuration:
  MinHeapFreeRatio          = 0
  MaxHeapFreeRatio          = 100
  MaxHeapSize                = 2147483648 (2048.0MB)
  NewSize                    = 88080384 (84.0MB)
  MaxNewSize                 = 715653120 (682.5MB)
  OldSize                    = 176160768 (168.0MB)
  NewRatio                   = 2
  SurvivorRatio              = 8
  MetaspaceSize              = 21807104 (20.796875MB)
  CompressedClassSpaceSize   = 1073741824 (1024.0MB)
  MaxMetaspaceSize           = 17592186044415 MB
  G1HeapRegionSize           = 0 (0.0MB)
```

```
2017-03-22 06:30:54,699 ERROR [main]: metadata.Hive (Hive.java:getTable(1118)) - Table emp_bucket not found: te
st.emp_bucket table not found
2017-03-22 06:30:54,700 ERROR [main]: ql.Driver (SessionState.java:printError(960)) - FAILED: SemanticException
org.apache.hadoop.hive.ql.metadata.InvalidTableException: Table not found emp_bucket
org.apache.hadoop.hive.ql.parse.SemanticException: org.apache.hadoop.hive.ql.metadata.InvalidTableException: Ta
ble not found emp_bucket
```

```
mysql> show grants;
+-----+
| Grants for hadoop@% |
+-----+
| GRANT ALL PRIVILEGES ON *.* TO 'hadoop'@'%' IDENTIFIED WITH GRANT OPTION |
| GRANT ALL PRIVILEGES ON `hive_db`.* TO 'hadoop'@'%' |
+-----+
```

```
[hadoop@edge1 cluster]$ hadoop fs -ls /user/hive/
Found 1 items
drwxrwxr-x - hadoop supergroup 0 2017-03-29 22:37 /user/hive/warehouse
[hadoop@edge1 cluster]$ hadoop fs -ls /tmp
Found 4 items
drwx----- - hadoop supergroup 0 2017-03-18 23:34 /tmp/hadoop-yarn
drwxr-xr-x - hadoop supergroup 0 2017-03-29 10:49 /tmp/hbase-hadoop
drwx-wx-wx - hadoop supergroup 0 2017-03-21 21:09 /tmp/hive
drwxrwxrwt - hadoop supergroup 0 2017-03-18 23:36 /tmp/logs
```

```
mysql> select * from VERSION;
+-----+-----+-----+
| VER_ID | SCHEMA_VERSION | VERSION_COMMENT |
+-----+-----+-----+
| 1 | 1.2.0 | Set by MetaStore hadoop@10.0.0.52 |
+-----+-----+-----+
1 row in set (0.00 sec)
```

```
[hadoop@edge1 cluster]$ lsof -Pan -p 14839 -i
COMMAND  PID  USER  FD  TYPE  DEVICE  SIZE/OFF  NODE NAME
java    14839  hadoop  390u  IPv4  42827      0t0  TCP  10.0.0.52:61606->10.0.0.104:3306 (ESTABLISHED)
java    14839  hadoop  391u  IPv4  44061      0t0  TCP  10.0.0.52:61608->10.0.0.104:3306 (ESTABLISHED)
java    14839  hadoop  392u  IPv4  44064      0t0  TCP  10.0.0.52:61614->10.0.0.104:3306 (ESTABLISHED)
java    14839  hadoop  393u  IPv4  44065      0t0  TCP  10.0.0.52:61616->10.0.0.104:3306 (ESTABLISHED)
```

```
[hadoop@edge1 cluster]$ zkServer.sh status
ZooKeeper JMX enabled by default
Using config: /opt/cluster/zoo/bin/../conf/zoo.cfg
Mode: follower
```

```
[hadoop@master1 log]$ lsof -Pan -p 18438 -i
COMMAND PID USER FD TYPE DEVICE SIZE/OFF NODE NAME
java 18438 hadoop 395u IPv6 906688 0t0 TCP 10.0.0.104:16000 (LISTEN)
java 18438 hadoop 439u IPv6 912196 0t0 TCP 10.0.0.104:57894->10.0.0.52:2181 (ESTABLISHED)
java 18438 hadoop 441u IPv6 912199 0t0 TCP *:16010 (LISTEN)
java 18438 hadoop 448u IPv6 910115 0t0 TCP 10.0.0.104:17252->10.0.0.4:2181 (ESTABLISHED)
java 18438 hadoop 449u IPv6 915480 0t0 TCP 10.0.0.104:15554->10.0.0.104:9000 (ESTABLISHED)
java 18438 hadoop 459u IPv6 915408 0t0 TCP 10.0.0.104:57996->10.0.0.52:2181 (ESTABLISHED)
java 18438 hadoop 463u IPv6 915411 0t0 TCP 10.0.0.104:57998->10.0.0.52:2181 (ESTABLISHED)
```

```
[hadoop@master1 log]$ tail -10f /opt/cluster/hbase/logs/hbase-hadoop-master-master1.cyrus.com.log
2017-04-03 05:57:21,826 INFO [master1:16000.activeMasterManager] util.FSUtils: Waiting for dfs to exit safe mode...
2017-04-03 05:57:31,829 INFO [master1:16000.activeMasterManager] util.FSUtils: Waiting for dfs to exit safe mode...
2017-04-03 05:57:31,829 INFO [master1:16000.activeMasterManager] util.FSUtils: Waiting for dfs to exit safe mode...
```

```
[hadoop@master1 log]$ hadoop fs -ls /hbase
Found 9 items
drwxr-xr-x - hadoop supergroup 0 2017-03-30 00:00 /hbase/.hbase-snapshot
drwxr-xr-x - hadoop supergroup 0 2017-03-29 23:28 /hbase/.tmp
drwxr-xr-x - hadoop supergroup 0 2017-04-02 23:41 /hbase/MasterProcWALs
drwxr-xr-x - hadoop supergroup 0 2017-03-29 12:11 /hbase/WALs
drwxr-xr-x - hadoop supergroup 0 2017-03-29 22:25 /hbase/archive
drwxr-xr-x - hadoop supergroup 0 2017-03-29 12:11 /hbase/data
-rw-r--r-- 1 hadoop supergroup 42 2017-03-29 12:11 /hbase/hbase.id
-rw-r--r-- 1 hadoop supergroup 7 2017-03-29 12:11 /hbase/hbase.version
drwxr-xr-x - hadoop supergroup 0 2017-04-02 23:22 /hbase/oldWALs
```

```
2017-04-03 06:02:56,968 INFO [master1:16000.activeMasterManager] master.ServerManager: Waiting for region servers count to settle; currently checked in 0, slept for 4511 ms, expecting minimum of 1, maximum of 2147483647, timeout of 4500 ms, interval of 1500 ms.
2017-04-03 06:02:58,471 INFO [master1:16000.activeMasterManager] master.ServerManager: Waiting for region servers count to settle; currently checked in 0, slept for 6015 ms, expecting minimum of 1, maximum of 2147483647, timeout of 4500 ms, interval of 1500 ms.
```

```
2017-04-03 06:15:15,753 INFO [master1:16000.activeMasterManager] master.AssignmentManager: Bulk assigning done
2017-04-03 06:15:15,756 INFO [master1:16000.activeMasterManager] master.AssignmentManager: Joined the cluster in 196ms, failover=false
2017-04-03 06:15:15,812 INFO [master1:16000.activeMasterManager] master.HMaster: Master has completed initialization
```

## Chapter 12: Security

```
[root@nn1 ~]# cryptsetup --verbose --verify-passphrase luksFormat /dev/sdb1
WARNING!
=====
This will overwrite data on /dev/sdb1 irrevocably.
Are you sure? (Type uppercase yes): YES
```

```
[root@nn1 ~]# cryptsetup luksDump /dev/sdb1
LUKS header information for /dev/sdb1

Version:          1
Cipher name:      aes
Cipher mode:      xts-plain64
Hash spec:        sha1
Payload offset:   4096
MK bits:          256
MK digest:        64 ac 54 cd 94 48 e4 e7 86 7c b6 49 43 92 3b 63 9b 25 9c 9a
MK salt:          98 d2 f9 5d 61 6c c0 34 19 92 59 87 54 82 b7 05
                  0c 7d d0 a4 8f 67 aa 6a d0 6a 1f c7 9c 9b 5b 88
MK iterations:    55625
UUID:             a25bfe03-541d-4e1a-a34b-7061c6141282

Key Slot 0: ENABLED
  Iterations:      218057
  Salt:            b5 d5 89 55 98 7e cf ab 6f 26 c8 76 ea 5a ba fa
                  fb 34 d6 15 d5 94 61 05 2f 34 58 cd 86 47 c5 46
  Key material offset: 8
  AF stripes:      4000
```

```
# dd if=/dev/urandom of=/root/keyfile bs=1024 count=4
# chmod 0400 /root/keyfile
# cryptsetup luksAddKey /dev/sdb1 /root/keyfile
```

```
/dev/mapper/sdb1      /data/namenode      ext4      defaults      0 0
/dev/mapper/sdc1      /data/namenode1    ext4      defaults      0 0
```

```
<property>
  <name>dfs.permissions.superusergroup</name>
  <value>hadoop</value>
</property>
```



```
[root@nn1 ~]# chown hadoop:hadoop /data
[root@nn1 ~]# chown -R hdfs:hadoop /data/namenode
[root@nn1 ~]# ls -ld /data/namenode
drwxr-xr-x 2 hdfs hadoop 4096 Apr 11 04:36 /data/namenode
```

```
[hdfs@nn1 ~]$ id
uid=1006(hdfs) gid=1006(hdfs) groups=1006(hdfs),1003(hadoop)
[hdfs@nn1 ~]$ ls -ld /data
drwxr-xr-x 5 hadoop hadoop 52 Oct 22 16:22 /data
[hdfs@nn1 ~]$ ls -ld /data/namenode
drwxr-xr-x 3 hdfs hdfs 4096 Apr 11 04:48 /data/namenode
[hdfs@nn1 ~]$ ls -ld /data/namenode/current
drwxrwxr-x 2 hdfs hdfs 4096 Apr 11 04:48 /data/namenode/current
```

```
[hdfs@nn1 ~]$ ps -ef | grep namenode
hdfs 1952 1 5 17:55 pts/0 00:00:05 /usr/java/latest/bin/java -Dproc namenode -Xmx1000m -Djava.net.preferIPv4Stack=true -Dhadoop.log.dir=/opt/cluster/hadoop-2.7.3/logs -Dhadoop.log.file=hadoop.log -Dhadoop.home.dir=/opt/cluster/hadoop-2.7.3 -Dhadoop.id.str=hdfs -Dhadoop.root.logger=INFO,console -Djava.library.path=/opt/cluster/hadoop-2.7.3/lib/native -Dhadoop.policy.file=hadoop-policy.xml -Djava.net.preferIPv4Stack=true -Djava.net.preferIPv4Stack=true -Djava.net.preferIPv4Stack=true -Dhadoop.log.dir=/opt/cluster/hadoop-2.7.3/logs -Dhadoop.p.log.file=hadoop-hdfs-namenode-nn1.cluster1.com.log -Dhadoop.home.dir=/opt/cluster/hadoop-2.7.3 -Dhadoop.id.str=hdfs -Dhadoop.root.logger=INFO,RFA -Djava.library.path=/opt/cluster/hadoop-2.7.3/lib/native -Dhadoop.policy.file=hadoop-policy.xml -Djava.net.preferIPv4Stack=true -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.password.file=/opt/cluster/hadoop/etc/hadoop/jmxremote.password -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.port=8004 -Dhadoop.security.logger=INFO,RFA -Dhdfs.audit.logger=INFO,NullAppender -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.password.file=/opt/cluster/hadoop/etc/hadoop/jmxremote.password -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.port=8004 -Dhadoop.security.logger=INFO,RFA -Dhdfs.audit.logger=INFO,NullAppender -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.password.file=/opt/cluster/hadoop/etc/hadoop/jmxremote.password -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false -Dcom.sun.management.jmxremote.port=8004 -Dhadoop.security.logger=INFO,RFA -Dhdfs.audit.logger=INFO,NullAppender -Dhadoop.security.logger=INFO,RFA org.apache.hadoop.hdfs.server.namenode.NameNode
```

```
[hdfs@nn1 ~]$ hadoop fs -put test /
[hdfs@nn1 ~]$ clear
[hdfs@nn1 ~]$ hadoop fs -ls /
Found 1 items
-rw-r--r-- 2 hdfs hadoop 20 2017-04-11 05:14 /test
```

```
[yarn@jt1 ~]$ yarn node -list
17/05/16 18:02:27 INFO client.RMProxy: Connecting to ResourceManager at jt1.cluster1.com/192.168.1.69:9003
Total Nodes:4
Node-Id Node-State Node-Http-Address Number-of-Running-Containers
dn1.cluster1.com:46217 RUNNING dn1.cluster1.com:8090 0
dn4.cluster1.com:54162 RUNNING dn4.cluster1.com:8090 0
dn2.cluster1.com:52476 RUNNING dn2.cluster1.com:8090 0
dn3.cluster1.com:59942 RUNNING dn3.cluster1.com:8090 0
```

```
[hdfs@nn1 ~]$ kms.sh start
Using CATALINA_BASE: /opt/cluster/hadoop/share/hadoop/kms/tomcat
Using CATALINA_HOME: /opt/cluster/hadoop/share/hadoop/kms/tomcat
Using CATALINA_TMPDIR: /opt/cluster/hadoop/share/hadoop/kms/tomcat/temp
Using JRE_HOME: /usr/java/latest
Using CLASSPATH: /opt/cluster/hadoop/share/hadoop/kms/tomcat/bin/bootstrap.jar
Using CATALINA_PID: /tmp/kms.pid
```

```
[hdfs@nn1 ~]$ hadoop key create key1
key1 has been successfully created with options Options{cipher='AES/CTR/NoPadding', bitLength=128
null', attributes=null}.
KMSPClientProvider[http://nn1.cluster1.com:16000/kms/v1/] has been updated.
```

```
[hdfs@nn1 ~]$ hadoop fs -mkdir /secure_zone
[hdfs@nn1 ~]$ hdfs crypto -createZone -keyName key1 -path /secure_zone
Added encryption zone /secure_zone
```

```
[hdfs@nn1 ~]$ hdfs crypto -listZones
/secure_zone key1
```

```
[hdfs@nn1 ~]$ hadoop fs -cp /test /secure_zone
[hdfs@nn1 ~]$ hadoop fs -ls /
Found 2 items
drwxr-xr-x - hdfs hadoop 0 2017-04-11 18:48 /secure_zone
-rw-r--r-- 2 hdfs hadoop 20 2017-04-11 18:47 /test
[hdfs@nn1 ~]$ hadoop fs -ls /secure_zone
Found 1 items
-rw-r--r-- 2 hdfs hadoop 20 2017-04-11 18:48 /secure_zone/test
```

```
[hdfs@nn1 ~]$ ls -la .kms.keystore.crc
-rwx----- 1 hdfs hdfs 20 Apr 11 18:31 .kms.keystore.crc
[hdfs@nn1 ~]$ ls -l kms.keystore
-rwx----- 1 hdfs hdfs 1116 Apr 11 18:31 kms.keystore
```

```
[hdfs@nn1 ~]$ kms.sh stop
Using CATALINA_BASE: /opt/cluster/hadoop/share/hadoop/kms/tomcat
Using CATALINA_HOME: /opt/cluster/hadoop/share/hadoop/kms/tomcat
Using CATALINA_TMPDIR: /opt/cluster/hadoop/share/hadoop/kms/tomcat/temp
Using JRE_HOME: /usr/java/latest
Using CLASSPATH: /opt/cluster/hadoop/share/hadoop/kms/tomcat/bin/b
Using CATALINA_PID: /tmp/kms.pid
[hdfs@nn1 ~]$ hadoop fs -cat /secure_zone/test
cat: Connection refused
```

```
[hdfs@nn1 ~]$ hdfs crypto -listZones
/secure_zone    key1
/secure_zone2  key2
```

```
[hdfs@nn1 ~]$ hadoop fs -cp /secure_zone/test /secure_zone2/
[hdfs@nn1 ~]$ hadoop fs -ls /secure_zone2/
Found 1 items
-rw-r--r--  2 hdfs hadoop      20 2017-04-11 19:22 /secure_zone2/test
```

```
[hadoop@nn1 keys]$ keytool -genkey -alias nn1.cluster1.com -keyalg rsa -ke
ysize 1024 -dname "CN=nn1.cluster1.com,OU=hadoop,O=Netxillon,L=Sydney,ST=N
SW,C=AU" -keypass hadoop@123 -keystore nn1-keystore.jks -storepass hadoop@
123
[hadoop@nn1 keys]$ ls -l
total 4
-rw-rw-r-- 1 hadoop hadoop 1410 Apr 11 22:10 nn1-keystore.jks
```

```
[hadoop@nn1 keys]$ for i in {nn1,dn1,dn2,dn3,dn4,jt1};do keytool -export
 -alias ${i}.cluster1.com -keystore ${i}-keystore.jks -rfc -file ${i}.cr
t -storepass hadoop@123; done
Certificate stored in file <nn1.crt>
Certificate stored in file <dn1.crt>
Certificate stored in file <dn2.crt>
Certificate stored in file <dn3.crt>
Certificate stored in file <dn4.crt>
Certificate stored in file <jt1.crt>
```

```
[hadoop@nn1 keys]$ for i in {nn1,dn1,dn2,dn3,dn4,jt1};do keytool -import
 -noprompt -alias ${i}.cluster1.com -file ${i}.crt -keystore ${i}-trusts
tore.jks -storepass hadoop@123; done
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
```



```
[hadoop@nn1 keys]$ for i in {nn1,dn1,dn2,dn3,dn4,jt1};do keytool -import
-noprompt -alias ${i}.cluster1.com -file ${i}.crt -keystore truststore.
jks -storepass hadoop@123; done
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
Certificate was added to keystore
```

```
[hadoop@nn1 keys]$ cp truststore.jks /opt/cluster/security/jks/
[hadoop@nn1 keys]$ cp nn1-truststore.jks /opt/cluster/security/certs/host-truststore.jks
[hadoop@nn1 keys]$ cp nn1-keystore.jks /opt/cluster/security/certs/host-keystore.jks
```

```
[hadoop@nn1 keys]$ for i in {dn1,dn2,dn3,dn4,jt1}; do scp truststore.jks ${i}.cluster1.com:/opt/c
luster/security/jks/; done
truststore.jks          100% 4058      4.0KB/s   00:00
truststore.jks          100% 4058      4.0KB/s   00:00
truststore.jks          100% 4058      4.0KB/s   00:00
truststore.jks          100% 4058      4.0KB/s   00:00
truststore.jks          100% 4058      4.0KB/s   00:00
[hadoop@nn1 keys]$ for i in {dn1,dn2,dn3,dn4,jt1}; do scp ${i}-truststore.jks ${i}.cluster1.com:/
opt/cluster/security/certs/host-truststore.jks; done
dn1-truststore.jks     100%   703      0.7KB/s   00:00
dn2-truststore.jks     100%   703      0.7KB/s   00:00
dn3-truststore.jks     100%   703      0.7KB/s   00:00
dn4-truststore.jks     100%   703      0.7KB/s   00:00
jt1-truststore.jks     100%   703      0.7KB/s   00:00
[hadoop@nn1 keys]$ for i in {dn1,dn2,dn3,dn4,jt1}; do scp ${i}-keystore.jks ${i}.cluster1.com:/op
t/cluster/security/certs/host-keystore.jks; done
dn1-keystore.jks       100%  1410      1.4KB/s   00:00
dn2-keystore.jks       100%  1410      1.4KB/s   00:00
dn3-keystore.jks       100%  1410      1.4KB/s   00:00
dn4-keystore.jks       100%  1408      1.4KB/s   00:00
jt1-keystore.jks       100%  1409      1.4KB/s   00:00
```

```
[root@repo ~]# clush -g all -b "ls -ld /opt/cluster/security"
-----
dn[1-4].cluster1.com,jt1.cluster1.com (5)
-----
drwxr-xr-x 4 hadoop hadoop 28 Apr 16 07:36 /opt/cluster/security
-----
nn1.cluster1.com
-----
drwxr-xr-x 4 hadoop hadoop 28 Apr 11 22:05 /opt/cluster/security
```



You are about to override how Firefox identifies this site.

**Legitimate banks, stores, and other public sites will not ask you to do this.**

### Server

Location:

[Get Certificate](#)

### Certificate Status

This site attempts to identify itself with invalid information.

[View...](#)

#### **Wrong Site**

The certificate belongs to a different site, which could mean that someone is trying to impersonate this site.

#### **Unknown Identity**

The certificate is not trusted because it hasn't been verified as issued by a trusted authority using a secure signature.

Permanently store this exception

[Confirm Security Exception](#)

[Cancel](#)



# Overview 'nn1.cluster1.com:9000' (active)

<b>Started:</b>	Wed Apr 12 00:35:06 IST 2017
<b>Version:</b>	2.7.3, rUnknown
<b>Compiled:</b>	2016-10-16T08:20Z by root from Unknown
<b>Cluster ID:</b>	CID-830d22f4-5bc4-4ad6-8685-72ced3b61f4a
<b>Block Pool ID:</b>	BP-370677777-192.168.1.70-1491866287092



# Datanode Information

## In operation

Node	Last contact	Admin State	Capacity	Used	Non DFS Used
dn2.cluster1.com:50010 (192.168.1.73:50010)	1	In Service	11.62 GB	136 KB	710.21 MB
dn1.cluster1.com:50010 (192.168.1.72:50010)	1	In Service	15.5 GB	136 KB	935.01 MB
dn3.cluster1.com:50010 (192.168.1.74:50010)	1	In Service	15.5 GB	136 KB	935.01 MB
dn4.cluster1.com:50010 (192.168.1.75:50010)	1	In Service	15.5 GB	136 KB	935.01 MB



## Nodes of the cluster

Logged in as: dr.who

- Cluster
- About
- Nodes
- Node Labels
- Applications
- NEW
- NEW\_SAVING
- SUBMITTED
- ACCEPTED
- RUNNING
- FINISHED
- FAILED
- KILLED
- Scheduler
- Tools

### Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Memory Reserved	VCores Used	VCores Total	VCores Reserved	Active Nodes	Decommissioned Nodes	Lost Nodes	Unhealthy Nodes	Rebooted Nodes
0	0	0	0	0	0 B	32 GB	0 B	0	32	0	4	0	0	0	0

### Scheduler Metrics

Scheduler Type	Scheduling Resource Type	Minimum Allocation	Maximum Allocation
Capacity Scheduler	[MEMORY]	<memory:1024, vCores:1>	<memory:8192, vCores:8>

Show 20 entries

Node Labels	Rack	Node State	Node Address	Node HTTP Address	Last health-update	Health-report	Containers	Mem Used	Mem Avail	VCores Used	VCores Avail	Version
	/default-rack	RUNNING	dn1.cluster1.com:46217	dn1.cluster1.com:8090	Tue May 16 18:06:10 +0530 2017		0	0 B	8 GB	0	8	2.7.3
	/default-rack	RUNNING	dn4.cluster1.com:54162	dn4.cluster1.com:8090	Tue May 16 18:06:08 +0530 2017		0	0 B	8 GB	0	8	2.7.3

192.168.1.70:50070/webhdfs/v1/?op=LISTSTATUS

JSON Raw Data Headers

Save Copy

```
FileStatuses:
  FileStatus:
    0:
      accessTime: 0
      blockSize: 0
      childrenNum: 1
      encBit: true
      fileId: 16494
      group: "hadoop"
      length: 0
      modificationTime: 1491916681841
      owner: "hdfs"
      pathSuffix: "secure_zone"
      permission: "755"
      replication: 0
      storagePolicy: 0
      type: "DIRECTORY"
```

```
zookeeper.ssl.keyStore.location="/opt/cluster/security/certs/host-keystore.jks"
zookeeper.ssl.keyStore.password="hadoop@123"
zookeeper.ssl.trustStore.location="/opt/cluster/security/certs/host-truststore.jks"
zookeeper.ssl.trustStore.password="hadoop@123"
```

```
[hadoop@nn1 conf]$ pwd
/opt/cluster/zoo/conf
[hadoop@nn1 conf]$ cat ../bin/zkServer.sh
#!/usr/bin/env bash

export SERVER_JVMFLAGS="
-Dzookeeper.serverCnxnFactory=org.apache.zookeeper.server.NettyServerCnxnFactory
-Dzookeeper.ssl.keyStore.location=/opt/cluster/security/certs/host-keystore.jks
-Dzookeeper.ssl.keyStore.password=hadoop@123
-Dzookeeper.ssl.trustStore.location=/opt/cluster/security/certs/host-truststore.jks
-Dzookeeper.ssl.trustStore.password=hadoop@123"
```

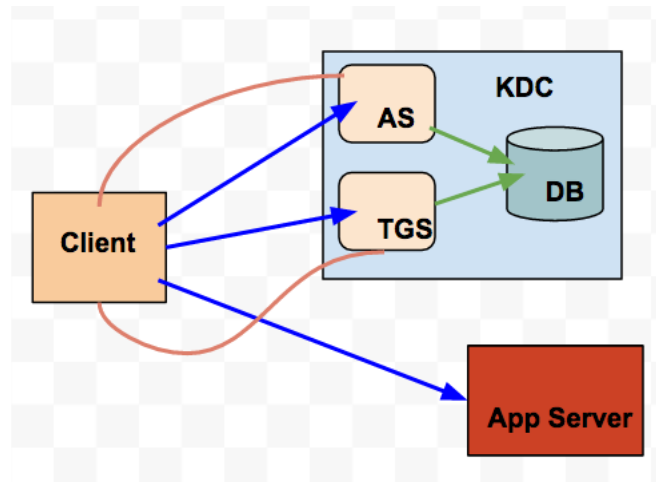
```
[hadoop@nn1 ~]$ jps
15183 QuorumPeerMain
15247 Jps
```

```
log4j.logger.SecurityLogger=INFO, RFAS
log4j.logger.org.apache.hadoop.hdfs.server.namenode.FSNamesystem.audit = INFO, RFAS
log4j.logger.org.apache.hadoop.mapred.AuditLogger = INFO, RFAS
log4j.logger.oozieaudit = INFO, RFAS
log4j.logger.SecurityLogger.org.apache.hadoop.hbase.security.access.AccessController= INFO, RFAS
```

```
$ ls -l ../..../logs/*audit*
hdfs 4348 Apr 11 20:23 ../..../logs/kms-audit.log
hadoop 0 Apr 11 05:11 ../..../logs/SecurityAuth-hdfs.audit
```

```
OK[op=DECRYPT_EEK, key=key1, user=hdfs, accessCount=1, interval=0ms]
OK[op=DECRYPT_EEK, key=key2, user=hdfs, accessCount=1, interval=0ms]
OK[op=DECRYPT_EEK, key=key1, user=hdfs, accessCount=1, interval=10247ms]
OK[op=DECRYPT_EEK, key=key2, user=hdfs, accessCount=1, interval=10059ms]
```

```
hadoop.mapreduce.jobsummary.logger=${hadoop.root.logger}
hadoop.mapreduce.jobsummary.log.file=hadoop-mapreduce.jobsummary.log
hadoop.mapreduce.jobsummary.log.maxfilesize=256MB
hadoop.mapreduce.jobsummary.log.maxbackupindex=20
log4j.appender.JSA=org.apache.log4j.RollingFileAppender
log4j.appender.JSA.File=${hadoop.log.dir}/${hadoop.mapreduce.jobsummary.log.file}
log4j.appender.JSA.MaxFileSize=${hadoop.mapreduce.jobsummary.log.maxfilesize}
log4j.appender.JSA.MaxBackupIndex=${hadoop.mapreduce.jobsummary.log.maxbackupindex}
log4j.appender.JSA.layout=org.apache.log4j.PatternLayout
log4j.appender.JSA.layout.ConversionPattern=%d{yy/MM/dd HH:mm:ss} %p %c{2}: %m%n
log4j.logger.org.apache.hadoop.mapred.JobInProgress$JobSummary=${hadoop.mapreduce.jobsummary.logger}
log4j.additivity.org.apache.hadoop.mapred.JobInProgress$JobSummary=false
```



```
[root@repo ~]# clush -g all -b "date"
-----
dn[1-4].cluster1.com,jt1.cluster1.com,nn1.cluster1.com (6)
-----
Mon Apr 17 04:40:33 IST 2017
```

```
[root@repo ~]# cat /var/kerberos/krb5kdc/kdc.conf
[kdcdefaults]
kdc_ports = 88
kdc_tcp_ports = 88

[realms]
CLUSTER1.COM = {
#master_key_type = aes256-cts
acl_file = /var/kerberos/krb5kdc/kadm5.acl
dict_file = /usr/share/dict/words
admin_keytab = /var/kerberos/krb5kdc/kadm5.keytab
supported_encetypes = aes256-cts:normal aes128-cts:normal des3-hmac-sha1:normal arcfour-hmac:normal
camellia256-cts:normal camellia128-cts:normal des-hmac-sha1:normal des-cbc-md5:normal des-cbc-crc:normal
}
}
```

```
[root@repo ~]# cat /var/kerberos/krb5kdc/kadm5.kadm5.acl kadm5.keytab
[root@repo ~]# cat /var/kerberos/krb5kdc/kadm5.acl
*/admin@CLUSTER1.COM *
```

```
kadmin.local: addprinc hadoop
WARNING: no policy specified for hadoop@CLUSTER1.COM;
Enter password for principal "hadoop@CLUSTER1.COM":
Re-enter password for principal "hadoop@CLUSTER1.COM":
Principal "hadoop@CLUSTER1.COM" created.
```



```
kadmin.local: listprincs
K/M@CLUSTER1.COM
hadoop@CLUSTER1.COM
hdfs@CLUSTER1.COM
kadmin/admin@CLUSTER1.COM
kadmin/changepw@CLUSTER1.COM
kadmin/repo.cluster1.com@CLUSTER1.COM
krbtgt/CLUSTER1.COM@CLUSTER1.COM
root/admin@CLUSTER1.COM
yarn@CLUSTER1.COM
```

```
kadmin.local: addprinc -randkey host/repo.cluster1.com
WARNING: no policy specified for host/repo.cluster1.com@CLUSTER1.COM
Principal "host/repo.cluster1.com@CLUSTER1.COM" created.
```

```
kadmin.local: ktadd host/repo.cluster1.com
Entry for principal host/repo.cluster1.com with kvno 2,
FILE:/etc/krb5.keytab.
Entry for principal host/repo.cluster1.com with kvno 2,
FILE:/etc/krb5.keytab.
Entry for principal host/repo.cluster1.com with kvno 2,
krb5.keytab.
```

```
[root@repo ~]# ktutil
ktutil: rkt /etc/krb5.keytab
ktutil: l
slot KVNO Principal
-----
1      2      host/repo.cluster1.com@CLUSTER1.COM
2      2      host/repo.cluster1.com@CLUSTER1.COM
3      2      host/repo.cluster1.com@CLUSTER1.COM
```



Authentication Configuration

<p>User Information</p> <p><input type="checkbox"/> Cache Information</p> <p><input type="checkbox"/> Use LDAP</p> <p><input type="checkbox"/> Use NIS</p> <p><input type="checkbox"/> Use IPAv2</p> <p><input type="checkbox"/> Use Winbind</p>	<p>Authentication</p> <p><input type="checkbox"/> Use MD5 Passwords</p> <p><input type="checkbox"/> Use Shadow Passwords</p> <p><input type="checkbox"/> Use LDAP Authentication</p> <p><input checked="" type="checkbox"/> Use Kerberos</p> <p><input type="checkbox"/> Use Fingerprint reader</p> <p><input type="checkbox"/> Use Winbind Authentication</p> <p><input checked="" type="checkbox"/> Local authorization is sufficient</p>
--	---

Kerberos Settings

Realm: CLUSTER1.COM  
 KDC: repo.cluster1.com  
 Admin Server: repo.cluster1.com

Use DNS to resolve hosts to realms  
 Use DNS to locate KDCs for realms

```

[root@repo ~]# su - hadoop
Last login: Fri Oct  7 06:49:33 IST 2016 on pts/0
[hadoop@repo ~]$ klist
klist: No credentials cache found (ticket cache KEYRING:persistent:1003:1003)
[hadoop@repo ~]$ kinit
Password for hadoop@CLUSTER1.COM:
[hadoop@repo ~]$ klist
Ticket cache: KEYRING:persistent:1003:1003
Default principal: hadoop@CLUSTER1.COM

Valid starting          Expires                Service principal
04/17/2017 06:09:34    04/18/2017 06:09:34    krbtgt/CLUSTER1.COM@CLUSTER1.COM
renew until 04/17/2017 06:09:34
  
```

```
[root@nn1 ~]# kadmin
Authenticating as principal root/admin@CLUSTER1.COM with password.
Password for root/admin@CLUSTER1.COM:
kadmin: ktadd host/nn1.cluster1.com
Entry for principal host/nn1.cluster1.com with kvno 5, encryption type aes256-cts-hmac-sha1-96 added to keytab FILE:/etc/krb5
.keytab.
Entry for principal host/nn1.cluster1.com with kvno 5, encryption type aes128-cts-hmac-sha1-96 added to keytab FILE:/etc/krb5
.keytab.
Entry for principal host/nn1.cluster1.com with kvno 5, encryption type des3-cbc-sha1 added to keytab FILE:/etc/krb5.keytab.
Entry for principal host/nn1.cluster1.com with kvno 5, encryption type arcfour-hmac added to keytab FILE:/etc/krb5.keytab.
Entry for principal host/nn1.cluster1.com with kvno 5, encryption type camellia256-cts-cmac added to keytab FILE:/etc/krb5.ke
ytab.
```

```
[hadoop@repo ~]$ ls -l .ssh/
total 4
-rw-r--r-- 1 hadoop hadoop 1341 Oct  7  2016 known_hosts
[hadoop@repo ~]$ ssh nn1.cluster1.com
Last login: Mon Apr 17 06:29:35 2017 from repo.cluster1.com
```

```
kadmin.local: listprincs
K/M@CLUSTER1.COM
hadoop@CLUSTER1.COM
hdfs@CLUSTER1.COM
host/dn1.cluster.com@CLUSTER1.COM
host/dn2.cluster.com@CLUSTER1.COM
host/dn3.cluster.com@CLUSTER1.COM
host/dn4.cluster.com@CLUSTER1.COM
host/jt1.cluster.com@CLUSTER1.COM
host/nn1.cluster1.com@CLUSTER1.COM
host/repo.cluster1.com@CLUSTER1.COM
```

```
kadmin.local: addprinc -randkey yarn/nn1.cluster1.com
WARNING: no policy specified for yarn/nn1.cluster1.com@C
Principal "yarn/nn1.cluster1.com@CLUSTER1.COM" created.
kadmin.local: addprinc -randkey mapred/nn1.cluster1.com
WARNING: no policy specified for mapred/nn1.cluster1.com
Principal "mapred/nn1.cluster1.com@CLUSTER1.COM" created
kadmin.local: addprinc -randkey HTTP/nn1.cluster1.com
WARNING: no policy specified for HTTP/nn1.cluster1.com@C
Principal "HTTP/nn1.cluster1.com@CLUSTER1.COM" created.
kadmin.local: addprinc -randkey hdfs/nn1.cluster1.com
WARNING: no policy specified for hdfs/nn1.cluster1.com@C
Principal "hdfs/nn1.cluster1.com@CLUSTER1.COM" created.
```

```
[hdfs@nn1 ~]$ hadoop fs -ls /
Apr 17 07:40:04 WARN ipc.Client: Exception encountered while connecting to the server :
javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException: No valid credentials
provided (Mechanism level: Failed to find any Kerberos tgt)]
```

```
[hdfs@nn1 ~]$ klist
Ticket cache: FILE:/tmp/krb5cc_1011
Default principal: hdfs@CLUSTER1.COM

Valid starting          Expires                Service principal
04/17/2017 07:45:29    04/18/2017 07:43:43  krbtgt/CLUSTER1.COM@CLUSTER1.COM
        renew until 04/17/2017 07:45:29

[hdfs@nn1 ~]$ hadoop fs -ls /
Found 4 items
drwxr-xr-x   - hdfs  hadoop           0 2017-04-11 18:48 /secure_zone
drwxr-xr-x   - hdfs  hadoop           0 2017-04-11 19:22 /secure_zone2
-rw-r--r--   2 hdfs  hadoop          19 2017-04-17 03:34 /test_audit
-rw-r--r--   2 hdfs  hadoop          20 2017-04-11 19:22 /tmp
```

```
#!/bin/bash

# Generate Hosts principals

for i in `cat dn_host_list`
do
    kadmin.local -q "addprinc -randkey host/$i"
    kadmin.local -q "addprinc -randkey http/$i"
    kadmin.local -q "addprinc -randkey dn/$i"
    kadmin.local -q "xst -norandkey -k dn.hdfs.keytab host/$i"
    kadmin.local -q "xst -norandkey -k dn.hdfs.keytab http/$i"
    kadmin.local -q "xst -norandkey -k dn.hdfs.keytab dn/$i"
done
```

```
#!/bin/bash
```

```
for k in `cat nn_host_list`  
do
```

```
    echo "$k"
```

```
    kadmin.local -q "addprinc -randkey host/$k"
```

```
    kadmin.local -q "addprinc -randkey http/$k"
```

```
    kadmin.local -q "addprinc -randkey nn/$k"
```

```
    kadmin.local -q "xst -norandkey -k nn.hdfs.keytab host/$k"
```

```
    kadmin.local -q "xst -norandkey -k nn.hdfs.keytab http/$k"
```

```
    kadmin.local -q "xst -norandkey -k nn.hdfs.keytab nn/$k"
```

```
    kadmin.local -q "xst -norandkey -k dn.hdfs.keytab host/$k"
```

```
    kadmin.local -q "xst -norandkey -k dn.hdfs.keytab http/$k"
```

```
    kadmin.local -q "xst -norandkey -k dn.hdfs.keytab nn/$k"
```

```
done
```

```
#!/bin/bash
```

```
for k in `cat user_host_list`  
do
```

```
    kadmin.local -q "xst -norandkey -k user.hdfs.keytab host/$k"
```

```
done
```

```
for p in `cat user_list`  
do
```

```
    kadmin.local -q "xst -norandkey -k user.hdfs.keytab $p"
```

```
done
```

```
# ls -l /usr/java/latest/jre/lib/security/

root root  2770 Oct  8  2013 blacklist
root root 82586 Oct  8  2013 cacerts
root root   158 Oct  8  2013 javafx.policy
root root  2254 Oct  8  2013 java.policy
root root 17677 Oct  8  2013 java.security
root root   98 Oct  8  2013 javaws.policy
root root  2865 Oct  8  2013 local_policy.jar
root root    0 Oct  8  2013 trusted.libraries
root root  2397 Oct  8  2013 US_export_policy.jar
```