



SNS COLLEGE OF TECHNOLOGY

(An Autonomous Institution)

Re-accredited by NAAC with A+ grade, Accredited by NBA(CSE, IT, ECE, EEE & Mechanical)
Approved by AICTE, New Delhi, Recognized by UGC, Affiliated to Anna University, Chennai



Department of MCA

Topic: HDFS – Monitoring & Maintenance

COURSE

16CA917

Big Data
Analytics

UNIT - III

Hadoop
Environment

CLASS

V Semester /
III MCA



Session Objectives




- ❑ Understand the routine administrative procedure followed in Hadoop
- ❑ Demonstrate how to add/remove node in Hadoop cluster



Monitoring



- ❑ Purpose of monitoring is to detect when the cluster is not providing the expected level of service
- ❑ Important to monitor: the namenodes (primary and secondary) and the jobtracker
- ❑ Cluster will be provided with extra capacity to manage other nodes





Logging



- ❑ All Hadoop daemons produce logfiles, which are useful for finding out what is happening in the system
- ❑ When debugging a problem, it is convenient to change the log level temporarily for a particular component in the system using log4j
- ❑ Hadoop daemons expose a web page that produces a thread dump for all running threads in the daemon's JVM



Metrics



- ❑ HDFS and MapReduce daemons collect information about events and measurements that are collectively known as metrics

- ❑ Metrics belong to a context,
 - dfs
 - mapred
 - rpc
 - Jvm



Metrics



- ❑ Java Management Extensions (JMX) is a standard Java API for monitoring and managing applications
- ❑ MBeans that expose the metrics in the “dfs” and “rpc”



Routine Administrative Procedures



Metadata backups:

- Keep multiple copies of different ages (hr/day/week/month) to protect against corruption
- A way to make backups is to write a script to periodically archive the secondary namenode's previous.checkpoint subdirectory to an offsite location

```
% hdfs dfsadmin -fetchImage fsimage.backup
```



Routine Administrative Procedures



Data backups

- Prioritize your data, highest priority is the data that cannot be regenerated, data is to regenerate: low priority
- distcp tool is ideal for making backups to other HDFS clusters
- HDFS allows administrators and users to take snapshots of the filesystem. A snapshot is a read-only copy of a filesystem subtree at a given point in time.



Routine Administrative Procedures



Filesystem check

- It is advisable to run HDFS's fsck tool regularly on the whole filesystem to proactively look for missing or corrupt blocks

Filesystem balancer

- Run balancer tool regularly to keep filesystem datanodes evenly balanced



Commissioning New node

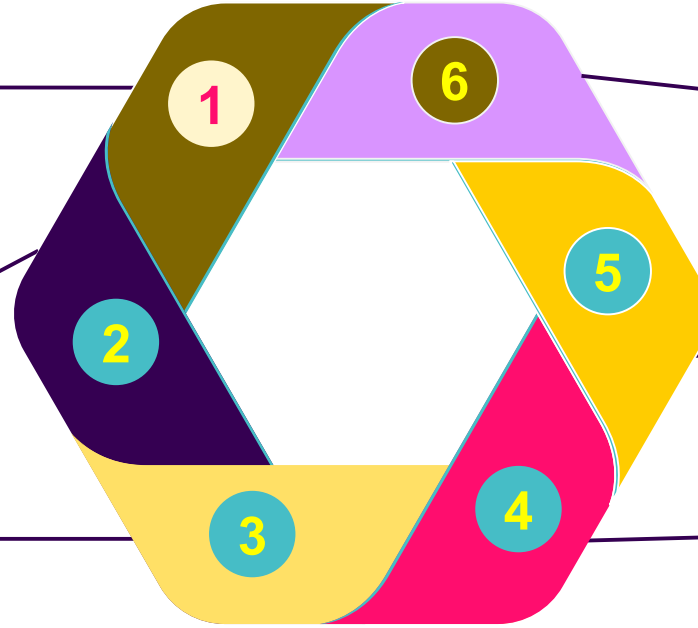


Add network addresses of the new nodes to the include file

Update namenode with new set of permitted datanodes using this command

```
% hadoop dfsadmin -refreshNodes
```

Update the slaves file with the new nodes, so that they are included in future operations the Hadoop control scripts



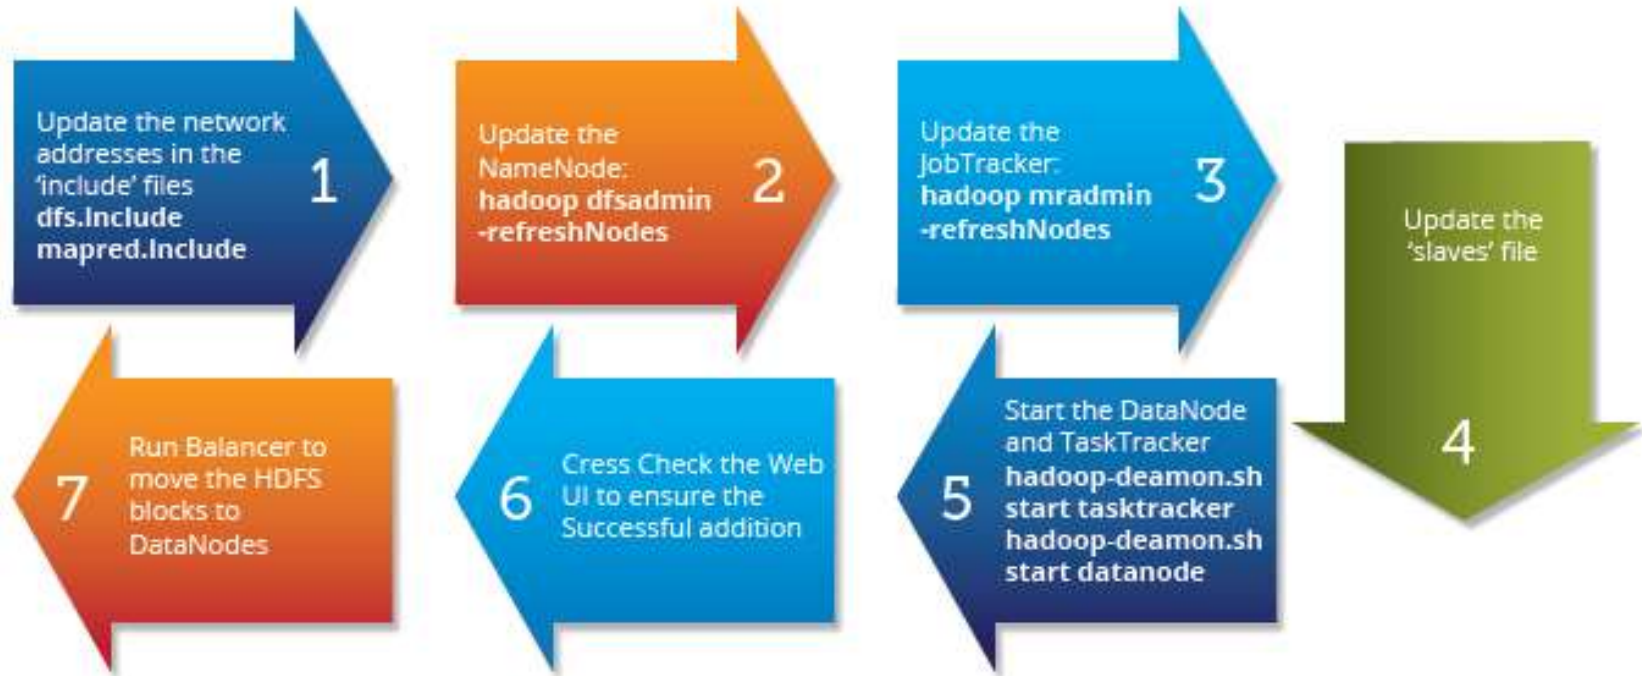
Check new datanodes & tasktrackers appear in the web

Restart the MapReduce cluster

Start the new datanodes



Commissioning New node





Commissioning New node



To add new nodes to the cluster:

1. Add the network addresses of the new nodes to the include file.

hdfs-site.xml

```
<property>
<name>dfs.hosts</name>
<value>/<hadoop-home>/conf/includes</value>
<final>>true</final>
</property>
```

mapred-site.xml

```
<property>
<name>mapred.hosts</name>
<value>/<hadoop-home>/conf/includes</value>
<final>>true</final>
</property>
```

[Datanodes](#) that are permitted to connect to the namenode are specified in a file whose name is specified by the dfs.hosts property.



Commissioning New node



Includes file resides on the namenode's local filesystem, and it contains a line for each datanode, specified by network address (as reported by the datanode; you can see what this is by looking at the namenode's web UI). If you need to specify multiple network addresses for a datanode, put them on one line, separated by whitespace.

eg :

slave01

slave02

slave03

.....

Similarly, tasktrackers that may connect to the jobtracker are specified in a file whose name is specified by the `mapred.hosts` property. In most cases, there is one shared file, referred to as the include file, that both `dfs.hosts` and `mapred.hosts` refer to, since nodes in the cluster run both datanode and tasktracker daemons.



Commissioning New node



2. Update the namenode with the new set of permitted datanodes using this command:

```
% hadoop dfsadmin -refreshNodes
```

3. Update the jobtracker with the new set of permitted tasktrackers using this command:

```
% hadoop mradmin -refreshNodes
```

4. Update the slaves file with the new nodes, so that they are included in future operations performed by the Hadoop control scripts.

5. Start the new datanodes and tasktrackers.

6. Check that the new datanodes and tasktrackers appear in the web UI.



Commissioning New node



2. Update the namenode with the new set of permitted datanodes using this command:

```
% hadoop dfsadmin -refreshNodes
```

3. Update the jobtracker with the new set of permitted tasktrackers using this command:

```
% hadoop mradmin -refreshNodes
```

4. Update the slaves file with the new nodes, so that they are included in future operations performed by the Hadoop control scripts.

5. Start the new datanodes and tasktrackers.

6. Check that the new datanodes and tasktrackers appear in the web UI.



Decommissioning New node



1

Add the network addresses of the nodes to be decommissioned to the exclude file. Do not update the include file at this point

2

Restart the MapReduce cluster to stop the task trackers on the nodes being decommissioned

3

Update the namenode with the new set of permitted datanodes, with this command:

```
%hadoop dfsadmin –  
refreshNodes
```




Decommissioning



4

Go to the web UI and check whether the admin state has changed to “Decommission In Progress” for the datanodes being decommissioned

5

When all the datanodes report their state as “Decommissioned,” then all the blocks have been replicated. Shut down the decommissioned nodes

6

Remove the nodes from the include file, and run:

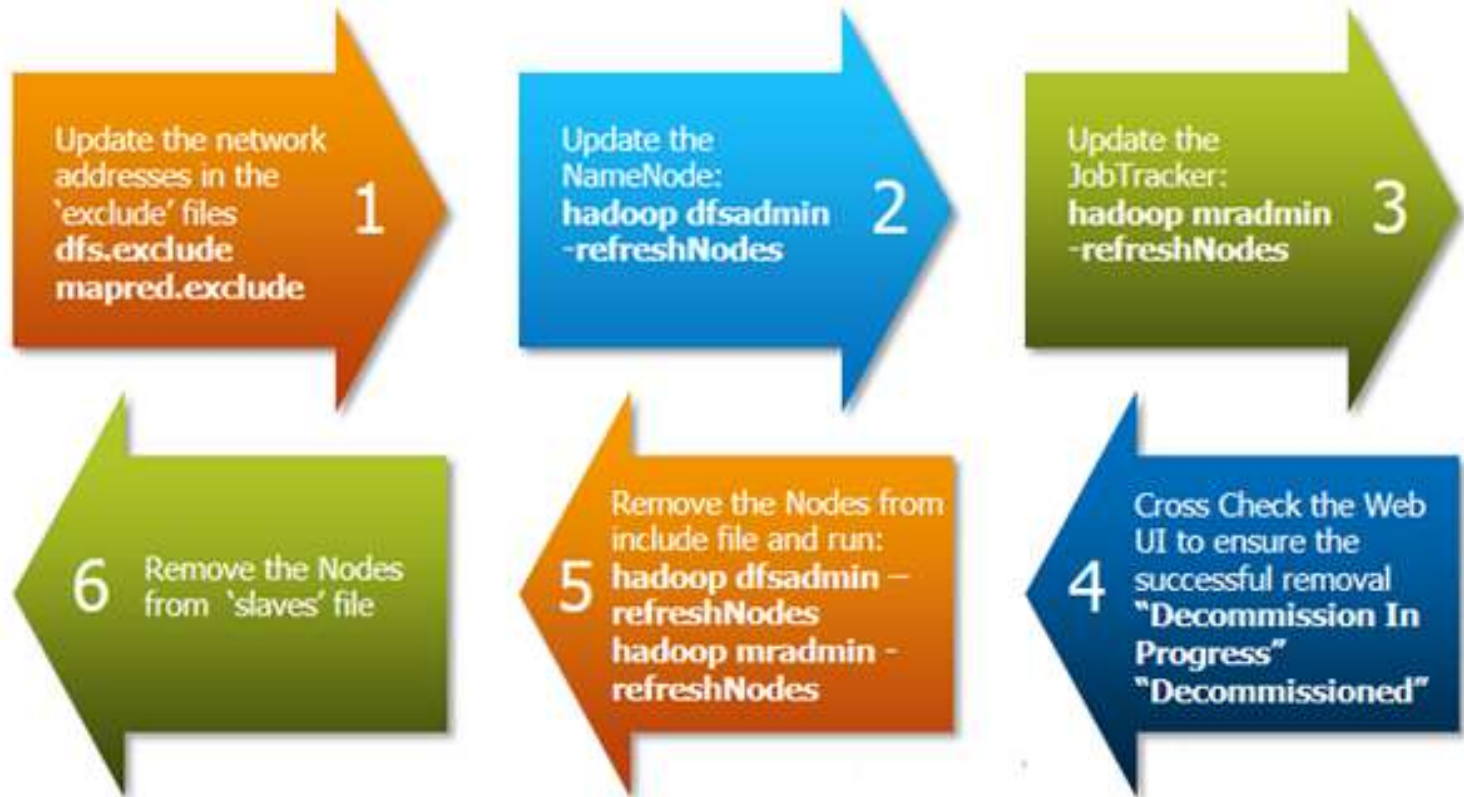
```
% hadoop  
dfsadmin -  
refreshNodes
```

7

Remove the nodes from the slaves file



Decommissioning New node





Decommissioning New node



To remove nodes from the cluster:

1. Add the network addresses of the nodes to be decommissioned to the exclude file. Do not update the include file at this point.

step by step process to decommission a DataNode in the cluster.

The first task is to update the 'exclude' files for both HDFS (hdfs-site.xml) and MapReduce (mapred-site.xml).

The 'exclude' file:

for jobtracker contains the list of hosts that should be excluded by the jobtracker. If the value is empty, no hosts are excluded.

for Namenode contains a list of hosts that are not permitted to connect to the Namenode. Here is the sample configuration for the exclude file in hdfs-site.xml and mapred-site.xml:

hdfs-site.xml

```
<property>
<name>dfs.hosts.exclude</name>
<value>/home/hadoop/excludes</value>
  <final>true</final>
</property>
```



Decommissioning New node



mapred-site.xml

```
<property>
<name>mapred.hosts.exclude</name>
<value>/home/hadoop/excludes</value>
<final>true</final>
</property>
```

Note: The full pathname of the files must be specified.

2. Update the namenode with the new set of permitted datanodes, using this command:

```
% hadoop dfsadmin -refreshNodes
```

3. Update the jobtracker with the new set of permitted tasktrackers using this command:

```
% hadoop mradmin -refreshNodes
```

4. Go to the web UI and check whether the admin state has changed to “Decommission In Progress” for the datanodes being decommissioned. They will start copying their blocks to other datanodes in the cluster.

5. When all the datanodes report their state as “Decommissioned,” all the blocks have been replicated. Shut down the decommissioned nodes.



Decommissioning New node



6. Remove the nodes from the include file, and run:

```
% hadoop dfsadmin -refreshNodes
```

```
% hadoop mradmin -refreshNodes
```

7. Remove the nodes from the slaves file.



Assessment



1. fsck tool is used check
 - A. Missing/corrupting blocks
 - B. Workload among nodes
 - C. Backup of meta data
 - D. Backup of data blocks
2. Include and excludes files for
 3. Logging
 - B. monitoring
 - C. Cluster Mgt
 - D. Backup
3. Purpose of metrics ?



References



Book

- ❑ Tom White, “ Hadoop: The Definitive Guide” Third Edition, O’reilly Media, 4th Edition, 2012

Web Resources

- ❑ https://www.tutorialspoint.com/map_reduce/map_reduce_hadoop_administration.htm#:~:text=HDFS%20administration%20includes%20monitoring%20the,nodes%2C%20application%20status%2C%20etc.
- ❑ <http://www.hadoopadmin.co.in/hadoop-administration-and-maintenance/>