

Hadoop

[Hadoop](#) is an Apache project which develops eponymous software for analyzing “big data” over many systems in a hadoop cluster. Because Hadoop clusters assume that they are dedicated to the purpose of running hadoop jobs, they don't interact well with the existing cluster system - because the two scheduling systems aren't aware of each other, they can both run jobs and cause nodes to become over occupied, reducing overall performance.

Fortunately, there is a piece of software called [Hadoop on Demand](#) (or HOD) which allows Hadoop to be run in preexisting cluster systems.

Hadoop on Demand (HOD)

HOD acts under the normal cluster system, creating a job for you that sets up a hadoop cluster. You can then run your hadoop jobs on this cluster. When you use HOD to create the Hadoop cluster (or allocate in HOD terminology), you configure how many nodes you want in your cluster, HOD then uses this information to reserve slots in the cluster system and so prevent over occupation. Using HOD therefore has 4 stages:

1. Configure HOD
2. Allocate the HOD cluster
3. Load your data into Hadoop and run your jobs
4. Stop (deallocate) the HOD cluster

Configure HOD

First you need to create a directory to hold the hadoop configuration and configure your environment to use the correct version of HOD:

```
[jbarber@submit ~]$ mkdir -p ~/hod/cluster
[jbarber@submit ~]$ mv hodrc ~/hod/
[jbarber@submit ~]$ module load hadoop/1.2.1
```

An example of the *hodrc* file which you should copy to the *~/hod/* directory follows:

```
[hod]
stream                = True
java-home              = /usr/lib/jvm/java-1.6.0-
openjdk-1.6.0.0.x86_64/jre
cluster                = batch
cluster-factor         = 1.8
xrs-port-range         = 32768-65536
debug                  = 3
allocate-wait-time     = 3600
original-dir           = /usr/local/hadoop-1.2.1
temp-dir               = /tmp/jbarber
log-dir                = /tmp/jbarber
```

```
[ringmaster]
register                = True
stream                 = False
http-port-range        = 8000-9000
xrs-port-range         = 32768-65536
debug                  = 3
max-master-failures    = 10
temp-dir               = /tmp/jbarber
log-dir                = /tmp/jbarber
work-dirs               = /tmp/jbarber/hod/1,/tmp/jbarber/hod/2

[hodring]
stream                 = False
register                = True
http-port-range        = 8000-9000
xrs-port-range         = 32768-65536
hadoop-port-range      = 50000-60000
debug                  = 3
java-home              = /usr/lib/jvm/java-1.6.0-
openjdk-1.6.0.0.x86_64/jre
temp-dir               = /tmp/jbarber
log-dir                = /tmp/jbarber

[resource_manager]
queue                  = batch
batch-home             = /usr
id                     = torque
log-dir                = /tmp/jbarber

[gridservice-mapred]
external               = False
tracker_port           = 8030
info_port              = 50080
pkgs                   = /usr/local/hadoop-1.2.1
log-dir                = /tmp/jbarber

[gridservice-hdfs]
external               = False
fs_port                = 8020
info_port              = 50070
pkgs                   = /usr/local/hadoop-1.2.1
log-dir                = /tmp/jbarber
```

You must change this example hodrc file to suit your requirements. At the least - you must change *jbarber* to your username. If you want to use a different version of the JVM you will need to change the *java-home* variable, and if you're not using Hadoop 1.2.1 you will also need to change the path */usr/local/hadoop-1.2.1* to point at the right location.

Allocate the HOD cluster

Next you create the Hadoop cluster with the HOD allocate command, giving the location of the HOD

configuration (`-c ~/hod/hodrc`), where to store the cluster information (`-d ~/hod/cluster`) and stating how many nodes you want in the cluster (here we ask for the minimum of 3 nodes with `-n 3`):

```
[jbarber@submit ~]$ hod -c ~/hod/hodrc allocate -d ~/hod/cluster -n 3
INFO - Cluster Id 74237.maui.grid.fe.up.pt
INFO - HDFS UI at http://avafat01.grid.fe.up.pt:53253
INFO - Mapred UI at http://avafat01.grid.fe.up.pt:54906
INFO - hadoop-site.xml at /homes/jbarber/hod/cluster
```

The command will tell you under what ID the HOD job is running, and the Hadoop cluster endpoint URLs. If you forget this information, you can find it later through the HOD `list` and `info` commands:

```
[jbarber@submit ~]$ hod -c ~/hod/hodrc list
INFO - alive 74237.maui.grid.fe.up.pt /homes/jbarber/hod/cluster
[jbarber@submit ~]$ hod -c ~/hod/hodrc list info /homes/jbarber/hod/cluster
INFO - HDFS UI at http://avafat01.grid.fe.up.pt:53253
INFO - Mapred UI at http://avafat01.grid.fe.up.pt:54906
INFO - Nodecount 3
INFO - Cluster Id 74237.maui.grid.fe.up.pt
INFO - hadoop-site.xml at /homes/jbarber/hod/cluster
```

Upload data and process it

You can now use the `hadoop` command and the configuration file created by HOD to interact with the hadoop cluster. First, you must tell the `hadoop` command which JVM to use, after that the full hadoop environment should be available:

```
[jbarber@submit ~]$ export JAVA_HOME=/usr/lib/jvm/java-1.6.0-
openjdk-1.6.0.0.x86_64/jre
[jbarber@submit ~]$ hadoop fs -conf ~/hod/cluster/hadoop-site.xml -ls /
Found 1 items
drwxr-xr-x - jbarber supergroup 0 2013-10-31 09:03 /mapredsystem
[jbarber@submit ~]$ hadoop fs -conf ~/hod/cluster/hadoop-site.xml -put
data.file /
[jbarber@submit ~]$ hadoop fs -conf ~/hod/cluster/hadoop-site.xml -ls /
Found 2 items
-rw-r--r-- 3 jbarber supergroup 38096663 2013-11-04 11:06 /data.file
drwxr-xr-x - jbarber supergroup 0 2013-11-04 11:00 /mapredsystem
```

Stop the cluster

The HOD command to stop the cluster is called `deallocate`. To use it you just need to provide the path to the cluster configuration:

```
[jbarber@submit ~]$ hod -c ~/hod/hodrc deallocate -d ~/hod/cluster
```

After this command completes, it may take a minute for the cluster job to complete, but you should make sure it changes from the R state to C, and if it doesn't, stop it manually with the `qdel` command:

```
[jbarber@submit ~]$ qstat
Job id          Name          User          Time Use S Queue
-----
74237.mai      HOD           jbarber       00:00:21 R batch
[jbarber@submit ~]$ qdel 74237.mai.grid.fe.up.pt
```

From:
<https://grid.fe.up.pt/dokuwiki/> - **GRID FEUP**

Permanent link:
<https://grid.fe.up.pt/dokuwiki/doku.php?id=documentation:software:hadoop>

Last update: **2016/05/25 16:34**

