

GangliaMetrics

Edit the file `hadoop-metrics.properties`, if it is not already present in your config directory.

The file should contain at least the following lines:

```
dfs.class=org.apache.hadoop.metrics.ganglia.GangliaContext
dfs.period=10
dfs.servers=<at:var at:name="GANGLIA" />:8649

mapred.class=org.apache.hadoop.metrics.ganglia.GangliaContext
mapred.period=10
mapred.servers=<at:var at:name="GANGLIA" />:8649

jvm.class=org.apache.hadoop.metrics.ganglia.GangliaContext
jvm.period=10
jvm.servers=<at:var at:name="GANGLIA" />:8649

rpc.class=org.apache.hadoop.metrics.ganglia.GangliaContext
rpc.period=10
rpc.servers=<at:var at:name="GANGLIA" />:8649
```

Replace with the hostname or IP of the ganglia endpoint. Depending on your Ganglia configuration, the host or IP you use might be:

- The value in the `udp_send_channel` you use in your `/etc/gmond.conf` (look for the line which says `mcast_join=<IP ADDRESS>`).
- The gmetad server
- Localhost

If you support multiple clusters, you may want to have Hadoop show up as a separate cluster. Add a data source lines for each cluster:
In the `gmetad.conf` file, add a `data_source` of the form

```
data_source "<at:var at:name="MyHadoopCluster" />" <at:var at:name="GANGLIA" />:8649
```

Source: <http://www.ibm.com/developerworks/wikis/display/WikiPtype/ganglia>

Be aware that versions 0.18.1, 0.19.0, and prior need to be patched in order to get Ganglia working; refer to [JIRA issue 3422](#).

Additionally, the Ganglia protocol change significantly between Ganglia 3.0 and Ganglia 3.1 (i.e., Ganglia 3.1 is not compatible with Ganglia 3.0 clients). This caused Hadoop to not work with Ganglia 3.1; there is a patch available for this, [HADOOP-4675](#). As of November 2010, this patch has been rolled into the mainline for 0.20.2 and later. To use the Ganglia 3.1 protocol in place of the 3.0, substitute `org.apache.hadoop.metrics.ganglia.GangliaContext31` for `org.apache.hadoop.metrics.ganglia.GangliaContext` in the `hadoop-metrics.properties` lines above.