


# hadoop-metrics.properties

The `hadoop-metrics.properties` files direct MapR where to output MRV1 job metrics and [service metrics](#) reports. You can output the reports to file (`FileContext`), to [Ganglia 3.1](#) (`MapRGangliaContext31`), or disable the output.

There are two `hadoop-metrics.properties` files:

- `/opt/mapr/hadoop/hadoop-0.2.0/conf/hadoop-metrics.properties` specifies output for MRV1 job metrics
- `/opt/mapr/conf/hadoop-metrics.properties` specifies output for MapR-specific services

The following table describes the parameters for each service in the `hadoop-metrics.properties` files.

Parameter	Description
<code>&lt;service&gt;.class</code>	The class that implements the interface responsible for sending the service metrics to the appropriate handler. When implementing a class that sends metrics to Ganglia, set this property to the class name.
<code>&lt;service&gt;.period</code>	The interval between 2 service metrics data exports to the appropriate interface. This is independent of how often are the metrics updated in the framework.
<code>&lt;service&gt;.fileName</code>	The path to the file where service metrics are exported when the <code>cldb.class</code> property is set to <code>FileContext</code> .
<code>&lt;service&gt;.servers</code>	The location of the gmon or gmeta that is aggregating metrics for this instance of the service, when the <code>cldb.class</code> property is set to <code>GangliaContext</code> . <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> By default, the CLDB and FileServer metrics are sent via unicast to the Ganglia gmon server running on localhost. To send the metrics directly to a Gmeta server, change the <code>cldb.servers</code> property to the hostname of the Gmeta server. To send the metrics to a multicast channel, change the <code>cldb.servers</code> property to the IP address of the multicast channel.</div>
<code>&lt;service&gt;.spooof</code>	Specifies whether the metrics being sent out from the server should be spoofed as coming from another server. All our fileserver metrics are also on cldb, but to make it appear to end users as if these properties were emitted by fileserver host, we spoof the metrics to Ganglia using this property. Currently only used for the FileServer service.

## Configure Metrics

1. Edit the appropriate `hadoop-metrics.properties` file.
  - a. For the MapR-FS and the CLDB edit `/opt/mapr/conf/hadoop-metrics.properties` on all CLDB nodes.
  - b. For MRV1 Metrics, edit `/opt/mapr/hadoop/hadoop-0.2.0/conf/hadoop-metrics.properties` on all metrics nodes.
2. In the sections specific to the service:
  - Un-comment the lines pertaining to the context to which you wish the service to send metrics.
  - Comment out the lines pertaining to other contexts.
3. Restart the service.

## Disable Metrics

Comment out the following parameters in the `/opt/mapr/hadoop/hadoop-0.2.0/conf/hadoop-metrics.properties` file to disable job metrics.

- `maprmepredvariant.class=com.mapr.job.mngmnt.hadoop.metrics.MaprRPCContext`
- `maprmepredvariant.period=10`
- `maprmapred.class=com.mapr.job.mngmnt.hadoop.metrics.MaprRPCContextFinal`
- `maprmapred.period=10`

 For information on how to enable or disable MapR-FS and CLDB metrics, see [service metrics](#).

## Examples

The `hadoop-metrics.properties` files are organized into sections for each service that provides metrics. Each section is divided into

subsections for the three contexts.

## MRv1 Job Metrics Example

The following file disables MRv1 job metrics in the `/opt/mapr/hadoop/hadoop-0.2.0/conf/hadoop-metrics.properties` file:



```
# Configuration of the "dfs" context for null
dfs.class=org.apache.hadoop.metrics.spi.NullContext
# Configuration of the "dfs" context for file
#dfs.class=org.apache.hadoop.metrics.file.FileContext
#dfs.period=10
#dfs.fileName=/tmp/dfsmetrics.log
# Configuration of the "dfs" context for ganglia
# Pick one: Ganglia 3.0 (former) or Ganglia 3.1 (latter)
# dfs.class=org.apache.hadoop.metrics.ganglia.GangliaContext
# dfs.class=org.apache.hadoop.metrics.ganglia.GangliaContext31
# dfs.period=10
# dfs.servers=localhost:8649

# Configuration of the "mapred" context for default
mapred.class=org.apache.hadoop.metrics.spi.MapRDefaultContext
mapred.period=30

# Configuration of the "mapred" context for file
#mapred.class=org.apache.hadoop.metrics.file.FileContext
#mapred.period=10
#mapred.fileName=/tmp/mrmetrics.log
# Configuration of the "mapred" context for ganglia
# Pick one: Ganglia 3.0 (former) or Ganglia 3.1 (latter)
# mapred.class=org.apache.hadoop.metrics.ganglia.GangliaContext
# mapred.class=org.apache.hadoop.metrics.ganglia.GangliaContext31
# mapred.period=10
# mapred.servers=localhost:8649

# Configuration of the "jvm" context for null
jvm.class=org.apache.hadoop.metrics.spi.NullContext
# Configuration of the "jvm" context for file
#jvm.class=org.apache.hadoop.metrics.file.FileContext
#jvm.period=10
#jvm.fileName=/tmp/jvmmetrics.log
# Configuration of the "jvm" context for ganglia
# jvm.class=org.apache.hadoop.metrics.ganglia.GangliaContext
# jvm.period=10
# jvm.servers=localhost:8649

# Configuration of the "ugi" context for null
ugi.class=org.apache.hadoop.metrics.spi.NullContext

# Configuration of the "fairscheduler" context for null
fairscheduler.class=org.apache.hadoop.metrics.spi.NullContext
# Configuration of the "fairscheduler" context for file
#fairscheduler.class=org.apache.hadoop.metrics.file.FileContext
#fairscheduler.period=10
#fairscheduler.fileName=/tmp/fairschedulermetrics.log
# Configuration of the "fairscheduler" context for ganglia
# fairscheduler.class=org.apache.hadoop.metrics.ganglia.GangliaContext
# fairscheduler.period=10
# fairscheduler.servers=localhost:8649

#maprmapredvariant.class=com.mapr.job.mngmnt.hadoop.metrics.MaprRPCContext
#maprmapredvariant.period=60
#maprmapred.class=com.mapr.job.mngmnt.hadoop.metrics.MaprRPCContextFinal
#maprmapred.period=60
```

## MapR-FS and CLDB Metrics Example

The following example enables the `MapRGangliaContext31` context for the MapR-FS and the CLDB in the `/opt/mapr/conf/hadoop-metrics.properties` file:

```
#CLDB metrics config - Pick one out of null,file or ganglia.
#Uncomment all properties in null, file or ganglia context, to send cldb metrics to
that context
# Configuration of the "cldb" context for null
#cldb.class=org.apache.hadoop.metrics.spi.NullContextWithUpdateThread
#cldb.period=10
# Configuration of the "cldb" context for file
#cldb.class=org.apache.hadoop.metrics.file.FileContext
#cldb.period=60
#cldb.fileName=/tmp/cldbmetrics.log
# Configuration of the "cldb" context for ganglia
cldb.class=com.mapr.fs.cldb.counters.MapRGangliaContext31
cldb.period=10
cldb.servers=localhost:8649
cldb.spoof=1
#FileServer metrics config - Pick one out of null,file or ganglia.
#Uncomment all properties in null, file or ganglia context, to send fileserver metrics
to that context
# Configuration of the "fileserver" context for null
#fileserver.class=org.apache.hadoop.metrics.spi.NullContextWithUpdateThread
#fileserver.period=10
# Configuration of the "fileserver" context for file
#fileserver.class=org.apache.hadoop.metrics.file.FileContext
#fileserver.period=60
#fileserver.fileName=/tmp/fsmetrics.log
# Configuration of the "fileserver" context for ganglia
fileserver.class=com.mapr.fs.cldb.counters.MapRGangliaContext31
fileserver.period=37
fileserver.servers=localhost:8649
fileserver.spoof=1
maprmapredvariant.class=com.mapr.job.mngmnt.hadoop.metrics.MaprRPCContext
maprmapredvariant.period=10
maprmapred.class=com.mapr.job.mngmnt.hadoop.metrics.MaprRPCContextFinal
maprmapred.period=10
```