

ManagementTools

Hadoop Cluster Management Tools

On a big cluster you don't want to have your phone page you every time a node goes down. The only individual machines you care about are the [NameNode](#), the Secondary [NameNode](#) and the [JobTracker](#). Worker nodes come and go. What matters there is the total cluster availability, the availability of the live data, and whether the rate of node failure is too high to get useful work done.

The other thing to be aware of is that the troublesome workers are not the dead ones; they are easy to detect; they don't report for duty. The troublesome ones are the nodes where the disk is playing up so badly that the system is really slow, so their work takes too long. Or their RAM isn't working properly so only 1GB of it appears there, and every job fails with memory problems. Or some strange motherboard/CPU/OS combination causes a machine to find race conditions in code where none surface elsewhere. That's what you need to identify: the troublemakers. Once found, you can set up Hadoop to blacklist nodes.

For detailed information, see Ed Capriolo's [presentation](#) from Hadoop World NYC 2009.

Nagios

There is support for Nagios in Hadoop.

Ganglia

There is support for Ganglia in Hadoop.

JMX Support

Hadoop has JMX support, so with the right JMX bridge for your chosen management tools, it should be possible to keep an eye on Hadoop from your favorite management console. For more on JMX and Hadoop, see Philip Zeyliger's [blog post](#) on the Cloudera blog.

JMX Bridging to Zenoss

Allen at [LinkedIn](#) says "We've working on getting our stats into Zenoss via the JMX connector and SNMP because Ganglia seems to have some fundamental issues (like grouping of hosts is a *client* side config). Note that Zenoss is available in both open source and commercial forms. We're using the commercial version, but the open source version would probably be just as good.

But that aside:

We're taking the approach of grid health by watching and monitoring the dead/live node count by scraping the NN and JT web pages. We also do daily fsck's, lsr's, and run a cut-down version of gridmix.

While monitoring individual nodes is useful in a pro-active sense, the bigger your grid gets, the less important it becomes"

Cacti

See Ed Capriolo's [blog post](#) about how About.com uses Cacti for monitoring their Hadoop cluster.

Cloudera Desktop

[Cloudera Desktop](#) includes a simple Cluster Health dashboard which displays node liveness information and other key data points about the cluster.