

NameNode

The [NameNode](#) is the centerpiece of an HDFS file system. It keeps the directory tree of all files in the file system, and tracks where across the cluster the file data is kept. It does not store the data of these files itself.

Client applications talk to the [NameNode](#) whenever they wish to locate a file, or when they want to add/copy/move/delete a file. The [NameNode](#) responds the successful requests by returning a list of relevant [DataNode](#) servers where the data lives.

The [NameNode](#) is a [Single Point of Failure](#) for the HDFS Cluster. HDFS is not currently a High Availability system. When the [NameNode](#) goes down, the file system goes offline. There is an optional [SecondaryNameNode](#) that can be hosted on a separate machine. It only creates checkpoints of the namespace by merging the edits file into the fsimage file and does not provide any real redundancy. Hadoop 0.21+ has a [BackupNameNode](#) that is part of a plan to have an HA name service, but it needs active contributions from the people who want it (i.e. you) to make it Highly Available.

It is essential to look after the [NameNode](#). Here are some recommendations from production use

- Use a good server with lots of RAM. The more RAM you have, the bigger the file system, or the smaller the block size.
- Use ECC RAM.
- On Java6u15 or later, run the server VM with compressed pointers `-XX:+UseCompressedOops` to cut the JVM heap size down.
- List more than one name node directory in the configuration, so that multiple copies of the file system meta-data will be stored. As long as the directories are on separate disks, a single disk failure will not corrupt the meta-data.
- Configure the [NameNode](#) to store one set of transaction logs on a separate disk from the image.
- Configure the [NameNode](#) to store another set of transaction logs to a network mounted disk.
- Monitor the disk space available to the [NameNode](#). If free space is getting low, add more storage.
- Do not host [DataNode](#), [JobTracker](#) or [TaskTracker](#) services on the same system.

If a [NameNode](#) does not start up, look at the [TroubleShooting](#) page.