

Quotas Guide

Table of contents

1 Overview.....	2
2 Name Quotas.....	2
3 Space Quotas.....	2
4 Administrative Commands.....	2
5 Reporting Command.....	3

1. Overview

The Hadoop Distributed File System (HDFS) allows the **administrator** to set quotas for the number of names used and the amount of space used for individual directories. Name quotas and space quotas operate independently, but the administration and implementation of the two types of quotas are closely parallel.

2. Name Quotas

The name quota is a hard limit on the number of file and directory names in the tree rooted at that directory. File and directory creations fail if the quota would be exceeded. Quotas stick with renamed directories; the rename operation fails if operation would result in a quota violation. The attempt to set a quota will still succeed even if the directory would be in violation of the new quota. A newly created directory has no associated quota. The largest quota is `Long.MaxValue`. A quota of one forces a directory to remain empty. (Yes, a directory counts against its own quota!)

Quotas are persistent with the `fsimage`. When starting, if the `fsimage` is immediately in violation of a quota (perhaps the `fsimage` was surreptitiously modified), a warning is printed for each of such violations. Setting or removing a quota creates a journal entry.

3. Space Quotas

The space quota is a hard limit on the number of bytes used by files in the tree rooted at that directory. Block allocations fail if the quota would not allow a full block to be written. Each replica of a block counts against the quota. Quotas stick with renamed directories; the rename operation fails if the operation would result in a quota violation. A newly created directory has no associated quota. The largest quota is `Long.MaxValue`. A quota of zero still permits files to be created, but no blocks can be added to the files. Directories don't use host file system space and don't count against the space quota. The host file system space used to save the file meta data is not counted against the quota. Quotas are charged at the intended replication factor for the file; changing the replication factor for a file will credit or debit quotas.

Quotas are persistent with the `fsimage`. When starting, if the `fsimage` is immediately in violation of a quota (perhaps the `fsimage` was surreptitiously modified), a warning is printed for each of such violations. Setting or removing a quota creates a journal entry.

4. Administrative Commands

Quotas are managed by a set of commands available only to the administrator.

- `dfsadmin -setQuota <N> <directory>...<directory>`
Set the name quota to be N for each directory. Best effort for each directory, with faults reported if N is not a positive long integer, the directory does not exist or it is a file, or the directory would immediately exceed the new quota.
- `dfsadmin -clrQuota <directory>...<director>`
Remove any name quota for each directory. Best effort for each directory, with faults reported if the directory does not exist or it is a file. It is not a fault if the directory has no quota.
- `dfsadmin -setSpaceQuota <N> <directory>...<directory>`
Set the space quota to be N bytes for each directory. This is a hard limit on total size of all the files under the directory tree. The space quota takes replication also into account, i.e. one GB of data with replication of 3 consumes 3GB of quota. N can also be specified with a binary prefix for convenience, for e.g. 50g for 50 gigabytes and 2t for 2 terabytes etc. Best effort for each directory, with faults reported if N is neither zero nor a positive integer, the directory does not exist or it is a file, or the directory would immediately exceed the new quota.
- `dfsadmin -clrSpaceQuota <directory>...<director>`
Remove any space quota for each directory. Best effort for each directory, with faults reported if the directory does not exist or it is a file. It is not a fault if the directory has no quota.

5. Reporting Command

An extension to the `count` command of the HDFS shell reports quota values and the current count of names and bytes in use.

- `fs -count -q <directory>...<directory>`
With the `-q` option, also report the name quota value set for each directory, the available name quota remaining, the space quota value set, and the available space quota remaining. If the directory does not have a quota set, the reported values are `none` and `inf`.