

Max number of files

[rodpuwell1](#) 4 posts since

Sep 14, 2008

I have a FAS3020 with 1Tb volume used as storage for an application that generates millions of small files - 30-40k each.

We have been getting an error message saying files are too large which we assume to mean there are too many.

We have increased to maximum number of files allowed as the volume is not full.

My questions:

1. How does ONTAP calculate the max number of files for a given volume size?
2. How many files can be added per volume above the default setting before you start to see an impact on performance?
3. How does having so many files affect performance when trying to do a backup?
4. Is there a best practice for data set containing millions of small files?

Thanks

Rod.

[rkaramchedu1](#) 96 posts since

Mar 9, 2008 1. **Re: Max number of files** Oct 18, 2008 12:07 PM

- Q. How does ONTAP calculate the max number of files for a given volume size?
A. By default, the volume will get 1 inode per every 32KB of disk space and the # of files that a volume can hold get arrived at using that number. The maximum number of inodes is limited to one inode per one block in the filesystem. (which is 1 inode per every 4KB). It is generally recommended to NOT go that low.

Max number of files

- Q. How many files can be added per volume above the default setting before you start to see an impact on performance?

A. I do not know if there has been any study that was done to measure the file system performance w.r.t to the number of files. Also note that the performance of a volume is not only dependent upon the # of files the volume has but also on the structure of the files. For example, you'll get a lot better performance if you distribute your files across multiple directories at the same level, as opposed to having all of them in 1 directory.

- Q. How does having so many files affect performance when trying to do a backup?

A. Very negatively. Your backup software will spend a LOT of time getting the file list and then actually doing the backup. Kick off a backup and examine the time difference between the time when it starts reading the file system and the time when it actually starts reading any bytes from the file system... that will tell you the story.

- Q. Is there a best practice for data set containing millions of small files?

A. I would try to organize the files in such a way that you have files distributed among multiple directories. Also, try not to have too deep of a directory structure.

Some other reading

http://now.netapp.com/NOW/knowledge/docs/ontap/rel73/html/ontap/cmdref/man1/na_maxfiles.1.htm

<https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs11850>

<https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs5094>

HTH

[cebulrdcis](#) 55 posts since

Apr 2, 2008 2. Re: **Max number of files** Oct 19, 2008 9:38 PM

Great answer, this is also helpful for me coz I also have a volume with so many small files on it.. Will read the provided link.

Max number of files

thanks,

tons

[rodpuwell1](#) 4 posts since

Sep 14, 2008 3. **Re: Max number of files** Oct 20, 2008 3:13 AM

Thanks dude. Great Answer, gives me the info I need.

[BrendonHiggins](#) 477 posts since

Mar 26, 2008 4. **Re: Max number of files** Oct 20, 2008 1:49 AM

I agree with everything which has been said and would also like to add that our backup application "HPData Protector" has a limit of 20,000,000 files per backup job. We found this out the 'hard way' once we were no longer able to backup the volume.