

Re: too many files....

## Re: too many files....

---

*Source:* <http://linux.derkeiler.com/Newsgroups/comp.os.linux.misc/2006-08/msg00432.html>

---

- *From:* Unruh <[unruh-spam@xxxxxxxxxxxxxxxx](mailto:unruh-spam@xxxxxxxxxxxxxxxx)>
  - *Date:* 21 Jul 2006 23:12:49 GMT
- 

Dances With Crows <[danSPANceswitTRAPhcrows@xxxxxxxx](mailto:danSPANceswitTRAPhcrows@xxxxxxxx)> writes:

On Fri, 21 Jul 2006 15:48:05 -0600, Douglas Mayne staggered into the Black Sun and said:

On Fri, 21 Jul 2006 14:28:05 -0700, CptDondo wrote:

I have a website with lots of pictures. At the moment, some 5,000+ and growing at the rate of about 300-500/month. The problem is that access to those files is getting slower and slower as the number of pictures gets larger.

Dividing the files into groups and directories is not an option; they are used in a website and all of the picture URLs are stored in an SQL database. Dividing them now would be a huge nightmare.....

You screwed up. Fix it now, or fix it later; the pain will be less if you fix it now. It's *\*never\** a good idea to try to pack every file into 1 directory if the number of files is going to grow and keep growing.

So... The pictures are stored on an ext3 filesystem. They range from small (10K) to moderate (1MB) in size. Is there some way to optimize this for faster access? Can I use a different filesystem?

Re: too many files....

Directory lookup times are faster for ReiserFS and XFS than they are for ext3, for various reasons. If you think you have to keep every file in 1 directory, and you don't want to fix your broken database/website, back up everything on this partition and mkreiserfs (or mkxfs) it. If using ReiserFS, make sure to use "-h tea" if you'll ever have more than about 200,000 files in 1 directory.

I simply do not believe that directory lookup is a significant time on any filesystem, even if he has 100,000 files ( which I do not believe. At 100K for size, that would be over 10GB. )

The hardware is a dual PIII 1Ghz HP server, 1 GB RAM, 5 SCSI 36 GB drives in a softRAID5 array

The hardware seems fine to me.

SoftRAID 5 requires more CPU resources than softRAID-1 for obvious reasons. That shouldn't be causing the problem that Dondo's reporting. One of the work machines has a hardRAID-5 of 72G 7200 RPM SCSI disks, and on 1 230G ext3 partition, it can open and read 1 100K file in a directory containing 20,000 files in about 0.1 seconds. (This machine is just a fileserver, though, not a web or database server.)

[Does] part of the problem of slow access involve some SQL lookup?

I would guess this is part of the problem. Make sure you're not doing something dumb in the queries being used and the server-side whatever (PHP? Ruby? CGI-Perl? \*yuck\* Java?) you're using.

Sounds much more likely— that he is doing something non-optimally.

Re: too many files....