

Re: Question about file system checks

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2008-03/msg00708.html>

- *From:* Alfred Perlstein <alfred@xxxxxxxxxxxx>
 - *Date:* Fri, 28 Mar 2008 15:01:55 -0700
-

* Ivan Voras <ivoras@xxxxxxxxxxxx> [080328 14:51] wrote:

Matthew Seaman wrote:

Ivan Voras wrote:

1. UFS+gjournal loses the least, but it's also the slowest.
2. UFS+SU had no truncated files or files of unexpected length
(apparently it just loses the file that would end up in this state)
3. XFS and JFS end up with a **huge** number of files that are truncated or of unexpected length (40%–50%!).
4. In no case has any of the above file systems gone completely corrupted or lost any of the files/directories not being updated.
5. ZFS on FreeBSD was the fastest, in the sense of creating the most files during this benchmark (though speed was not the target for this benchmark so this is a low-quality observation), closely followed by JFS and XFS.
6. ZFS crashed the kernel at least once.

Hmmm... in many ways a corrupt or truncated file is a worse outcome than a completely missing file — at least if the file has gone away you know you've got to do something to fix it. A damaged file could end up silently causing weird behavioural effects and (by the law of natural cussedness) it is almost bound not to be tracked down until the day after the last good copy on the backup tapes gets overwritten...

How do the different filesystems compare if you total all lost, damaged

Re: Question about file system checks

or truncated files?

The only things that happen are that XFS and JFS get disproportionately bad numbers and that ext3 gets almost identically bad results with UFS+SU. Overall ratios remain approximately the same.

To put this into perspective, for total "bad" files this means that, e.g. UFS+SU created 20000 files, of which 750 were in some way "bad", and ZFS created 46000 files, of which 900 were bad (so percentage is in favour of ZFS). JFS created 43000 files of which 20000 were of wrong size, but only 45 were completely lost. How bad this is depends, of course on what is done with the file system.

A big surprise for me was that Windows' NTFS did very good, though it was the slowest in most other tests (which are smarter and probably use fsync a lot), it managed to create 32000 files and have only 121 "bad" in some way.

I know this sounds pretty awful, but honestly any file modified within an hour and not fsync'd being "lost" is not really a bad thing.

It's pretty much "the unix way" that only fsync'd files/directories or file modified more than several minutes ago are safe.

--

- Alfred Perlstein

freebsd-stable@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-stable>

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@xxxxxxxxxxx"