

Re: Sidetracked: why gjournal over soft-updates (Was: Re: UFS2 Journaling implementation detail)

Re: Sidetracked: why gjournal over soft-updates (Was: Re: UFS2 Journaling implementation detail)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/questions/2008-04/msg01241.html>

- *From:* "Ivan Voras" <ivoras@xxxxxxxxxxxxx>
 - *Date:* Sat, 19 Apr 2008 22:02:02 +0200
-

2008/4/19 Mel <fbsd.questions@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx>:

On Friday 18 April 2008 17:40:04 Ivan Voras wrote:

>> 5. "Some UFS implementations avoid journaling and
>> instead implement soft updates: they order their
>> writes in such a way that the on-disk file system is
>> never inconsistent, or that the only inconsistency
>> that can be created in the event of a crash is a
>> storage leak. To recover from these leaks, the free
>> space map is reconciled against a full walk of the
>> file system at next mount." –
>> (http://en.wikipedia.org/wiki/Journaling_file_system)
>>
>> So the disadvantage of Soft Update is it is necessary
>> to run fsck after reboot in event of a crash or power
>> failure?
>
> Yes. The advantage is that practically, the data is as safe as with
> journalling.

I've been following this with interest, however it's still not clear to me,
why I'd want a journaling filesystem, because:

- 1) If you have soft-updates the data is as safe as with journal
- 2) If you have soft-updates fsck will run in the background
- 3) Soft-updates don't require disk space.

So...other than "journaling filesystems are cool", what's the real advantage?

gjournal is a kind of hack and it's only usable for avoiding fscks.
It's actually slower than softupdates for high-demand disk loads.
"Real" journalling file systems would have the following benefits, as
replies to the points from your post:

- 1) Soft-updates were created in a different time, with different
requirements than modern hard drives (especially desktop hard drives)
can deliver. Especially, SU requires that data it once sends to the

Re: Sidetracked: why gjournal over soft-updates (Was: Re: UFS2 Journaling implementation detail)

Re: Sidetracked: why gjournal over soft-updates (Was: Re: UFS2 Journaling implementation detail)

drive gets written immediately, not cached by the drive. Modern desktop drives don't do that so journaling is today actually safer than SU (because it requires smaller, compact data to be written to the journal, instead of scattered bits of data across the disk drive like SU does). There are reasons why there's only one single soft-updates file system in existence today while there are dozens of journaled file systems. Soft-updates is infinitely better than nothing, but at the very least it needs to be seriously updated to be useful for professional purposes today.

2) Background fsck is considered unsafe by many high-end users of FreeBSD, especially with large drives (many terabytes). One of the reasons for that is because it relies on file system snapshots which have in the past behaved quirky. Another reason to avoid background fsck-s is that some applications require all their disk IO bandwidth immediately on boot, and background fsck may "steal" the bandwidth while it's running. On the other hand, for desktop-sized drives and low-demanding purposes, background fsck works ok.

3) "Real" journalling file systems have journals sized in the order of dozens of megabytes (e.g. 10 MB - 50 MB) because they don't need to journal every single IO operation like gjournal does. You'll agree that those sizes are trivial for today's drives.

For details, you can search the mailing list archives - what I've said is nothing original and the topic has been talked about a lot.

freebsd-questions@xxxxxxxxxxx mailing list

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

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

Re: Sidetracked: why gjournal over soft-updates (Was: Re: UFS2 Journaling implementation detail)