

## Re: Best practices for using gjournal with gmirror?

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/hackers/2007-01/msg00220.html>

---

- *From:* Pawel Jakub Dawidek <[pjd@xxxxxxxxxxxx](mailto:pjd@xxxxxxxxxxxx)>
  - *Date:* Fri, 12 Jan 2007 19:57:22 +0100
- 

On Wed, Jan 10, 2007 at 11:21:01PM -0500, John Nielsen wrote:

I have a few questions for pjd (or anyone else) about using gjournal, particularly when used with gmirror.

1) I'm running 6-STABLE and plan to test with gjournal6\_20061030.patch (from the mailing list; updated version of 20061024 that applies cleanly). Is there a better/newer version for -STABLE that I should use instead?

There probably should be a newer version as there were some minor changes after I committed the code to HEAD. I'll try to create a new patch during the weekend.

2) When using gjournal and for a gmirror volume, does the journal need to be mirrored as well to maintain redundancy? If so, when storing the journal on the same physical disks as the mirror, is it better to mirror at the slice level (journal and fs on different partitions in the same mirror) or at the partition level (journal and fs each have their own mirror) or does it matter?

The problem with mirroring each partition/slice separately is that when you have a crash, on boot, gmirror will start to rebuild all partitions at once, which may be problematic. On the other hand, when you mirror each partition/slice separately, and some partitions weren't modified in last few seconds before the crash, gmirror will not resync them on boot, so not entire disk will be synchronized.

When you run gjournal on top of gmirror/raid3 there is no need for resync after a crash, so basically all cons against mirroring the whole disks and against mirroring partitions are no longer true. Both configurations will work the same. In that case I'd suggest mirroring the whole disks, because when one of your disks dies, you may just replace it and be down with it. If you mirror partitions separately, you first have to create partitions and insert each of them into their mirrors, which is more complex than simple 'gmirror insert foo newdisk'.

Re: Best practices for using gjournal with gmirror?

3) I remember reading where pjd said that gjournal plus gmirror or graid3 would eliminate the need to re-sync the array after a crash. While clearly a design goal, is that actually the case with the version of the patch mentioned above? If so, are any config changes needed or will it just happen automatically?

No, you need to:

```
# gmirror configure -F <mirror_name>
```

4) In the same vein as 3)--does a gjournal volume need to be fsck'ed after a crash? If not, will it just work (e.g. fsck -p sees that the filesystem is clean) or does it need to be disabled somehow?

Gjournalled file system has to be fscked, but only to handle orphaned files. Such fsck on multiterabyte provider takes seconds, not hours.

5) Finally, how dangerous is this code? I realize it's experimental and only plan to use it with data that has recent backups, but how much should I worry about it blowing up my system or corrupting my files?

I'm using it in production, my customer using it in production on large number of FreeBSD servers and I also have heard already many success stories, BUT I still consider the code to be experimental.

---

Pawel Jakub Dawidek <http://www.wheel.pl>  
pjd@xxxxxxxxxxx <http://www.FreeBSD.org>  
FreeBSD committer Am I Evil? Yes, I Am!

**Attachment:** [pgpT0LaTLWVhE.pgp](#)

*Description:* PGP signature