

## Re: Preventing ext3 fsck at boot?

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/RedHat/2007-09/msg00245.html>

---

- *From:* "Sandor W. Sklar" <[ssklar@xxxxxxxxxxxxx](mailto:ssklar@xxxxxxxxxxxxx)>
  - *Date:* Sat, 29 Sep 2007 10:03:19 -0700
- 

On Sep 28, 2007, at 10:34 PM, Mike Kearey wrote:

Sandor W. Sklar wrote:

I've got a number of large EXT3 filesystems (2-8 TB each), presented via dual-path Fibre HBAs via SAN switches from several Nexsan SATAbeast arrays, to a number of systems running RHEL4.

Now, you're just showing off :) Imagine 2 -8TB's 20 years ago:

<http://sd4.sd-lj.si/digg/20yago.jpg> 1GB 20 years ago compared to a flash available now.

Oh, I wish I was just showing off! I never thought I'd miss the days where all of my storage was presented in 4 GB luns from a "top-of-the-line" Symmetrix, connected with lots of really long SCSI cables under the floor. :-) Our last purchase of Nexsan SATAbeasts included drives that were 1 TB in size; I shudder to think of the rebuild time when one drive in the RAID set fails.

The question of whether EXT3 is the right filesystem to be using for this is probably best saved for another email (but I'd love to hear about better options; I'm relatively new to Linux, compared to AIX and Solaris.)

ext3 is best used on a RHEL4 system because it's what we develop, test and support. That is a very important consideration. Note that this does not mean it's the best one on a technical and theoretical or performance standpoint.

## Re: Preventing ext3 fsck at boot?

That is an interesting point, and one that I didn't consider. All of our RHEL systems are built from a local Satellite Server, but we have bought a few "retail" licenses, for the purposes of support. So, can I take it that you're stating that if we were to have a problem with an XFS, or Reiser filesystem, and opened a support case with it, we might experience some issues? That is an important point, so thanks ... that does help inform our decision.

My main problem is that when we reboot these servers for scheduled maintenance (or for any reason), odds are pretty good that I'm going to get the (dreaded) ...

`/dev/nsvg/lvol0` has gone 182 days without being checked, check forced.

... message, and then my downtime is extended by 2–3 hours while the system does its fsck (and usually finds 0 problems.)

So, my questions are:

– The man page for `tune2fs` says that this can be disabled with the `"-c"` option, but recommends strongly against it. Is it really such a bad thing to disable, if I'm using EXT3 (with the journaling that makes it "3" instead of "2")? I've used JFS/JFS2 for years on AIX, and UFS journaling on Solaris, and neither seems to want to force an fsck just because some arbitrary time period has past since it last checked.

ext2 is Linux extended filesystem 2. ext3 is Linux extended filesystem 3

The major difference between the two is the journal capability in ext3. ext3 filesystem can be mounted as ext2 BTW, backwards compatibility is good.

Sure, its good for many people in many cases; irrelevant for us, though. I need, in order of importance, (a) reliability, (b) performance, and (c) ease of administrative tasks.

– If the consensus is that it would be ok to disable these checks, what is the proper syntax? I tried:

```
# tune2fs -c0 /dev/mapper/nsvg-lvol0
tune2fs 1.35 (28-Feb-2004)
Setting maximal mount count to -1
```

... but that didn't work. Looking for some practical advice and recommendations, here, please!

Re: Preventing ext3 fsck at boot?

We have the kbase article [http://kbase.redhat.com/faq/FAQ\\_80\\_5779.shtm](http://kbase.redhat.com/faq/FAQ_80_5779.shtm)  
It says that the filesystem needs to be unmount, but I am not completely sure that is required.

Ah, thanks for that URL. The man page wasn't clear on whether or not the filesystem needed to be unmounted (at least, to my bleary eyes, it wasn't.)

Anyway, verify it is set :

```
# dumpe2fs /dev/mapper/myvg-rootvol |grep Max
dumpe2fs 1.39 (29-May-2006)
Maximum mount count: -1
```

That should be enough. I believe RHEL5 comes delivered with Maximum mount count set this way for the root filesystem.

In my opinion when you have more 800GB for a filesystem you are well and truly at a point where an fsck is a waste of time compared to a clean mkfs and restore from backup. So take the dire warnings from the tune2fs manual as something that was valid and relevant 5 years ago or more and not quite applicable to your situation :)

Thanks very much for your help. Generally, I agree with your opinion (but I'd prefer to never have to either fsck or restore filesystems this big!)

—s—

—

Sandor W. Sklar  
Unix Systems Administrator  
Stanford University Libraries & Academic Information Resources (SULAIR)  
Digital Libraries Systems & Services (DLSS)

—

redhat-list mailing list  
unsubscribe <mailto:redhat-list-request@xxxxxxxxxx?subject=unsubscribe>  
<https://www.redhat.com/mailman/listinfo/redhat-list>