## Subject: Re: OpenVZ 7 containers crashing with ext4 errors Posted by nathan.brownrice on Mon, 24 Aug 2020 16:48:39 GMT

View Forum Message <> Reply to Message

Hello All, we're having this same issue as well.

The issue is that overnight, a VPSs filesystem will go read-only. We've seen this on no fewer than 5-10 different VPSs since switching to ovz7 in the last year. In some cases the filesystem can be repaired using the normal recovery methods (I.E.

https://virtuozzosupport.force.com/s/article/000014682 and

https://inertz.org/container-corruption-easy-repair-using-fsck/), but sometimes things are irrecoverable and we have to restore the entire container from backup. This is a pretty big deal.

The first thing we noticed was, in /var/log/messages on the host machine, this is happening right when the VPS goes read-only (notice the similar timestamps to the OP's issue):

```
Aug 21 02:00:04 ovz7-3-taos pcompact[29446]: {"operation":"pcompactStart", "uuid":"fa79f45c-5f32-4b6c-8ce5-9d4a012e43c8", "disk_id":0, "task_id":"290cde3e-e52f-4e29-9f51-a9db197887eb", "ploop_size":98078, "image_size":93758, "data_size":34888, "balloon_size":280802, "rate":60.0, "config_dry":0, "config_threhshold":10}

Aug 21 02:00:11 ovz7-3-taos pcompact[29446]: {"operation":"pcompactFinish", "uuid":"fa79f45c-5f32-4b6c-8ce5-9d4a012e43c8", "disk_id":0, "task_id":290cde3e-e52f-4e29-9f51-a9db197887eb", "was_compacted":1, "ploop_size":98078, "stats_before": {"image_size":93758, "data_size":34888, "balloon_size":280802}, "stats_after": {"image_size":93758, "data_size":34888, "balloon_size":280802}, "time_spent":"7.016s", "result":-1}
```

The next thing we noticed, after seeing the above error, is that the pcompact.log has the following:

```
2020-08-21T02:00:04-0600 pcompact: Inspect fa79f45c-5f32-4b6c-8ce5-9d4a012e43c8
2020-08-21T02:00:04-0600 pcompact : Inspect
/vz/private/fa79f45c-5f32-4b6c-8ce5-9d4a012e43c8/root.hdd/DiskDescriptor.xml
2020-08-21T02:00:04-0600 pcompact : ploop=98078MB image=93758MB data=34888MB
balloon=280802MB
2020-08-21T02:00:04-0600 pcompact : Rate: 60.0 (threshold=10)
2020-08-21T02:00:04-0600 pcompact : Start compacting (to free 53965MB)
2020-08-21T02:00:04-0600 : Start defrag dev=/dev/ploop43779p1
mnt=/vz/root/fa79f45c-5f32-4b6c-8ce5-9d4a012e43c8 blocksize=2048
2020-08-21T02:00:11-0600 : Error in wait pid (balloon.c:962): The /usr/sbin/e4defrag2 process
failed with code 1
2020-08-21T02:00:11-0600 : /usr/sbin/e4defrag2 exited with error
2020-08-21T02:00:11-0600: Trying to find free extents bigger than 0 bytes granularity=1048576
2020-08-21T02:00:11-0600 : Error in ploop_trim (balloon.c:892): Can't trim file system:
Input/output error
2020-08-21T02:00:11-0600 pcompact : ploop=98078MB image=93758MB data=34888MB
balloon=280802MB
2020-08-21T02:00:11-0600 pcompact : Stats: uuid=fa79f45c-5f32-4b6c-8ce5-9d4a012e43c8
```

ploop\_size=98078MB image\_size\_before=93758MB image\_size\_after=93758MB compaction\_time=7.016s type=online 2020-08-21T02:00:11-0600 pcompact : End compacting

This is basically identical to what's being discussed in this thread. We've just spun up a new host machine with a fresh OS install, and it looks like the newest ISO still has the old kernel version (vz7.151.14), so we've applied the patch as discussed here.

What I'd like to discuss:

- 1) Others that have had this same issue, and have applied the kernel update, did this fix your issues?
- 2) We have several other production host machines, and it's going to take a lot of moving things around before we can safely kernel update them. We're working on this, but in the meantime is there a way to ensure this doesn't happen?

It looks like the initial error is happening during pcompact defrag, which we see can be disabled as per https:// docs.openvz.org/openvz\_command\_line\_reference.webhelp/\_pcomp act\_conf.html . Perhaps this would prevent the issue from happening if this were to be temporarily disabled until we can get the kernels updated. Or, perhaps we could disable pcompact altogether. Any thoughts or suggestions on this?

Thanks for the wonderful software and the great community behind it!