
Subject: OpenVZ 7 containers crashing with ext4 errors
Posted by [allan.talver](#) on Fri, 03 Jul 2020 15:06:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

We have been running OpenVZ 6 in our environments for several years. The platform has been stable and predictable. Recently we have started evaluating OpenVZ 7 as the replacement. In most aspects, OpenVZ 7 has proven to be good and suitable for our purpose. However, recently we began to experience seemingly random crashes with symptoms pointing to ext4 filesystem and ploop. When the crash happens, virtual container is left with its disk in read only state. Restart is not successful due to errors present in the filesystem. After running fsck manually, container is able to start and we have not experienced any data loss. However, even without data loss, such events reduce the confidence to run production workloads with critical data on these servers.

In total we have now had 4 such events. First 3 were Ubuntu 16.04 containers that got migrated from OpenVZ 6 to OpenVZ 7. Before migration with `ovztransfer.sh`, the server disks got converted from `simfs` to `ploop` with `vzctl convert`. Initially we thought that the issue might be something in our migration procedure or something specific to Ubuntu 16.04 operating system, because no Ubuntu 18.04 server (created fresh on OpenVZ 7) had crashed. However, two days ago the container affected by the last crash was 18.04 which never was migrated from OpenVZ 6 (although it has been migrated between OpenVz 7 nodes).

Another thing we have noticed is that the crash seems to be happening roughly at the same time when `pcompact` is running on the hardware node. And also, 2 out of 3 containers that had been migrated from OpenVZ 6, crashed the next night right after the migration.

Are these errors something that the community has seen before and could help us explain?

In all cases, the log output in hardware node `dmesg` has been similar and as follows:

```
[2020-05-29 02:02:20] WARNING: CPU: 12 PID: 317821 at fs/ext4/ext4_jbd2.c:266
__ext4_handle_dirty_metadata+0x1c2/0x220 [ext4]
[2020-05-29 02:02:20] Modules linked in: nfsv3 nfs_acl ip6table_mangle nf_log_ipv4
nf_log_common xt_LOG nfsv4 dns_resolver nfs lockd grace fscache xt_multiport
nf_conntrack_ipv4 nf_defrag_ipv4 xt_conntrack xt_comment binfmt_misc xt_CHECKSUM
iptable_mangle ip6t_REJECT nf_reject_ipv6 ipt_REJECT nf_reject_ipv4 tun 8021q garp mrp
devlink ip6table_filter ip6_tables iptable_filter bonding ebtfilter ebt_arp ebtables sunrpc
iTCO_wdt iTCO_vendor_support sb_edac intel_powerclamp coretemp intel_rapl iosf_mbi
kvm_intel kvm irqbypass crc32_pclmul ghash_clmulni_intel aesni_intel lrw gf128mul glue_helper
ablk_helper cryptd pcspkr i2c_i801 joydev mei_me lpc_ich mei sg ioatdma wmi ipmi_si
ipmi_devintf ipmi_msghandler acpi_power_meter acpi_pad pcc_cpufreq ip_vs nf_conntrack
libcrc32c br_netfilter veth overlay ip6_vzprivnet
[2020-05-29 02:02:20] ip6_vznetstat ip_vznetstat ip_vzprivnet vziolimit vzevent vzlist vzstat
vznetstat vznetdev vzmon vzdev bridge stp llc ip_tables ext4 mbcache jbd2 sd_mod crc_t10dif
crct10dif_generic ast i2c_algo_bit drm_kms_helper syscopyarea sysfillrect sysimgblt fb_sys_fops
ttm ixgbe ahci drm libahci libata megaraid_sas crct10dif_pclmul crct10dif_common crc32c_intel
mdio ptp pps_core drm_panel_orientation_quirks dca dm_mirror dm_region_hash dm_log
dm_mod pio_kaio pio_nfs pio_direct pfmt_raw pfmt_ploop1 ploop
[2020-05-29 02:02:20] CPU: 12 PID: 317821 Comm: e4defrag2 ve: 0 Kdump: loaded Not tainted
```

3.10.0-1062.12.1.vz7.131.10 #1 131.10

[2020-05-29 02:02:20] Hardware name: Supermicro SYS-1028R-WC1RT/X10DRW-iT, BIOS 2.0a 07/26/2016

[2020-05-29 02:02:20] Call Trace:

```
[2020-05-29 02:02:20] [<ffffff81baebc7>] dump_stack+0x19/0x1b
[2020-05-29 02:02:20] [<ffffff8149bdc8>] __warn+0xd8/0x100
[2020-05-29 02:02:20] [<ffffff8149bf0d>] warn_slowpath_null+0x1d/0x20
[2020-05-29 02:02:20] [<ffffffc047cd82>] __ext4_handle_dirty_metadata+0x1c2/0x220 [ext4]
[2020-05-29 02:02:20] [<ffffffc0474f44>] ext4_ext_split+0x304/0x9a0 [ext4]
[2020-05-29 02:02:20] [<ffffffc0476dfd>] ext4_ext_insert_extent+0x7bd/0x8d0 [ext4]
[2020-05-29 02:02:20] [<ffffffc0479a7f>] ext4_ext_map_blocks+0x5cf/0xf60 [ext4]
[2020-05-29 02:02:20] [<ffffffc0446676>] ext4_map_blocks+0x136/0x6b0 [ext4]
[2020-05-29 02:02:20] [<ffffffc047496c>] ? ext4_alloc_file_blocks.isra.36+0xbc/0x2f0 [ext4]
[2020-05-29 02:02:20] [<ffffffc047498f>] ext4_alloc_file_blocks.isra.36+0xdf/0x2f0 [ext4]
[2020-05-29 02:02:20] [<ffffffc047b6bd>] ext4_fallocate+0x15d/0x990 [ext4]
[2020-05-29 02:02:20] [<ffffff8166cb78>] ? __sb_start_write+0x58/0x120
[2020-05-29 02:02:20] [<ffffff81666c72>] vfs_fallocate+0x142/0x1e0
[2020-05-29 02:02:20] [<ffffff81667cdb>] SyS_fallocate+0x5b/0xa0
[2020-05-29 02:02:20] [<ffffff81bc1fde>] system_call_fastpath+0x25/0x2a
[2020-05-29 02:02:20] ---[ end trace cf8fe0ecbf57efcc ]---
[2020-05-29 02:02:20] EXT4-fs: ext4_ext_split:1139: aborting transaction: error 28 in
__ext4_handle_dirty_metadata
[2020-05-29 02:02:20] EXT4-fs error (device ploop52327p1): ext4_ext_split:1139: inode
#325519: block 6002577: comm e4defrag2: journal_dirty_metadata failed: handle type 3 started at
line 4741, credits 8/0, errcode -28
[2020-05-29 02:02:20] Aborting journal on device ploop52327p1-8.
[2020-05-29 02:02:20] EXT4-fs (ploop52327p1): Remounting filesystem read-only
[2020-05-29 02:02:20] EXT4-fs error (device ploop52327p1) in ext4_free_blocks:4915: Journal
has aborted
[2020-05-29 02:02:20] EXT4-fs error (device ploop52327p1) in ext4_free_blocks:4915: Journal
has aborted
[2020-05-29 02:02:20] EXT4-fs error (device ploop52327p1) in ext4_reserve_inode_write:5360:
Journal has aborted
[2020-05-29 02:02:20] EXT4-fs error (device ploop52327p1) in ext4_alloc_file_blocks:4753: error
28
```

Thanks! :p

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [khorenko](#) on Fri, 03 Jul 2020 16:07:31 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

1) update, vz7 update 14 has quite a number of ploop-related fixes (kernel vz7.151.14).

2) if you face this issue again on vz7 u14, please file a bug at bugs.openvz.org
(which full logs, not just a snippet)

3)

Quote:

```
EXT4-fs: ext4_ext_split:1139: aborting transaction: error 28 in __ext4_handle_dirty_metadata  
"error 28" means space shortage
```

Quote:

```
#define ENOSPC 28 /* No space left on device */  
So please check it ploop usage is close to it's size.
```

4) i don't know if you have messages like

Quote:

```
[2040002.704309] Purging lru entry from extent tree for inode 356516155 (map_size=50243  
ratio=12612%)
```

```
[2040002.704318] max_extent_map_pages=16384 is too low for  
ploop_io_images_size=7914092756992 bytes
```

if you do, increase the following parameter, say by 10 times
`/sys/module/pio_direct/parameters/max_extent_map_pages`

and check if it helps.

The increase can be done on the fly, no reboot required.

But first of all - update!

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Wed, 22 Jul 2020 11:45:35 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello

Thank you for the reply! According to your suggestion we updated all our existing OpenVZ 7 nodes to the 14 version over the past two weeks. We have not yet experienced any crashes on updated nodes (one crash did happen on 11 July on a node that had not yet been updated).

As it was pointed out, it seems that disk space could be a contributing factor to the issue. But I can assure that disks of these failing containers are definitely not full. Some have very low utilisation (around 30%). We have noticed another behaviour which we see is possibly related to the crashes we have experienced, and also points to issues with not enough disk space available. While pcompact is running, some virtual containers show extreme changes in disk utilisation. Usually the disk suddenly shows as full and goes down to normal several times during the time pcompact runs. One example of pcompact.log output while one of such vps' is being compacted:

```
2020-07-22T02:00:12+0200 pcompact : Inspect 7a81d5ef-9a70-4a20-bb57-cf38f45b2926
2020-07-22T02:00:12+0200 pcompact : Inspect /vz/private/4001/root.hdd/DiskDescriptor.xml
2020-07-22T02:00:12+0200 pcompact : ploop=107520MB image=39805MB data=20760MB
balloon=0MB
2020-07-22T02:00:12+0200 pcompact : Rate: 17.7 (threshold=10)
2020-07-22T02:00:12+0200 pcompact : Start compacting (to free 13669MB)
2020-07-22T02:00:12+0200 : Start defrag dev=/dev/ploop12981p1 mnt=/vz/root/4001
blocksize=2048
2020-07-22T02:09:48+0200 : Trying to find free extents bigger than 0 bytes granularity=1048576
2020-07-22T02:09:49+0200 pcompact : ploop=107520MB image=29687MB data=20767MB
balloon=0MB
2020-07-22T02:09:49+0200 pcompact : Stats: uuid=7a81d5ef-9a70-4a20-bb57-cf38f45b2926
ploop_size=107520MB image_size_before=39805MB image_size_after=29687MB
compaction_time=577.227s type=online
2020-07-22T02:09:49+0200 pcompact : End compacting
```

And then we have one container where disk usage stays near 100% (but fluctuating) for 2 hours until pcompact times out (I tried attaching a screenshot, but got an error that "Attachment is too big", even if the file was quite small.). Ploop.log shows:

```
2020-07-22T02:00:01+0200 pcompact : Inspect 240e4613-e12c-46b1-bc06-d001b12463c8
2020-07-22T02:00:01+0200 pcompact : Inspect /vz/private/4116/root.hdd/DiskDescriptor.xml
2020-07-22T02:00:01+0200 pcompact : ploop=261120MB image=116695MB data=87568MB
balloon=0MB
2020-07-22T02:00:01+0200 pcompact : Rate: 11.2 (threshold=10)
2020-07-22T02:00:01+0200 pcompact : Start compacting (to free 16070MB)
2020-07-22T02:00:01+0200 : Start defrag dev=/dev/ploop48627p1 mnt=/vz/root/4116
blocksize=2048
2020-07-22T04:00:21+0200 : Error in wait_pid (balloon.c:967): The /usr/sbin/e4defrag2 process
killed by signal 15
2020-07-22T04:00:21+0200 : /usr/sbin/e4defrag2 exited with error
2020-07-22T04:00:21+0200 : Trying to find free extents bigger than 0 bytes granularity=1048576
2020-07-22T04:00:23+0200 pcompact : ploop=261120MB image=100782MB data=87487MB
balloon=0MB
2020-07-22T04:00:23+0200 pcompact : Stats: uuid=240e4613-e12c-46b1-bc06-d001b12463c8
ploop_size=261120MB image_size_before=116695MB image_size_after=100782MB
compaction_time=7221.741s type=online
2020-07-22T04:00:23+0200 pcompact : End compacting
```

This node is running a MySQL server which shows errors during these 2 hours (different errors pointing to disk being full). Eventually MySQL crashes.

We'll continue to monitor and report back how it goes. But has anyone experienced such fluctuations in disk utilisation while pcompact is running? Is it somehow expected? How to get around applications failing due to disk showing as full (even when it is actually not). Worth mentioning that all these issues described in this post happen on newly created Ubuntu 18.04

containers (as opposed to my initial post where issues were mostly related to 16.04 containers migrated from OpenVZ 6).

Thanks! :)

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Thu, 23 Jul 2020 08:06:59 GMT

[View Forum Message](#) <> [Reply to Message](#)

And today we had another crash. Hardware node with latest OpenVZ 7 version. Ubuntu 18.04 container.

pcompact.log

```
2020-07-22T02:00:02+0200 : Start defrag dev=/dev/ploop43013p1 mnt=/vz/root/4064
blocksize=2048
2020-07-22T02:01:07+0200 : Trying to find free extents bigger than 0 bytes granularity=1048576
2020-07-22T02:01:10+0200 pcompact : ploop=35840MB image=11098MB data=10739MB
balloon=0MB
2020-07-22T02:01:10+0200 pcompact : Stats: uuid=b78116fc-c179-482b-9ab5-c813d3a895a1
ploop_size=35840MB image_size_before=21123MB image_size_after=11098MB
compaction_time=68.751s type=online
2020-07-22T02:01:10+0200 pcompact : End compacting
```

dmesg

```
[Wed Jul 22 02:02:06 2020] -----[ cut here ]-----
[Wed Jul 22 02:02:06 2020] WARNING: CPU: 25 PID: 567333 at fs/ext4/ext4_jbd2.c:266
__ext4_handle_dirty_metadata+0x1c2/0x220 [ext4]
[Wed Jul 22 02:02:06 2020] Modules linked in: binfmt_misc nf_contrack_ipv4 nf_defrag_ipv4
xt_comment xt_multiport xt_contrack nfsv4 dns_resolver nfsd nfs_auth_rpcgss nfs_acl lockd
grace fscache xt_CHECKSUM iptable_mangle ip6t_REJECT nf_reject_ipv6 ipt_REJECT
nf_reject_ipv4 tun 8021q garp mrp devlink ip6table_filter ip6_tables iptable_filter bonding
ehtable_filter ebt_among ebtables sunrpc iTCO_wdt iTCO_vendor_support sb_edac
intel_powerclamp coretemp intel_rapl iosf_mbi kvm_intel kvm irqbypass crc32_pclmul
ghash_clmulni_intel aesni_intel lrw gf128mul glue_helper ablk_helper cryptd pcspkr ipc_ich
i2c_i801 joydev mei_me sg mei ioatdma wmi ipmi_si ipmi_devintf ipmi_msghandler acpi_pad
acpi_power_meter ip_vs nf_contrack libcrc32c br_netfilter veth overlay ip6_vzprivnet
ip6_vznetstat ip_vznetstat ip_vzprivnet vziolimit
[Wed Jul 22 02:02:06 2020] vzevent vzlist vzstat vznetstat vznetdev vzmon vzdev bridge stp llc
ip_tables ext4 mbcache jbd2 sd_mod crc_t10dif crct10dif_generic ast drm_kms_helper
syscopyarea sysfillrect sysimgblt fb_sys_fops ttm drm ahci libahci igb libata crct10dif_pclmul
crct10dif_common crc32c_intel megaraid_sas ptp pps_core dca i2c_algo_bit
drm_panel_orientation_quirks nfit libnvdimm dm_mirror dm_region_hash dm_log dm_mod
pio_kaio pio_nfs pio_direct pfmt_raw pfmt_ploop1 ploop
[Wed Jul 22 02:02:06 2020] CPU: 25 PID: 567333 Comm: e4defrag2 ve: 0 Kdump: loaded Not
tainted 3.10.0-1127.8.2.vz7.151.14 #1 151.14
```

[Wed Jul 22 02:02:06 2020] Hardware name: Supermicro Super Server/X10DRL-i, BIOS 3.1
06/07/2018
[Wed Jul 22 02:02:06 2020] Call Trace:
[Wed Jul 22 02:02:06 2020] [<ffffffff881b67f1>] dump_stack+0x19/0x1b
[Wed Jul 22 02:02:06 2020] [<ffffffff87a9d168>] __warn+0xd8/0x100
[Wed Jul 22 02:02:06 2020] [<ffffffff87a9d2ad>] warn_slowpath_null+0x1d/0x20
[Wed Jul 22 02:02:06 2020] [<ffffffffffc042a382>] __ext4_handle_dirty_metadata+0x1c2/0x220
[ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc04224b4>] ext4_ext_split+0x304/0x9a0 [ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc042436d>] ext4_ext_insert_extent+0x7bd/0x8d0 [ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc042700f>] ext4_ext_map_blocks+0x5cf/0xfc0 [ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc03f3ae6>] ext4_map_blocks+0x136/0x6b0 [ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc0421ecc>] ? ext4_alloc_file_blocks.isra.37+0xbc/0x300
[ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc0421eef>] ext4_alloc_file_blocks.isra.37+0xdf/0x300 [ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffffc0428c9d>] ext4_fallocate+0x15d/0x9b0 [ext4]
[Wed Jul 22 02:02:06 2020] [<ffffffffff87c72428>] ? __sb_start_write+0x58/0x120
[Wed Jul 22 02:02:06 2020] [<ffffffffff87c6c612>] vfs_fallocate+0x142/0x1e0
[Wed Jul 22 02:02:06 2020] [<ffffffffff87c6d66b>] SyS_fallocate+0x5b/0xa0
[Wed Jul 22 02:02:06 2020] [<ffffffffff881c9fd2>] system_call_fastpath+0x25/0x2a
[Wed Jul 22 02:02:06 2020] ---[end trace 00dd6f1b1f749469]---
[Wed Jul 22 02:02:06 2020] EXT4-fs: ext4_ext_split:1139: aborting transaction: error 28 in
__ext4_handle_dirty_metadata
[Wed Jul 22 02:02:06 2020] EXT4-fs error (device ploop43013p1): ext4_ext_split:1139: inode
#407808: block 2308152: comm e4defrag2: journal_dirty_metadata failed: handle type 3 started at
line 4741, credits 8/0, errcode -28
[Wed Jul 22 02:02:06 2020] Aborting journal on device ploop43013p1-8.
[Wed Jul 22 02:02:06 2020] EXT4-fs (ploop43013p1): Remounting filesystem read-only
[Wed Jul 22 02:02:06 2020] EXT4-fs error (device ploop43013p1) in ext4_free_blocks:4933:
Journal has aborted
[Wed Jul 22 02:02:06 2020] EXT4-fs error (device ploop43013p1) in ext4_free_blocks:4933:
Journal has aborted
[Wed Jul 22 02:02:06 2020] EXT4-fs (ploop43013p1): ext4_writespages: jbd2_start: 1023 pages,
ino 824451; err -30
[Wed Jul 22 02:02:06 2020] EXT4-fs error (device ploop43013p1) in
ext4_reserve_inode_write:5358: Journal has aborted
[Wed Jul 22 02:02:06 2020] EXT4-fs error (device ploop43013p1) in ext4_alloc_file_blocks:4753:
error 28
[Thu Jul 23 02:03:43 2020] EXT4-fs (ploop43013p1): error count since last fsck: 5
[Thu Jul 23 02:03:43 2020] EXT4-fs (ploop43013p1): initial error at time 1595376065:
ext4_ext_split:1139: inode 407808: block 2308152
[Thu Jul 23 02:03:43 2020] EXT4-fs (ploop43013p1): last error at time 1595376065:
ext4_alloc_file_blocks:4753: inode 407808: block 2308152

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [ktkhai](#) on Fri, 24 Jul 2020 10:56:04 GMT

[View Forum Message](#) <> [Reply to Message](#)

Please, try this kpatch module as a fix for kernel 3.10.0-1062.12.1.vz7.131.10:

#<http://fe.virtuozzo.com/602a9723c49782e0a8142a5b6fac65f3/kpatch-p.ko>

Installation is simple:

```
$kpatch load kpatch-p.ko
```

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Fri, 24 Jul 2020 14:56:21 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote:Please, try this kpatch module as a fix for kernel 3.10.0-1062.12.1.vz7.131.10:

#<http://fe.virtuozzo.com/602a9723c49782e0a8142a5b6fac65f3/kpatch-p.ko>

Installation is simple:

```
$kpatch load kpatch-p.ko
```

Thanks, but our recent issues have all happened on nodes that already have the latest kernel 3.10.0-1127.8.2.vz7.151.14

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [ktkhai](#) on Fri, 24 Jul 2020 15:27:34 GMT

[View Forum Message](#) <> [Reply to Message](#)

Then, here is kpatch for 3.10.0-1127.8.2.vz7.151.14:

<http://fe.virtuozzo.com/ad49bab0fd1a6daf3a9fa54e86159931/kpatch-p.ko>

Please, check that problem becomes disappeared with it.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Wed, 29 Jul 2020 14:58:48 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quote:Then, here is kpatch for 3.10.0-1127.8.2.vz7.151.14:

<http://fe.virtuozzo.com/ad49bab0fd1a6daf3a9fa54e86159931/kpatch-p.ko>

Please, check that problem becomes disappeared with it.

Thanks! But before applying it, is it possible to somehow check what the patch is supposed to change and to verify that this patch is legitimate and that the integrity of the file has not been compromised? Some signature/checksum system perhaps?

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [ktkhai](#) on Wed, 29 Jul 2020 15:24:49 GMT

[View Forum Message](#) <> [Reply to Message](#)

#md5sum kpatch-p.ko

8ec9997bc4a65c1a0b0a65d4950187a6 kpatch-p.ko

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [ktkhai](#) on Mon, 03 Aug 2020 10:19:21 GMT

[View Forum Message](#) <> [Reply to Message](#)

Any feedback on this? Does this fix your problem?
