
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] [<ffffffff81baebc7>] dump_stack+0x19/0x1b
[2020-05-29 02:02:20] [<ffffffff8149bdc8>] __warn+0xd8/0x100
[2020-05-29 02:02:20] [<ffffffff8149bf0d>] warn_slowpath_null+0x1d/0x20
[2020-05-29 02:02:20] [<ffffffffc047cd82>] __ext4_handle_dirty_metadata+0x1c2/0x220 [ext4]
[2020-05-29 02:02:20] [<ffffffffc0474f44>] ext4_ext_split+0x304/0x9a0 [ext4]
[2020-05-29 02:02:20] [<ffffffffc0476dfd>] ext4_ext_insert_extent+0x7bd/0x8d0 [ext4]
[2020-05-29 02:02:20] [<ffffffffc0479a7f>] ext4_ext_map_blocks+0x5cf/0xf60 [ext4]
[2020-05-29 02:02:20] [<ffffffffc0446676>] ext4_map_blocks+0x136/0x6b0 [ext4]
[2020-05-29 02:02:20] [<ffffffffc047496c>] ? ext4_alloc_file_blocks.isra.36+0xbc/0x2f0 [ext4]
[2020-05-29 02:02:20] [<ffffffffc047498f>] ext4_alloc_file_blocks.isra.36+0xdf/0x2f0 [ext4]
[2020-05-29 02:02:20] [<ffffffffc047b6bd>] ext4_fallocate+0x15d/0x990 [ext4]
[2020-05-29 02:02:20] [<ffffffff8166cb78>] ? __sb_start_write+0x58/0x120
[2020-05-29 02:02:20] [<ffffffff81666c72>] vfs_fallocate+0x142/0x1e0
[2020-05-29 02:02:20] [<ffffffff81667cdb>] SyS_fallocate+0x5b/0xa0
[2020-05-29 02:02:20] [<ffffffff81bc1fde>] 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!

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 if 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_conntrack_ipv4 nf_defrag_ipv4
xt_comment xt_multiport xt_conntrack 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
ebtable_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 lpc_ich
i2c_i801 joydev mei_me sg mei ioatdma wmi ipmi_si ipmi_devintf ipmi_msghandler acpi_pad
acpi_power_meter ip_vs nf_conntrack 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] [`<ffffffffc042a382>`] `__ext4_handle_dirty_metadata+0x1c2/0x220` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc04224b4>`] `ext4_ext_split+0x304/0x9a0` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc042436d>`] `ext4_ext_insert_extent+0x7bd/0x8d0` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc042700f>`] `ext4_ext_map_blocks+0x5cf/0xfc0` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc03f3ae6>`] `ext4_map_blocks+0x136/0x6b0` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc0421ecc>`] `? ext4_alloc_file_blocks.isra.37+0xbc/0x300` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc0421eef>`] `ext4_alloc_file_blocks.isra.37+0xdf/0x300` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffffc0428c9d>`] `ext4_fallocate+0x15d/0x9b0` [ext4]

[Wed Jul 22 02:02:06 2020] [`<ffffffff87c72428>`] `? __sb_start_write+0x58/0x120`

[Wed Jul 22 02:02:06 2020] [`<ffffffff87c6c612>`] `vfs_fallocate+0x142/0x1e0`

[Wed Jul 22 02:02:06 2020] [`<ffffffff87c6d66b>`] `Sys_fallocate+0x5b/0xa0`

[Wed Jul 22 02:02:06 2020] [`<ffffffff881c9fd2>`] `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_writepages: 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?

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [noc.r](#) on Thu, 20 Aug 2020 12:34:46 GMT

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

Hey there. I'm not OP.

We are experiencing similar problems while using kernel 3.10.0-1127.8.2.vz7.151.14:

```
2020-08-20T02:00:47+0200 pcompact : Inspect d35c0c6f-0b08-4541-82a0-d5a611c49ae2
2020-08-20T02:00:47+0200 pcompact : Inspect /vz/private/2660/root.hdd/DiskDescriptor.xml
2020-08-20T02:00:47+0200 pcompact : ploop=31258MB image=32907MB data=26247MB
balloon=3302MB
2020-08-20T02:00:47+0200 pcompact : Rate: 21.3 (threshold=10)
2020-08-20T02:00:47+0200 pcompact : Start compacting (to free 5097MB)
2020-08-20T02:00:47+0200 : Start defrag dev=/dev/ploop31385p1 mnt=/vz/root/2660
blocksize=2048
2020-08-20T02:00:56+0200 : Trying to find free extents bigger than 0 bytes granularity=512
2020-08-20T02:00:56+0200 pcompact : ploop=31258MB image=29563MB data=27143MB
balloon=3302MB
2020-08-20T02:00:56+0200 pcompact : Stats: uuid=d35c0c6f-0b08-4541-82a0-d5a611c49ae2
ploop_size=31258MB image_size_before=32907MB image_size_after=29563MB
compaction_time=9.058s type=online
2020-08-20T02:00:56+0200 pcompact : End compacting
```

So far the problem always seems to be related to images being higher than the total ploop size.

Is the patch on this thread supposed to solve these issues?

Any other thing which should be looked at?

Best regards and thank you,

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [kthai](#) on Thu, 20 Aug 2020 12:53:19 GMT

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

Quote: Is the patch on this thread supposed to solve these issues?

In case of you met the same message in dmesg as Allan, the patch should fix your problem.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [kthai](#) on Thu, 20 Aug 2020 12:54:49 GMT

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

Hi, Allan! Any news? Has patch fixed your problems?

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [noc.r](#) on Thu, 20 Aug 2020 15:04:11 GMT

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

kthai wrote on Thu, 20 August 2020 12:53: Quote: Is the patch on this thread supposed to solve these issues?

In case of you met the same message in dmesg as Allan, the patch should fix your problem.

I've tried to use the module you've linked but I'm finding some trouble.

```
# md5sum kpatch.ko
8ec9997bc4a65c1a0b0a65d4950187a6 kpatch.ko
```

```
# uname -r
3.10.0-1127.8.2.vz7.151.14
```

```
kpatch load kpatch.ko
loading core module: /usr/lib/modules/3.10.0-1127.8.2.vz7.151.14/extra/kpatch/kpatch.ko
insmod: ERROR: could not insert module
/usr/lib/modules/3.10.0-1127.8.2.vz7.151.14/extra/kpatch/kpatch.ko: Unknown symbol in module
```

kpatch: failed to load core module

Could you please double check whether that's the correct module for this kernel version?

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [ktkhai](#) on Fri, 21 Aug 2020 11:33:46 GMT

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

Try modprobe kpatch before this. Please, check you downloaded correct kpatch-p.ko (there are two different for two kernel versions in this BUG).

```
#uname -a
```

```
Linux xxxxx 3.10.0-1127.8.2.vz7.151.14 #1 SMP Tue Jun 9 12:58:54 MSK 2020 x86_64 x86_64  
x86_64 GNU/Linux
```

```
# modprobe kpatch
```

```
# wget http://fe.virtuozzo.com/ad49bab0fd1a6daf3a9fa54e86159931/kpatch-p.ko
```

```
# kpatch load ./kpatch-p.ko
```

```
loaded core module
```

```
loading patch module: ./kpatch-p.ko
```

```
# lsmod | grep kpatch
```

```
kpatch_p          26959  1
```

```
kpatch           56653  1 kpatch_p
```

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [noc.r](#) on Mon, 24 Aug 2020 10:45:13 GMT

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

ktkhai wrote on Fri, 21 August 2020 11:33Try modprobe kpatch before this. Please, check you downloaded correct kpatch-p.ko (there are two different for two kernel versions in this BUG).

```
#uname -a
```

```
Linux xxxxx 3.10.0-1127.8.2.vz7.151.14 #1 SMP Tue Jun 9 12:58:54 MSK 2020 x86_64 x86_64  
x86_64 GNU/Linux
```

```
# modprobe kpatch
```

```
# wget http://fe.virtuozzo.com/ad49bab0fd1a6daf3a9fa54e86159931/kpatch-p.ko
```

```
# kpatch load ./kpatch-p.ko
```

```
loaded core module
```

```
loading patch module: ./kpatch-p.ko
```

```
# lsmod | grep kpatch
```

```
kpatch_p          26959  1
```

```
kpatch           56653  1 kpatch_p
```

I'm using the correct version posted and I've checked that it was matching the md5sum. I'm using

the kernel version 3.10.0-1127.8.2.vz7.151.14 which seems the one the patch was posted for.

I did notice that the kpatch kmod is not loading on my system, which is weird.

I've investigated this a little bit and realised that the kpatch kmod is generated by the "kpatch-kmod" package at the "virtuozzlinux-base" repository, but apparently its bindings are not for this kernel version:

--> Resolución de dependencias finalizada

Removing kpatch-kmod-15.2-0.3.3-62.20160301.vl7.x86_64 since corresponding vzkernel package is not installed.

Do you have any idea as to why this is happening?

Am I forced to compile the kpatch kmod on my own?

Regards,

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [ktkhai](#) on Mon, 24 Aug 2020 11:16:28 GMT

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

Oh, sure, please, get kpatch-mod from here

<http://fe.virtuozzo.com/30917b643cb8f77fc3d6214d144a47f8/kpa>

tch-kmod-151.14-0.5.0-4.vl7.x86_64.rpm

md5sum kpatch-kmod-151.14-0.5.0-4.vl7.x86_64.rpm

5cfc0c589e03293486a13dba4c8d0a6 kpatch-kmod-151.14-0.5.0-4.vl7.x86_64.rpm

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [noc.r](#) on Mon, 24 Aug 2020 14:11:19 GMT

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

ktkhai wrote on Mon, 24 August 2020 11:16Oh, sure, please, get kpatch-mod from here

<http://fe.virtuozzo.com/30917b643cb8f77fc3d6214d144a47f8/kpa>

tch-kmod-151.14-0.5.0-4.vl7.x86_64.rpm

md5sum kpatch-kmod-151.14-0.5.0-4.vl7.x86_64.rpm

5cfc0c589e03293486a13dba4c8d0a6 kpatch-kmod-151.14-0.5.0-4.vl7.x86_64.rpm

Thank you very much. That seemed to work.

I have a few questions, if you dont mind and know the answer:

- * Why is the kpatch-kmod package not available for my kernel? is there some kind of issue on my end?
- * Is it supposed to be available for my kernel?
- * Finally, any way to test that the patch works other than wait for the issues to appear?

Thank you very much,

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_threshold":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/_pcompact_conf.html](https://docs.openvz.org/openvz_command_line_reference.webhelp/_pcompact_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!

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [eshatokhin](#) on Tue, 25 Aug 2020 17:50:32 GMT

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

noc.r wrote on Mon, 24 August 2020 14:11 Why is the kpatch-kmod package not available for my kernel? is there some kind of issue on my end?

No, this is expected.

kpatch-kmod packages are currently used for ReadyKernel only (<https://readykernel.com/faq>), which is available for Virtuozzo but not for OpenVZ. The packages are kept in a separate

repository for that reason.

noc.r wrote on Mon, 24 August 2020 14:11Is it supposed to be available for my kernel?
No, kpatch-kmod packages are not supposed to be available in OpenVZ. Still, we can provide these packages when needed for debugging, etc., like it was in this case.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [khorenko](#) on Wed, 26 Aug 2020 11:46:07 GMT

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

noc.r wrote on Mon, 24 August 2020 17:11I have a few questions, if you dont mind and know the answer:

- * Why is the kpatch-kmod package not available for my kernel? is there some kind of issue on my end?
- * Is it supposed to be available for my kernel?
- * Finally, any way to test that the patch works other than wait for the issues to appear?

Well, Zhenya has covered first 2 questions, AFAIS,
so the only last question left here.

And - no, i do not see any simple way for verification because of

1) the kpatch patch contain 2 mainstream commits:

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=812c0cab2c0dfad977605dbadf9148490ca5d93f>

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=4134f5c88dcd>

And those patches verification is not so easy.

But most important reason 2:

- 2) Even if we verify those patches to fix the problem they are intended to fix,
this does not mean, it will 100% fix your exact problem.
We think so, but this is not 100% proven.

So, please let us know if you experience issues when the kpatch patch is loaded.

Thank you.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [khorenko](#) on Wed, 26 Aug 2020 12:02:23 GMT

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

nathan.brownrice wrote on Mon, 24 August 2020 19:48What I'd like to discuss:

1) Others that have had this same issue, and have applied the kernel update, did this fix your issues?

We would be also very interesting in the feedback on the patch.

nathan.brownrice wrote on Mon, 24 August 2020 19:48:22 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.

You can install the kpatch patch from this thread on one node and check if it helps - it does not require Node reboot.

nathan.brownrice wrote on Mon, 24 August 2020 19:48:22 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/_pcompact_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?

May be you are right and pcompact significantly increases the chances the issue to trigger. But the truth is - until you (or someone else if you 100% sure you have exactly the same issue) test the patch, you cannot be sure the issue is ever fixed, so you can wait for "updated kernels" forever. i mean - something will be fixed in newer kernels, but it easily could be not a fix for your particular issue.

So i really suggest to install the kpatch patch and check if it helps. After that you will know you (and we) are on the right way at least.

BTW, if someone prefers full kernel instead of kpatch patch installation, here you are:
<http://fe.virtuozzo.com/a3001136c32272a6889092b16af03f64/>

This kernel contains a dozen of patches comparing to vz7.151.14 kernel, but all of them are important, most of them are already released as ReadyKernel patches, so this kernel can be considered as a stable one.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors
Posted by [noc.r](#) on Wed, 26 Aug 2020 12:24:29 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thank you all for the detailed answers.

I'll deploy the patch and get back to you with feedback.

As this is kind of an odd an arbitrary issue --it only happens every once in a while-- it will be

difficult to confirm and only time will tell. I'll get back to you regardless.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Tue, 01 Sep 2020 11:00:31 GMT

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

Hi,

Sorry for the delay, but I was away from work for some time. Meanwhile our team has stress tested OpenVZ 7 nodes by creating lot of random reads and writes on the ploop devices and then triggering pcompact. And we have been able to establish the pattern that with unpatched kernel we now regularly see the issue happening. Our planned next step was to apply the patch and see if the pattern will change. Unfortunately the link provided to patch for the 3.10.0-1127.8.2.vz7.151.14 kernel does not work anymore.

Could you please provide the link again. We could then apply the patch shortly and get back with the feedback about the status.

Best regards,

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [kthkhai](#) on Tue, 01 Sep 2020 12:04:06 GMT

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

<http://fe.virtuozzo.com/ae04f66cb5e1398aae2d1522a8095496/kpatch-p.ko>

\$md5sum kpatch-p.ko

d3e330b20f20c431887f93ff747846ea kpatch-p.ko

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [noc.r](#) on Thu, 03 Sep 2020 12:37:36 GMT

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

I have a few machines that have an updated kernel: 3.10.0-1127.10.1.vz7.162.9

Could we please have the kpatch-kmod and the kpatch.ko for this kernel so that we can try the patch?

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [eshatokhin](#) on Thu, 03 Sep 2020 13:00:55 GMT

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

noc.r wrote on Thu, 03 September 2020 12:37I have a few machines that have an updated kernel: 3.10.0-1127.10.1.vz7.162.9

Could we please have the kpatch-kmod and the kpatch.ko for this kernel so that we can try the patch?

Unfortunately, not yet. Currently, live patches are only built for the kernels from Virtuozzo releases and hotfixes. Kernel 3.10.0-1127.10.1.vz7.162.9 is not yet supported there.

I can try to build the modules for that particular kernel manually though, perhaps, in a day or two.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [eshatokhin](#) on Fri, 04 Sep 2020 16:59:41 GMT

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

http://fe.virtuozzo.com/eb9978e1071ab0c10993c2fd62dea865/kpatch-kmod-162.9-0.5.0-4.vl7.x86_64.rpm

<http://fe.virtuozzo.com/eb9978e1071ab0c10993c2fd62dea865/kpatch-p-162.9.ko>

sha1sum kpatch-kmod-162.9-0.5.0-4.vl7.x86_64.rpm kpatch-p-162.9.ko:

377a1549212e89cccc0055c1cab84115b60bf919 kpatch-kmod-162.9-0.5.0-4.vl7.x86_64.rpm

ae4c8e7d18de45723e4b186b64f30c7b6f41a2b7 kpatch-p-162.9.ko

I have only checked that the patch module loads OK, so, please, use it with caution.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Fri, 11 Sep 2020 14:08:50 GMT

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

Hello,

Wanted to give a short update and ask a couple of questions.

First of all, last night one of the virtual containers on a non-patched host went into read-only. However, different from the previous cases is that pcompact was not involved in this case. We actually have pcompact cron disabled and we trigger it manually on the nodes that are in our test sample. The errors in messages log were:

Sep 11 00:34:00 server-n697 kernel: bash (639778): drop_caches: 3

Sep 11 00:34:36 server-n697 systemd: Started Session c185489 of user root.

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_free_blocks:4933: Out of memory

Sep 11 00:34:39 server-n697 kernel: Aborting journal on device ploop35478p1-8.

Sep 11 00:34:39 server-n697 kernel: EXT4-fs (ploop35478p1): Remounting filesystem read-only

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_ext_remove_space:3073: IO failure

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_ext_truncate:4692: Journal has aborted

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_reserve_inode_write:5358: Journal has aborted

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_truncate:4145: Journal has aborted

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_reserve_inode_write:5358: Journal has aborted

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_orphan_del:2731: Journal has aborted

Sep 11 00:34:39 server-n697 kernel: EXT4-fs error (device ploop35478p1) in ext4_reserve_inode_write:5358: Journal has aborted

Sep 11 00:34:39 server-n697 kernel: EXT4-fs (ploop35478p1): ext4_writepages: jbd2_start: 0 pages, ino 661650; err -30

Sep 11 00:34:41 server-n697 kernel: dd invoked oom-killer: gfp_mask=0xd0, order=0, oom_score_adj=0

Sep 11 00:34:41 server-n697 kernel: dd cpuset=3988 mems_allowed=0

Sep 11 00:34:41 server-n697 kernel: CPU: 6 PID: 273822 Comm: dd ve: 3988 Kdump: loaded Not tainted 3.10.0-1127.8.2.vz7.151.14 #1 151.14

Sep 11 00:34:41 server-n697 kernel: Hardware name: Supermicro X9DRW/X9DRW, BIOS 3.0c 03/24/2014

Sep 11 00:34:41 server-n697 kernel: Call Trace:

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb95b67f1>] dump_stack+0x19/0x1b

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb95b0fc6>] dump_header+0x90/0x229

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb8fd7076>] ? find_lock_task_mm+0x56/0xc0

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb8fd7dad>] oom_kill_process+0x47d/0x640

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb90040fe>] ? get_task_oom_score_adj+0xee/0x100

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb8fd7213>] ? oom_badness+0x133/0x1e0

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb905f509>] mem_cgroup_oom_synchronize+0x4b9/0x510

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb8fd84c3>] pagefault_out_of_memory+0x13/0x50

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb95af06d>] mm_fault_error+0x6a/0x157

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb95c49a1>] __do_page_fault+0x491/0x500

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb95c4a45>] do_page_fault+0x35/0x90

Sep 11 00:34:41 server-n697 kernel: [<ffffffffffb95c0778>] page_fault+0x28/0x30

Sep 11 00:34:41 server-n697 kernel: Task in /machine.slice/3988 killed as a result of limit of /machine.slice/3988

Sep 11 00:34:41 server-n697 kernel: memory: usage 4192232kB, limit 4194304kB, failcnt 13429301

Sep 11 00:34:41 server-n697 kernel: memory+swap: usage 4325376kB, limit 4325376kB, failcnt 31733768989

Sep 11 00:34:41 server-n697 kernel: kmem: usage 36356kB, limit 9007199254740988kB, failcnt 0

Sep 11 00:34:41 server-n697 kernel: Memory cgroup stats for /machine.slice/3988:

rss_huge:479232KB mapped_file:33780KB shmem:84KB slab_unreclaimable:9776KB

swap:133144KB cache:3459360KB rss:696520KB slab_reclaimable:10664KB

inactive_anon:199280KB active_anon:497324KB inactive_file:597760KB active_file:2861360KB

unevictable:0KB

Sep 11 00:34:41 server-n697 kernel: Memory cgroup out of memory: Kill process 671560

(mysqld) score 87 or sacrifice child

Sep 11 00:34:41 server-n697 kernel: Killed process 581188 (mysqld) in VE "3988", UID 116, total-vm:2820740kB, anon-rss:376664kB, file-rss:0kB, shmem-rss:0kB

Sep 11 00:34:53 server-n697 systemd: Started Session c185490 of user root.

Sep 11 00:35:01 server-n697 kernel: bash (639778): drop_caches: 3

Please note that the container was running the script that was creating and deleting random files. The script we use to generate read-write activity for testing purposes. Of course, the script also failed after filesystem was switched to read-only mode.

Would you think that the cause of this error is related to the same bug that the kernel patch should fix? Meaning that pcompact is not part of the problem, but just happens to amplify the issue.

Secondly, we see that while pcompact is running, the virtual container disk utilisation fluctuates rapidly and the disk also comes 100% full several time. Is that an expected behaviour during a pcompact run? As a result, we have seen some applications (like MySQL) throwing errors saying that it's unable to write to disk. Feels like it could cause data corruption.

Thirdly, worth noting that we have not yet seen a container switching to read-only on a node that has the kernel patch applied. But even on that node, the disk was showing as full when pcompact ran, so the concern regarding potential data corruption applies.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [noc.r](#) on Mon, 28 Sep 2020 14:12:26 GMT

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

Hey there. Wanted to post an update and say that patched machines didn't have any read only error on ~4 weeks, so I'm assuming that the patch is working.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [nathan.brownrice](#) on Sat, 03 Oct 2020 18:13:22 GMT

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

Looks like this link is stale again, can we get a fresh one please?

http://fe.virtuozzo.com/ae04f66cb5e1398aae2d1522a8095496/kpa_tch-p.ko

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [eshatokhin](#) on Sun, 04 Oct 2020 15:33:35 GMT

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

nathan.brownrice wrote on Sat, 03 October 2020 18:13 Looks like this link is stale again, can we get a fresh one please?

<http://fe.virtuozzo.com/0ee3b7b7a249d06134b4eaf7ef50f3d8/>

The link will expire in 30 days.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Mon, 05 Oct 2020 05:56:49 GMT

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

One important and related question. Is it reasonable to expect that the fix will be part of any future kernels after 3.10.0-1127.8.2.vz7.151.14? So that if one day we have an installer image with newer kernel or if we should update the kernels of existing servers, the issue will not reappear?

Thanks,
Allan

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [eshatokhin](#) on Mon, 05 Oct 2020 10:22:56 GMT

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

allan.talver wrote on Mon, 05 October 2020 05:56: Is it reasonable to expect that the fix will be part of any future kernels after 3.10.0-1127.8.2.vz7.151.14?

The patches were included into kernel 3.10.0-1127.18.2.vz7.163.2. If they actually fix your issue, updating kernel to that or a newer version (when it is available) could help.

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [allan.talver](#) on Wed, 14 Oct 2020 08:02:49 GMT

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

Hi all,

A short update and a few questions:

- * We have now roughly 50 nodes running with vz7.151.14 kernel and the patch and we have not had containers turning into read-only mode. I think this is good.

- * However, we still have another issue related to pcompact and, as we have increased the number of nodes on vz7, it also is causing us more trouble. The issue is that while pcompact and defrag are running, the disk space utilisation of the container starts fluctuating rapidly and during that period, it is also shown as 100% full several times. In turn, when the disk shows as full, applications are unable to write to the disk and it causes unexpected errors and application crashes (for example MySQL and Redis instances with configured with persistence). Are we the only ones experiencing this behaviour? Is there a way to avoid this? For now we have stopped pcompact again because it has caused several production incidents.

- * A question related to the previous point, but what I would like to address separately. What are the downsides if we don't run pcompact at all? I think one obvious outcome is that we won't be able to release unused space in the container disk images back to the host operating system. But we don't see it as a big issue because we normally do not overcommit hardware node disk space anyway (meaning that total size of container disks is below the size of disk space on the host)

node). In this scenario, could we just leave pcompact disabled?

* We noticed that now there is new kernel version vz7.158.8 which is getting installed on the newer nodes. Of course the patch that has been shared in this thread here is not usable for this version. But I believe that the kernel fix is not yet in that version of the kernel. Is it possible to have another patch for vz7.158.8?

Thanks and best regards,
Allan

Subject: Re: OpenVZ 7 containers crashing with ext4 errors

Posted by [eshatokhin](#) on Wed, 14 Oct 2020 08:27:46 GMT

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

allan.talver wrote on Wed, 14 October 2020 08:02

* We noticed that now there is new kernel version vz7.158.8 which is getting installed on the newer nodes. Of course the patch that has been shared in this thread here is not usable for this version. But I believe that the kernel fix is not yet in that version of the kernel. Is it possible to have another patch for vz7.158.8?

I have checked - the kernel vz7.158.8 also has the needed fixes for ext4, same as the in-development vz7.163.x series. Live patches with these fixes are not needed there.
