

Import 'protective MBR' GPT disk images like SUSE or openSUSE to LVM-Thin with disabled zeroing cause corruption

Date: 2022-04-14
modified: 2022-04-15
tags: Linux, Proxmox, SUSE, openSUSE, LVM. GPT
description: SUSE / openSUSE KVM Images - Timeout / Issue on boot '/dev/disk/by-uuid'
category: Linux
slug: import_protective-mbr_gpt_disk_images_like_suse_or_opensuse_to_lvm-thin_with_disabled_zeroing_cause_corruption
Author: Dominik Wombacher
lang: en
transid: import_protective-mbr_gpt_disk_images_like_suse_or_opensuse_to_lvm-thin_with_disabled_zeroing_cause_corruption
Status: published

Any available SUSE and openSUSE KVM Image I tried, ended in a VM that stuck on boot with the message "a start job is running for /dev/disk/by-uuid" or dropped into the dracut emergency shell after a timeout.

No partitions could be found, only "/dev/sda" was visible from inside the VM, gdisk dropped a lot of warnings and errors as well:

```
Caution! After loading partitions, the CRC doesn't check out!  
Warning! Main and backup partition tables differ! Use the 'c' and 'e' options  
on the recovery & transformation menu to examine the two tables.  
  
Warning! One or more CRCs don't match. You should repair the disk!  
Main header: OK  
Backup header: OK  
Main partition table: ERROR  
Backup partition table: ERROR  
  
Partition table scan:  
  MBR: protective  
  BSD: not present  
  APM: not present  
  GPT: damaged  
  
*****  
Caution: Found protective or hybrid MBR and corrupt GPT. Using GPT, but disk  
verification and recovery are STRONGLY recommended.  
*****  
  
Warning! Secondary partition table overlaps the last partition by  
18315034227491254276 blocks!  
You will need to delete this partition or resize it in another utility.
```

I tried it on a Proxmox VE Host with a LVM-Thin Datastore, a closer look at the LVM-Thin Device of one of the affected VMs confirmed the missing partitions:

```
fdisk -l /dev/mapper/pve-vm--9181--disk--0
Disk /dev/mapper/pve-vm--9181--disk--0: 24 GiB, 25769803776 bytes, 50331648 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 524288 bytes / 524288 bytes
Disklabel type: dos
Disk identifier: 0x406b6f28

Device                               Boot Start      End  Sectors Size Id Type
/dev/mapper/pve-vm--9181--disk--0-part1 1 50331647 50331647 24G ee GPT

Partition 1 does not start on physical sector boundary.
```

I chased that issue down for weeks and almost gave up because I didn't find any clue, it just felt like no one ever faced a similar issue before. After a lot of different tests on multiple systems and configurations, my LVM-Thin Storage seemed to cause the issue and by accident I finally found a hint in the [Proxmox Forum](#) (Archive: [\[1\]](#), [\[2\]](#)).

Indeed, no idea why but I disabled zeroing on the LV during conversion to Thin:

```
lvconvert --type thin-pool pve/lv_test -Zn -y
```

The official [Proxmox Documentation](#) (Archive: [\[1\]](#), [\[2\]](#)) doesn't mention the Parameter `-Zn`, obviously because of very good reasons.

Tests with a new LVM-Thin with enabled zeroing (default) confirmed that all issues I faced earlier seem to be caused by that. A new VM based on the same `qcow2` image as before is now starting without any issue, delay or timeout, just as expected, also the partition table looks way better:

```
fdisk -l /dev/mapper/pve-vm--9181--disk--0
Disk /dev/mapper/pve-vm--9181--disk--0: 24 GiB, 25769803776 bytes, 50331648 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 65536 bytes / 65536 bytes
Disklabel type: gpt
Disk identifier: 0D13903E-0408-4963-9C0D-A455B37C6062

Device                               Start      End  Sectors Size Type
/dev/mapper/pve-vm--9181--disk--0-part1 2048      6143   4096 2M BIOS boot
/dev/mapper/pve-vm--9181--disk--0-part2 6144     73727  67584 33M EFI System
/dev/mapper/pve-vm--9181--disk--0-part3 73728 50331614 50257887 24G Linux filesystem
```

I installed some VMs "normally" by booting a .iso files, also importing `qcow2` images with a "flat" / "oldschool" Partition Layout, like the Rocky Linux Generic Cloud Image, which comes with a single XFS formatted MBR Partition, was working without the above problems.

So it looks like that's something specific to disk images with 'protective MBR' GPT and EFI Partitions. So far only SUSE and openSUSE Images seem to use such a Layout by default.

To be perfectly honest, I didn't have a fully technical explanation why disabled zeroing on LVM-Thin has such an impact when importing PMBR disks. I can only assume that during the LV creation and `qcow2` import as well as conversion to `raw` there is some messing around with the first few sectors.

Fact is: I enabled zeroing on the LVM-Thin LV, everything got back to normal and the pre-build images behave now as they are supposed to.