

ceph-volume - Bug #23581

ceph-volume should use mount -o discard when underlying device is using VDO

04/06/2018 07:15 PM - David Galloway

Status: Resolved	% Done: 0%
Priority: Immediate	
Assignee:	
Category:	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport:	Pull request ID:
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	

Description

```
[root@reesi004 ~]# lsblk /dev/mapper/vdo_sda
NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
vdo_sda              253:12  0   3.6T  0 vdo
└─vg_sda-lv_sda     253:13  0   3.6T  0 lvm

[root@reesi004 ~]# mount | grep osd
tmpfs on /var/lib/ceph/osd/ceph-55 type tmpfs (rw,relatime,seclabel)

# ceph-volume lvm list

===== osd.55 =====

[ db]      /dev/journals/sda

type                db
osd id               55
cluster fsid         dd5d3a03-ba02-48a5-97ec-869b5443fb7b
cluster name         ceph
osd fsid             19e83348-bc84-4516-9a78-fbac863ca3f7
db device            /dev/journals/sda
encrypted            0
db uuid              Y7JcqI-YQ2A-CetU-rqwA-Fp48-yT7h-mWjG07
cephx lockbox secret
block uuid           EutNis-Gb1K-D281-dAMa-tHtD-6nYW-qiOCXv
block device         /dev/vg_sda/lv_sda
crush device class   None

[block]     /dev/vg_sda/lv_sda

type                block
osd id               55
cluster fsid         dd5d3a03-ba02-48a5-97ec-869b5443fb7b
cluster name         ceph
osd fsid             19e83348-bc84-4516-9a78-fbac863ca3f7
db device            /dev/journals/sda
encrypted            0
db uuid              Y7JcqI-YQ2A-CetU-rqwA-Fp48-yT7h-mWjG07
cephx lockbox secret
block uuid           EutNis-Gb1K-D281-dAMa-tHtD-6nYW-qiOCXv
block device         /dev/vg_sda/lv_sda
crush device class   None
```

https://bugzilla.redhat.com/show_bug.cgi?id=1563794#c2

History

#1 - 04/06/2018 07:16 PM - David Galloway

- Description updated

#2 - 04/09/2018 09:08 PM - Alfredo Deza

- Priority changed from Normal to Immediate

#3 - 04/09/2018 09:43 PM - Alfredo Deza

This only affects filestore, bluestore does not need any of the special mounts since it is using tmpfs.

#4 - 04/16/2018 05:20 PM - Alfredo Deza

PR at <https://github.com/ceph/ceph/pull/21451>

#5 - 04/16/2018 05:20 PM - Alfredo Deza

- Status changed from New to Fix Under Review

#6 - 04/18/2018 03:17 PM - Alfredo Deza

- Status changed from Fix Under Review to Resolved

merged commit 06c0933 into master

#7 - 05/02/2018 10:27 AM - Sergey Ponomarev

Alfredo Deza wrote:

This only affects filestore, bluestore does not need any of the special mounts since it is using tmpfs.

Please answer me, how can enable block discarding VDO if top of OSD -> Bluestore (not Filestore)?
I see counter "block map discard required" don't decrease and always growing up

#8 - 05/09/2018 11:51 AM - Sergey Ponomarev

Sergey Ponomarev wrote:

Alfredo Deza wrote:

This only affects filestore, bluestore does not need any of the special mounts since it is using tmpfs.

Please answer me, how can enable block discarding VDO if top of OSD -> Bluestore (not Filestore)?
I see counter "block map discard required" don't decrease and always growing up

I don't understand, for Ceph 12.2.5 in chain VDO+Bluestore free space don't discard and reclaim?
This work only for VDO+Filestore?

#9 - 05/16/2018 12:09 PM - Alfredo Deza

With bluestore, we don't mount devices so no discard flags can be enabled on devices. Bluestore does its own discard internally. More details on how that is done in this PR <https://github.com/ceph/ceph/pull/14727>