

## bluestore - Feature #20801

### ability to rebuild BlueStore WAL journals is missing

07/27/2017 05:38 PM - Dmitry Smirnov

<b>Status:</b> Rejected	<b>% Done:</b> 0%
<b>Priority:</b> Normal	
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b> v12.1.0	
<b>Source:</b> Community (user)	<b>Reviewed:</b>
<b>Tags:</b>	<b>Affected Versions:</b> v12.0.0
<b>Backport:</b>	<b>Pull request ID:</b>
<b>Description</b>	
<p>We have a six-node CEPH cluster v12.0.1 with BlueStore back-end, all RocksDB WAL journals were located on Flashtec NVRAM PCIe cards. Due to power outage and not yet clarified hardware issue we have lost all WAL partitions (5 partitions for 5 SAS OSD in each server). Data disks are not affected and looks healthy, OSD tree map is healthy as well, but the entire cluster is down. Is it any possibility to rebuild those journals from scratch? I tried to create GPT partitions with correct WAL GUID code (5CE17FCE-4087-4169-B7FF-056CC58473F9), same size as original and respective Partition UUID (taken from /var/lib/ceph/osd/ceph-X/block.wal_uuid). But it seems not enough, some data should be written to that new partition as well.</p> <pre>/usr/sbin/ceph-disk --verbose activate-block /dev/nvme0n1p1</pre> <p>get_dm_uuid: get_dm_uuid /dev/nvme0n1p1 uuid path is /sys/dev/block/259:2/dm/uuid command: Running command: /usr/sbin/blkid -o udev -p /dev/nvme0n1p1 command: Running command: /usr/bin/ceph-osd --get-device-fsid /dev/nvme0n1p1 get_space_osd_uuid: Block /dev/nvme0n1p1 has OSD UUID 00000000-0000-0000-0000-000000000000 main_activate_space: activate: OSD device not present, not starting, yet</p>	

#### History

##### #1 - 11/29/2017 11:07 PM - Sage Weil

- Project changed from Ceph to bluestore
- Category deleted (OSD)

##### #2 - 02/20/2018 02:29 PM - Sage Weil

- Status changed from New to Rejected

The wal or journal is an integral part of the store. The data store cannot be reconstructed without it.