

Ceph - Bug #19489

ceph-disk: failing to activate osd with multipath

04/04/2017 03:02 PM - Matt Stroud

| | | | |
|------------------------|----------|---------------------------|------------|
| Status: | Resolved | Start date: | 04/04/2017 |
| Priority: | Normal | Due date: | |
| Assignee: | | % Done: | 0% |
| Category: | | Estimated time: | 0.00 hour |
| Target version: | | Spent time: | 0.00 hour |
| Source: | | Severity: | 2 - major |
| Tags: | | Reviewed: | |
| Backport: | jewel | Affected Versions: | |
| Regression: | No | ceph-qa-suite: | |

Description

We are trying to setup a new cluster using multipath disks. I'm able to prepare the osd just fine, but it fails to activate. Here is the output:

```
[root@mon01 ceph-config]# ceph-deploy osd activate osd01:mapper/mpatha
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (1.5.37): /usr/bin/ceph-deploy osd activate osd01:mapper/mpatha
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username : None
[ceph_deploy.cli][INFO ] verbose : False
[ceph_deploy.cli][INFO ] overwrite_conf : False
[ceph_deploy.cli][INFO ] subcommand : activate
[ceph_deploy.cli][INFO ] quiet : False
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at 0x170e710>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] func : <function osd at 0x1701b90>
[ceph_deploy.cli][INFO ] ceph_conf : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.cli][INFO ] disk : [('osd01', '/dev/mapper/mpatha', None)]
[ceph_deploy.osd][DEBUG ] Activating cluster ceph disks osd01:/dev/mapper/mpatha:
[osd01][DEBUG ] connected to host: osd01
[osd01][DEBUG ] detect platform information from remote host
[osd01][DEBUG ] detect machine type
[osd01][DEBUG ] find the location of an executable
[ceph_deploy.osd][INFO ] Distro info: CentOS Linux 7.3.1611 Core
[ceph_deploy.osd][DEBUG ] activating host osd01 disk /dev/mapper/mpatha
[ceph_deploy.osd][DEBUG ] will use init type: systemd
[osd01][DEBUG ] find the location of an executable
[osd01][INFO ] Running command: /usr/sbin/ceph-disk -v activate --mark-init systemd --mount /dev/mapper/mpatha
[osd01][WARNIN] main_activate: path = /dev/mapper/mpatha
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:3/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:3/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] command: Running command: /sbin/blkid -p -s TYPE -o value -- /dev/mapper/mpatha
[osd01][WARNIN] Traceback (most recent call last):
[osd01][WARNIN] File "/usr/sbin/ceph-disk", line 9, in <module>
```

```

[osd01][WARNIN] load_entry_point('ceph-disk==1.0.0', 'console_scripts', 'ceph-disk')()
[osd01][WARNIN] File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 5047, in run
[osd01][WARNIN] main(sys.argv[1:])
[osd01][WARNIN] File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 4998, in main
[osd01][WARNIN] args.func(args)
[osd01][WARNIN] File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 3357, in main_activate
[osd01][WARNIN] reactivate=args.reactivate,
[osd01][WARNIN] File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 3067, in mount_activate
[osd01][WARNIN] e,
[osd01][WARNIN] ceph_disk.main.FilesystemTypeError: Cannot discover filesystem type: device /dev/mapper/mpatha: Line is truncated:
[osd01][ERROR ] RuntimeError: command returned non-zero exit status: 1
[ceph_deploy][ERROR ] RuntimeError: Failed to execute command: /usr/sbin/ceph-disk -v activate --mark-init systemd --mount /dev/mapper/mpatha

```

I have tried switching between using root and the ceph user. With the ceph user I get permission denied while trying to activate. I'm trying to use Jewel due to its long term support, however I'm going to give kraken a try after reporting this.

Related issues:

Copied to Ceph - Backport #20837: jewel: ceph-disk: failing to activate osd w...

Resolved

History

#1 - 04/04/2017 03:11 PM - Matt Stroud

Sorry I pasted incorrect output:

```

[root@mon01 ceph-config]# ceph-deploy osd activate osd01:mapper/mpathal
[ceph_deploy.conf][DEBUG ] found configuration file at: /root/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (1.5.37): /usr/bin/ceph-deploy osd activate osd01:mapper/mpathal
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username : None
[ceph_deploy.cli][INFO ] verbose : False
[ceph_deploy.cli][INFO ] overwrite_conf : False
[ceph_deploy.cli][INFO ] subcommand : activate
[ceph_deploy.cli][INFO ] quiet : False
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at 0x7fb
d52801758>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] func : <function osd at 0x7fbd527f4b90>
[ceph_deploy.cli][INFO ] ceph_conf : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.cli][INFO ] disk : [('osd01', '/dev/mapper/mpathal', None)]
[ceph_deploy.osd][DEBUG ] Activating cluster ceph disks osd01:/dev/mapper/mpathal:
[osd01][DEBUG ] connected to host: osd01
[osd01][DEBUG ] detect platform information from remote host
[osd01][DEBUG ] detect machine type
[osd01][DEBUG ] find the location of an executable
[ceph_deploy.osd][INFO ] Distro info: CentOS Linux 7.3.1611 Core
[ceph_deploy.osd][DEBUG ] activating host osd01 disk /dev/mapper/mpathal
[ceph_deploy.osd][DEBUG ] will use init type: systemd
[osd01][DEBUG ] find the location of an executable
[osd01][INFO ] Running command: /usr/sbin/ceph-disk -v activate --mark-init systemd --mount /dev/mapper/mpathal
[osd01][WARNIN] main_activate: path = /dev/mapper/mpathal
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathal uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathal uuid is part1-mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathal uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathal uuid is part1-mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] command: Running command: /usr/sbin/blkid -o udev -p /dev/mapper/mpathal
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathal uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathal uuid is part1-mpath-360060e80074e840000304e8400004000

```

```

[osd01][WARNIN]
[osd01][WARNIN] command: Running command: /sbin/blkid -p -s TYPE -o value -- /dev/mapper/mpathal
[osd01][WARNIN] command: Running command: /usr/bin/ceph-conf --cluster=ceph --name=osd. --lookup osd_mount_opt
ions_xfs
[osd01][WARNIN] command: Running command: /usr/bin/ceph-conf --cluster=ceph --name=osd. --lookup osd_fs_mount_
options_xfs
[osd01][WARNIN] mount: Mounting /dev/mapper/mpathal on /var/lib/ceph/tmp/mnt.EKR3Pp with options noatime,inode
64
[osd01][WARNIN] command_check_call: Running command: /usr/bin/mount -t xfs -o noatime,inode64 -- /dev/mapper/m
pathal /var/lib/ceph/tmp/mnt.EKR3Pp
[osd01][WARNIN] command: Running command: /usr/sbin/restorecon /var/lib/ceph/tmp/mnt.EKR3Pp
[osd01][WARNIN] activate: Cluster uuid is fb991e48-c425-4f82-a70e-5ce748ae186b
[osd01][WARNIN] command: Running command: /usr/bin/ceph-osc --cluster=ceph --show-config-value=fsid
[osd01][WARNIN] activate: Cluster name is ceph
[osd01][WARNIN] activate: OSD uuid is 924f6191-152c-4d1d-b920-9e7237467573
[osd01][WARNIN] activate: OSD id is 1
[osd01][WARNIN] activate: Initializing OSD...
[osd01][WARNIN] command_check_call: Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd
--keyring /var/lib/ceph/bootstrap-osd/ceph.keyring mon getmap -o /var/lib/ceph/tmp/mnt.EKR3Pp/activate.monmap
[osd01][WARNIN] got monmap epoch 2
[osd01][WARNIN] command: Running command: /usr/bin/timeout 300 ceph-osc --cluster ceph --mkfs --mkkey -i 1 --m
onmap /var/lib/ceph/tmp/mnt.EKR3Pp/activate.monmap --osd-data /var/lib/ceph/tmp/mnt.EKR3Pp --osd-journal /var/
lib/ceph/tmp/mnt.EKR3Pp/journal --osd-uuid 924f6191-152c-4d1d-b920-9e7237467573 --keyring /var/lib/ceph/tmp/mn
t.EKR3Pp/keyring --setuser ceph --setgroup ceph
[osd01][WARNIN] mount_activate: Failed to activate
[osd01][WARNIN] unmount: Unmounting /var/lib/ceph/tmp/mnt.EKR3Pp
[osd01][WARNIN] command_check_call: Running command: /bin/umount -- /var/lib/ceph/tmp/mnt.EKR3Pp
[osd01][WARNIN] Traceback (most recent call last):
[osd01][WARNIN]   File "/usr/sbin/ceph-disk", line 9, in <module>
[osd01][WARNIN]     load_entry_point('ceph-disk==1.0.0', 'console_scripts', 'ceph-disk')()
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 5230, in run
[osd01][WARNIN]     main(sys.argv[1:])
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 5181, in main
[osd01][WARNIN]     args.func(args)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 3532, in main_activate
[osd01][WARNIN]     reactivate=args.reactivate,
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 3289, in mount_activate
[osd01][WARNIN]     (osd_id, cluster) = activate(path, activate_key_template, init)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 3465, in activate
[osd01][WARNIN]     keyring=keyring,
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 2927, in mkfs
[osd01][WARNIN]     '--setgroup', get_ceph_group(),
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 2874, in ceph_osd_mkfs
[osd01][WARNIN]     raise Error('%s failed : %s' % (str(arguments), error))
[osd01][WARNIN] ceph_disk.main.Error: Error: ['ceph-osc', '--cluster', 'ceph', '--mkfs', '--mkkey', '-i', u'1'
, '--monmap', '/var/lib/ceph/tmp/mnt.EKR3Pp/activate.monmap', '--osd-data', '/var/lib/ceph/tmp/mnt.EKR3Pp', '-
osd-journal', '/var/lib/ceph/tmp/mnt.EKR3Pp/journal', '--osd-uuid', u'924f6191-152c-4d1d-b920-9e7237467573',
'--keyring', '/var/lib/ceph/tmp/mnt.EKR3Pp/keyring', '--setuser', 'ceph', '--setgroup', 'ceph'] failed : 2017-
04-04 09:10:58.437244 7f045872c940 -1 filestore(/var/lib/ceph/tmp/mnt.EKR3Pp) mkjournal error creating journal
on /var/lib/ceph/tmp/mnt.EKR3Pp/journal: (13) Permission denied
[osd01][WARNIN] 2017-04-04 09:10:58.437284 7f045872c940 -1 OSD::mkfs: ObjectStore::mkfs failed with error -13
[osd01][WARNIN] 2017-04-04 09:10:58.437395 7f045872c940 -1 ** ERROR: error creating empty object store in /va
r/lib/ceph/tmp/mnt.EKR3Pp: (13) Permission denied
[osd01][WARNIN]
[osd01][ERROR ] RuntimeError: command returned non-zero exit status: 1
[ceph_deploy][ERROR ] RuntimeError: Failed to execute command: /usr/sbin/ceph-disk -v activate --mark-init sys
temd --mount /dev/mapper/mpathal

```

More issues:

```
[ceph@mon01 ceph-config]$ ceph-deploy osd prepare osd01:mapper/mpathc:mapper/mpatha
[ceph_deploy.conf][DEBUG ] found configuration file at: /home/ceph/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (1.5.37): /bin/ceph-deploy osd prepare osd01:mapper/mpathc:mapper/mpatha
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username : None
[ceph_deploy.cli][INFO ] disk : [('osd01', '/dev/mapper/mpathc', '/dev/mapper/mpatha')]
[ceph_deploy.cli][INFO ] dmccrypt : False
[ceph_deploy.cli][INFO ] verbose : False
[ceph_deploy.cli][INFO ] bluestore : None
[ceph_deploy.cli][INFO ] overwrite_conf : False
[ceph_deploy.cli][INFO ] subcommand : prepare
[ceph_deploy.cli][INFO ] dmccrypt_key_dir : /etc/ceph/dmccrypt-keys
[ceph_deploy.cli][INFO ] quiet : False
[ceph_deploy.cli][INFO ] cd_conf : <ceph_deploy.conf.cephdeploy.Conf instance at 0x175e710>
[ceph_deploy.cli][INFO ] cluster : ceph
[ceph_deploy.cli][INFO ] fs_type : xfs
[ceph_deploy.cli][INFO ] func : <function osd at 0x1751b90>
[ceph_deploy.cli][INFO ] ceph_conf : None
[ceph_deploy.cli][INFO ] default_release : False
[ceph_deploy.cli][INFO ] zap_disk : False
[ceph_deploy.osd][DEBUG ] Preparing cluster ceph disks osd01:/dev/mapper/mpathc:/dev/mapper/mpatha
[osd01][DEBUG ] connection detected need for sudo
[osd01][DEBUG ] connected to host: osd01
[osd01][DEBUG ] detect platform information from remote host
[osd01][DEBUG ] detect machine type
[osd01][DEBUG ] find the location of an executable
[ceph_deploy.osd][INFO ] Distro info: CentOS Linux 7.3.1611 Core
[ceph_deploy.osd][DEBUG ] Deploying osd to osd01
[osd01][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[osd01][WARNIN] osd keyring does not exist yet, creating one
[osd01][DEBUG ] create a keyring file
[ceph_deploy.osd][DEBUG ] Preparing host osd01 disk /dev/mapper/mpathc journal /dev/mapper/mpatha activate False
[osd01][DEBUG ] find the location of an executable
[osd01][INFO ] Running command: sudo /usr/sbin/ceph-disk -v prepare --cluster ceph --fs-type xfs -- /dev/mapper/mpathc /dev/mapper/mpatha
[osd01][WARNIN] command: Running command: /usr/bin/ceph-osd --cluster=ceph --show-config-value=fsid
[osd01][WARNIN] command: Running command: /usr/bin/ceph-osd --check-allows-journal -i 0 --cluster ceph --setuser ceph --setgroup ceph
[osd01][WARNIN] command: Running command: /usr/bin/ceph-osd --check-wants-journal -i 0 --cluster ceph --setuser ceph --setgroup ceph
[osd01][WARNIN] command: Running command: /usr/bin/ceph-osd --check-needs-journal -i 0 --cluster ceph --setuser ceph --setgroup ceph
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid is mpath-360060e80074e840000304e8400004002
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid is mpath-360060e80074e840000304e8400004002
[osd01][WARNIN]
[osd01][WARNIN] command: Running command: /usr/bin/ceph-osd --cluster=ceph --show-config-value=osd_journal_size
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid is mpath-360060e80074e840000304e8400004002
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid is mpath-360060e80074e840000304e8400004002
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpathc uuid is mpath-360060e80074e840000304e8400004002
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/dm-4 uuid path is /sys/dev/block/253:4/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/dm-4 uuid is mpath-360060e80074e840000304e8400004002
[osd01][WARNIN]
[osd01][WARNIN] command: Running command: /usr/bin/ceph-conf --cluster=ceph --name=osd. --lookup osd_mkfs_options_xfs
[osd01][WARNIN] command: Running command: /usr/bin/ceph-conf --cluster=ceph --name=osd. --lookup osd_fs_mkfs_options_xfs
[osd01][WARNIN] command: Running command: /usr/bin/ceph-conf --cluster=ceph --name=osd. --lookup osd_mount_opt
```

```

ions_xfs
[osd01][WARNIN] command: Running command: /usr/bin/ceph-conf --cluster=ceph --name=osd. --lookup osd_fs_mount_
options_xfs
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:2/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:2/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] prepare_device: OSD will not be hot-swappable if journal is not the same device as the osd dat
a
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:2/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:2/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] ptype_tobe_for_name: name = journal
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid path is /sys/dev/block/253:2/dm/uuid
[osd01][WARNIN] get_dm_uuid: get_dm_uuid /dev/mapper/mpatha uuid is mpath-360060e80074e840000304e8400004000
[osd01][WARNIN]
[osd01][WARNIN] command: Running command: /sbin/parted --machine -- /dev/mapper/mpatha print
[osd01][WARNIN] get_free_partition_index: get_free_partition_index: analyzing BYT;
[osd01][WARNIN] /dev/mapper/mpatha:2199GB:dm:512:512:gpt:Linux device-mapper (multipath)::
[osd01][WARNIN]
[osd01][WARNIN] Traceback (most recent call last):
[osd01][WARNIN]   File "/usr/sbin/ceph-disk", line 9, in <module>
[osd01][WARNIN]     load_entry_point('ceph-disk==1.0.0', 'console_scripts', 'ceph-disk')()
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 5047, in run
[osd01][WARNIN]     main(sys.argv[1:])
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 4998, in main
[osd01][WARNIN]     args.func(args)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 1812, in main
[osd01][WARNIN]     Prepare.factory(args).prepare()
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 1801, in prepare
[osd01][WARNIN]     self.prepare_locked()
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 1832, in prepare_locked
[osd01][WARNIN]     self.data.prepare(self.journal)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 2494, in prepare
[osd01][WARNIN]     self.prepare_device(*to_prepare_list)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 2670, in prepare_device
[osd01][WARNIN]     to_prepare.prepare()
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 2003, in prepare
[osd01][WARNIN]     self.prepare_device()
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 2095, in prepare_device
[osd01][WARNIN]     num=num)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 1531, in create_partition
[osd01][WARNIN]     num = get_free_partition_index(dev=self.path)
[osd01][WARNIN]   File "/usr/lib/python2.7/site-packages/ceph_disk/main.py", line 1380, in get_free_partition_
index
[osd01][WARNIN]     raise Error('parted output expected to contain ' + dev + ': ' + lines)
[osd01][WARNIN] ceph_disk.main.Error: Error: parted output expected to contain /dev/mapper/mpatha: BYT;
[osd01][WARNIN] /dev/mapper/mpatha:2199GB:dm:512:512:gpt:Linux device-mapper (multipath)::
[osd01][WARNIN]
[osd01][ERROR ] RuntimeError: command returned non-zero exit status: 1
[ceph_deploy.osd][ERROR ] Failed to execute command: /usr/sbin/ceph-disk -v prepare --cluster ceph --fs-type x
fs -- /dev/mapper/mpathc /dev/mapper/mpatha
[ceph_deploy][ERROR ] GenericError: Failed to create 1 OSDs

```

#3 - 04/04/2017 08:26 PM - Nathan Cutler

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

#4 - 04/04/2017 08:26 PM - Nathan Cutler

- Related to Feature #11881: ceph-disk support for multipath added

#5 - 04/04/2017 08:31 PM - Loic Dachary

- Project changed from devops to Ceph
- Subject changed from Failing to activate osd to ceph-disk: failing to activate osd with multipath

#6 - 04/04/2017 08:37 PM - Loic Dachary

Are you able to workaround the problem with chown ceph:ceph on the relevant device ? I'm not suggesting this is the right thing to do, just trying to assert where the problem is ;-)

#7 - 04/04/2017 09:48 PM - Matt Stroud

Loic Dachary wrote:

Are you able to workaround the problem with chown ceph:ceph on the relevant device ? I'm not suggesting this is the right thing to do, just trying to assert where the problem is ;-)

The fix I came up with was in /lib/udev/rules.d/60-ceph-by-parttypeuuid.rules line number 29 I changed the following:

```
ENV{ID_PART_ENTRY_SCHEME}=="gpt", ENV{ID_PART_ENTRY_TYPE}=="?*", ENV{ID_PART_ENTRY_UUID}=="?*", SYMLINK+="disk/by-parttypeuuid/${env{ID_PART_ENTRY_TYPE}}.${env{ID_PART_ENTRY_UUID}}"
```

to

```
ENV{ID_PART_ENTRY_SCHEME}=="gpt", ENV{ID_PART_ENTRY_TYPE}=="?*", ENV{ID_PART_ENTRY_UUID}=="?*", SYMLINK+="disk/by-parttypeuuid/${env{ID_PART_ENTRY_TYPE}}.${env{ID_PART_ENTRY_UUID}}", SYMLINK+="disk/by-partuuid/${env{ID_PART_ENTRY_UUID}}"
```

#8 - 04/17/2017 02:38 PM - Matt Stroud

Any updates to this? I wouldn't expect this kind of behavior for a fresh install.

#9 - 07/04/2017 07:14 PM - David Disseldorp

- Assignee set to David Disseldorp

Matt Stroud wrote:

Loic Dachary wrote:

Are you able to workaround the problem with chown ceph:ceph on the relevant device ? I'm not suggesting this is the right thing to do, just trying to assert where the problem is ;-)

The fix I came up with was in /lib/udev/rules.d/60-ceph-by-parttypeuuuid.rules line number 29 I changed the following:

```
[...]  
to  
[...]
```

Thanks for the patch, Matt. I think the by-partuuid symlink should be created by 60-persistent-storage.rules, which is part of upstream udev. Please try the following change to 60-persistent-storage.rules instead:

```
--- 60-persistent-storage.rules 2017-07-04 21:07:57.136282571 +0200  
+++ 60-persistent-storage.rules.new 2017-07-04 21:08:52.168990967 +0200  
@@ -6,7 +6,7 @@  
 ACTION=="remove", GOTO="persistent_storage_end"  
  
 SUBSYSTEM!="block", GOTO="persistent_storage_end"  
-KERNEL!="loop*|mmcblk*[0-9]|msblk*[0-9]|mspblk*[0-9]|nvme*|hd*|sd*|sr*|vd*|xvd*|bcache*|cciss*|dasd*|ubd*|scm  
*|pmem*", GOTO="persistent_storage_end"  
+KERNEL!="loop*|mmcblk*[0-9]|msblk*[0-9]|mspblk*[0-9]|nvme*|hd*|sd*|sr*|vd*|xvd*|bcache*|cciss*|dasd*|ubd*|scm  
*|pmem*|dm*", GOTO="persistent_storage_end"  
  
# ignore partitions that span the entire disk  
TEST=="whole_disk", GOTO="persistent_storage_end"
```

Works for me... That said, there are still a couple of other minor issues that I'm still chasing up:

- the sd devices that correspond to individual paths for the dm device are triggering the ceph-disk activate udev rules
- + for the case of the journal, even ceph-disk suppress-activate still results in an error

#10 - 07/05/2017 03:27 PM - David Disseldorp

David Disseldorp wrote:

...

Thanks for the patch, Matt. I think the by-partuuid symlink should be created by 60-persistent-storage.rules, which is part of upstream udev. Please try the following change to 60-persistent-storage.rules instead:

[...]

Works for me... That said, there are still a couple of other minor issues that I'm still chasing up:

- the sd devices that correspond to individual paths for the dm device are triggering the ceph-disk activate udev rules
- + for the case of the journal, even ceph-disk suppress-activate still results in an error

Following further local testing, I submitted the udev 60-persistent-storage.rules fix upstream (<https://lists.freedesktop.org/archives/systemd-devel/2017-July/039167.html>), and proposed a fix for the ceph-disk suppress-activate journal issue via <https://github.com/ceph/ceph/pull/16123> .

#11 - 07/10/2017 11:40 AM - David Disseldorp

David Disseldorp wrote:

...

Following further local testing, I submitted the udev 60-persistent-storage.rules fix upstream (<https://lists.freedesktop.org/archives/systemd-devel/2017-July/039167.html>), and proposed a fix for the ceph-disk suppress-activate journal issue via <https://github.com/ceph/ceph/pull/16123> .

The udev dm by-partuuid symlink functionality has gone into upstream 13-dm-disk.rules (provided by lvm2/device-mapper), rather than 60-persistent-storage.rules:

<https://sourceware.org/git/?p=lvm2.git;a=commit;h=c48149cf80c6582c2369bc7f8a33d794021d9dae>

IMO this bug can be closed once <https://github.com/ceph/ceph/pull/16123> has been merged.

#12 - 07/10/2017 07:16 PM - Nathan Cutler

- Status changed from New to Need Review
- Target version deleted (v10.2.7)

- Backport set to jewel

David, I guess it might take some time before distros (especially the stable versions like Trusty) roll out the updated 13-dm-disk.rules . . . is this safe to backport to jewel immediately?

#13 - 07/10/2017 11:11 PM - David Disseldorp

Nathan Cutler wrote:

David, I guess it might take some time before distros (especially the stable versions like Trusty) roll out the updated 13-dm-disk.rules . . . is this safe to backport to jewel immediately?

<https://github.com/ceph/ceph/pull/16123> is an independent change. It's safe to go in without the udev changes.

#14 - 07/26/2017 04:13 AM - Kefu Chai

- Status changed from Need Review to Resolved

instead of <https://lists.freedesktop.org/archives/systemd-devel/2017-July/039167.html>, the lvm2 upstream has a fix <https://sourceware.org/git/?p=lvm2.git;a=commit;h=c48149cf80c6582c2369bc7f8a33d794021d9dae>. so we should just wait until the lvm2 change is included by the downstream distros.

#15 - 07/26/2017 04:13 AM - Kefu Chai

- Status changed from Resolved to Pending Backport

- Assignee deleted (David Disseldorp)

#16 - 07/28/2017 06:18 AM - Nathan Cutler

- Copied to Backport #20837: jewel: ceph-disk: failing to activate osd with multipath added

#17 - 09/07/2017 08:22 AM - Nathan Cutler

- Status changed from Pending Backport to Resolved