

Ceph - Bug #21498

ceph-disk omits "--runtime" when enabling ceph-osd@\$ID.service units for device-backed OSDs

09/22/2017 03:30 AM - Carl Xiong

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category: ceph cli	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport: jewel,luminous	Pull request ID:
Regression: Yes	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	

Description

Ceph version 10.2.9

In a new deployment, after rebooting the nodes following error can be found in ceph-osd log files

```
/var/log/ceph/ceph-osd.17.log:2017-08-01 08:52:05.824620 7f77c42b8b80 -1 ** ERROR: unable to open
OSD superblock on /var/lib/ceph/osd/ceph-17: (2) No such file or directory
/var/log/ceph/ceph-osd.3.log:2017-08-01 08:52:05.824323 7f46e90efb80 -1 ** ERROR: unable to open
OSD superblock on /var/lib/ceph/osd/ceph-3: (2) No such file or directory
/var/log/ceph/ceph-osd.5.log:2017-08-01 08:52:05.823254 7fcb53149b80 -1 ** ERROR: unable to open
OSD superblock on /var/lib/ceph/osd/ceph-5: (2) No such file or directory
/var/log/ceph/ceph-osd.9.log:2017-08-01 08:52:05.820776 7f86322bab80 -1 ** ERROR: unable to open
OSD superblock on /var/lib/ceph/osd/ceph-9: (2) No such file or directory
```

It turns out that ceph-disk is creating ceph-osd.target "wants" symlinks under /etc/systemd instead of /run/systemd:

```
# ls -l /etc/systemd/system/ceph-osd.target.wants
total 0
lrwxrwxrwx 1 root root 41 Sep 18 14:26 ceph-osd@0.service -> /usr/lib/systemd/system/ceph-osd@.ser
vice
lrwxrwxrwx 1 root root 41 Sep 18 14:26 ceph-osd@1.service -> /usr/lib/systemd/system/ceph-osd@.ser
vice
lrwxrwxrwx 1 root root 41 Sep 18 14:27 ceph-osd@2.service -> /usr/lib/systemd/system/ceph-osd@.ser
vice
# ls -l /run/systemd/system/ceph-osd.target.wants
ls: cannot access '/run/systemd/system/ceph-osd.target.wants': No such file or directory.
```

After some debugging, I believe the commit <https://github.com/ceph/ceph/commit/9b77b16b888b8efbf7d50d333e1880a6ec70d87a>, has introduced a regression by using 'is_mounted' to check whether a path is mounted in 'systemd_start', while 'is_mounted' is supposedly used to check a device.

A potential fix might be <https://github.com/SUSE/ceph/pull/151>

Related issues:

Copied to Ceph - Backport #21504: ceph-disk omits "--runtime" when enabling c...	Resolved
Copied to Ceph - Backport #21522: jewel: ceph-disk omits "--runtime" when ena...	Resolved

History

#1 - 09/22/2017 04:18 AM - Carl Xiong

<https://github.com/ceph/ceph/pull/17904>

#2 - 09/22/2017 05:28 AM - Kefu Chai

- Description updated

#3 - 09/22/2017 05:57 AM - Kefu Chai

- Status changed from New to Fix Under Review

- Backport set to jewel,luminous

#4 - 09/22/2017 03:08 PM - Kefu Chai

- Status changed from Fix Under Review to Pending Backport

#5 - 09/22/2017 03:16 PM - Kefu Chai

- Copied to Backport #21504: ceph-disk omits "--runtime" when enabling ceph-osd@\$ID.service units for device-backed OSDs added

#6 - 09/22/2017 08:33 PM - Nathan Cutler

- Copied to Backport #21522: jewel: ceph-disk omits "--runtime" when enabling ceph-osd@\$ID.service units for device-backed OSDs added

#7 - 10/11/2017 02:38 AM - Kefu Chai

- Status changed from Pending Backport to Resolved