

fs - Bug #40182

luminous: pybind: luminous volume client breaks against nautilus cluster

06/06/2019 07:22 AM - Jan Fajerski

Status:	Resolved	Start date:	
Priority:	High	Due date:	
Assignee:	Jan Fajerski	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v12.2.13	Affected Versions:	
Source:	Community (dev)	ceph-qa-suite:	
Tags:		Component(FS):	VolumeClient
Backport:		Labels (FS):	Manila
Regression:	No	Pull request ID:	28445
Severity:	3 - minor		
Reviewed:			

Description

Due to the removal of the 'ceph mds dump' command in nautilus, a luminous ceph_volume_client does not work against a nautilus cluster. This breaks some version combinations of openstack cloud and ceph.

Here's a log fragment of manila:

```
2019-05-23 09:56:50.763 INFO manila.share.drivers.cephfs.driver [req-34c1c009-cf00-48e8-ab3a-e19ea4bc8df8 None None] [CEPHFSNATIVE1] Ceph client found, connecting...
2019-05-23 09:56:50.802 INFO ceph_volume_client [req-34c1c009-cf00-48e8-ab3a-e19ea4bc8df8 None None] evict clients with auth_name=manila
2019-05-23 09:56:50.872 ERROR manila.share.manager [req-34c1c009-cf00-48e8-ab3a-e19ea4bc8df8 None None] Error encountered during initialization of driver CephFSDriver@192.168.24.218@cephfsnative1:
Error: command is obsolete; please check usage and/or man page
2019-05-23 09:56:50.872 TRACE manila.share.manager Traceback (most recent call last):
2019-05-23 09:56:50.872 TRACE manila.share.manager File "/usr/lib/python2.7/site-packages/manila/share/manager.py", line 305, in _driver_setup
2019-05-23 09:56:50.872 TRACE manila.share.manager self.driver.do_setup(ctxt)
2019-05-23 09:56:50.872 TRACE manila.share.manager File "/usr/lib/python2.7/site-packages/manila/share/drivers/cephfs/driver.py", line 144, in do_setup
2019-05-23 09:56:50.872 TRACE manila.share.manager ceph_vol_client=self.volume_client)
2019-05-23 09:56:50.872 TRACE manila.share.manager File "/usr/lib/python2.7/site-packages/manila/share/drivers/cephfs/driver.py", line 216, in volume_client
2019-05-23 09:56:50.872 TRACE manila.share.manager self._volume_client.connect(premount_evict=premount_evict)
2019-05-23 09:56:50.872 TRACE manila.share.manager File "/usr/lib/python2.7/site-packages/ceph_volume_client.py", line 474, in connect
2019-05-23 09:56:50.872 TRACE manila.share.manager self.evict(premount_evict)
2019-05-23 09:56:50.872 TRACE manila.share.manager File "/usr/lib/python2.7/site-packages/ceph_volume_client.py", line 399, in evict
2019-05-23 09:56:50.872 TRACE manila.share.manager mds_map = self._rados_command("mds dump", {
})
2019-05-23 09:56:50.872 TRACE manila.share.manager File "/usr/lib/python2.7/site-packages/ceph_volume_client.py", line 1340, in _rados_command
2019-05-23 09:56:50.872 TRACE manila.share.manager raise rados.Error(outs)
2019-05-23 09:56:50.872 TRACE manila.share.manager Error: command is obsolete; please check usage and/or man page
```

https://github.com/ceph/ceph/compare/luminous...jan-f:ceph_volume_client-version-aware implements a potential fix.

Another option would be to simply replace the 'mds dump' calls with 'fs dump', since the latter is present in luminous. Not sure if that breaks something else though.

Related issues:

Related to fs - Bug #22524: NameError: global name 'get_mds_map' is not defined	Resolved	12/21/2017
Related to fs - Bug #22038: ceph-volume-client: rados.Error: command not known	Resolved	11/03/2017

History**#1 - 06/06/2019 05:21 PM - Patrick Donnelly**

- Subject changed from *luminous volume client breaks against nautilus cluster* to *pybind: luminous volume client breaks against nautilus cluster*
- Assignee set to *Rishabh Dave*
- Priority changed from *Normal* to *High*
- Target version set to *v15.0.0*
- Start date deleted (*06/06/2019*)
- Source set to *Community (dev)*
- Backport set to *nautilus,mimic,luminous*

I think adopting `fs dump` instead of `mds dump` is the right thing to do.

#2 - 06/07/2019 07:20 AM - Jan Fajerski

proposed fix: <https://github.com/ceph/ceph/pull/28445>

#3 - 06/07/2019 07:17 PM - Patrick Donnelly

- Subject changed from *pybind: luminous volume client breaks against nautilus cluster* to *luminous: pybind: luminous volume client breaks against nautilus cluster*
- Status changed from *New* to *Need Review*
- Assignee changed from *Rishabh Dave* to *Jan Fajerski*
- Target version changed from *v15.0.0* to *v12.2.13*
- Backport deleted (*nautilus,mimic,luminous*)
- Pull request ID set to *28445*

Let's treat this as a backport. Please cherry-pick the commits from here:

<https://github.com/ceph/ceph/pull/17266/commits/cbbdd0da7d40e4e5def5cc0b9a9250348e71019f>
<https://github.com/ceph/ceph/pull/19633>

#4 - 06/11/2019 09:56 AM - Jan Fajerski

Patrick Donnelly wrote:

Let's treat this as a backport. Please cherry-pick the commits from here:

<https://github.com/ceph/ceph/pull/17266/commits/cbbdd0da7d40e4e5def5cc0b9a9250348e71019f>
<https://github.com/ceph/ceph/pull/19633>

This won't be sufficient as <https://github.com/ceph/ceph/pull/17266/commits/cbbdd0da7d40e4e5def5cc0b9a9250348e71019f> leaves the ceph fs `add_data_pool` call broken. I'm happy to add all necessary backports, but imho this is needlessly complex. Hence the simple commit.

#5 - 06/11/2019 06:04 PM - Patrick Donnelly

Jan Fajerski wrote:

Patrick Donnelly wrote:

Let's treat this as a backport. Please cherry-pick the commits from here:

<https://github.com/ceph/ceph/pull/17266/commits/cbbdd0da7d40e4e5def5cc0b9a9250348e71019f>

<https://github.com/ceph/ceph/pull/19633>

This won't be sufficient as <https://github.com/ceph/ceph/pull/17266/commits/cbbdd0da7d40e4e5def5cc0b9a9250348e71019f> leaves the ceph fs add_data_pool call broken. I'm happy to add all necessary backports, but imho this is needlessly complex. Hence the simple commit.

I understand it's more complex but we have rules about backports so it's possible to check commit history to see if something has been backported. Please just add this ticket too:

<http://tracker.ceph.com/issues/22038>

Are there any other issues?

#6 - 06/12/2019 01:05 PM - Jan Fajerski

Patrick Donnelly wrote:

Are there any other issues?

A couple more. PR is updated.

#7 - 06/12/2019 01:18 PM - Nathan Cutler

- Related to Bug #22524: *NameError: global name 'get_mds_map' is not defined* added

#8 - 06/12/2019 01:18 PM - Nathan Cutler

- Related to Bug #22038: ceph-volume-client: rados.Error: command not known added

#9 - 06/25/2019 04:25 PM - Yuri Weinstein

Jan Fajerski wrote:

proposed fix: <https://github.com/ceph/ceph/pull/28445>

merged

#10 - 06/26/2019 05:07 AM - Patrick Donnelly

- Status changed from Need Review to Resolved