

CephFS - Bug #44276

pybind/mgr/volumes: cleanup stale connection hang

02/24/2020 06:58 PM - Patrick Donnelly

Status:	Resolved	% Done:	0%
Priority:	Normal		
Assignee:	Venky Shankar		
Category:			
Target version:	v16.0.0		
Source:	Development	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	octopus,nautilus	Component(FS):	mgr/volumes
Regression:	No	Labels (FS):	
Severity:	3 - minor	Pull request ID:	35496
Reviewed:		Crash signature:	

Description

```
2020-02-24 13:59:12.933 7fa5f7aa5700 4 mgr[volumes] scanning for idle connections..
2020-02-24 13:59:12.933 7fa5f7aa5700 4 mgr[volumes] cleaning up connection for 'cephfs'
2020-02-24 13:59:12.933 7fa5f7aa5700 20 mgr[volumes] self.fs_id=43, fs_id=43
2020-02-24 13:59:12.933 7fa5f7aa5700 4 mgr[volumes] disconnecting from cephfs 'cephfs'
2020-02-24 13:59:12.933 7fa5f7aa5700 1 -- 172.21.15.16:0/1494282980 --> [v2:172.21.15.61:6832/333
244281,v1:172.21.15.61:6833/333244281] -- client_session(request_close seq 11) v3 -- 0x560b4af0dd4
0 con 0x560b4b415400
```

From:
/ceph/teuthology-archive/vshankar-2020-02-24_12:33:54-fs-wip-vshankar-testing-testing-basic-smithi/4798102/remote/smithi016/log/ceph-mgr.x.log.gz

We don't see 7fa5f7aa5700 again for the rest of the log. It's stuck in disconnecting the handle or an exception (as Venky mentioned just in comment 2).

Later on, the subvolume group command hangs:

```
2020-02-24 13:59:31.481 7fa6032fc700 1 -- [v2:172.21.15.16:6800/25634,v1:172.21.15.16:6801/25634]
<== client.31354 172.21.15.61:0/1231806590 1 ==== command(tid 0: {"prefix": "fs subvolume create"
, "vol_name": "cephfs", "sub_name": "subvolume_6955", "target": ["mgr", ""], "size": 20971520}) v1
==== 150+0+0 (crc 0 0 0) 0x560b4ae290e0 con 0x560b4cd1d800
2020-02-24 13:59:31.481 7fa6032fc700 4 mgr.server _handle_command decoded 5
2020-02-24 13:59:31.481 7fa6032fc700 4 mgr.server _handle_command prefix=fs subvolume create
2020-02-24 13:59:31.482 7fa6032fc700 20 is_capable service=py module=volumes command=fs subvolume
create read write addr - on cap allow *
2020-02-24 13:59:31.482 7fa6032fc700 20 allow so far , doing grant allow *
2020-02-24 13:59:31.482 7fa6032fc700 20 allow all
2020-02-24 13:59:31.482 7fa6032fc700 10 mgr.server _allowed_command client.admin capable
2020-02-24 13:59:31.482 7fa6032fc700 0 log_channel(audit) log [DBG] : from='client.31354 -' entit
y='client.admin' cmd=[{"prefix": "fs subvolume create", "vol_name": "cephfs", "sub_name": "subvolu
me_6955", "target": ["mgr", ""], "size": 20971520}]: dispatch
2020-02-24 13:59:31.482 7fa6032fc700 10 mgr.server _handle_command passing through 5
2020-02-24 13:59:31.482 7fa603afd700 20 mgr Gil Switched to new thread state 0x560b4b16a2c0
```

This is testing with the backport of <https://github.com/ceph/ceph/pull/33413>

Related issues:

Related to CephFS - Bug #44207: mgr/volumes: deadlock when trying to purge la...	Resolved
Duplicated by CephFS - Bug #44281: pybind/mgr/volumes: cleanup stale connecti...	Duplicate
Copied to CephFS - Backport #46388: nautilus: pybind/mgr/volumes: cleanup sta...	Resolved
Copied to CephFS - Backport #46389: octopus: pybind/mgr/volumes: cleanup stal...	Resolved

History**#1 - 02/24/2020 06:58 PM - Patrick Donnelly**

- Project changed from mgr to CephFS
- Status changed from New to In Progress
- Component(FS) mgr/volumes added

#2 - 02/25/2020 03:52 AM - Venky Shankar

Patrick Donnelly wrote:

[...]

From:

/ceph/teuthology-archive/vshankar-2020-02-24_12:33:54-fs-wip-vshankar-testing-testing-basic-smithi/4798102/remote/smithi016/log/ceph-mgr.x.log.gz

This is testing with the backport of <https://github.com/ceph/ceph/pull/33413>

the perceived hang may be due to an unhandled (+ not logged) exception. I'll take a look.

#3 - 02/25/2020 03:57 AM - Patrick Donnelly

- Description updated

#4 - 02/25/2020 03:59 AM - Patrick Donnelly

- Subject changed from nautilus: pybind/mgr/volumes: apparent deadlock to pybind/mgr/volumes: cleanup stale connection hang
- Target version changed from v14.2.8 to v15.0.0
- Backport set to nautilus

Copying from [#44281](#):

```
2020-02-24T20:00:50.686+0000 7fa06a524700 0 [volumes] [INFO] [volumes.fs.operations.volume] scanning for idle
connections..
2020-02-24T20:00:50.686+0000 7fa06a524700 0 [volumes] [INFO] [volumes.fs.operations.volume] cleaning up conne
ction for 'cephfs'
2020-02-24T20:00:50.686+0000 7fa06a524700 0 [volumes] [DEBUG] [volumes.fs.operations.volume] self.fs_id=50, f
s_id=50
2020-02-24T20:00:50.686+0000 7fa06a524700 0 [volumes] [INFO] [volumes.fs.operations.volume] disconnecting fro
m cephfs 'cephfs'
2020-02-24T20:00:50.686+0000 7fa06a524700 1 -- 172.21.15.172:0/1073184588 --> [v2:172.21.15.31:6832/416293353
3,v1:172.21.15.31:6833/4162933533] -- client_session(request_close seq 538) v4 -- 0x55f6d16dff80 con 0x55f6d18
cac00
```

From:
/ceph/teuthology-archive/vshankar-2020-02-24_18:41:47-fs-wip-vshankar-testing-2020-02-24-102202-testing-basic-smithi/4798643/remote/smithi172/
log/ceph-mgr.y.log.gz

See also other failures in: http://pulpito.ceph.com/vshankar-2020-02-24_18:41:47-fs-wip-vshankar-testing-2020-02-24-102202-testing-basic-smithi/

Because this cleanup method holds the ConnectionPool lock, no other subvolume method will make progress.

#5 - 02/25/2020 04:00 AM - Patrick Donnelly

- Duplicated by Bug #44281: pybind/mgr/volumes: cleanup stale connection hang added

#6 - 02/25/2020 04:00 AM - Patrick Donnelly

- Related to Bug #44207: mgr/volumes: deadlock when trying to purge large number of trash entries added

#7 - 02/25/2020 02:29 PM - Venky Shankar

It's the connection cleanup thread that's waiting on cephfs shutdown() call after initiating a session disconnect:

```
2020-02-25 09:24:51.854 7f85ed4bd700 4 mgr[volumes] disconnecting from cephfs 'cephfs'
2020-02-25 09:24:51.854 7f85ed4bd700 1 -- 172.21.15.121:0/959936343 --> [v2:172.21.15.121:6834/987573718,v1:172.21.15.121:6835/987573718] -- client_session(request_close seq 11) v3 -- 0x56404a91bd40 con 0x56404ac9bc00
```

Ideally, there should have been messages exchanged between the manager and mds for connection termination:

```
2020-02-25 09:08:51.835 7f85ed4bd700 4 mgr[volumes] disconnecting from cephfs 'cephfs'
2020-02-25 09:08:51.835 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 --> [v2:172.21.15.121:6834/1785088858,v1:172.21.15.121:6835/1785088858] -- client_session(request_close seq 120) v3 -- 0x56404ae68480 con 0x56405498ec00
2020-02-25 09:08:51.839 7f859dbde700 1 -- 172.21.15.121:0/593434532 <== mds.0 v2:172.21.15.121:6834/1785088858 314 ==== client_session(close) v1 ==== 28+0+0 (crc 0 0 0) 0x56404ae68480 con 0x56405498ec00
2020-02-25 09:08:51.839 7f859dbde700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6834/1785088858,v1:172.21.15.121:6835/1785088858] conn(0x56405498ec00 msgr2=0x56405495f600 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=0).mark_down
2020-02-25 09:08:51.839 7f859dbde700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6834/1785088858,v1:172.21.15.121:6835/1785088858] conn(0x56405498ec00 0x56405495f600 crc :-1 s=READY pgs=16 cs=0 l=0 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.46:6808/26678,v1:172.21.15.46:6809/26678] conn(0x56405498f000 msgr2=0x564054960100 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.46:6808/26678,v1:172.21.15.46:6809/26678] conn(0x56405498f000 0x564054960100 crc :-1 s=READY pgs=70 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.46:6800/26680,v1:172.21.15.46:6801/26680] conn(0x56404cde5c00 msgr2=0x564054979700 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.46:6800/26680,v1:172.21.15.46:6801/26680] conn(0x56404cde5c00 0x564054979700 crc :-1 s=READY pgs=56 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.46:6816/26679,v1:172.21.15.46:6817/26679] conn(0x56404ce3fc00 msgr2=0x564054978100 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.46:6816/26679,v1:172.21.15.46:6817/26679] conn(0x56404ce3fc00 0x564054978100 crc :-1 s=READY pgs=56 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f86161a6700 1 -- 172.21.15.121:0/593434532 reap_dead start
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6802/27401,v1:172.21.15.121:6803/27401] conn(0x56404cd54800 msgr2=0x564054979180 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6802/27401,v1:172.21.15.121:6803/27401] conn(0x56404cd54800 0x564054979180 crc :-1 s=READY pgs=62 cs=0 l=1 rx=0 tx=0).stop
```

```

2020-02-25 09:08:51.840 7f86161a6700 1 -- 172.21.15.121:0/593434532 reap_dead start
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6810/27402,v1:172.21.15.121:6811/27402] conn(0x564054962800 msgr2=0x564054a04000 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6810/27402,v1:172.21.15.121:6811/27402] conn(0x564054962800 0x564054a04000 crc :-1 s=READY pgs=60 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6817/27403,v1:172.21.15.121:6821/27403] conn(0x56404c74b000 msgr2=0x564054978c00 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6817/27403,v1:172.21.15.121:6821/27403] conn(0x56404c74b000 0x564054978c00 crc :-1 s=READY pgs=59 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6813/27404,v1:172.21.15.121:6815/27404] conn(0x56404ce3e000 msgr2=0x564054978680 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6813/27404,v1:172.21.15.121:6815/27404] conn(0x56404ce3e000 0x564054978680 crc :-1 s=READY pgs=56 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:3300/0,v1:172.21.15.121:6789/0] conn(0x5640556c2400 msgr2=0x56405495fb80 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=1).mark_down
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:3300/0,v1:172.21.15.121:6789/0] conn(0x5640556c2400 0x56405495fb80 crc :-1 s=READY pgs=3142 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 shutdown_connections
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6810/27402,v1:172.21.15.121:6811/27402] conn(0x564054962800 0x564054a04000 unknown :-1 s=CLOSED pgs=60 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6813/27404,v1:172.21.15.121:6815/27404] conn(0x56404ce3e000 0x564054978680 unknown :-1 s=CLOSED pgs=56 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:3300/0,v1:172.21.15.121:6789/0] conn(0x5640556c2400 0x56405495fb80 unknown :-1 s=CLOSED pgs=3142 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 --2- 172.21.15.121:0/593434532 >> [v2:172.21.15.121:6817/27403,v1:172.21.15.121:6821/27403] conn(0x56404c74b000 0x564054978c00 unknown :-1 s=CLOSED pgs=59 cs=0 l=1 rx=0 tx=0).stop
2020-02-25 09:08:51.840 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 shutdown_connections
2020-02-25 09:08:51.841 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 wait complete.
2020-02-25 09:08:51.841 7f85ed4bd700 1 -- 172.21.15.121:0/593434532 >> 172.21.15.121:0/593434532 conn(0x564054963000 msgr2=0x564054959000 unknown :-1 s=STATE_NONE l=0).mark_down

```

The cleanup thread holds the conn pool lock so commands are blocked until shutdown returns. There's no exception logged and cephfs shutdown is a blocking call, so, pretty much everything in mgr/volumes comes to a standstill.

On the MDS side, mds.b (to which the client has a session), respawned:
(note that this is *just* before the client initiates session close)

```

2020-02-25 09:24:50.993 7f3582f29700 10 mds.b my gid is 14360
2020-02-25 09:24:50.993 7f3582f29700 10 mds.b map says I am mds.-1.-1 state null
2020-02-25 09:24:50.993 7f3582f29700 10 mds.b msgr says i am [v2:172.21.15.121:6834/987573718,v1:172.21.15.121:6835/987573718]
2020-02-25 09:24:50.993 7f3582f29700 1 mds.b Map removed me (mds.-1 gid:14360) from cluster due to lost contact; respawning
2020-02-25 09:24:50.993 7f3582f29700 1 mds.b respawn!

```

Post respawn, the only log entries for client "172.21.15.121:0/959936343" are below which is just after the dump of events.

```

2020-02-25 09:24:50.993 7f3585f2f700 1 -- [v2:172.21.15.121:6834/987573718,v1:172.21.15.121:6835/987573718] >
> 172.21.15.121:0/959936343 conn(0x5613c0689400 msgr2=0x5613c0451700 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=
0).read_bulk peer close file descriptor 35
2020-02-25 09:24:50.993 7f3585f2f700 1 -- [v2:172.21.15.121:6834/987573718,v1:172.21.15.121:6835/987573718] >
> 172.21.15.121:0/959936343 conn(0x5613c0689400 msgr2=0x5613c0451700 crc :-1 s=STATE_CONNECTION_ESTABLISHED l=
0).read_until read failed
2020-02-25 09:24:50.993 7f3585f2f700 1 --2- [v2:172.21.15.121:6834/987573718,v1:172.21.15.121:6835/987573718]
>> 172.21.15.121:0/959936343 conn(0x5613c0689400 0x5613c0451700 crc :-1 s=READY pgs=4 cs=0 l=0 rx=0 tx=0).han
dle_read_frame_preamble_main read frame length and tag failed r=-1 ((1) Operation not permitted)
2020-02-25 09:24:50.993 7f3585f2f700 1 --2- [v2:172.21.15.121:6834/987573718,v1:172.21.15.121:6835/987573718]
>> 172.21.15.121:0/959936343 conn(0x5613c0689400 0x5613c0451700 unknown :-1 s=READY pgs=4 cs=0 l=0 rx=0 tx=0)
._fault with nothing to send, going to standby

```

mds.b goes through boot sequence boot->standby->creating->active and enters normal operation. mds.b never saw "request_close" message from this client.

#8 - 02/27/2020 04:24 AM - Venky Shankar

Patrick,

This issue is present in master too:

http://pulpito.ceph.com/vshankar-2020-02-26_17:46:30-fs-wip-vshankar-testing-2020-02-26-093932-testing-basic-smithi/

Didn't see in the master run that I ran couple of days ago -- that might be because I scheduled less number of repeat jobs. This is with "-N 10".

So, it's not only in nautilus as I mentioned yesterday. That's good as I don't have to dig into what's in master but not in nautilus.

#9 - 04/16/2020 01:35 PM - Venky Shankar

- *Priority changed from Immediate to Normal*

lowering priority since happens (most I guess) on libcephfs shutdown (in teuthology).

#10 - 05/05/2020 05:16 PM - Patrick Donnelly

- *Target version changed from v15.0.0 to v16.0.0*

#11 - 07/06/2020 11:52 PM - Patrick Donnelly

- *Status changed from In Progress to Pending Backport*

- *Backport changed from nautilus to octopus,nautilus*

- *Pull request ID set to 35496*

#12 - 07/07/2020 11:12 AM - Nathan Cutler

- *Copied to Backport #46388: nautilus: pybind/mgr/volumes: cleanup stale connection hang added*

#13 - 07/07/2020 11:12 AM - Nathan Cutler

- *Copied to Backport #46389: octopus: pybind/mgr/volumes: cleanup stale connection hang added*

#14 - 08/26/2020 09:17 AM - Nathan Cutler

- *Status changed from Pending Backport to Resolved*

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".