

RADOS - Bug #36183

[objecter] client socket failure leads to hung connection

09/25/2018 05:50 PM - Jason Dillaman

Status:	Resolved	Start date:	09/25/2018
Priority:	High	Due date:	
Assignee:	Jason Dillaman	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	luminous,mimic	ceph-qa-suite:	
Regression:	No	Component(RADOS):	Objecter
Severity:	3 - minor	Pull request ID:	

Description

During an rbd-mirror thrash test run, the process failed to shut down cleanly because it was stuck in an librados read operation for over 10 minutes. The OSD logs show that the client attempted to connect but disconnected due to the inject socket failure option. There were no further attempts for the connection to re-connect after the initial failure.

```
2018-09-25 00:09:14.227 7f82794b1700 1 -- 172.21.15.31:6808/12957 >> - conn(0x23a27480 legacy :6808 s=ACCEPTING pgs=0 cs=0 l=0).send_server_banner sd=57 172.21.15.31:33298/0
2018-09-25 00:09:14.227 7f82794b1700 1 -- 172.21.15.31:6808/12957 >> 172.21.15.31:33298/0 conn(0x23a27480 legacy :6808 s=STATE_CONNECTION_ESTABLISHED l=0).read_bulk peer close file descriptor 57
2018-09-25 00:09:14.227 7f82794b1700 1 -- 172.21.15.31:6808/12957 >> 172.21.15.31:33298/0 conn(0x23a27480 legacy :6808 s=STATE_CONNECTION_ESTABLISHED l=0).read_until read failed
2018-09-25 00:09:14.227 7f82794b1700 1 -- 172.21.15.31:6808/12957 >> - conn(0x23a27480 legacy :6808 s=ACCEPTING pgs=0 cs=0 l=0).handle_client_banner read peer banner and addr failed
```

The client log only shows that it injected a socket failure:

```
2018-09-25 00:09:14.227 7efbd8d99700 0 -- 172.21.15.31:47216/4186053723 >> 172.21.15.31:6808/12957 conn(0x2dfa850 legacy :-1 s=STATE_CONNECTION_ESTABLISHED l=1)._try_send injecting socket failure
```

... but there is no corresponding log from Objecter's "ms_handle_reset" callback. In Objecter, adding a "sleep(1)" between "messenger->connect_to_osd(osdmap->get_addrs(osd))," and "s->con->set_priv(RefCountedPtr{s});" can reproduce the issue if the connection is dropped during the handshake.

http://qa-proxy.ceph.com/teuthology/jdillaman-2018-09-24_18:10:36-rbd-wip-rbd-mirror-distro-basic-smithi/3066290/teuthology.log

Related issues:

Copied to RADOS - Backport #36295: luminous: [objecter] client socket failure...	Resolved
Copied to RADOS - Backport #36296: mimic: [objecter] client socket failure le...	Resolved

History

#1 - 09/25/2018 05:51 PM - Jason Dillaman

- Description updated

#2 - 09/25/2018 06:19 PM - Jason Dillaman

- Status changed from In Progress to Need Review

PR: <https://github.com/ceph/ceph/pull/24276>

#3 - 09/25/2018 06:20 PM - Jason Dillaman

- Subject changed from *[async msgr] client socket failure leads to hung connection* to *[objecter] client socket failure leads to hung connection*

#4 - 09/25/2018 06:20 PM - Jason Dillaman

- Component(RADOS) Objecter added

#5 - 10/02/2018 10:46 PM - Neha Ojha

- Status changed from *Need Review* to *Pending Backport*

#6 - 10/03/2018 08:09 AM - Nathan Cutler

- Copied to Backport #36295: *luminous: [objecter] client socket failure leads to hung connection* added

#7 - 10/03/2018 08:09 AM - Nathan Cutler

- Copied to Backport #36296: *mimic: [objecter] client socket failure leads to hung connection* added

#8 - 10/26/2018 02:30 AM - Nathan Cutler

- Status changed from *Pending Backport* to *Resolved*