

CephFS - Bug #46887

kceph: testing branch: hang in workunit by 1/2 clients during tree export

08/10/2020 08:38 PM - Patrick Donnelly

Status:	Need More Info	% Done:	0%
Priority:	Urgent		
Assignee:	Patrick Donnelly		
Category:			
Target version:	v17.0.0		
Source:	Q/A	ceph-qa-suite:	
Tags:		Component(FS):	MDS, kceph
Backport:	pacific,octopus,nautilus	Labels (FS):	qa-failure
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature (v1):	
Reviewed:		Crash signature (v2):	
Affected Versions:			

Description

```
2020-08-08T08:52:51.696 INFO:teuthology.orchestra.run.smithi026:> sudo adjust-ulimits ceph-coverage /home/ubuntu/ceph-test/archive/coverage nsenter --net=/var/run/netns/ceph-ns--home-ubuntu-ceph-test-mnt.1 /bin/mount -t ceph :/ /home/ubuntu/ceph-test/mnt.1 -v -o name=1,norequire_active_mds,conf=/etc/ceph/ceph.conf
2020-08-08T08:52:51.823 INFO:teuthology.orchestra.run.smithi026.stdout:parsing options: rw,name=1,norequire_active_mds,conf=/etc/ceph/ceph.conf
2020-08-08T08:52:51.823 INFO:teuthology.orchestra.run.smithi026.stdout:mount.ceph: options "name=1,norequire_active_mds" will pass to kernel.
2020-08-08T08:52:51.857 INFO:teuthology.orchestra.run.smithi026.stdout:parsing options: rw,name=1,norequire_active_mds,conf=/etc/ceph/ceph.conf
2020-08-08T08:52:51.857 INFO:teuthology.orchestra.run.smithi026.stdout:mount.ceph: options "name=1,norequire_active_mds" will pass to kernel.
2020-08-08T08:52:51.863 INFO:teuthology.orchestra.run:Running command with timeout 300
2020-08-08T08:52:51.864 INFO:teuthology.orchestra.run.smithi026:> sudo chmod 1777 /home/ubuntu/ceph-test/mnt.1
...
2020-08-08T08:54:16.377 INFO:teuthology.orchestra.run.smithi026:workunit test fs/norstats/kernel_untar_tar.sh> mkdir -p -- /home/ubuntu/ceph-test/mnt.1/client.1/tmp && cd -- /home/ubuntu/ceph-test/mnt.1/client.1/tmp && CEPH_CLI_TEST_DUP_COMMAND=1 CEPH_REF=34905f0a75aeb20a442c8aedc05af7a94d2a3b8a TESTDIR="/home/ubuntu/ceph-test" CEPH_ARGS="--cluster ceph" CEPH_ID="1" PATH=$PATH:/usr/sbin CEPH_BASE=/home/ubuntu/ceph-test/clone.client.1 CEPH_ROOT=/home/ubuntu/ceph-test/clone.client.1 adjust-ulimits ceph-coverage /home/ubuntu/ceph-test/archive/coverage timeout 3h /home/ubuntu/ceph-test/clone.client.1/qa/workunits/fs/norstats/kernel_untar_tar.sh
...
2020-08-08T08:54:28.862 INFO:tasks.workunit.client.1.smithi026.stdout:linux-4.0.5/Documentation/zorro.txt
...
2020-08-08T11:54:16.444 DEBUG:teuthology.orchestra.run:got remote process result: 124
2020-08-08T11:54:16.459 INFO:tasks.workunit:Stopping ['fs/norstats'] on client.1...
2020-08-08T11:54:16.460 INFO:teuthology.orchestra.run.smithi026:> sudo rm -rf -- /home/ubuntu/ceph-test/workunits.list.client.1 /home/ubuntu/ceph-test/clone.client.1
2020-08-08T11:54:17.180 ERROR:teuthology.run_tasks:Saw exception from tasks.
Traceback (most recent call last):
  File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/run_tasks.py", line 90
, in run_tasks
    manager = run_one_task(taskname, ctx=ctx, config=config)
  File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/run_tasks.py", line 69
, in run_one_task
    return task(**kwargs)
  File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20200808.001303/qa/tasks/w
```

```

rkunit.py", line 138, in task
    cleanup=cleanup)
File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20200808.001303/qa/tasks/wo
rkunit.py", line 288, in _spawn_on_all_clients
    timeout=timeout)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 84,
in __exit__
    for result in self:
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 98,
in __next__
    resurrect_traceback(result)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 30,
in resurrect_traceback
    raise exc.exc_info[1]
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 23,
in capture_traceback
    return func(*args, **kwargs)
File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20200808.001303/qa/tasks/wo
rkunit.py", line 413, in _run_tests
    label="workunit test {workunit}".format(workunit=workunit)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/remote.py",
line 204, in run
    r = self._runner(client=self.ssh, name=self.shortname, **kwargs)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 446, in run
    r.wait()
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 160, in wait
    self._raise_for_status()
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 182, in _raise_for_status
    node=self.hostname, label=self.label
teuthology.exceptions.CommandFailedError: Command failed (workunit test fs/norstats/kernel_untar_t
ar.sh) on smithi026 with status 124: 'mkdir -p -- /home/ubuntu/cephptest/mnt.1/client.1/tmp && cd -
- /home/ubuntu/cephptest/mnt.1/client.1/tmp && CEPH_CLI_TEST_DUP_COMMAND=1 CEPH_REF=34905f0a75aeb20
a442c8aedc05af7a94d2a3b8a TESTDIR="/home/ubuntu/cephptest" CEPH_ARGS="--cluster ceph" CEPH_ID="1" P
ATH=$PATH:/usr/sbin CEPH_BASE=/home/ubuntu/cephptest/clone.client.1 CEPH_ROOT=/home/ubuntu/cephtest
/clone.client.1 adjust-ulimits ceph-coverage /home/ubuntu/cephptest/archive/coverage timeout 3h /ho
me/ubuntu/cephptest/clone.client.1/qa/workunits/fs/norstats/kernel_untar_tar.sh'
2020-08-08T11:54:17.317 ERROR:teuthology.run_tasks: Sentry event: http://sentry.ceph.com/sepia/teu
thology/?q=79b9d3d9d1db43cc9c9c59952f1da44c
Traceback (most recent call last):
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/run_tasks.py", line 90
, in run_tasks
    manager = run_one_task(taskname, ctx=ctx, config=config)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/run_tasks.py", line 69
, in run_one_task
    return task(**kwargs)
File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20200808.001303/qa/tasks/wo
rkunit.py", line 138, in task
    cleanup=cleanup)
File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20200808.001303/qa/tasks/wo
rkunit.py", line 288, in _spawn_on_all_clients
    timeout=timeout)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 84,
in __exit__
    for result in self:
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 98,
in __next__
    resurrect_traceback(result)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 30,
in resurrect_traceback
    raise exc.exc_info[1]
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 23,
in capture_traceback
    return func(*args, **kwargs)
File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20200808.001303/qa/tasks/wo

```

```
rkunit.py", line 413, in _run_tests
    label="workunit test {workunit}".format(workunit=workunit)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/remote.py",
line 204, in run
    r = self._runner(client=self.ssh, name=self.shortname, **kwargs)
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 446, in run
    r.wait()
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 160, in wait
    self._raise_for_status()
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 182, in _raise_for_status
    node=self.hostname, label=self.label
teuthology.exceptions.CommandFailedError: Command failed (workunit test fs/norstats/kernel_untar_t
ar.sh) on smithi026 with status 124: 'mkdir -p -- /home/ubuntu/ceph_test/mnt.1/client.1/tmp && cd -
- /home/ubuntu/ceph_test/mnt.1/client.1/tmp && CEPH_CLI_TEST_DUP_COMMAND=1 CEPH_REF=34905f0a75aeb20
a442c8aedc05af7a94d2a3b8a TESTDIR="/home/ubuntu/ceph_test" CEPH_ARGS="--cluster ceph" CEPH_ID="1" P
ATH=$PATH:/usr/sbin CEPH_BASE=/home/ubuntu/ceph_test/clone.client.1 CEPH_ROOT=/home/ubuntu/ceph_test
/clone.client.1 adjust-ulimits ceph-coverage /home/ubuntu/ceph_test/archive/coverage timeout 3h /ho
me/ubuntu/ceph_test/clone.client.1/qa/workunits/fs/norstats/kernel_untar_tar.sh'
```

From:
/ceph/teuthology-archive/pdonnell-2020-08-08_02:21:01-multimds-wip-pdonnell-testing-20200808.001303-distro-basic-smithi/5319230/teuthology.log

The hung client appears to be client.4412. From kernel log:

```
2020-08-08T08:52:47.591880+00:00 smithi026 kernel: libceph: mon2 (1)172.21.15.200:6790 session est
ablished
2020-08-08T08:52:47.599882+00:00 smithi026 kernel: libceph: client4391 fsid 34ef1224-5cb7-4eac-88c
0-ed7c101900ca
2020-08-08T08:52:51.576198+00:00 smithi026 kernel: IPv6: ADDRCONF(NETDEV_CHANGE): brx.1: link beco
mes ready
2020-08-08T08:52:51.576286+00:00 smithi026 kernel: IPv6: ADDRCONF(NETDEV_CHANGE): veth0: link beco
mes ready
2020-08-08T08:52:51.620289+00:00 smithi026 kernel: ceph-brx: port 2(brx.1) entered blocking state
2020-08-08T08:52:51.620502+00:00 smithi026 kernel: ceph-brx: port 2(brx.1) entered disabled state
2020-08-08T08:52:51.625716+00:00 smithi026 kernel: device brx.1 entered promiscuous mode
2020-08-08T08:52:51.637811+00:00 smithi026 kernel: ceph-brx: port 2(brx.1) entered blocking state
2020-08-08T08:52:51.637916+00:00 smithi026 kernel: ceph-brx: port 2(brx.1) entered forwarding stat
e
2020-08-08T08:52:51.843897+00:00 smithi026 kernel: libceph: mon2 (1)172.21.15.200:6790 session est
ablished
2020-08-08T08:52:51.844026+00:00 smithi026 kernel: libceph: client4412 fsid 34ef1224-5cb7-4eac-88c
0-ed7c101900ca
2020-08-08T11:54:18.793937+00:00 smithi026 kernel: ceph-brx: port 2(brx.1) entered disabled state
2020-08-08T11:54:18.830062+00:00 smithi026 kernel: device brx.1 left promiscuous mode
2020-08-08T11:54:18.830231+00:00 smithi026 kernel: ceph-brx: port 2(brx.1) entered disabled state
2020-08-08T11:54:18.969493+00:00 smithi026 kernel: device brx.0 left promiscuous mode
2020-08-08T11:54:18.969581+00:00 smithi026 kernel: ceph-brx: port 1(brx.0) entered disabled state
```

From:
/ceph/teuthology-archive/pdonnell-2020-08-08_02:21:01-multimds-wip-pdonnell-testing-20200808.001303-distro-basic-smithi/5319230/remote/smithi026/syslog/kern.log.gz

The last thing the MDS saw that client mount do was:

```
2020-08-08T08:54:29.748+0000 7f6fc50c1700 7 mds.0.server reply_client_request 0 ((0) Success) cli
ent_request(client.4412:4029 create #0x1000000005e2/zorro.txt 2020-08-08T08:54:28.857992+0000 calle
r_uid=1000, caller_gid=1227{ }) v2
```

Then the client gets notified of exports:

```
2020-08-08T08:54:31.033+0000 7f6fc50c1700 10 mds.0.6 send_message_client_counted client.4412 seq 4
049 client_caps(export ino 0x1000000003ea 1 seq 0 caps=- dirty=- wanted=- follows 0 size 0/0 mtime
0.000000) v11
...
```

Client makes no progress from there.

History

#1 - 08/17/2020 01:42 PM - Patrick Donnelly

- Status changed from New to Need More Info
- Assignee set to Patrick Donnelly

Would be good to adjust the qa code to fetch the stack if the process hangs. Get /sys/debug/fs/ceph files as well.

#2 - 10/09/2020 06:48 PM - Patrick Donnelly

- Labels (FS) qa-failure added

#3 - 01/15/2021 10:44 PM - Patrick Donnelly

- Target version changed from v16.0.0 to v17.0.0
- Backport set to *pacific,octopus,nautilus*