

Orchestrator - Bug #45235

cephadm: mons are not properly undeployed

04/23/2020 03:07 PM - Sebastian Wagner

Status: Can't reproduce	% Done: 0%
Priority: Low	
Assignee:	
Category: cephadm	
Target version:	
Source:	Reviewed:
Tags:	Affected Versions:
Backport:	ceph-qa-suite:
Regression: No	Pull request ID:
Severity: 3 - minor	Crash signature:
Description	
<pre>\$ ceph orch daemon rm mon.ubuntu --force</pre>	
<pre>Removed mon.ubuntu from host 'ubuntu'</pre>	
<pre>\$ ceph orch ps</pre>	
<pre>No daemons reported</pre>	
<p>suddenly, everything hangs:</p>	
<pre>\$ ceph orch apply mon --placement 'ubuntu:127.0.0.1'</pre>	
<pre>^CCluster connection aborted</pre>	
<pre>\$ ceph -s</pre>	
<pre>^CCluster connection aborted</pre>	
<p>gdb revealed: cephadm hangs in set_store:</p>	
<pre>Thread 40 (Thread 0x7f8917ec4700 (LWP 23648)):</pre>	
<pre>#0 0x00007f894179b9f3 in futex_wait_cancelable (private=<optimized out>, expected=0, futex_word=0x7f8917ec0400) at ../sysdeps/unix/sysv/linux/futex-internal.h:88</pre>	
<pre>#1 __pthread_cond_wait_common (abstime=0x0, mutex=0x7f8917ec03b0, cond=0x7f8917ec03d8) at pthread_cond_wait.c:502</pre>	
<pre>#2 __pthread_cond_wait (cond=0x7f8917ec03d8, mutex=0x7f8917ec03b0) at pthread_cond_wait.c:655</pre>	
<pre>#3 0x00007f8941274f0c in std::condition_variable::wait(std::unique_lock<std::mutex>&) () from /usr/lib/x86_64-linux-gnu/libstdc++.so.6</pre>	
<pre>#4 0x0000561b001cd1ee in std::condition_variable::wait<C_SaferCond::wait()::{lambda()#1}>(std::unique_lock<std::mutex>&, C_SaferCond::wait()::{lambda()#1}) (this=0x7f8917ec03d8, __lock=..., __p=..) at /usr/include/c++/9/condition_variable:101</pre>	
<pre>#5 0x0000561b001c91ae in C_SaferCond::wait (this=0x7f8917ec03a8) at /home/sebastian/Repos/ceph/src/common/Cond.h:100</pre>	
<pre>#6 0x0000561b001c9a82 in Command::wait (this=0x7f8917ec03a0) at /home/sebastian/Repos/ceph/src/mgr/MgrContext.h:39</pre>	
<pre>#7 0x0000561b001c188a in ActivePyModules::set_store (this=0x561b03e5e000, module_name="cephadm", key="host.ubuntu", val=...) at /home/sebastian/Repos/ceph/src/mgr/ActivePyModules.cc:679</pre>	
<pre>#8 0x0000561b001e738a in ceph_store_set (self=0x7f89217685e8,</pre>	
<pre>args=('host.ubuntu', '{"daemons": {}, "devices": [{"rejected_reasons": ["locked"], "available": false, "path": "/dev/nvme0n1", "sys_api": {"removable": "0", "ro": "0", "vendor": "", "model": "KXG502NV512G NVME TOSHIBA 512GB", "rev": "", "sas_address": "", "sas_device_handle": "", "support_discard": "512", "rotational": "0", "nr_requests": "1023", "scheduler_mode": "none", "partitions":</pre>	

```

{"nvme0nlp3": {"start": "1288192", "sectors": "998926336", "sectorsize": 512, "size": 51145028403
2.0, "human_readable_size": "476.33 GB", "holders": []}, "nvme0nlp1": {"start": "2048", "sectors":
"1024000", "sectorsize": 512, "size": 524288000.0, "human_readable_size": "500.00 MB", "holders":
[]}, "nvme0nlp2": {"start": "1026048", "sectors": "262144", "sectorsize": 512, "size": 134217728.
0, "human_readable_size": "128.00 MB", "holders": []}}, "sectors": 0, "sectorsize": "512", "size":
512110190592.0, "human_readable_size": "476.94 GB", "path": "/dev/nvme0n1", "locked": 1}, "lvs":
[], "human_readable_type": "ssd", "device_id":'}) at /home/sebastian/Repos/ceph/src/mgr/BaseMgrMod
ule.cc:511
#9 0x00007f8941e3077b in PyCFunction_Call (func=<unknown at remote 0x7f8917ec0400>, args=<unknown
at remote 0x80>, kwds=0x0) at ../Objects/methodobject.c:103
#10 0x00000000000000002 in ?? ()
#11 0x00000000000000000 in ?? ()

```

monmap still lists the old mon:

```

$ ceph daemon mon.a mon_status
  "monmap": {
    "epoch": 2,
    "fsid": "a9df56ad-29b0-4bf4-8e47-a35c7657b332",
    "modified": "2020-04-23T13:26:56.292827Z",
    "created": "2020-04-23T11:14:26.191548Z",
    "min_mon_release": 15,
    "min_mon_release_name": "octopus",
    "features": {
      "persistent": [
        "kraken",
        "luminous",
        "mimic",
        "osdmap-prune",
        "nautilus",
        "octopus"
      ],
      "optional": []
    },
    "mons": [
      {
        "rank": 0,
        "name": "a",
        "public_addr": {
          "addrvec": [
            {
              "type": "v2",
              "addr": "172.17.0.1:40843",
              "nonce": 0
            },
            {
              "type": "v1",
              "addr": "172.17.0.1:40844",
              "nonce": 0
            }
          ]
        },
        "addr": "172.17.0.1:40844/0",
        "public_addr": "172.17.0.1:40844/0",
        "priority": 0,
        "weight": 0
      },
      {
        "rank": 1,
        "name": "ubuntu",
        "public_addr": {
          "addrvec": [
            {
              "type": "v2",

```

```
        "addr": "127.0.0.1:3300",
        "nonce": 0
      },
      {
        "type": "v1",
        "addr": "127.0.0.1:6789",
        "nonce": 0
      }
    ]
  },
  "addr": "127.0.0.1:6789/0",
  "public_addr": "127.0.0.1:6789/0",
  "priority": 0,
  "weight": 0
}
]
```

Related issues:

Related to Orchestrator - Bug #45167: cephadm: mons are not properly deployed

New

History

#1 - 04/23/2020 03:09 PM - Sebastian Wagner

- Description updated

#2 - 06/08/2020 01:24 PM - Sebastian Wagner

- Related to Bug #45167: cephadm: mons are not properly deployed added

#3 - 07/22/2020 11:53 AM - Sebastian Wagner

- Priority changed from Normal to Low

#4 - 07/31/2020 06:11 PM - Sebastian Wagner

- Status changed from New to Can't reproduce