

## RADOS - Bug #23360

### call to 'ceph osd erasure-code-profile set' asserts the monitors

03/14/2018 02:37 PM - Sebastian Wagner

<b>Status:</b> Duplicate	<b>% Done:</b> 0%
<b>Priority:</b> Normal	<b>Spent time:</b> 0.00 hour
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b> Development	<b>Affected Versions:</b>
<b>Tags:</b>	<b>ceph-qa-suite:</b>
<b>Backport:</b>	<b>Component(RADOS):</b>
<b>Regression:</b> No	<b>Pull request ID:</b>
<b>Severity:</b> 3 - minor	<b>Crash signature:</b>
<b>Reviewed:</b>	
<b>Description</b> I've attached `thread apply all bt` mixed with `thread apply all py-bt`  Threads 38 35 34 32 and 31 are waiting for futex 0x55a285204640  Thread 37 waits in File "/src/pybind/mgr/mgr_module.py", line 71, in wait self.ev.wait()  AFAICT, all other threads are not part of this deadlock.	
<b>Related issues:</b> Duplicates RADOS - Bug #23345: `ceph osd erasure-code-profile set` crashes th... <b>Resolved</b> <b>03/13/2018</b>	

#### History

##### #1 - 03/14/2018 02:48 PM - Ricardo Dias

Could you point to the code, or provide a small python example, that triggers this deadlock?

##### #2 - 03/14/2018 02:52 PM - Sebastian Wagner

The send\_command() function visible in this traceback is:

<https://github.com/ceph/ceph/pull/20865/files#diff-188b91d966a54c045089a507b218e586R108>

The code of this deadlock including erasure\_code\_profile.py is here:

<https://github.com/sebastian-philipp/ceph/commit/3444cf5f2c22a35bf411ecb566f7fa1ba5670bc7>

##### #3 - 03/14/2018 03:08 PM - Sebastian Wagner

Hm. quite possible that this is in fact not a classic deadlock.

Turns out, the `ceph` command line tool is also broken:

```
$ ceph osd erasure-code-profile set threetwo k=3 m=2
```

```
^CError EINTR: Interrupted!
```

Environment:

Fresh build on git commit on the latest master (619d435a) with a vstart.sh cluster (ceph version 13.0.1-3034)

#### #4 - 03/14/2018 04:15 PM - Sebastian Wagner

- Project changed from mgr to RADOS
- Category deleted (python interface)

Found the cause of this. From the mon.a.log:

```
5> 2018-03-14 16:55:26.470 7f61aa02a700 10 log_client will send 2018-03-14 16:55:26.473781 mon.a mon.0 192.168.178.29:40706/0 93 : audit [INF] from='client.154151 -' entity='client.admin' cmd=[{"profile": [{"k="3", "m="3"}], "prefix": "osd erasure-code-profile set", "name": "myprofile"}]: dispatch
... snip ...
/home/sebastian/Repos/ceph/src/mon/OSDMonitor.cc: 5501: FAILED assert((*erasure_code_profile_map).count("plugin"))
```

```
ceph version 13.0.1-3034-g619d435a71 (619d435a71f002571bda0c71dd26d75deaf480c3) mimic (dev)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x137) [0x7f61b722c9fb]
2: (OSDMonitor::parse_erasure_code_profile(std::vector<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>, const&, std::map<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>, std::less<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>>, std::allocator<char>>>>, std::allocator<std::pair<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>> const, std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>>*, std::ostream*)+0x145) [0x5584109774ef]
3: (OSDMonitor::prepare_command_impl(boost::intrusive_ptr<MonOpRequest>, std::map<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, boost::variant<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, bool, long, double, std::vector<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>, std::allocator<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>>, std::allocator<char>>>>, std::vector<long, std::allocator<long>>>, std::vector<double, std::allocator<double>>>, std::less<void>, std::allocator<std::pair<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>> const, boost::variant<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, bool, long, double, std::vector<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>, std::allocator<char>>>>, std::allocator<std::__cxx11::basic_string<char, std::char_traits<char>, std::allocator<char>>>>, std::vector<
```

#### #5 - 03/14/2018 04:16 PM - Sebastian Wagner

A proper fix would be to provide a proper error message in OSDMonitor::parse\_erasure\_code\_profile instead of asserting, when the list of ecp plugins is empty.

**#6 - 03/14/2018 04:49 PM - Joao Eduardo Luis**

- Status changed from New to Duplicate

duplicate of <http://tracker.ceph.com/issues/23345>

**#7 - 03/14/2018 04:49 PM - Joao Eduardo Luis**

- Duplicates Bug #23345: `ceph osd erasure-code-profile set` crashes the monitors on vstart clusters added

**#8 - 03/14/2018 04:54 PM - Joao Eduardo Luis**

- Subject changed from Python Deadlock in mgr\_module.CommandResult#wait to call to 'ceph osd erasure-code-profile set' asserts the monitors

## Files

---

deadlock.txt	651 KB	03/14/2018	Sebastian Wagner
--------------	--------	------------	------------------