

Ceph - Bug #16297

Monitor die if moncommand without "prefix" item

06/14/2016 03:21 PM - Xiaoxi Chen

Status: Resolved	% Done: 0%
Priority: Urgent	Spent time: 0.00 hour
Assignee: Joao Eduardo Luis	
Category: Monitor	
Target version:	
Source: other	Reviewed:
Tags:	Affected Versions: v0.94.6
Backport: jewel,hammer	ceph-qa-suite:
Regression: No	Pull request ID:
Severity: 2 - major	Crash signature:
Description we are using cluster.mon_command() in rados.py to do with *prefix*="osd pool stats",format="json", monitor will assert Recheck the code we wrongly make "prefix" to "perfix" :(, and anything other than "prefix"(say "abc") will kill the monitor. This is dangerous as anyone has access to the rados can take down whole cluster. [Bug seem in hammer 0.94.6, not tested against other version	
Related issues: Copied to Ceph - Backport #16549: jewel: Monitor die if moncommand without "p... Resolved Copied to Ceph - Backport #16550: hammer:Monitor die if moncommand without "p... Resolved	

History

#1 - 06/14/2016 03:49 PM - Xiaoxi Chen

- Release set to hammer
- Release set to infernalis
- Release set to jewel

as this part of code doesn't change , should also affect all version

#2 - 06/14/2016 03:50 PM - Xiaoxi Chen

```
2016-06-14 05:31:58.676295 7fa378a5c700 0 mon.lvs2b02c-1mgr@2(peon) e3 handle_command mon_command({"prefix": "osd pool s
tats", "format": "json"}v0)v1
2016-06-14 05:31:58.721433 7fa378a5c700 -1 ** Caught signal (Segmentation fault) *
in thread 7fa378a5c700
```

```
ceph version 0.94.6 (e832001feaf8c176593e0325c8298e3f16dfb403)
1: /usr/bin/ceph-mon() [0x9acf9a]
2: (()+0x10340) [0x7fa38182c340]
3: (std::string::assign(std::string const&)+0x10) [0x7fa380631480]
4: (Monitor::handle_command(MMonCommand*)+0xaeef) [0x5cba6f]
5: (Monitor::dispatch(MonSession*, Message*, bool)+0xf9) [0x5cf649]
6: (Monitor::_ms_dispatch(Message*)+0x1a6) [0x5d02c6]
7: (Monitor::ms_dispatch(Message*)+0x23) [0x5ee073]
8: (DispatchQueue::entry()+0x649) [0x928839]
9: (DispatchQueue::DispatchThread::entry()+0xd) [0x7c8eed]
10: (()+0x8182) [0x7fa381824182]
11: (clone()+0x6d) [0x7fa37fd8f47d]
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```

#3 - 06/14/2016 05:07 PM - Joao Eduardo Luis

- Category set to Monitor

- Status changed from New to In Progress

#4 - 06/14/2016 06:48 PM - Joao Eduardo Luis

I can confirm current master also suffers from this. I'm assuming both infernalis and jewel also suffer from it.

#5 - 06/14/2016 06:53 PM - Ji You

Joao Luis wrote:

I can confirm current master also suffers from this. I'm assuming both infernalis and jewel also suffer from it.

patch submitted: <https://github.com/ceph/ceph/pull/9700>

#6 - 06/14/2016 10:15 PM - Joao Eduardo Luis

Ji You wrote:

Joao Luis wrote:

I can confirm current master also suffers from this. I'm assuming both infernalis and jewel also suffer from it.

patch submitted: <https://github.com/ceph/ceph/pull/9700>

Ji You, in the future, if you see the ticket is assigned to someone and is in progress, please check with the person it is assigned to before submitting a patch. It reduces the amount of duplicate work.

I will comment on your pull request in a bit.

#7 - 06/15/2016 02:01 AM - Ji You

Joao Luis wrote:

Ji You wrote:

Joao Luis wrote:

I can confirm current master also suffers from this. I'm assuming both infernalis and jewel also suffer from it.

patch submitted: <https://github.com/ceph/ceph/pull/9700>

Ji You, in the future, if you see the ticket is assigned to someone and is in progress, please check with the person it is assigned to before submitting a patch. It reduces the amount of duplicate work.

I will comment on your pull request in a bit.

Very sorry to not obey the process. My big mistake.

Yesterday deep night in china time, when I find this bug in our production, after communication with XiaoXi, Chen. I was trying to write a patch for this issue.

XiaoXi, Chen writes this bug issue on tracker.

My mistake is taking this bug is assigned to XiaoXi, Chen. So I submitted this patch.

Sorry for my careless. Everyone in my team is against me to do this kind of thing. Because this not follow regular working process, it's too bad to do this with careless.

And, many thanks for your valuable comments. That helps me a lot.

#8 - 06/29/2016 03:57 PM - Ken Dreyer

Red Hat product security has assigned CVE-2016-5009 to this issue today.

#9 - 06/30/2016 03:12 AM - Sage Weil

- Status changed from *In Progress* to *Pending Backport*

- Backport set to *jewel,hammer*

#10 - 06/30/2016 03:29 AM - Xiaoxi Chen

- Copied to Backport #16549: *jewel: Monitor die if moncommand without "prefix" item added*

#11 - 06/30/2016 03:30 AM - Xiaoxi Chen

- Copied to Backport #16550: hammer:Monitor die if moncommand without "prefix" item added

#12 - 08/08/2016 08:42 AM - Loic Dachary

- Status changed from Pending Backport to Resolved