

Ceph - Bug #20745

Error on create-initial with --cluster

07/22/2017 03:35 PM - Oscar Segarra

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:	Monitor		
Target version:	v12.1.0		
Source:	Community (user)	Reviewed:	
Tags:		Affected Versions:	v12.1.0
Backport:		ceph-qa-suite:	ceph-deploy
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	

Description

```
[vdicceph@vdicnode01 ceph]$ ceph-deploy --cluster vdicmgmt --username vdicceph mon create-initial
[ceph_deploy.conf][DEBUG ] found configuration file at: /home/vdicceph/.cephdeploy.conf
[ceph_deploy.cli][INFO ] Invoked (1.5.38): /bin/ceph-deploy --cluster vdicmgmt --username vdicceph mon create-initial
[ceph_deploy.cli][INFO ] ceph-deploy options:
[ceph_deploy.cli][INFO ] username          : vdicceph
[ceph_deploy.cli][INFO ] verbose           : False
[ceph_deploy.cli][INFO ] overwrite_conf    : False
[ceph_deploy.cli][INFO ] subcommand        : create-initial
[ceph_deploy.cli][INFO ] quiet             : False
[ceph_deploy.cli][INFO ] cd_conf           : <ceph_deploy.conf.cephdeploy.Conf instance at 0x1b4bab8>
[ceph_deploy.cli][INFO ] cluster           : vdicmgmt
[ceph_deploy.cli][INFO ] func              : <function mon at 0x1b3d668>
[ceph_deploy.cli][INFO ] ceph_conf         : None
[ceph_deploy.cli][INFO ] default_release   : False
[ceph_deploy.cli][INFO ] keyrings          : None
[ceph_deploy.mon][DEBUG ] Deploying mon, cluster vdicmgmt hosts vdicnode01
[ceph_deploy.mon][DEBUG ] detecting platform for host vdicnode01 ...
[vdicnode01][DEBUG ] connection detected need for sudo
[vdicnode01][DEBUG ] connected to host: vdicceph@vdicnode01
[vdicnode01][DEBUG ] detect platform information from remote host
[vdicnode01][DEBUG ] detect machine type
[vdicnode01][DEBUG ] find the location of an executable
[ceph_deploy.mon][INFO ] distro info: CentOS Linux 7.3.1611 Core
[vdicnode01][DEBUG ] determining if provided host has same hostname in remote
[vdicnode01][DEBUG ] get remote short hostname
[vdicnode01][DEBUG ] deploying mon to vdicnode01
[vdicnode01][DEBUG ] get remote short hostname
[vdicnode01][DEBUG ] remote hostname: vdicnode01
[vdicnode01][DEBUG ] write cluster configuration to /etc/ceph/{cluster}.conf
[vdicnode01][DEBUG ] create the mon path if it does not exist
[vdicnode01][DEBUG ] checking for done path: /var/lib/ceph/mon/vdicmgmt-vdicnode01/done
[vdicnode01][DEBUG ] done path does not exist: /var/lib/ceph/mon/vdicmgmt-vdicnode01/done
[vdicnode01][INFO ] creating keyring file: /var/lib/ceph/tmp/vdicmgmt-vdicnode01.mon.keyring
[vdicnode01][DEBUG ] create the monitor keyring file
[vdicnode01][INFO ] Running command: sudo ceph-mon --cluster vdicmgmt --mkfs -i vdicnode01 --keyring
/var/lib/ceph/tmp/vdicmgmt-vdicnode01.mon.keyring --setuser 167 --setgroup 167
[vdicnode01][INFO ] unlinking keyring file /var/lib/ceph/tmp/vdicmgmt-vdicnode01.mon.keyring
[vdicnode01][DEBUG ] create a done file to avoid re-doing the mon deployment
[vdicnode01][DEBUG ] create the init path if it does not exist
[vdicnode01][INFO ] Running command: sudo systemctl enable ceph.target
[vdicnode01][INFO ] Running command: sudo systemctl enable ceph-mon@vdicnode01
[vdicnode01][WARNIN] Created symlink from /etc/systemd/system/ceph-mon.target.wants/ceph-mon@vdicnode01.service to
/usr/lib/systemd/system/ceph-mon@.service.
[vdicnode01][INFO ] Running command: sudo systemctl start ceph-mon@vdicnode01
```

```

[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[vdicnode01][WARNIN] monitor: mon.vdicnode01, might not be running yet
[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[vdicnode01][WARNIN] monitor vdicnode01 does not exist in monmap
[ceph_deploy.mon][INFO ] processing monitor mon.vdicnode01
[vdicnode01][DEBUG ] connection detected need for sudo
[vdicnode01][DEBUG ] connected to host: vdicceph@vdicnode01
[vdicnode01][DEBUG ] detect platform information from remote host
[vdicnode01][DEBUG ] detect machine type
[vdicnode01][DEBUG ] find the location of an executable
[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[ceph_deploy.mon][WARNIN] mon.vdicnode01 monitor is not yet in quorum, tries left: 5
[ceph_deploy.mon][WARNIN] waiting 5 seconds before retrying
[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[ceph_deploy.mon][WARNIN] mon.vdicnode01 monitor is not yet in quorum, tries left: 4
[ceph_deploy.mon][WARNIN] waiting 10 seconds before retrying
[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[ceph_deploy.mon][WARNIN] mon.vdicnode01 monitor is not yet in quorum, tries left: 3
[ceph_deploy.mon][WARNIN] waiting 10 seconds before retrying
[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[ceph_deploy.mon][WARNIN] mon.vdicnode01 monitor is not yet in quorum, tries left: 2
[ceph_deploy.mon][WARNIN] waiting 15 seconds before retrying
[vdicnode01][INFO ] Running command: sudo ceph --cluster=vdicmgmt --admin-daemon
/var/run/ceph/vdicmgmt-mon.vdicnode01.asok mon_status
[vdicnode01][ERROR ] admin_socket: exception getting command descriptions: [Errno 2] No such file or directory
[ceph_deploy.mon][WARNIN] mon.vdicnode01 monitor is not yet in quorum, tries left: 1
[ceph_deploy.mon][WARNIN] waiting 20 seconds before retrying
[ceph_deploy.mon][ERROR ] Some monitors have still not reached quorum:
[ceph_deploy.mon][ERROR ] vdicnode01

```

History

#1 - 07/24/2017 05:28 PM - Vasu Kulkarni

Have you followed all steps properly, The nightly tests I see on centos 7.3 dont show any issues.

#2 - 07/25/2017 09:04 AM - Oscar Segarra

Yes... there is not so much history about creating a ceph cluster...

```

--> as root
echo "vdicceph ALL = (root) NOPASSWD:ALL" | tee /etc/sudoers.d/vdicceph
echo "Defaults:vdicceph ""!""requiretty" | tee -a /etc/sudoers.d/vdicceph
chmod 0440 /etc/sudoers.d/vdicceph

```

```

--> as user vdicceph:
ceph-deploy --cluster vdicmgmt --username vdicceph new vdicnode01 --cluster-network 192.168.100.0/24 --public-network 192.168.100.0/24
--> What finishes successfully
ceph-deploy --cluster vdicmgmt --username vdicceph mon create-initial

```

Note that the primary IP of the host vdicnode01 is 192.168.2.101. Network 192.168.100.0/24 is the private network.

Thanks a lot.

#3 - 07/26/2017 03:03 AM - Sage Weil

- Status changed from New to Resolved

Please avoid using the --cluster option when deploy the cluster itself. The support is prereserved for clients to attached to other clusters, but we're trying to avoid naming clusters on the server side (it just makes life harder for the operator, especially since we've transitioned to systemd).

#4 - 07/27/2017 09:25 PM - Oscar Segarra

Hi,

Just to clarify,

--> Please avoid using the --cluster option when deploy the cluster itself.

Why might I avoid using --cluster in the "new" command isn't it? <-- It is not explained in documentation and the --cluster argument in the ceph-deploy new command works perfectly.

--> The support is prereserved for clients to attached to other clusters, but we're trying to avoid naming clusters on the server side (it just makes life harder for the operator, especially since we've transitioned to systemd).

Is there any known problem when the admin node is the first monitor node?

Thanks a lot.

#5 - 08/08/2017 06:04 PM - Sage Weil

Oscar Segarra wrote:

Hi,

Just to clarify,

--> Please avoid using the --cluster option when deploy the cluster itself.

Why might I avoid using --cluster in the "new" command isn't it? <-- It is not explained in documentation and the --cluster argument in the ceph-deploy new command works perfectly.

Setting cluster != ceph for daemons is very poorly tested and documented, and likely to break. For example, you have to manually edit /etc/sysconfig/ceph to set CLUSTER=whatever in order for systemd to manage your daemons.

I thought Alfredo removed --cluster support but I guess not; he probably just removed the ceph-disk support. In any case, please don't use it.

--> The support is prereserved for clients to attached to other clusters, but we're trying to avoid naming clusters on the server side (it just makes life harder for the operator, especially since we've transitioned to systemd).

Is there any known problem when the admin node is the first monitor node?

Nothing comes to mind!

#6 - 08/15/2017 04:39 PM - Oscar Segarra

Well Sage,

The issue has been closed as "Resolved" but I think the usage of the supported and documented "cluster" clause that is not tested and does not work properly might be reviewed and fixed or mark it as "experimental" or whatever...

Thanks a lot.