

fs - Bug #17105

multimds: allow_multimds not required when max_mds is set in ceph.conf at startup

08/23/2016 06:02 PM - Patrick Donnelly

Status:	Resolved	Start date:	08/23/2016
Priority:	Normal	Due date:	
Assignee:	Patrick Donnelly	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Affected Versions:	
Source:	Development	ceph-qa-suite:	
Tags:		Component(FS):	MDSMonitor
Backport:	jewel	Labels (FS):	multimds
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Reviewed:			

Description

Problem:

```
$ vstart.sh -d -l --short -n -o 'max_mds = 3'
[...]
$ ceph status
  cluster 24600546-44d2-41ad-85ba-996b7af1fd73
  health HEALTH_OK
  monmap e1: 3 mons at {a=127.0.0.1:6789/0,b=127.0.0.1:6790/0,c=127.0.0.1:6791/0}
    election epoch 6, quorum 0,1,2 a,b,c
  fsmap e7: 3/3/3 up {0=a=up:active,1=c=up:creating,2=b=up:active}
  osdmap e10: 3 osds: 3 up, 3 in
    flags sortbitwise,require_jewel_osds,require_kraken_osds
  pgmap v13: 24 pgs, 3 pools, 0 bytes data, 0 objects
    550 GB used, 114 GB / 664 GB avail
    24 active+clean
$ ceph mds set_max_mds 4
Error EINVAL: multi-MDS clusters are not enabled; set 'allow_multimds' to enable
```

The last command is not the problem. It's just showing what we currently do when trying to set max_mds when the flag is unset. The "issue" is that setting max_mds in the ceph.conf enables multimds but the CEPH_MDSMAP_ALLOW_MULTIMDS is not also set.

Either we should:

1. Enable allow_multimds by setting the CEPH_MDSMAP_ALLOW_MULTIMDS flag so the administrator can modify max_mds without redundantly setting the flag (because multimds is already enabled even though the flag is unset).
2. Require a new config option setting allow_multimds in addition to max_mds and reject max_mds otherwise (silently?).

Related issues:

Copied to fs - Backport #17264: jewel: multimds: allow_multimds not required ...

Resolved

History

#1 - 08/23/2016 06:23 PM - Greg Farnum

I think we want to force users to set multi-mds flags explicitly, not implicitly via the initial config. I'm fine with it being silent if being loud is inconvenient.

#2 - 08/24/2016 11:41 PM - Patrick Donnelly

- Status changed from New to Need Review

PR: <https://github.com/ceph/ceph/pull/10848>

#3 - 08/29/2016 09:23 PM - Patrick Donnelly

New PR: <https://github.com/ceph/ceph/pull/10914>

#4 - 09/06/2016 01:43 PM - John Spray

- Status changed from Need Review to Pending Backport

#5 - 09/06/2016 04:07 PM - Patrick Donnelly

- Release set to jewel

Backport: <https://github.com/ceph/ceph/pull/10997>

#6 - 09/12/2016 07:58 AM - Loic Dachary

- Backport set to jewel

#7 - 09/12/2016 07:59 AM - Loic Dachary

- Copied to Backport #17264: jewel: multimds: allow_multimds not required when max_mds is set in ceph.conf at startup added

#8 - 10/13/2016 03:51 PM - Loic Dachary

- Status changed from Pending Backport to Resolved

#9 - 03/09/2019 12:31 AM - Patrick Donnelly

- Category deleted (90)

- Labels (FS) multimds added