

fs - Bug #41728

mds: hang during fragmentdir

09/09/2019 06:29 PM - Nathan Fish

Status: Can't reproduce	% Done: 0%
Priority: High	
Assignee: Zheng Yan	
Category: Correctness/Safety	
Target version:	
Source: Community (user)	Affected Versions: v14.2.3
Tags:	ceph-qa-suite:
Backport: nautilus,mimic	Component(FS): MDS, osdc
Regression:	Labels (FS):
Severity: 2 - major	Pull request ID:
Reviewed:	Crash signature:

Description

When doing a parallel cp, the active MDS on the CephFS hung on a fragmentdir op.
It might be this bug: <http://lists.ceph.com/pipermail/ceph-users-ceph.com/2019-May/034785.html>

I was able to repair the cluster by failing the active MDS, unmounting and remounting the client, then failing the new active MDS. I increased "mds bal split rd" to 250000 and "mds bal split wr" to 100000 in the hopes this would reduce the number of fragmentdir ops, then restarted the cp. The issue occurred again an hour or two in.

Active MDS ops:
ceph daemon mds.dc-3558-422-C ops
<https://termbin.com/c3sm>

ceph health detail
<https://termbin.com/xx4k>

All machines are Ubuntu 18.04 with hwe kernel 5.0, and Nautilus 14.2.3.

Related issues:

Related to fs - Bug #41434: mds: infinite loop in Locker::file_update_finish()	Resolved
--	-----------------

History

#1 - 09/11/2019 05:47 PM - Patrick Donnelly

- Subject changed from MDS hangs during fragmentdir to mds: hang during fragmentdir
- Priority changed from Normal to High
- Target version set to v15.0.0
- Start date deleted (09/09/2019)
- Backport set to nautilus,mimic
- Component(FS) MDS added

#2 - 09/16/2019 01:56 PM - Zheng Yan

- Component(FS) osdc added

Nathan Fish wrote:

When doing a parallel cp, the active MDS on the CephFS hung on a fragmentdir op.
It might be this bug: <http://lists.ceph.com/pipermail/ceph-users-ceph.com/2019-May/034785.html>

Above bug is fixed by <https://github.com/ceph/ceph/pull/29902>

#3 - 09/16/2019 02:28 PM - Nathan Fish

Thanks!

#4 - 09/16/2019 07:06 PM - Patrick Donnelly

- Related to Bug #41434: *mds: infinite loop in Locker::file_update_finish()* added

#5 - 09/16/2019 07:06 PM - Patrick Donnelly

- Status changed from *New* to *Need More Info*

- Assignee set to *Zheng Yan*

#6 - 01/23/2020 01:32 AM - Patrick Donnelly

- Target version deleted (*v15.0.0*)

#7 - 03/30/2020 09:19 AM - Zheng Yan

- Status changed from *Need More Info* to *Can't reproduce*