

Ceph - Bug #21613

backfill cancelation makes target crash; now triggered by recovery preemption

10/01/2017 07:26 PM - Sage Weil

Status:	Resolved	% Done:	0%
Priority:	Immediate	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	luminous	Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

if backfill is in progress and we cancel (previous due to unfound, now due to preemption), we send a MBackfillReserve REJECT message to the backfill target

```
2017-09-29 19:49:28.319666 7fba27681700 10 osd.6 pg_epoch: 751 pg[2.14( v 606'828 (241'472,606'828
] local-lis/les=747/749 n=2 ec=57/17 lis/c 747/688 les/c/f 749/690/0 747/747/484) [6,3,0]/[6,3,4]
r=0 lpr=747 pi=[688,747)/1 bft=0 crt=606'828 lcod 605'826 mlcod 0'0 active+remapped+backfilling sn
aptrimq=[149~1,226~1,285~1,296~3]] state<Started/Primary/Active/Backfilling>: defer backfill, retr
y delay 0
2017-09-29 19:49:28.319694 7fba27681700 1 -- 172.21.15.187:6814/32046 --> 172.21.15.201:6809/1489
5 -- MBackfillReserve REJECT pgid: 2.14, query_epoch: 751 v3 -- ?+0 0x5627e019d200 con 0x5627dbfa
0f40
```

REJECT is meant to be sent from requestee to requester, not requester to cancel at requestee. the backfill target does not like this message:

```
2017-09-29 19:49:28.346586 7f7b95f9b700 1 -- 172.21.15.201:6809/14895 <== osd.6 172.21.15.187:681
4/32046 4478 ==== MBackfillReserve REJECT pgid: 2.14, query_epoch: 751 v3 ==== 30+0+0 (2776089898
0 0) 0x55fb09669d40 con 0x55fb09f1e7e0
2017-09-29 19:49:28.346590 7f7b95f9b700 20 osd.0 751 OSD::ms_dispatch: MBackfillReserve REJECT pg
id: 2.14, query_epoch: 751 v3
2017-09-29 19:49:28.346593 7f7b95f9b700 20 osd.0 751 _dispatch 0x55fb09669d40 MBackfillReserve REJ
ECT pgid: 2.14, query_epoch: 751 v3
...
2017-09-29 19:49:28.346638 7f7b8a701700 10 osd.0 pg_epoch: 751 pg[2.14( v 606'828 (241'472,606'828
] local-lis/les=747/749 n=2 ec=57/17 lis/c 747/688 les/c/f 749/690/0 747/747/484) [6,3,0]/[6,3,4]
r=-1 lpr=749 pi=[688,747)/1 luod=0'0 crt=606'828 lcod 0'0 active] handle_peering_event: epoch_sent
: 751 epoch_requested: 751 RemoteReservationRejected
2017-09-29 19:49:28.346657 7f7b8a701700 5 osd.0 pg_epoch: 751 pg[2.14( v 606'828 (241'472,606'828
] local-lis/les=747/749 n=2 ec=57/17 lis/c 747/688 les/c/f 749/690/0 747/747/484) [6,3,0]/[6,3,4]
r=-1 lpr=749 pi=[688,747)/1 luod=0'0 crt=606'828 lcod 0'0 active] exit Started/ReplicaActive/RepNo
tRecovering 0.232928 4 0.000064
2017-09-29 19:49:28.346668 7f7b8a701700 5 osd.0 pg_epoch: 751 pg[2.14( v 606'828 (241'472,606'828
] local-lis/les=747/749 n=2 ec=57/17 lis/c 747/688 les/c/f 749/690/0 747/747/484) [6,3,0]/[6,3,4]
r=-1 lpr=749 pi=[688,747)/1 luod=0'0 crt=606'828 lcod 0'0 active] exit Started/ReplicaActive 0.896
000 0 0.000000
2017-09-29 19:49:28.347184 7f7b8a701700 5 osd.0 pg_epoch: 751 pg[2.14( v 606'828 (241'472,606'828
] local-lis/les=747/749 n=2 ec=57/17 lis/c 747/688 les/c/f 749/690/0 747/747/484) [6,3,0]/[6,3,4]
r=-1 lpr=749 pi=[688,747)/1 luod=0'0 crt=606'828 lcod 0'0 active] exit Started 1.070936 0 0.000000
2017-09-29 19:49:28.347192 7f7b8a701700 5 osd.0 pg_epoch: 751 pg[2.14( v 606'828 (241'472,606'828
```

```
] local-lis/les=747/749 n=2 ec=57/17 lis/c 747/688 les/c/f 749/690/0 747/747/484) [6,3,0]/[6,3,4]
r=-1 lpr=749 pi=[688,747)/1 luod=0'0 crt=606'828 lcod 0'0 active] enter Crashed
```

/a/sage-2017-09-29_18:35:33-rados-wip-sage-testing-2017-09-29-1154-distro-basic-smithi/1686969

Related issues:

Copied to Ceph - Backport #22780: luminous: backfill cancelation makes target...

Resolved

History

#1 - 10/01/2017 08:11 PM - Sage Weil

- Status changed from 12 to Fix Under Review

<https://github.com/ceph/ceph/pull/18070>

#2 - 10/04/2017 01:31 PM - Sage Weil

/a/sage-2017-10-03_21:58:15-rados-wip-sage-testing-2017-10-03-1358-distro-basic-smithi/1700060

#3 - 10/06/2017 01:27 PM - Sage Weil

- Status changed from Fix Under Review to Resolved

#4 - 01/23/2018 09:50 PM - David Zafman

- Status changed from Resolved to Pending Backport

#5 - 01/23/2018 09:51 PM - David Zafman

- Copied to Backport #22780: luminous: backfill cancelation makes target crash; now triggered by recovery preemption added

#6 - 01/23/2018 11:29 PM - David Zafman

- Status changed from Pending Backport to Resolved