

Linux kernel client - Bug #22014

krbd: memory allocation related deadlock (cloned image)

11/03/2017 06:42 AM - Zheng Yan

Status:	Resolved	Start date:	11/03/2017
Priority:	Urgent	Due date:	
Assignee:	Ilya Dryomov	% Done:	0%
Category:	rbd	Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Severity:	3 - minor
Tags:		Reviewed:	
Backport:		Affected Versions:	
Regression:	No	ceph-qa-suite:	
Description			
calltrace of ceph-msgr [7507093.943117] [<fffffffa075dfc9>] jbd2_log_wait_commit+0xa9/0x130 [jbd2] [7507093.943147] [<ffffff81218761>] evict+0xb1/0x190 [7507093.943160] [<ffffff81192c97>] shrink_slab.part.40+0x1f7/0x3c0 [7507093.943173] [<ffffff8118a8a2>] __alloc_pages_nodemask+0x582/0xab0 [7507093.943196] [<ffffffa0748835>] rbd_img_obj_request_submit+0x295/0x4b0 [rbd] [7507093.943216] [<ffffffa06f6dbc>] handle_reply+0x3fc/0x810 [libceph] [7507093.943237] [<ffffff810ae4c2>] ? sched_clock_cpu+0x72/0xa0 [7507093.943254] [<ffffff81096804>] worker_thread+0x114/0x470 [7507093.943262] [<ffffff8109c270>] ? kthread_park+0x60/0x60 calltrace of jbd2 [<ffffff8130a0e4>] blkdev_issue_discard+0x1f4/0x250 [<ffffffa07b3fdb>] ext4_free_data_callback+0x20b/0x2c0 [ext4] [<ffffffa078f6e3>] ext4_journal_commit_callback+0x83/0xc0 [ext4] [<ffffffa075a442>] jbd2_journal_commit_transaction+0x12e2/0x1940 [jbd2] [<ffffffa075e48a>] kjournald2+0xca/0x260 [jbd2] [<ffffff8109c348>] kthread+0xd8/0xf0 [<ffffff816ee90f>] ret_from_fork+0x3f/0x70 [<ffffffffffffff>] 0xffffffffffff			

History

#1 - 11/03/2017 06:45 AM - Zheng Yan

happened on layered image

#2 - 11/03/2017 10:34 AM - Ilya Dryomov

This particular one is easy to fix, I'll patch it next week.

#3 - 11/06/2017 10:28 AM - Ilya Dryomov

- Subject changed from *krbd: memory allocation related deadlock* to *krbd: memory allocation related deadlock (cloned image)*

- Category set to *rbd*

- Status changed from *New* to *In Progress*

#4 - 11/06/2017 10:30 AM - Ilya Dryomov

- Priority changed from *Normal* to *Urgent*

#5 - 11/09/2017 08:41 AM - Ilya Dryomov

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

[PATCH] rbd: use GFP_NOIO for parent stat and data requests

#6 - 11/11/2017 09:32 AM - Ilya Dryomov

- *Status changed from Pending Upstream to Pending Backport*

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=1e37f2f84680fa7f8394fd444b6928e334495ccc> in 4.14

#7 - 11/27/2017 09:28 AM - Ilya Dryomov

- *Status changed from Pending Backport to Resolved*

In 3.18.81, 4.4.98, 4.9.62 and 4.13.13.