Linux kernel client - Bug #1769

osd_client: susceptibility to low memory deadlocks

11/30/2011 11:28 AM - Anonymous

Status:	New	**% Done:**	0%
Priority:	Low	**Spent time:**	0.00 hour
Assignee:			
Category:			
Target version:			
Source:		**Reviewed:**	
Tags:		**Affected Versions:**	
Backport:		**ceph-qa-suite:**	
Regression:	No	**Crash signature (v1):**	
Severity:	3 - minor	**Crash signature (v2):**	

Description

We could be trying to flush the cache in order to free up memory, and find ourselves unable to allocate a ceph_osd or socket to (re)establish the connection through which we would flush our cache and free up some memory ... giving rise to a memory deadlock.

It would be good if OSD connection (re)establishment could be performed without additional memory allocation.
Exactly how is not clear, but Sage opines that we might pin old connections open, and then try to finish all old requests before starting new ones.

He also observes that there are probably several other such susceptibilities that should be recognized and addressed.

History

#1 - 12/02/2011 11:55 AM - Sage Weil

- Project changed from Ceph to Linux kernel client

#2 - 01/03/2012 10:52 AM - Sage Weil

- translation missing: en.field_position set to 27

#3 - 01/03/2012 10:52 AM - Sage Weil

- translation missing: en.field_position deleted (32)

- translation missing: en.field_position set to 35