

CephFS - Bug #21807

mds: trims all unpinned dentries when memory limit is reached

10/14/2017 08:46 PM - Patrick Donnelly

Status:	Resolved	% Done:	0%
Priority:	Immediate		
Assignee:	Patrick Donnelly		
Category:			
Target version:			
Source:	Development	ceph-qa-suite:	
Tags:		Component(FS):	MDS
Backport:	luminous	Labels (FS):	
Regression:	No	Pull request ID:	
Severity:	2 - major	Crash signature (v1):	
Reviewed:		Crash signature (v2):	
Affected Versions:			
Description Generally dentries are pinned by the client cache so this was easy to miss in testing. Bug is here: https://github.com/ceph/ceph/blob/5a96d19fc0cf2a3a04641cb2fa608f08b91279a2/src/mds/MDCache.cc#L6461 count will underflow is it is 0 (count=0 indicates to trim until the cache is no longer toofull). This bug is particularly bad for large directory workloads where the mjaority of dentries are unpinned because clients are only reading a small selectoin of dentries from each dirfrag. https://bugzilla.redhat.com/show_bug.cgi?id=1502178			
Related issues: Copied to CephFS - Backport #21810: luminous: mds: trims all unpinned dentrie... <div>Resolved</div>			

History

#1 - 10/14/2017 08:49 PM - Patrick Donnelly

- Status changed from New to Fix Under Review

<https://github.com/ceph/ceph/pull/18309>

#2 - 10/16/2017 12:03 AM - Patrick Donnelly

- Copied to Backport #21810: luminous: mds: trims all unpinned dentries when memory limit is reached added

#3 - 10/16/2017 12:07 AM - Patrick Donnelly

- Status changed from Fix Under Review to Pending Backport

#4 - 10/29/2017 06:14 PM - Patrick Donnelly

- Status changed from Pending Backport to Resolved