

Ceph - Bug #758

dd-truncate elsewhere-dd again is slow

02/02/2011 11:29 AM - Greg Farnum

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Greg Farnum		
Category:			
Target version:	v0.25		
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Description			
<p>Jim Schutt has reported some troubling behavior apparently involved with client caps surrounding distributed (but not simultaneous) use of files. See the "Odd "data used" reporting behavior by ceph -w" thread (http://www.spinics.net/lists/ceph-devel/msg01222.html).</p> <p>After some effort I was able to reproduce his issue with dd-truncate-dd in an environment using vstart.sh (1 of each daemon), the kernel client (via UML) doing the dd, and cfuse doing the truncate. This did not occur if cfuse was used to perform the dd and the kernel client performed the truncate. I eventually discovered an issue with loner assignment in the MDS that made the kclient-based dd succeed, although obviously this loner assignment could not be the sole cause.</p> <p>Jim reported back that this didn't seem to alleviate his problem at all. Now we need to fix it.</p>			

History

#1 - 02/08/2011 10:21 AM - Greg Farnum

- Status changed from New to Resolved

From Jim:

Also FWIW, I'm currently unable to reproduce the slow write after truncate issue, using stable branch on server side and 2.6.38-rc4 + ceph-client master branch (commit 1afeac8b0) on client side. I'm not sure what was going on with my above report; sorry for the noise.