

Ceph - Bug #9217

wrong value of osd::op_w_latency

08/25/2014 02:40 AM - Xinxin Shu

Status:	Resolved	Start date:	08/25/2014
Priority:	Normal	Due date:	
Assignee:	Xinxin Shu	% Done:	0%
Category:	OSD	Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	Community (dev)	Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Description			
<p>when i use fio to test rbd performance (4kb random write) , i found that osd::op_w_latency which i got from perf counter is much larger than the latency reported by fio, after i digged into it and found that, with XFS , after data is written to journal , osd sends a reply to client to imply that the write request has completed , it updates the osd::op_w_latency perf counter in ReplicatedPG::eval_repop function , when data is written to filestore, the osd::op_w_latency perf counter updates again, so op_w_latency perf counter records the apply latency , not the commit latency , is this the defined behavior or a bug?</p>			

Associated revisions

Revision f3bf2468 - 08/26/2014 06:19 AM - xinxin shu

fix wrong value of op_w_latency perf counter

Fixes: #9217

Signed-off-by: xinxin shu <xinxin.shu@intel.com>

History

#1 - 08/25/2014 02:43 AM - Xinxin Shu

- Category set to OSD

#2 - 08/25/2014 09:36 AM - Sage Weil

- Source changed from other to Community (dev)

I think it's a bug:

```
osd_plb.add_time_avg(l_osd_op_w_rlat, "op_w_rlat"); // client write readable/applied latency
osd_plb.add_time_avg(l_osd_op_w_lat, "op_w_latency"); // client write latency
```

#3 - 08/25/2014 05:39 PM - Xinxin Shu

- Assignee set to Xinxin Shu

#4 - 08/25/2014 08:49 PM - Xinxin Shu

- Status changed from New to Need Review

#5 - 08/25/2014 08:50 PM - Xinxin Shu

<https://github.com/ceph/ceph/pull/2325>

#6 - 08/26/2014 07:50 AM - Sage Weil

- Status changed from Need Review to Resolved