RADOS - Bug #26998

IOPS churn with "osd op queue" = "mclock_opclass" or "mclock_client"

08/22/2018 02:13 AM - Kefu Chai

Status:	Resolved	**% Done:**	0%
Priority:	Normal	**Spent time:**	0.00 hour
Assignee:	Kefu Chai		
Category:	Performance/Resource Usage		
Target version:			
Source:		**Affected Versions:**	
Tags:		**ceph-qa-suite:**	
Backport:		**Component(RADOS):**	OSD
Regression:	No	**Pull request ID:**	
Severity:	3 - minor	**Crash signature (v1):**	
Reviewed:		**Crash signature (v2):**	

Description

for more details on this issue, please refer to https://github.com/ceph/dmclock/pull/58 . in short, if "osd op queue" is "mclock_opclass" or "mclock_client" on OSD, the IOPS churns around expected number.

History

#1 - 08/22/2018 02:14 AM - Kefu Chai

- Assignee set to Kefu Chai

- https://github.com/ceph/dmclock/pull/58
- https://github.com/ceph/ceph/pull/23643

#2 - 08/22/2018 02:56 AM - Kefu Chai

- Status changed from New to 17

#3 - 08/23/2018 11:23 AM - Kefu Chai

- Status changed from 17 to Resolved