

## Ceph - Bug #891

### osd: fix last\_epoch\_started updates

03/15/2011 09:28 PM - Sage Weil

<b>Status:</b>	Resolved	<b>Start date:</b>	03/15/2011
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Sage Weil	<b>% Done:</b>	0%
<b>Category:</b>	OSD	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v0.26	<b>Spent time:</b>	2.00 hours
<b>Source:</b>		<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>		<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor		
<b>Description</b>			
<p>last_epoch_started is used to bound how far back in time we query other OSDs in order to recovery PG state (this is the prior_set). Currently we update last_epoch_started in PG::activate() and broadcast to replicas, but we do this before any state we recovered hits disk. This means we might hear about a last_epoch_started X even though the nodes in X (primary OR replicas) crashed before committing that recovered state to disk.</p> <p>Instead, we should activate all peers, everyone commit, send acks back to the primary saying "yes, I have committed the recovered pg info and log for this interval", and only then, once all replicas have done so, update last_epoch_started, rebroadcast to replicas, and queue for disk.</p>			
<b>Related issues:</b>			
Related to Ceph - Bug #865: osd: mark pg clean only after purging strays		<b>Won't Fix</b>	<b>03/04/2011</b>

### History

#### #1 - 03/15/2011 10:36 PM - Sage Weil

- Status changed from New to Resolved
- Assignee set to Sage Weil

[844093fbf7cc47add68e5e8fcac52c6ec875dcc3](#)

#### #2 - 03/20/2011 09:25 PM - Sage Weil

- translation missing: en.field\_story\_points set to 2
- translation missing: en.field\_position set to 555

#### #3 - 03/28/2011 10:48 AM - Sage Weil

- translation missing: en.field\_position deleted (575)
- translation missing: en.field\_position set to 574

#### #4 - 03/28/2011 10:50 AM - Sage Weil

- translation missing: en.field\_story\_points changed from 2 to 3
- translation missing: en.field\_position deleted (574)
- translation missing: en.field\_position set to 574