

## Ceph - Bug #55

### osd: fix transition from snaps -> no snaps -> snaps

04/21/2010 03:28 PM - Sage Weil

<b>Status:</b>	Resolved	<b>Start date:</b>	04/21/2010
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Yehuda Sadeh	<b>% Done:</b>	0%
<b>Category:</b>	OSD	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v0.23	<b>Spent time:</b>	0.00 hour
<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>Crash signature:</b>	
<b>Description</b>			
After filestore btrfs snaps = 1 -> 0, we may have leftover snap dirs. Later, if we turn them back on again, we'll roll back to a horribly ancient state.			
<b>Related issues:</b>			
Blocked by Ceph - Feature #359: osd: use new btrfs snapshot ioctls		<b>Resolved</b>	<b>08/17/2010</b>

#### History

##### #1 - 06/02/2010 12:14 PM - Sage Weil

- Target version changed from v0.21 to v0.22

##### #2 - 09/09/2010 12:23 PM - Sage Weil

- Target version changed from v0.22 to v0.23

##### #3 - 10/22/2010 11:42 AM - Sage Weil

I think all we need to do is look at current/commit\_op\_seq. If it is greater than the newest snap, than that snap is stale, because someone went and committed to current/.

An optional arg --osd-use-stale-snap bool would override that. That way you could still manually roll back to an older snap (by, say, deleting/renaming the newest) if you needed to for some reason.

##### #4 - 10/22/2010 04:28 PM - Sage Weil

- Assignee set to Yehuda Sadeh

##### #5 - 10/27/2010 01:21 PM - Yehuda Sadeh

- Status changed from New to Resolved

Fixed with [c1d078160a454c92fea899659d506e0b0ab7d92b](https://github.com/c1d078160a454c92fea899659d506e0b0ab7d92b).

##### #6 - 11/10/2010 09:20 AM - Sage Weil

- translation missing: en.field\_position deleted (317)

- translation missing: en.field\_position set to 183