

## CephFS - Feature #19109

### Use data pool's 'df' for statfs instead of global stats, if there is only one data pool

02/28/2017 01:37 PM - John Spray

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Immediate	
<b>Assignee:</b> Douglas Fuller	
<b>Category:</b> Administration/Usability	
<b>Target version:</b> v12.1.0	
<b>Source:</b> other	<b>Affected Versions:</b>
<b>Tags:</b>	<b>Component(FS):</b> Client
<b>Backport:</b> luminous	<b>Labels (FS):</b>
<b>Reviewed:</b>	<b>Pull request ID:</b>
<b>Description</b>	
<p>The client sends a MStatfs to the mon to get the info for a statfs system call. Currently the mon gives it the global usage information of the cluster.</p> <p>If a filesystem has exactly one data pool, it should be straightforward to modify this to pass a pool ID and respond to statfs with pool-specific stats instead.</p> <p>In the case of filesystems with multiple data pools, we need to fall back to the global stats. Doing a sum of the per-pool stats would probably not make sense because the pools could use overlapping sets of OSDs and we wouldn't want to double-report space.</p> <p>This would require updates to the mon, fuse client, and kernel client.</p> <p>Opened from: <a href="https://bugzilla.redhat.com/show_bug.cgi?id=1427512">https://bugzilla.redhat.com/show_bug.cgi?id=1427512</a></p>	
<b>Related issues:</b>	
Related to Linux kernel client - Subtask #20669: add support for filtered sta...	<b>Resolved</b> 07/18/2017
Duplicated by CephFS - Bug #21734: mount client shows total capacity of clust...	<b>Duplicate</b> 10/09/2017
Duplicated by CephFS - Bug #22091: statfs get wrong fs size	<b>Duplicate</b> 11/09/2017

## History

### #1 - 03/17/2017 08:45 AM - Ramana Raja

- Description updated

### #2 - 07/12/2017 04:25 AM - Patrick Donnelly

- Assignee set to Douglas Fuller

- Priority changed from Normal to Immediate

- Target version set to v12.1.0

- Source set to other

Doug, please take this one.

### #3 - 07/13/2017 04:08 PM - Douglas Fuller

What value for free space should we give in this case? If it's the global free space, it might be misleading to specify only the space used by a single data pool.

### #4 - 07/14/2017 12:19 AM - John Spray

You'd give the available space for that pool (i.e. how many bytes can they write before it becomes full). Same as in the "ceph df" output.

**#5 - 07/17/2017 09:31 PM - Douglas Fuller**

- Status changed from New to Fix Under Review

<https://github.com/ceph/ceph/pull/16378>

**#6 - 07/18/2017 02:48 PM - Douglas Fuller**

- Related to Subtask #20669: add support for filtered statfs in kernel client added

**#7 - 08/08/2017 05:35 PM - Patrick Donnelly**

Partially resolved by: <https://github.com/ceph/ceph/commit/eabe6626141df3f1b253c880aa6cb852c8b7ac1d>

**#8 - 08/17/2017 06:14 PM - Patrick Donnelly**

- Status changed from Fix Under Review to Pending Backport

- Backport set to luminous

Waiting for

<https://github.com/ceph/ceph-client/commit/b7f94d6a95dfe2399476de1e0d0a7c15c01611d0>

to be merged upstream.

In the meantime, we need a backport of <https://github.com/ceph/ceph/pull/16378> to luminous.

**#9 - 08/17/2017 06:22 PM - Douglas Fuller**

There's no need to wait for the kernel client since the message encoding is versioned. This has already been merged into luminous. (4aef4eab99c63d6ed84968e7ac6861ca081e737c, 76fc99a0a5112847ecc0eb7de8730c1a56c5930c, aeeeb43da90b61f8dad83251991088d667b7be47)

**#10 - 08/17/2017 06:41 PM - Patrick Donnelly**

- Status changed from Pending Backport to Resolved

Oh, oops. I forgot I merged this into luminous. Thanks Doug.

**#11 - 11/01/2017 06:38 PM - Patrick Donnelly**

- Duplicated by Bug #21734: mount client shows total capacity of cluster but not of a pool added

**#12 - 11/15/2017 05:52 PM - Patrick Donnelly**

- Duplicated by Bug #22091: statfs get wrong fs size added