

## CephFS - Bug #21745

### mds: MDBalancer using total (all time) request count in load statistics

10/10/2017 09:47 AM - John Spray

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> High	
<b>Assignee:</b> Zheng Yan	
<b>Category:</b>	
<b>Target version:</b> v13.0.0	
<b>Source:</b> Community (dev)	<b>Affected Versions:</b> v12.2.5
<b>Tags:</b> balancer	<b>ceph-qa-suite:</b>
<b>Backport:</b> luminous	<b>Component(FS):</b> MDS
<b>Regression:</b> No	<b>Labels (FS):</b> multimds
<b>Severity:</b> 2 - major	<b>Pull request ID:</b>
<b>Reviewed:</b>	<b>Crash signature:</b>
<b>Description</b>	
This was pointed out by Xiaoxi Chen	
The get_req_rate() function is returning the value of l_mds_request, which is a counter.	
This is then used in the load calculation in MDBalancer, resulting in crazy high values like this:	
<pre>2017-10-09 10:05:09.128325 7fc899748700 0 mds.0.bal mds.1 mdsload&lt;[0,0 0]/[0,0 0], req 3.11991e+07, hr 0, qlen 0, cpu 0.12&gt; = 3.11991e+07 ~ 15711.7</pre>	
<b>Related issues:</b>	
Copied to CephFS - Backport #23671: luminous: mds: MDBalancer using total (al...	<b>Resolved</b>

## History

### #1 - 10/11/2017 02:40 AM - Xiaoxi Chen

although it is simple to add last\_timestamp and last\_reqcount so that we can get an average TPS, but TPS may fluctuate a lot, which may result dirfrag ping-pong between multi mdss.

We probably need longer(configurable?) average for high fluctuate value like q\_len and req\_rate.

### #2 - 04/09/2018 08:48 PM - Patrick Donnelly

- Subject changed from MDBalancer using total (all time) request count in load statistics to mds: MDBalancer using total (all time) request count in load statistics

- Category deleted (90)

- Assignee set to Zheng Yan

- Target version set to v13.0.0

- Source set to Community (dev)

- Tags set to multimds,balancer

- Severity changed from 3 - minor to 2 - major

- Affected Versions v12.2.5 added

- Component(FS) MDS added

<https://github.com/ceph/ceph/pull/19220/commits/fb8d07772ffd3b061d2752c6b3375f6cb187be4b>

Zheng, please amend the above commit that it fixes this issue.

**#3 - 04/10/2018 01:22 AM - Zheng Yan**

- Status changed from New to Fix Under Review

<https://github.com/ceph/ceph/pull/19220/commits/e9689c1ff7e75394298c0e86aa9ed4e703391c3e>

**#4 - 04/11/2018 05:14 AM - Patrick Donnelly**

- Status changed from Fix Under Review to Pending Backport

- Tags changed from multimds,balancer to balancer

- Labels (FS) multimds added

**#5 - 04/12/2018 01:35 AM - Nathan Cutler**

- Copied to Backport #23671: luminous: mds: MDBalancer using total (all time) request count in load statistics added

**#6 - 06/13/2018 05:15 PM - Nathan Cutler**

- Status changed from Pending Backport to Resolved