

## RADOS - Bug #24768

### rgw workload makes osd memory explode

07/04/2018 07:17 PM - Sage Weil

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Urgent	<b>Spent time:</b> 0.00 hour
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b>	<b>Affected Versions:</b>
<b>Tags:</b>	<b>ceph-qa-suite:</b>
<b>Backport:</b> mimic,luminous,jewel	<b>Component(RADOS):</b>
<b>Regression:</b> No	<b>Pull request ID:</b>
<b>Severity:</b> 3 - minor	<b>Crash signature:</b>
<b>Reviewed:</b>	

#### Description

From ML,

On 07/03/2018 05:55 AM, Sage Weil wrote:

> On Fri, 29 Jun 2018, Aleksei Gutikov wrote:

> > Throughput is 100% the same, just sliced into bigger chunks (rados objects).

> > And this throughput is not high, less than single object per second. And

> > memory stay occupied even after writing stopped.

> >

> > Currently I'm sure that is side effect of sharing buffer::raw object among

> > different buffer::ptr objects.

> >

> > Please, have a look into this dump of ObjectContext::attr\_cache of one of

> > context in PrimaryLogPG::object\_contexts, made after uploading single 4M

> > object into S3.

> > Notice "\_user.rgw.idtag" and "\_user.rgw.tail\_tag" xattrs, both 44 bytes

> > length, holding 4194304 bytes buffer::raw object (nref=2).

>

> That is the smoking gun! What version is this?

Particularly this dump from 12.2.2

But issue was also reproducible for 12.2.5 and master.

#### Related issues:

Copied to RADOS - Backport #24805: mimic: rgw workload makes osd memory explode

**Resolved**

Copied to RADOS - Backport #24806: luminous: rgw workload makes osd memory ex...

**Resolved**

Copied to RADOS - Backport #24847: jewel: rgw workload makes osd memory explode

**Resolved**

#### History

#1 - 07/05/2018 12:44 PM - Sage Weil

- Status changed from 12 to Fix Under Review

<https://github.com/ceph/ceph/pull/22858>

#2 - 07/06/2018 06:44 PM - Sage Weil

- Status changed from Fix Under Review to Pending Backport

#3 - 07/06/2018 09:06 PM - Nathan Cutler

- Copied to Backport #24805: *mimic: rgw workload makes osd memory explode added*

**#4 - 07/06/2018 09:06 PM - Nathan Cutler**

- Copied to Backport #24806: *luminous: rgw workload makes osd memory explode added*

**#5 - 07/10/2018 09:33 AM - Kefu Chai**

jewel backport: <https://github.com/ceph/ceph/pull/22959>

i knew that jewel is (almost) EOL. just in case anyone is still stuck with it.

**#6 - 07/10/2018 09:54 AM - Nathan Cutler**

- Backport changed from *mimic,luminous* to *mimic,luminous,jewel*

**#7 - 07/10/2018 09:54 AM - Nathan Cutler**

- Copied to Backport #24847: *jewel: rgw workload makes osd memory explode added*

**#8 - 10/16/2018 08:51 AM - Kefu Chai**

- Status changed from *Pending Backport* to *Resolved*