

rgw - Bug #47074

rgw: bucket radoslist looping endlessly while processing a bucket

08/21/2020 07:20 PM - J. Eric Ivancich

Status:	New	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	J. Eric Ivancich		
Category:			
Target version:	v16.0.0		
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			
Description			
This was reported to via irc by Manuel Rios Fernandez. This tracker is an official channel exchange information.			

History

#1 - 08/21/2020 07:29 PM - J. Eric Ivancich

Manuel, you asked what I would need to help debug.

Please run:

```
radosgw-admin bucket radoslist --bucket=<bucketname> --debug-rgw=20 > radoslist-out.txt 2> radoslist-err.txt
```

and monitor the tail of "radoslist-err.txt" for objects being processed, choose an object, and grep the file for it. When you see the same object going through the exact same processing at, say, 10 different time instances, you can stop the process and attach "radoslist-err.txt". NOTE: I don't mean 10 mentions of the same object, I mean the equivalent sequence of log lines happening at 10 different instances; I hope that's clear.

Additionally:

```
radosgw-admin bi list --bucket=<bucketname> > bilist-out.txt
```

and attach "bilist-out.txt".

Let me know if you have any questions.

#2 - 08/27/2020 09:17 PM - J. Eric Ivancich

Manuel Rios Fernandez provided me with some logs to examine. Because they contain information that is not public, they were not attached to this tracker.

Those logs do not show a bug, however.

If we grep radoslist-err.txt for "start_after", we'll see that the value of start_after increases each time. That indicates we're making forward progress through the bucket and that we're not stuck in a loop.

If we run the following sequence of commands:

```
grep start_after radoslist-err.txt | sed -e 's/.*start_after="//' -e 's/", prefix=.*//' > sa1.txt
LANG=C sort sa1.txt > sa2.txt
diff sa1.txt sa2.txt
```

we'll see that the start_after values are sorted in the log and increase. Using 'uniq' we can also see that no values are repeated.

#3 - 08/27/2020 09:18 PM - J. Eric Ivancich

Manuel Rios Fernandez: I know you told me the answer to this via irc, but would you add the version of ceph you're running to this tracker? Thanks!

#4 - 08/27/2020 09:20 PM - Manuel Rios

Im running:

```
ceph version 14.2.11 (f7fdb2f52131f54b891a2ec99d8205561242cdaf) nautilus (stable)
```

#5 - 09/03/2020 02:33 PM - Casey Bodley

- Subject changed from *rgw: orphan list looping endlessly while processing a bucket* to *rgw: bucket radoslist looping endlessly while processing a bucket*

#6 - 09/10/2020 02:23 PM - Casey Bodley

Eric Ivancich wrote:

Manuel Rios Fernandez provided me with some logs to examine. Because they contain information that is not public, they were not attached to this tracker.

Those logs do not show a bug, however.

If we grep radoslist-err.txt for "start_after", we'll see that the value of start_after increases each time. That indicates we're making forward progress through the bucket and that we're not stuck in a loop.

If we run the following sequence of commands:
[...]

we'll see that the start_after values are sorted in the log and increase. Using 'uniq' we can also see that no values are repeated.

Do we still suspect there's an infinite loop here?

#7 - 12/08/2020 06:47 PM - James McEwan

I'm experiencing the same issue in my environment. I recently attempted to run the 'rgw-orphan-list' tool against our cluster (octopus 15.2.7) to identify any orphans and noticed that the 'radosgw-admin bucket radoslist' command appeared to be stuck in a loop.

I saw in the 'radosgw-admin-XXXXXX.intermediate' output file the same sequence of objects looping repeatedly, and the command would not progress.

I ran the command 'radosgw-admin bucket radoslist --bucket=<bucketname> --debug-rgw=20 > radoslist-out.txt 2> radoslist-err.txt' in order to identify what was happening,

In the radoslist-err.txt file, I see entries similar to the following, which just repeat until I cancel the command:

```
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=345930 stripe_ofs=345930 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=360383 stripe_ofs=360383 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=302704 stripe_ofs=302704 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=381794 stripe_ofs=381794 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=334269 stripe_ofs=334269 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=378182 stripe_ofs=378182 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=325041 stripe_ofs=325041 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=295118 stripe_ofs=295118 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=315823 stripe_ofs=315823 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=346112 stripe_ofs=346112 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=303128 stripe_ofs=303128 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=364011 stripe_ofs=364011 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): result: ofs=331137 stripe_ofs=331137 part_ofs=0
rule->part_size=0
2020-12-04T12:37:22.227+0100 7fc405f79080 20 RGWObjManifest::operator++(): rule->part_size=0 rules.size()=1
```

Just before these entries begin, I see the following:

```
2020-12-04T12:37:22.203+0100 7fc405f79080 20 RGWRados::Bucket::List::list_objects_ordered INFO end of outer loop, truncated=1, count=4,
attempt=9
2020-12-04T12:37:22.203+0100 7fc405f79080 20 RGWRadosList::do_incomplete_multipart processing incomplete multipart entry
```

```
RGWMultipartUploadEntry{ obj.key="_multipart_/halon_2_settings_2020-08-28.json.2~TD-WUPWhOicUKkz4NTxZcxuEI3b_lqC.meta"
mp=RGWMPObj:{ prefix="/halon_2_settings_2020-08-28.json.2~TD-WUPWhOicUKkz4NTxZcxuEI3b_lqC",
meta="/halon_2_settings_2020-08-28.json.2~TD-WUPWhOicUKkz4NTxZcxuEI3b_lqC.meta" } }
2020-12-04T12:37:22.203+0100 7fc405f79080 20 RGWRadosList::do_incomplete_multipart processing incomplete multipart entry
RGWMultipartUploadEntry{ obj.key="_multipart_gl-events_0_2020-08-28.json.2~OQWHwayxwHwUMLteZWQZWRnsdgT5CI2.meta"
mp=RGWMPObj:{ prefix="gl-events_0_2020-08-28.json.2~OQWHwayxwHwUMLteZWQZWRnsdgT5CI2",
meta="gl-events_0_2020-08-28.json.2~OQWHwayxwHwUMLteZWQZWRnsdgT5CI2.meta" } }
2020-12-04T12:37:22.203+0100 7fc405f79080 20 RGWRadosList::do_incomplete_multipart processing incomplete multipart entry
RGWMultipartUploadEntry{ obj.key="_multipart_gl-events_0_2020-08-28.json.2~zz-6rqcuHGmSPYTY2IRueMP5HIFOL6t.meta" mp=RGWMPObj:{
prefix="gl-events_0_2020-08-28.json.2~zz-6rqcuHGmSPYTY2IRueMP5HIFOL6t",
meta="gl-events_0_2020-08-28.json.2~zz-6rqcuHGmSPYTY2IRueMP5HIFOL6t.meta" } }
2020-12-04T12:37:22.203+0100 7fc405f79080 20 RGWRadosList::do_incomplete_multipart processing incomplete multipart entry
RGWMultipartUploadEntry{ obj.key="_multipart_gl-events_1_2020-08-28.json.2~-WCK4AJE1k7or-KN6qTePxfIvHL1Nlf.meta" mp=RGWMPObj:{
prefix="gl-events_1_2020-08-28.json.2~-WCK4AJE1k7or-KN6qTePxfIvHL1Nlf",
meta="gl-events_1_2020-08-28.json.2~-WCK4AJE1k7or-KN6qTePxfIvHL1Nlf.meta" } }
```

Can you see anything here that would confirm a bug?