

Ceph - Bug #9008

Objecter: pg listing can deadlock when throttling is in use

08/05/2014 02:44 AM - Guang Yang

Status:	Resolved	% Done:	0%
Priority:	Urgent	Spent time:	0.00 hour
Assignee:	Nathan Cutler		
Category:			
Target version:			
Source:	Community (user)	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	firefly	Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

In our Ceph cluster (with radosgw), we found that occasionally the processing threads hang forever and eventually hang all thread pool workers.

1> radosgw.log

The following log entries repeats forever even there is no load on the cluster.

```
2014-08-05 09:31:16.986181 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd21dc700' had timed out after 600
2014-08-05 09:31:16.986183 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd2bdd700' had timed out after 600
2014-08-05 09:31:16.986185 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd3fdf700' had timed out after 600
2014-08-05 09:31:16.986186 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd49e0700' had timed out after 600
2014-08-05 09:31:16.986188 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd53e1700' had timed out after 600
2014-08-05 09:31:16.986190 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd5de2700' had timed out after 600
2014-08-05 09:31:16.986191 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd67e3700' had timed out after 600
2014-08-05 09:31:16.986193 7f4d824ee700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f4cd71e4700' had timed out after 600
```

2> GDB backtrace

Taking one thread for example, here is the GDB backtrace:

```
(gdb) thread 1078
[Switching to thread 1078 (Thread 0x7f4cb0fa7700 (LWP 14399))]#0 0x00000032b740b43c in pthread_cond_wait@@GLIBC_2.3.2 () from /lib64/libpthread.so.0
(gdb) bt
#0 0x00000032b740b43c in pthread_cond_wait@@GLIBC_2.3.2 () from /lib64/libpthread.so.0
#1 0x00007f4d84e913e3 in Cond::Wait (this=0x7f4cb0fa4a70, mutex=...) at ./common/Cond.h:55
#2 0x00007f4d84e88f4b in librados::IoCtxImpl::operate_read (this=0x7f4af402a440, oid=..., o=0x7f4af402b740, pbl=0x7f4cb0fa4b80, flags=0) at librados/IoCtxImpl.cc:569
#3 0x00007f4d84e5c35f in librados::IoCtx::operate (this=0x7f4cb0fa4c80, oid=<value optimized out>, o=0x7f4cb0fa4c90, pbl=0x7f4cb0fa4b80) at librados/librados.cc:1039
#4 0x00000000052cae8 in RGWRados::obj_stat (this=0x1744010, ctx=<value optimized out>, obj=...,
```

```
psize=0x7f4af403aa80, pmtime=0x7f4af403aa88, epoch=0x7f4af403aa90, attrs=0x7f4af403ae10, first_chunk=0x0, objv_tracker=0x0)
  at rgw/rgw_rados.cc:5132
#5 0x000000000054f73c in RGWCache<RGWRados>::obj_stat (this=0x1744010, ctx=0x7f4cb0fa6670, obj=..
., psize=0x7f4af403aa80, pmtime=0x7f4af403aa88, pepoch=0x7f4af403aa90, attrs=0x7f4af403ae10, first
_chunk=0x0, objv_tracker=0x0)
  at rgw/rgw_cache.h:475
#6 0x0000000000532de4 in RGWRados::get_obj_state (this=0x1744010, rctx=0x7f4cb0fa6670, obj=..., s
tate=0x7f4cb0fa5810, objv_tracker=0x0) at rgw/rgw_rados.cc:3852
#7 0x0000000000534d64 in RGWRados::prepare_atomic_for_write_impl (this=0x1744010, rctx=<value opt
imized out>, obj=<value optimized out>, op=..., pstate=0x7f4cb0fa5810, reset_obj=true, ptag=0x7f4a
f4073648) at rgw/rgw_rados.cc:3978
#8 0x0000000000536c37 in RGWRados::put_obj_meta_impl (this=0x1744010, ctx=0x7f4cb0fa6670, obj=...
, size=1131685, mtime=0x7f4af4019920, attrs=std::map with 4 elements = {...}, category=RGW_OBJ_CAT
EGORY_MAIN, flags=1, rmatrs=0x0,
  data=0x7f4af40734e8, manifest=0x7f4af4073720, ptag=0x7f4af4073648, remove_objs=0x0, modify_ver
sion=false, objv_tracker=0x0, set_mtime=0, bucket_owner="flickrtest1") at rgw/rgw_rados.cc:2783
#9 0x000000000055412f in RGWCache<RGWRados>::put_obj_meta_impl (this=0x1744010, ctx=0x7f4cb0fa667
0, obj=..., size=1131685, mtime=0x7f4af4019920, attrs=std::map with 4 elements = {...}, category=R
GW_OBJ_CATEGORY_MAIN, flags=1,
  rmatrs=0x0, data=0x7f4af40734e8, manifest=0x7f4af4073720, ptag=0x7f4af4073648, remove_objs=0x
0, modify_version=false, objv_tracker=0x0, set_mtime=0, owner="flickrtest1") at rgw/rgw_cache.h:41
5
#10 0x0000000000541d28 in RGWRados::put_obj_meta (this=<value optimized out>, ctx=<value optimized
 out>, obj=<value optimized out>, size=<value optimized out>, attrs=<value optimized out>, categor
y=<value optimized out>, flags=1,
  params=...) at rgw/rgw_rados.h:1460
#11 0x000000000052e9f6 in RGWPutObjProcessor_Atomic::do_complete (this=0x7f4af4073490, etag=<value
 optimized out>, mtime=<value optimized out>, set_mtime=0, attrs=std::map with 4 elements = {...})
  at rgw/rgw_rados.cc:1264
#12 0x0000000000518809 in RGWPutObjProcessor::complete (this=0x7f4af4073490, etag=<value optimized
 out>, mtime=<value optimized out>, set_mtime=<value optimized out>, attrs=<value optimized out>)
  at rgw/rgw_rados.cc:986
#13 0x0000000000562d56 in RGWPutObj::execute (this=0x7f4af40197a0) at rgw/rgw_op.cc:1628
#14 0x00000000004cb522 in process_request (store=0x1744010, rest=<value optimized out>, req=0x7f4c
f0597650, client_io=0x7f4cb0fa6840, olog=0x0) at rgw/rgw_main.cc:600
#15 0x00000000004cc12a in RGWFCGXProcess::handle_request (this=0x179d830, r=<value optimized out>)
  at rgw/rgw_main.cc:631
#16 0x00000000004cf687 in RGWProcess::RGWWQ::_process (this=0x179dad0, req=0x7f4cf0597650) at rgw/
rgw_main.cc:200
#17 0x00007f4d84f1d031 in ThreadPool::worker (this=0x179d898, wt=0x7f4cf000afd0) at common/WorkQue
ue.cc:125
#18 0x00007f4d84f20070 in ThreadPool::WorkThread::entry (this=<value optimized out>) at common/Wor
kQueue.h:317
#19 0x00000032b7407851 in start_thread () from /lib64/libpthread.so.0
#20 0x00000032b70e890d in clone () from /lib64/libc.so.6
```

More information: * ceph version 0.80.1 (a38fe1169b6d2ac98b427334c12d7cf81f809b74) * platform: RHEL6

Related issues:

Copied to Ceph - Backport #12007: Objecter: pg listing can deadlock when thro...

Resolved

08/05/2014

History

#1 - 08/05/2014 08:28 AM - Sage Weil

please query the admin socket for the process like so:

```
ceph daemon /var/run/ceph/ceph-client.*.asok objecter_requests
```

to see the in flight requests. That will tell you which request(s) are blocked and with which OSD they are outstanding. Then you can check the OSD with

```
ceph daemon /var/run/ceph/ceph-osd.NNN.asok dump_ops_in_flight
```

to see what their state is there. I'm guessing the request is hung on teh OSD side of things...

#2 - 08/05/2014 08:28 AM - Sage Weil

- Status changed from New to Need More Info

#3 - 08/05/2014 06:15 PM - Guang Yang

I'm guessing the request is hung on teh OSD side of things...

Thanks Sage. Sadly after radosgw daemon restarting, I am not able to reproduce it for now (I will keep monitoring).

Even with in-flight OP, shouldn't they be timeout after some time without response? From radosgw's log, it seems it just hangs forever.

Following are logs for the same thread:

```
2014-08-05 06:09:35.448230 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:09:40.448390 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:09:45.448586 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:09:50.448730 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:09:55.448888 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:00.449061 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:05.449275 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:10.449521 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:15.449688 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:20.449893 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:25.450096 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:30.450232 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:35.450404 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:40.450660 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:45.450877 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:50.451132 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:10:55.451361 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:00.451537 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:05.451785 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:10.451987 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:15.452205 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:20.452401 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:25.452626 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:30.452823 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:35.453083 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
```

ad timed out after 600
2014-08-05 06:11:40.453248 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:45.453470 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:50.453713 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:11:55.453945 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:00.454124 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:05.454291 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:10.454522 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:15.454707 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:20.454907 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:25.455115 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:30.455417 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:35.455658 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:40.455947 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:45.456107 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:50.456318 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:12:55.456575 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:00.456801 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:05.457052 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:10.457330 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:15.457591 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:20.457815 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:25.458048 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:30.458224 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:35.458459 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600
2014-08-05 06:13:40.458695 7faadee30700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7fa9dc7f3700' h
ad timed out after 600

#4 - 09/03/2014 06:40 AM - Guang Yang

```
@radosgw
$ ceph daemon /var/run/ceph/ceph-client.*.asok objecter_requests { "ops": [ { "tid": 15582485,
"pg": "3.ed2ef9fb",
"osd": 146,
"object_id": "default.7291.42_2546732403_de301a716f_o.jpg",
"object_locator": "@3",
"target_object_id": "default.7291.42_2546732403_de301a716f_o.jpg",
"target_object_locator": "@3",
"paused": 0,
"used_replica": 0,
"precalc_pgid": 0,
"last_sent": "2014-09-03 12:48:31.135795",
"attempts": 1,
"snapid": "head",
"snap_context": "0=[]",
"mtime": "0.000000",
"osd_ops": [
"getxattrs",
"stat",
"read 0~5898240"]},
]...

```

```
@osd146
$ ceph daemon /var/run/ceph/ceph-osd.NNN.asok dump_ops_in_flight { "num_ops": 0,
"ops": []}

```

And currently all radosgw threads hang in the same way..

#5 - 09/03/2014 08:19 AM - Sage Weil

Guang Yang wrote:

```
@radosgw
$ ceph daemon /var/run/ceph/ceph-client.*.asok objecter_requests { "ops": [ { "tid": 15582485,
"pg": "3.ed2ef9fb",
"osd": 146,
"object_id": "default.7291.42_2546732403_de301a716f_o.jpg",
"object_locator": "@3",
"target_object_id": "default.7291.42_2546732403_de301a716f_o.jpg",
"target_object_locator": "@3",
"paused": 0,
"used_replica": 0,
"precalc_pgid": 0,
"last_sent": "2014-09-03 12:48:31.135795",
"attempts": 1,
"snapid": "head",
"snap_context": "0=[]",

```

```
"mtime": "0.000000",
"osd_ops": [
  "getxattrs",
  "stat",
  "read 0~5898240"],
]...
```

```
@osd146
$ ceph daemon /var/run/ceph/ceph-osd.NNN.asok dump_ops_in_flight { "num_ops": 0,
"ops": []}
```

And currently all radosgw threads hang in the same way..

Hmm, and this is reproducible? It may be the throttling in the msgr layer.. debug ms = 20 might help, if you can reproduce this situation. You could also ceph pg <pgid> query to confirm that the pg is peered and active. From this hung state, if you ceph osd down 146, will things then repeer and recover?

#6 - 09/04/2014 05:28 PM - Guang Yang

Sage Weil wrote:

Hmm, and this is reproducible? It may be the throttling in the msgr layer.. debug ms = 20 might help, if you can reproduce this situation. You could also ceph pg <pgid> query to confirm that the pg is peered and active. From this hung state, if you ceph osd down 146, will things then repeer and recover?

1. Yeah, it is very reproducible on v0.80.4, however, we didn't see this pattern before (prior to v0.80).
2. I can confirm that all PGs are active+clean during the time.
3. Let me increase the debug ms to 20 and check.

Thanks!

#7 - 09/10/2014 07:17 AM - Guang Yang

While in the progress of testing with debug ms = 20, I have a couple of questions:

- 1> What is the risk of adding timeout for radosgw op? My understanding is that the change should be around librados::IoCtxImpl::operate_* to add a timeout for the conditional wait, is that correct?
- 2> Is it possible it was caused by network jitter (NIC error, etc.)? I assume not since there is explicitly message ack handling at messenger layer, but it would be nice to get confirmed?
- 3> By throttling you mentioned previously, can you elaborate what type of throttling might cause this behavior?

Thanks,
Guang

#8 - 09/10/2014 07:18 AM - Guang Yang

Another hang we observed with debug ms = 1 at radosgw side, it is confirmed that the `osd_op` does not have its peer `osd_op_reply` which hang the thread.

#9 - 09/10/2014 07:47 AM - Guang Yang

- File `bug9008.log` added

radosgw side log before hang.

#10 - 09/11/2014 06:08 AM - Guang Yang

I finally got some logs which revealed that the `osd_op_reply` message had been received by Pipe, but it stayed in `DispatchQueue` and there was no dispatcher trying to dispatch it (and the dispatcher thread was throttled).

Radosgw WQ thread logs:

```
2014-09-09 14:46:24.857510 7f657ad0a700 20 -- 10.214.140.210:0/1064300 submit_message osd_op(client.29459.0:174471939 [cmpxattr user.rgw.idtag (22) op 1 mode 1,create 0~0,delete,setxattr user.rgw.idtag (23),writefull 0~2784877,setxattr user.rgw.manifest (466),setxattr user.rgw.acl (127),setxattr user.rgw.content_type (25),setxattr user.rgw.etag (33),setxattr user.rgw.x-amz-meta-origin (56)] 3.3000ebaa ondisk+write e523) v4 remote, 10.214.141.21:6805/33707, have pipe.
2014-09-09 14:56:25.939603 7f66b0219700 1 heartbeat_map is_healthy 'RGWProcess::m_tp thread 0x7f657ad0a700' had timed out after 600
```

reader/writer pipe logs:

```
2014-09-09 14:46:24.857650 7f63b9608700 20 -- 10.214.140.210:0/1064300 >> 10.214.141.21:6805/33707 pipe(0x7f669c148b00 sd=325 :35922 s=2 pgs=1474 cs=1 l=1 c=0x7f669c13b530).writer encoding 212897 features 17592186044415 0x7f647c02ac50 osd_op(client.29459.0:174471939 default.29329.25_7633301044_eeb7d3d282_o.jpg [cmpxattr user.rgw.idtag (22) op 1 mode 1,create 0~0,delete,setxattr user.rgw.idtag (23),writefull 0~2784877,setxattr user.rgw.manifest (466),setxattr user.rgw.acl (127),setxattr user.rgw.content_type (25),setxattr user.rgw.etag (33),setxattr user.rgw.x-amz-meta-origin (56)] 3.3000ebaa ondisk+write e523) v4
2014-09-09 14:46:24.857476 7f657ad0a700 1 -- 10.214.140.210:0/1064300 --> 10.214.141.21:6805/33707 -- osd_op(client.29459.0:174471939 default.29329.25_7633301044_eeb7d3d282_o.jpg [cmpxattr user.rgw.idtag (22) op 1 mode 1,create 0~0,delete,setxattr user.rgw.idtag (23),writefull 0~2784877,setxattr user.rgw.manifest (466),setxattr user.rgw.acl (127),setxattr user.rgw.content_type (25),setxattr user.rgw.etag (33),setxattr user.rgw.x-amz-meta-origin (56)] 3.3000ebaa ondisk+write e523) v4 -- ?+0 0x7f647c02ac50 con 0x7f669c13b530
2014-09-09 14:46:25.304788 7f63b48bb700 10 -- 10.214.140.210:0/1064300 >> 10.214.141.21:6805/33707 pipe(0x7f669c148b00 sd=325 :35922 s=2 pgs=1474 cs=1 l=1 c=0x7f669c13b530).reader got message 235178 0x7f64d80fb410 osd_op_reply(174471939 default.29329.25_7633301044_eeb7d3d282_o.jpg [cmpxattr (22) op 1 mode 1,create 0~0,delete,setxattr (23),writefull 0~2784877,setxattr (466),setxattr (127),setxattr (25),setxattr (33),setxattr (56)] v523'20313 uv20313 ondisk = 0) v6
2014-09-09 14:46:25.304821 7f63b48bb700 20 -- 10.214.140.210:0/1064300 queue 0x7f64d80fb410 prio 127
```

And continue searching 'done calling dispatch on 0x7f64d80fb410', nothing could be found.

So that we can see that even the message was received, it failed to be dispatched and thus not able to deliver to upper layer (radosgw wq thread).

Another pattern I saw was that as those `osd_op_reply` didn't get dispatched, the objecter throttle hit its limit (since it the tokens were release upon dispatch), lots of following radosgw wq thread got hung with another pattern (waiting on throttler).

However the dispatcher thread was hung there as well

```
2014-09-09 14:46:24.984694 7f66aee17700 1 -- 10.214.140.210:0/1064300 <== osd.107 10.214.140.25:6810/3812 540
```

```
190 ==== osd_op_reply(174471909 [pgls start_epoch 0] v0'0 uv0 ondisk = 1) v6 ==== 167+0+44 (1640671103 0 1390
81063) 0x7f64f80bd3c0 con 0x7f669c0c46e0
2014-09-09 14:46:24.984742 7f66aee17700 2 throttle(objecter_bytes 0x204e468) _wait waiting...
```

After looking at the code, I failed to find the reason why dispatcher will need to wait for throttle budget... It looks like it should only release budget, or did I miss anything obvious?

#11 - 09/11/2014 01:57 PM - Greg Farnum

Well, the dispatcher doesn't normally take budget directly, but it could be doing something else farther down the call chain that requires sending a new message. Do you have more complete logs you can share for somebody to look at?

I went over it briefly with Yehuda and we didn't see any obvious way that this should be able to deadlock, but that's my guess. You could test it by disabling objecter throttling and seeing if the rgw hangs disappear.

#12 - 09/11/2014 02:34 PM - Greg Farnum

- Subject changed from *rgw: processing thread hangs forever if the pending OP does not get replied to Objecter: pg listing can deadlock when throttling is in use*

- Status changed from *Need More Info* to *12*

Okay, this is because our object listing code is incorrect (in both Firefly and Giant-to-be). A pgls response has a context which **directly** triggers a new pgls message to go out — but if the Objecter doesn't have any available throttle, it blocks. And this happens in the dispatch thread which calls `handle_osd_op_reply` (in Firefly that's the main dispatch loop; in Giant it's the Pipe reader thread).

We might just be able to throw the callback into the Finisher thread instead — although then the Finisher thread can block, which isn't exactly ideal either. :/ The other option is to somehow have pg listing replies and outgoing continuation messages not drop and take op budget...

#13 - 09/11/2014 02:42 PM - Sage Weil

my vote is to make the pgls continuation hold onto existing budget (and not take new budget). is that feasible?

#14 - 09/15/2014 04:52 AM - Guang Yang

Please help to review: <https://github.com/ceph/ceph/pull/2489>

#15 - 09/23/2014 09:25 AM - Ian Colle

- Project changed from *rgw* to *Ceph*

#16 - 09/23/2014 01:27 PM - Samuel Just

- Status changed from *12* to *7*

#17 - 09/29/2014 02:11 PM - Samuel Just

- Priority changed from *High* to *Urgent*

#18 - 10/07/2014 11:28 AM - Samuel Just

- Status changed from *7* to *Pending Backport*

#19 - 12/19/2014 04:05 PM - Chris Armstrong

Hi folks, we're running into this on Giant. Is there any estimate as to when this'll be fixed in a maintenance release? Thanks!

#20 - 06/14/2015 07:50 AM - Nathan Cutler

- Assignee set to *Nathan Cutler*

- Backport set to *firefly*

- Regression set to *No*

#21 - 09/04/2015 03:19 PM - Loïc Dachary

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

Files

bug9008.log	155 KB	09/10/2014	Guang Yang
-------------	--------	------------	------------