

Ceph - Bug #1151

OSD: CephxClientHandler::handle_response

06/09/2011 03:06 AM - Wido den Hollander

Status:	Resolved	Start date:	06/09/2011
Priority:	Normal	Due date:	
Assignee:	Sage Weil	% Done:	0%
Category:	cephx	Estimated time:	0.00 hour
Target version:	v0.30	Spent time:	1.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

I just saw a bunch of OSD's going down with:

```
(gdb) bt
#0 0x00007f609f8e97bb in raise () from /lib/libpthread.so.0
#1 0x00000000063d443 in reraise_fatal (signum=4000) at common/signal.cc:61
#2 0x00000000063e55b in handle_fatal_signal (signum=6) at common/signal.cc:108
#3 <signal handler called>
#4 0x00007f609e4b9a75 in raise () from /lib/libc.so.6
#5 0x00007f609e4bd5c0 in abort () from /lib/libc.so.6
#6 0x00007f609ed6f8e5 in __gnu_cxx::__verbose_terminate_handler() () from /usr/lib/libstdc++.so.6
#7 0x00007f609ed6dd16 in ?? () from /usr/lib/libstdc++.so.6
#8 0x00007f609ed6dd43 in std::terminate() () from /usr/lib/libstdc++.so.6
#9 0x00007f609ed6de3e in __cxa_throw () from /usr/lib/libstdc++.so.6
#10 0x0000000006077e1 in ceph::__ceph_assert_fail (assertion=<value optimized out>, file=<value optimized out>, line=<value optimized out>, func=0x6c29a0 "virtual int CephxClientHandler::handle_response(int, ceph::buffer::list::iterator&)" at common/assert.cc:50)
#11 0x000000000657b09 in CephxClientHandler::handle_response (this=0x91b630, ret=-1627771328, ind ata=...) at auth/cephx/CephxClientHandler.cc:162
#12 0x00000000062fa58 in MonClient::handle_auth (this=0x7ff1a1b5fe0, m=0x6793200) at mon/MonClient.cc:438
#13 0x000000000630ebc in MonClient::ms_dispatch (this=0x7ff1a1b5fe0, m=0x6793200) at mon/MonClient.cc:272
#14 0x0000000006245a3 in Messenger::ms_deliver_dispatch (this=0x1c9b000) at msg/Messenger.h:98
#15 SimpleMessenger::dispatch_entry (this=0x1c9b000) at msg/SimpleMessenger.cc:353
#16 0x00000000048fd9c in SimpleMessenger::DispatchThread::entry (this=0x1c9b488) at msg/SimpleMessenger.h:544
#17 0x00007f609f8e09ca in start_thread () from /lib/libpthread.so.0
#18 0x00007f609e56c70d in clone () from /lib/libc.so.6
#19 0x0000000000000000 in ?? ()
```

The logging was on:

```
debug osd = 20
debug ms = 1
```

My logs showed:

```
Jun 8 18:39:13 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_boot
```

```

Jun  8 18:39:13 atom2 osd.8[4000]: 7f6094259700 osd8 41918 assuming cluster_addr ip matches clien
t_addr
Jun  8 18:39:13 atom2 osd.8[4000]: 7f6094259700 osd8 41918 assuming hb_addr ip matches cluster_ad
dr
Jun  8 18:39:13 atom2 osd.8[4000]: 7f6094259700 osd8 41918 client_addr [2a00:f10:113:1:225:90ff:f
e33:49f2]:6806/3999, cluster_addr [2a00:f10:113:1:225:90ff:fe33:49f2]:6807/3999, hb addr [2a00:f10
:113:1:225:90ff:fe33:49f2]:6808/3999
Jun  8 18:39:13 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_alive up_thru currently 41870 want
0
Jun  8 18:39:13 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_pg_stats
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6092a56700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 m
ark_down 0x1c941dc0 -- 0x2016da00
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6092a56700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 -
-> [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- auth(proto 0 26 bytes) v1 -- ?+0 0x6793600 con 0x
1c941c80
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6092154700 osd8 41918 OSD::ms_get_authorizer type=mon
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 ms_handle_connect on mon
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_boot
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 assuming cluster_addr ip matches clien
t_addr
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 assuming hb_addr ip matches cluster_ad
dr
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 client_addr [2a00:f10:113:1:225:90ff:f
e33:49f2]:6806/3999, cluster_addr [2a00:f10:113:1:225:90ff:fe33:49f2]:6807/3999, hb addr [2a00:f10
:113:1:225:90ff:fe33:49f2]:6808/3999
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_alive up_thru currently 41870 want
0
Jun  8 18:39:16 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_pg_stats
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6092a56700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 m
ark_down 0x1c941c80 -- 0x2016d280
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 <
== mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 1 ==== auth_reply(proto 2 0 Success) v1 ==== 33
+0+0 (3727614903 0 0) 0x6793800 con 0x1c941c80
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6092a56700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 -
-> [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- auth(proto 0 26 bytes) v1 -- ?+0 0x6793200 con 0x
1c941b40
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 -
-> [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- auth(proto 2 32 bytes) v1 -- ?+0 0x6793800 con 0x
1c941b40
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6092053700 osd8 41918 OSD::ms_get_authorizer type=mon
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 ms_handle_connect on mon
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_boot
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 assuming cluster_addr ip matches clien
t_addr
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 assuming hb_addr ip matches cluster_ad
dr
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 client_addr [2a00:f10:113:1:225:90ff:f
e33:49f2]:6806/3999, cluster_addr [2a00:f10:113:1:225:90ff:fe33:49f2]:6807/3999, hb addr [2a00:f10
:113:1:225:90ff:fe33:49f2]:6808/3999
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_alive up_thru currently 41870 want
0
Jun  8 18:39:19 atom2 osd.8[4000]: 7f6094259700 osd8 41918 send_pg_stats
Jun  8 18:39:21 atom2 osd.8[4000]: 7f6094259700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/3999 <
== mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 1 ==== auth_reply(proto 2 0 Success) v1 ==== 33
+0+0 (3639190872 0 0) 0x6793200 con 0x1c941b40
Jun  8 18:39:21 atom2 osd.8[4000]: 7f6094259700 cephx client: unknown request_type 64001
Jun  8 18:39:21 atom2 osd.8[4000]: auth/cephx/CephxClientHandler.cc: In function 'virtual int Ceph
xClientHandler::handle_response(int, ceph::buffer::list::iterator&)', in thread '0x7f6094259700'#0
12auth/cephx/CephxClientHandler.cc: 162: FAILED assert(0)
Jun  8 18:39:21 atom2 osd.8[4000]: ceph version 0.28.2-260-ge2c808a (commit:e2c808aea97ced6e9a55b
143116b66d344f72c0b)#012 1: (CephxClientHandler::handle_response(int, ceph::buffer::list::iterator
&)+0x369) [0x657b09]#012 2: (MonClient::handle_auth(MAuthReply*)+0xb8) [0x62fa58]#012 3: (MonClie
nt::ms_dispatch(Message*)+0x26c) [0x630ebc]#012 4: (SimpleMessenger::dispatch_entry()+0x893) [0x624
5a3]#012 5: (SimpleMessenger::DispatchThread::entry()+0x1c) [0x48fd9c]#012 6: ((()+0x69ca) [0x7f609
f8e09ca]#012 7: (clone()+0x6d) [0x7f609e56c70d]
Jun  8 18:39:21 atom2 osd.8[4000]: ceph version 0.28.2-260-ge2c808a (commit:e2c808aea97ced6e9a55b
143116b66d344f72c0b)#012 1: (CephxClientHandler::handle_response(int, ceph::buffer::list::iterator

```

```
&)+0x369) [0x657b09]#012 2: (MonClient::handle_auth(MAuthReply*)+0xb8) [0x62fa58]#012 3: (MonClient::ms_dispatch(Message*)+0x26c) [0x630ebc]#012 4: (SimpleMessenger::dispatch_entry()+0x893) [0x6245a3]#012 5: (SimpleMessenger::DispatchThread::entry()+0x1c) [0x48fd9c]#012 6: (()+0x69ca) [0x7f609f8e09ca]#012 7: (clone()+0x6d) [0x7f609e56c70d]
Jun  8 18:39:21 atom2 osd.8[4000]: *** Caught signal (Aborted) **#012 in thread 0x7f6094259700
Jun  8 18:39:21 atom2 osd.8[4000]: ceph version 0.28.2-260-ge2c808a (commit:e2c808aea97ced6e9a55b143116b66d344f72c0b)#012 1: /usr/bin/cosd() [0x63e33e]#012 2: (()+0xf8f0) [0x7f609f8e98f0]#012 3: (gsignal()+0x35) [0x7f609e4b9a75]#012 4: (abort()+0x180) [0x7f609e4bd5c0]#012 5: (__gnu_cxx::__verbose_terminate_handler()+0x115) [0x7f609ed6f8e5]#012 6: (()+0xcad16) [0x7f609ed6dd16]#012 7: (()+0xcad43) [0x7f609ed6dd43]#012 8: (()+0xcae3e) [0x7f609ed6de3e]#012 9: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x371) [0x6077e1]#012 10: (CephxClientHandler::handle_response(int, ceph::buffer::list::iterator&)+0x369) [0x657b09]#012 11: (MonClient::handle_auth(MAuthReply*)+0xb8) [0x62fa58]#012 12: (MonClient::ms_dispatch(Message*)+0x26c) [0x630ebc]#012 13: (SimpleMessenger::dispatch_entry()+0x893) [0x6245a3]#012 14: (SimpleMessenger::DispatchThread::entry()+0x1c) [0x48fd9c]#012 15: (()+0x69ca) [0x7f609f8e09ca]#012 16: (clone()+0x6d) [0x7f609e56c70d]
```

I'm still seeing my monitor eating more and more memory and going OOM, this might be related to it?

Associated revisions

Revision 35663dee - 06/13/2011 05:04 PM - Sage Weil

monc: protect msg source check with monc_lock

Push the locking up into ms_dispatch, so that we protect the source check and don't race with session reopens. Also simplifies things a bit.

Fixes: #1151

Signed-off-by: Sage Weil <sage@newdream.net>

History

#1 - 06/09/2011 09:01 AM - Sage Weil

- Assignee set to Sage Weil

- Target version set to v0.30

Can you attach the full osd log? It looks like there are two threads doing the auth handshake.

#2 - 06/09/2011 09:56 AM - Sage Weil

Also, tell us more about this cmon memory thing! :)

#3 - 06/10/2011 10:21 AM - Wido den Hollander

The full OSD log is a bit large to attach, so I've left it on my server at the office.

You can access it by SSH'ing to [root@logger.ceph.widodh.nl](ssh://root@logger.ceph.widodh.nl) and from there on 'ssh noisy', there you'll find the log in /var/log/remote, the filename is osd.8.ms_crash.log

#4 - 06/13/2011 10:04 AM - Sage Weil

- Status changed from New to Resolved

Fixed locking in mon_client, [35663dee2f7128df642594d0ebdf16640fe1f29a](https://github.com/ceph/ceph/pull/35663dee2f7128df642594d0ebdf16640fe1f29a)