

Ceph - Bug #1033

osd: CephxClientHandler::handle_response

04/29/2011 02:31 AM - Wido den Hollander

Status:	Resolved	Start date:	04/29/2011
Priority:	Normal	Due date:	
Assignee:	Sage Weil	% Done:	0%
Category:	cephx	Estimated time:	0.00 hour
Target version:	v0.28	Spent time:	1.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	

Description

On one of my OSD's I noticed the following crash:

```
(gdb) bt
#0 0x00007f328d4fd7bb in raise () from /lib/libpthread.so.0
#1 0x0000000000613243 in reraise_fatal (signum=2020) at common/signal.cc:63
#2 0x00000000006142db in handle_fatal_signal (signum=6) at common/signal.cc:110
#3 <signal handler called>
#4 0x00007f328c0cda75 in raise () from /lib/libc.so.6
#5 0x00007f328c0d15c0 in abort () from /lib/libc.so.6
#6 0x00007f328c9838e5 in __gnu_cxx::__verbose_terminate_handler() () from /usr/lib/libstdc++.so.6
#7 0x00007f328c981d16 in ?? () from /usr/lib/libstdc++.so.6
#8 0x00007f328c981d43 in std::terminate() () from /usr/lib/libstdc++.so.6
#9 0x00007f328c981e3e in __cxa_throw () from /usr/lib/libstdc++.so.6
#10 0x00000000005f999a in ceph::__ceph_assert_fail (assertion=<value optimized out>, file=<value optimized out>, line=<value optimized out>,
    func=0x6523c0 "virtual int CephxClientHandler::handle_response(int, ceph::buffer::list::iterator&)" ) at common/assert.cc:86
#11 0x0000000000628ce3 in CephxClientHandler::handle_response (this=0x8a07b0, ret=-1933873600, ind
    ata=...) at auth/cephx/CephxClientHandler.cc:162
#12 0x00000000006065df in MonClient::handle_auth (this=0x7ffff4164370, m=0x17a78600) at mon/MonClient.cc:397
#13 0x0000000000609563 in MonClient::ms_dispatch (this=0x7ffff4164370, m=0x17a78600) at mon/MonClient.cc:242
#14 0x0000000000474ffa in Messenger::ms_deliver_dispatch (this=0x27d5000) at msg/Messenger.h:98
#15 SimpleMessenger::dispatch_entry (this=0x27d5000) at msg/SimpleMessenger.cc:352
#16 0x000000000046a48c in SimpleMessenger::DispatchThread::entry (this=0x27d5488) at msg/SimpleMessenger.h:533
#17 0x00007f328d4f49ca in start_thread () from /lib/libpthread.so.0
#18 0x00007f328c18070d in clone () from /lib/libc.so.6
#19 0x0000000000000000 in ?? ()
```

```
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want 36392
```

Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392

```
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:c
f64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:f
e32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10
:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want
36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:c
f64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:f
e32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10
:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want
36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:c
f64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:f
e32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10
:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want
36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:c
f64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:f
e32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10
:113:1:225:90ff:fe32:cf64]:6813/2019
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want
36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 <
== mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 1 ==== auth_reply(proto 2 0 Success) v1 ==== 33
+0+0 (3471581752 0 0) 0x188d4400 con 0x1ef2ec80
Apr 28 17:58:26 atom1 osd.4[2020]: 7f328256d700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 -
-> mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- auth(proto 2 128 bytes) v1 -- ?+0 0x188d4400
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 <
== mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 2 ==== auth_reply(proto 2 0 Success) v1 ==== 22
5+0+0 (1529762884 0 0) 0x188d4400 con 0x1ef2ec80
Apr 28 17:58:27 atom1 osd.4[2020]: 7f3280d6a700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 m
ark_down [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- 0x22e2a280
Apr 28 17:58:27 atom1 osd.4[2020]: 7f3280d6a700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 -
-> mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- auth(proto 0 26 bytes) v1 -- ?+0 0x17a78600
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 -
-> mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 -- auth(proto 2 128 bytes) v1 -- ?+0 0x188d4400
Apr 28 17:58:27 atom1 osd.4[2020]: 7f3278e31700 osd4 36561 OSD:ms_get_authorizer type=mon
Apr 28 17:58:27 atom1 osd.4[2020]: 2011-04-28 17:56:57.162086 osd4 [2a00:f10:113:1:225:90ff:fe32:c
f64]:6800/2019 1 : [WRN] map e36561 wrongly marked me down or wrong addr
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 osd4 36561 ms_handle_connect on mon
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_boot
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 osd4 36561 client_addr [2a00:f10:113:1:225:90ff:f
e32:cf64]:6800/2019, cluster_addr [2a00:f10:113:1:225:90ff:fe32:cf64]:6812/2019, hb addr [2a00:f10
:113:1:225:90ff:fe32:cf64]:6813/2019
```

```
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive up_thru currently 35256 want 36392
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_alive want 36392
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 osd4 36561 send_pg_stats
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 -- [2a00:f10:113:1:225:90ff:fe32:cf64]:6800/2019 <
== mon0 [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0 1 ==== auth_reply(proto 2 0 Success) v1 ==== 33
+0+0 (3796173676 0 0) 0x17a78600 con 0x1955d140
Apr 28 17:58:27 atom1 osd.4[2020]: 7f328256d700 cephx client: unknown request_type 55041
Apr 28 17:58:27 atom1 osd.4[2020]: auth/cephx/CephxClientHandler.cc: In function 'virtual int Ceph
xClientHandler::handle_response(int, ceph::buffer::list::iterator&)', in thread '0x7f328256d700'#0
12auth/cephx/CephxClientHandler.cc: 162: FAILED assert(0)
Apr 28 17:58:27 atom1 osd.4[2020]: ceph version 0.26-304-gfa5382d (commit:fa5382d3b389bbbbe1b7805
1f77ebd981bc54222)#012 1: (CephxClientHandler::handle_response(int, ceph::buffer::list::iterator&
)+0x323) [0x628ce3]#012 2: (MonClient::handle_auth(MAuthReply*)+0xaf) [0x6065df]#012 3: (MonClient:
:ms_dispatch(Message*)+0x143) [0x609563]#012 4: (SimpleMessenger::dispatch_entry()+0x7ea) [0x474ffa]
#012 5: (SimpleMessenger::DispatchThread::entry()+0x1c) [0x46a48c]#012 6: (()+0x69ca) [0x7f328d4
f49ca]#012 7: (clone()+0x6d) [0x7f328c18070d]
Apr 28 17:58:27 atom1 osd.4[2020]: ceph version 0.26-304-gfa5382d (commit:fa5382d3b389bbbbe1b7805
1f77ebd981bc54222)#012 1: (CephxClientHandler::handle_response(int, ceph::buffer::list::iterator&
)+0x323) [0x628ce3]#012 2: (MonClient::handle_auth(MAuthReply*)+0xaf) [0x6065df]#012 3: (MonClient:
:ms_dispatch(Message*)+0x143) [0x609563]#012 4: (SimpleMessenger::dispatch_entry()+0x7ea) [0x474ffa]
#012 5: (SimpleMessenger::DispatchThread::entry()+0x1c) [0x46a48c]#012 6: (()+0x69ca) [0x7f328d4
f49ca]#012 7: (clone()+0x6d) [0x7f328c18070d]
Apr 28 17:58:27 atom1 osd.4[2020]: *** Caught signal (Aborted) ***#012 in thread 0x7f328256d700
Apr 28 17:58:27 atom1 osd.4[2020]: ceph version 0.26-304-gfa5382d (commit:fa5382d3b389bbbbe1b7805
1f77ebd981bc54222)#012 1: /usr/bin/cosd() [0x6140be]#012 2: (()+0xf8f0) [0x7f328d4fd8f0]#012 3: (g
signal()+0x35) [0x7f328c0cda75]#012 4: (abort()+0x180) [0x7f328c0d15c0]#012 5: (__gnu_cxx::__verbo
se_terminate_handler()+0x115) [0x7f328c9838e5]#012 6: (()+0xcad16) [0x7f328c981d16]#012 7: (()+0xc
ad43) [0x7f328c981d43]#012 8: (()+0xcae3e) [0x7f328c981e3e]#012 9: (ceph::__ceph_assert_fail(char
const*, char const*, int, char const*)+0x36a) [0x5f999a]#012 10: (CephxClientHandler::handle_respo
nse(int, ceph::buffer::list::iterator&)+0x323) [0x628ce3]#012 11: (MonClient::handle_auth(MAuthRep
ly*)+0xaf) [0x6065df]#012 12: (MonClient::ms_dispatch(Message*)+0x143) [0x609563]#012 13: (SimpleM
essenger::dispatch_entry()+0x7ea) [0x474ffa]#012 14: (SimpleMessenger::DispatchThread::entry()+0x1
c) [0x46a48c]#012 15: (()+0x69ca) [0x7f328d4f49ca]#012 16: (clone()+0x6d) [0x7f328c18070d]
```

Seems like the monitor send something the OSD did not understood?

Associated revisions

Revision de640d85 - 05/10/2011 11:54 PM - Sage Weil

monclient: maintain explicit session connection; ignore stray messages

Maintain an explicit Connection handle to send messages and mark_down old monitor connections. Ignore any incoming message that is not part of that session. This fixes problems with incoming messages that race with session restarts.

Fixes: #1033

Reported-by: Wido den Hollander <wido@widodh.nl>

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

History

#1 - 05/02/2011 02:54 PM - Sage Weil

- Target version set to v0.27.1

#2 - 05/02/2011 06:13 PM - Sage Weil

- translation missing: en.field_position set to 2

#3 - 05/10/2011 04:36 PM - Sage Weil

- Status changed from New to Resolved

- Assignee set to Sage Weil

- Target version changed from v0.27.1 to v0.28

[de640d85fa3e0e5e5a31704eab5a8714a1ffe867](https://bugzilla.redhat.com/show_bug.cgi?id=714141)