CephFS - Bug #49725

client: crashed in cct->_conf.get_val() in Client::start_tick_thread()

03/11/2021 01:11 AM - Xiubo Li

Status:	Resolved	**% Done:**	0%
Priority:	Urgent		
Assignee:	Xiubo Li		
Category:			
Target version:	v17.0.0		
Source:	Development	**ceph-qa-suite:**	fs
Tags:		**Component(FS):**	libcephfs
Backport:	pacific	**Labels (FS):**	crash
Regression:	No	**Pull request ID:**	40028
Severity:	3 - minor	**Crash signature (v1):**	
Reviewed:		**Crash signature (v2):**	
Affected Versions:			

Description

The call trace:

```
    1 -- Logs begin at Mon 2021-02-08 09:26:45 CST, end at Wed 2021-03-10 16:22:02 CST. --
    2 Mar 10 16:22:02 lxbceph1 chronyd[1260]: Source 50.205.244.20 replaced with 2600:3c00::f03c:91ff:fe05:b640
    3 Mar 10 16:11:56 lxbceph1 systemd[1]: systemd-coredump@7917-219584-0.service: Succeeded.
    4 Mar 10 16:11:56 lxbceph1 systemd-coredump[219585]: Process 147255 (ceph_test_libce) of user 0 dumped core.
    5
    6                                                   Stack trace of thread 147716:
    7                                                   #0  0x000015293132b47a __memcmp_avx2_movbe (libc.so.6)
    8                                                   #1  0x0000152935a3e498 _ZNSt11char_traitsIcE7compareEPKcS2_m (libceph-common.so.2      )
    9                                                   #2  0x0000152935a8befc _ZNKSt17basic_string_viewIcSt11char_traitsIcEE7compareES2_       (libceph-common.so.2)
   10                                                   #3  0x0000152935bb478a _ZStltIcSt11char_traitsIcEEbSt17basic_string_viewIT_T0_ES5      _ (libceph-common.so.2)
   11                                                   #4  0x0000152935bb2fbd _ZNKSt4lessISt17basic_string_viewIcSt11char_traitsIcEEEclE      RKS3_S6_ (libceph-common.so.2)
   12                                                   #5  0x0000152935bb2f08 _ZNKSt8_Rb_treeISt17basic_string_viewIcSt11char_traitsIcEE      St4pairIKS3_RK6OptionESt10_Select1stIS9_ESt4lessIS3_ESaIS9_EE14_M_lower_boundEPKSt13_Rb_tree_nodeIS9_EPKSt18_Rb_tree_node_baseRS5_ (      libceph-common.so.2)
   13                                                   #6  0x0000152935bb0a36 _ZNKSt8_Rb_treeISt17basic_string_viewIcSt11char_traitsIcEE      St4pairIKS3_RK6OptionESt10_Select1stIS9_ESt4lessIS3_ESaIS9_EE4findERS5_ (libceph-common.so.2)
   14                                                   #7  0x0000152935badc21 _ZNKSt3mapISt17basic_string_viewIcSt11char_traitsIcEERK6Op      tionSt4lessIS3_ESaISt4pairIKS3_S6_EEE4findERSA_ (libceph-common.so.2)
   15                                                   #8  0x0000152935c3b333 _ZNK11md_config_t11find_optionESt17basic_string_viewIcSt11      char_traitsIcEE (libceph-common.so.2)
   16                                                   #9  0x0000152935c40f56 _ZNK11md_config_t8_get_valERK12ConfigValuesSt17basic_strin      g_viewIcSt11char_traitsIcEEPN5boost9container12small_vectorISt4pairIPK6OptionPKNS7_7variantINS7_5blankEJNSt7__cxx1112basic_stringIcS      5_SaIcEEEmldb13entity_addr_t16entity_addrvec_tNSt6chrono8durationIlSt5ratioILl1ELl1EEEENSN_IlSO_ILl1ELl1000EEEENSB_6size_tE6uuid_dEE      EELm4EvvEEPSo (libceph-common.so.2)
   17                                                   #10 0x0000152935c40e8b _ZNK11md_config_t15get_val_genericB5cxx11ERK12ConfigValues      St17basic_string_viewIcSt11char_traitsIcEE (libceph-common.so.2)
   18                                                   #11 0x0000152933f5dcdb _ZNK11md_config_t7
```

```
get_valINSt6chrono8durationIlSt5ratioILl      1ELl1EEEEEEKT_RK12ConfigValuesSt17basic_string_viewI
cSt11char_traitsIcEE (libcephfs.so.2)
   19                                                 #12 0x0000152933f41dc7 _ZNK4ceph6common11
ConfigProxy7get_valINSt6chrono8durationI      lSt5ratioILl1ELl1EEEEEEKT_St17basic_string_viewIcSt1
1char_traitsIcEE (libcephfs.so.2)
   20                                                 #13 0x0000152933ec30a9 _ZZN6Client17start
_tick_threadEvENKUlvE_clEv (libcephfs.so      .2)
   21                                                 #14 0x0000152933f12afd __invoke_impl<void
, Client::start_tick_thread()::<lambda()      > > (libcephfs.so.2)
   22                                                 #15 0x0000152933f12737 __invoke<Client::s
tart_tick_thread()::<lambda()> > (libcep      hfs.so.2)
   23                                                 #16 0x0000152933f14ec0 _M_invoke<0> (libc
ephfs.so.2)
   24                                                 #17 0x0000152933f14e96 operator() (libcep
hfs.so.2)
   25                                                 #18 0x0000152933f14e7a _M_run (libcephfs.
so.2)
   26                                                 #19 0x0000152931bf0ba3 execute_native_thr
ead_routine (libstdc++.so.6)
   27                                                 #20 0x0000152933bad2de start_thread (libp
thread.so.0)
   28                                                 #21 0x00001529312cd133 __clone (libc.so.6
)
   29
   30                                                 Stack trace of thread 147708:
   31                                                 #0  0x0000152933bb37ca futex_abstimed_wai
t_cancelable (libpthread.so.0)
   32                                                 #1  0x000015293406973c __gthread_cond_tim
edwait (libcephfs.so.2)
   33                                                 #2  0x00001529340c72ca _ZNSt18condition_v
ariable17__wait_until_implINSt6chrono8du      rationIlSt5ratioILl1ELl1000000000EEEEEESt9cv_statusR
St11unique_lockISt5mutexERKNS1_10time_pointINS1_3_V212system_clockET_EE (libceph      fs.so.2)
   34                                                 #3  0x00001529340b650c _ZNSt18condition_v
ariable10wait_untilIN4ceph17coarse_mono_      clockENSt6chrono8durationImSt5ratioILl1ELl1000000000
EEEEEESt9cv_statusRSt11unique_lockISt5mutexERKNS3_10time_pointIT_T0_EE (libcephf      s.so.2)
   35                                                 #4  0x000015293409f895 _ZN4ceph5timerINS_
17coarse_mono_clockEE12timer_threadEv (l      ibcephfs.so.2)
   36                                                 #5  0x00001529340b66dc _ZSt13__invoke_imp
lIvMN4ceph5timerINS0_17coarse_mono_clock      EEEFvvEPS3_JEET_St21__invoke_memfun_derefOT0_OT1_DpO
T2_ (libcephfs.so.2)
   37                                                 #6  0x000015293409f93b _ZSt8__invokeIMN4c
eph5timerINS0_17coarse_mono_clockEEEFvvE      JPS3_EENSt15__invoke_resultIT_JDpT0_EE4typeEOS8_DpOS
9_ (libcephfs.so.2)
   38                                                 #7  0x000015293411574f _ZNSt6thread8_Invo
kerISt5tupleIJMN4ceph5timerINS2_17coarse      _mono_clockEEEFvvEPS5_EEE9_M_invokeIJLm0ELm1EEEEDTcl
8__invokespcl10_S_declvalIXT_EEEEESt12_Index_tupleIJXspT_EEE (libcephfs.so.2)
   39                                                 #8  0x000015293411286e _ZNSt6thread8_Invo
kerISt5tupleIJMN4ceph5timerINS2_17coarse      _mono_clockEEEFvvEPS5_EEEclEv (libcephfs.so.2)
   40                                                 #9  0x000015293410e184 _ZNSt6thread11_Sta
te_implINS_8_InvokerISt5tupleIJMN4ceph5t      imerINS3_17coarse_mono_clockEEEFvvEPS6_EEEEE6_M_runE
v (libcephfs.so.2)
   41                                                 #10 0x0000152931bf0ba3 execute_native_thr
ead_routine (libstdc++.so.6)
   42                                                 #11 0x0000152933bad2de start_thread (libp
thread.so.0)
   43                                                 #12 0x00001529312cd133 __clone (libc.so.6
)
   44
```

This could be easier to reproduce with my inode lock patches. And have tried it without the inode lock patches it still could happen after running it hours.

Maybe there should be one dedicate lock in Config classh to protect the `schema` map. In the Client class we can protect all the `cct->_conf` everywhere by `Client_lock`, but the `cct->_conf` still could be access or change out side the Client, which is not under

Client class control/scope.

I will try to fix it.

Related issues:

Copied to CephFS - Backport #49854: pacific: client: crashed in cct->_conf.ge...	**Resolved**

History

#1 - 03/11/2021 02:23 AM - Xiubo Li

- Priority changed from Normal to Urgent

#2 - 03/11/2021 10:43 AM - Xiubo Li

- Pull request ID set to 40028

#3 - 03/11/2021 02:14 PM - Xiubo Li

- Status changed from New to Fix Under Review

- Assignee set to Xiubo Li

#4 - 03/12/2021 12:46 AM - Xiubo Li

With the upstream code, I can reproduce it around 10 time by running 8 hours at night.

#5 - 03/17/2021 03:25 AM - Patrick Donnelly

- Status changed from Fix Under Review to Pending Backport

- Target version set to v17.0.0

- Source set to Development

- Backport set to pacific

#6 - 03/17/2021 03:30 AM - Backport Bot

- Copied to Backport #49854: pacific: client: crashed in cct->_conf.get_val() in Client::start_tick_thread() added

#7 - 04/02/2021 07:06 AM - Loïc Dachary

- Status changed from Pending Backport to Resolved

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".