

CephFS - Bug #13443

Ceph-fuse won't start correctly when the option log_max_new in ceph.conf set to zero

10/10/2015 06:28 AM - Wenjun Huang

<div>Status:Resolved</div> <div>Priority:Normal</div> <div>Assignee:</div> <div>Category:</div> <div>Target version:</div> <div>Source:Community (dev)</div> <div>Tags:</div> <div>Backport:infernalis</div> <div>Regression:No</div> <div>Severity:2 - major</div> <div>Reviewed:</div> <div>Affected Versions:</div>	<div>% Done:0%</div> <div>ceph-qa-suite:fs</div> <div>Component(FS):</div> <div>Labels (FS):</div> <div>Pull request ID:</div> <div>Crash signature (v1):</div> <div>Crash signature (v2):</div>
<div>Description</div> <div>We found a bug in ceph-fuse when we set log_max_new = 0 in ceph.conf, which will lead ceph-fuse process hang when it is started. We set this option, because we want to minimize the memory used by logging subsystem. After checking the related code, we found something incorrect in main methods within ceph_fuse.cc. The suspicious code snippet is below:</div> <div><pre>if (childpid == 0) { common_init_finish(g_ceph_context); //cout << "child, mounting" << std::endl; ::close(fd[0]); if (restart_log) g_ceph_context->_log->start();</pre></div> <div>The hang event happens when executing common_init_finish(g_ceph_context), and in AdminSocket::init it hits</div> <div><pre>ldout(m_cct, 5) << "init " << path << dendl;</pre></div> <div>This will depend on log subsystem working properly. But till then, the Log thread haven't been started in the child process. So, I think we should start the log thread before the common_init_finish method. As far as I know to fix this issue is to bring g_ceph_context->_log->start(); before common_init_finish(g_ceph_context);</div>	
<div>Related issues:</div> <div>Copied to CephFS - Backport #13593: Ceph-fuse won't start correctly when the ...</div> <div>Resolved</div>	

Associated revisions

Revision f2763085 - 10/10/2015 06:46 AM - wenjunhuang

ceph-fuse.cc: While starting ceph-fuse, start the log thread first

<http://tracker.ceph.com/issues/13443> Fixes: #13443

Signed-off-by: Wenjun Huang <wenjunhuang@tencent.com>

Revision 34d0b692 - 10/27/2015 02:30 PM - wenjunhuang

ceph-fuse.cc: While starting ceph-fuse, start the log thread first

<http://tracker.ceph.com/issues/13443> Fixes: #13443

Signed-off-by: Wenjun Huang <wenjunhuang@tencent.com>
(cherry picked from commit f2763085754462610730a23bb5652237714abc2a)

History

#1 - 10/10/2015 06:50 AM - Wenjun Huang

I have made a pull request here <https://github.com/ceph/ceph/pull/6224>

#2 - 10/12/2015 06:30 AM - Nathan Cutler

- Status changed from New to Fix Under Review
- Source changed from other to Community (dev)

#3 - 10/25/2015 08:09 PM - Sage Weil

- Status changed from Fix Under Review to Pending Backport
- Backport set to infernalis

#4 - 10/25/2015 09:06 PM - Greg Farnum

- Project changed from rgw to CephFS

#5 - 10/26/2015 06:18 AM - Nathan Cutler

- Copied to Backport #13593: Ceph-fuse won't start correctly when the option log_max_new in ceph.conf set to zero added

#6 - 02/08/2016 05:06 AM - Loïc Dachary

- Status changed from Pending Backport to Resolved