Linux kernel client - Bug #1871

ceph_client: crash after running xfstests 002

01/03/2012 11:46 AM - Alex Elder

Status:	Resolved	**Start date:**	01/03/2012
Priority:	Normal	**Due date:**	
Assignee:	Alex Elder	**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:	v3.3	**Spent time:**	0.00 hour
Source:		**Reviewed:**	
Tags:		**Affected Versions:**	
Backport:		**ceph-qa-suite:**	
Regression:	No	**Crash signature:**	
Severity:	3 - minor		

Description

Running xfstests under UML against a ceph filesystem, I get a
client crash due to dereferencing a null pointer in ceph_d_prune()
after running (only) test 002 has completed. The crash is
happening during the unmount of the filesystem. My ceph volume
has 2 OSD's and 1 MDS, and I have 3 monitors. Note that I am
running both ceph and the UML on the same system--a HT quad core
Intel Core i7-2600 @ 3.40 GHz.

ceph: master branch,
a1252463 librados: take lock in rollback
ceph-client: for-linus branch,
a4d46363ce ceph: disable use of dcache for readdir etc.
xfstests: master branch,
e219e1cb5 275: add write and reserve test
I have added a few small changes to xfstests so it can be
run with a ceph target filesystem. The filesystem mount
is something like:
192.168.122.1:/test_dir mounted on /mnt/test_dir.12345

I am using cephx authentication.

Dumped stack trace will be below.

Test 002 simply creates a file (like "foo.1") in the target
filesystem, then makes 19 hard links to that same file (using
commands such as "ln foo.1 foo.20"). After each link gets made,
the link count for the original file is checked to ensure it is
the right value. In the second phase of the test, the links are
deleted in reverse order, starting with "foo.20" and ending with
"foo.1". Again, before removing each link, the link count
is checked to ensure it is correct.

When testing is done, the filesystem gets unmounted by a wrapper
script. It is at this point that the kernel panic occurs.

Below is the stack trace from the UML panic.

...
Ran: 001 002
Passed all 2 tests
/
[3707.450000]
[3707.450000] Modules linked in:
[3707.450000] Pid: 4713, comm: umount Not tainted 3.1.0-09588-ga4d4636-dirty

```
[ 3707.450000] RIP: 0033:[<00000000603f9e3c>]
[ 3707.450000] RSP: 000000009e443d10  EFLAGS: 00010206
[ 3707.450000] RAX: 0000000000000000 RBX: 0000000098fa2f18 RCX: 0000000000000000
[ 3707.450000] RDX: 0000000098fa2fd0 RSI: 000000009eb9dd40 RDI: 0000000098fa2f18
[ 3707.450000] RBP: 000000009e443d20 R08: 0000000000100000 R09: 0000000000100000
[ 3707.450000] R10: 0000000000000000 R11: 00007f8ac02ac403 R12: 0000000060585bf0
[ 3707.450000] R13: 000000009e79ff08 R14: 0000000000000000 R15: 000000009e443eb8
[ 3707.450000] Call Trace:
[ 3707.450000] 606f9618:  [<6001b919>] segv+0x229/0x23b
[ 3707.450000] 606f9628:  [<6003a114>] __do_softirq+0x12b/0x141
[ 3707.450000] 606f96e8:  [<6001b994>] segv_handler+0x69/0x6f
[ 3707.450000] 606f9708:  [<60018e76>] sigio_handler+0x58/0x5d
[ 3707.450000] 606f9728:  [<60028f15>] sig_handler_common+0x84/0x98
[ 3707.450000] 606f97b0:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f97d0:  [<6006ae91>] pagevec_lru_move_fn+0x2f/0xc5
[ 3707.450000] 606f9828:  [<600164e4>] _einittext+0x2108/0x34b4
[ 3707.450000] 606f9838:  [<600157c8>] _einittext+0x13ec/0x34b4
[ 3707.450000] 606f9918:  [<600164e4>] _einittext+0x2108/0x34b4
[ 3707.450000] 606f9a58:  [<60029005>] sig_handler+0x2d/0x38
[ 3707.450000] 606f9a78:  [<60028c3b>] hard_handler+0x6b/0x9d
[ 3707.450000] 606f9b68:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000]
[ 3707.450000] Kernel panic - not syncing: Kernel mode fault at addr 0x0, ip 0x603f9e3c
[ 3707.450000] Call Trace:
[ 3707.450000] 606f9500:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f9518:  [<605314f8>] panic+0xdb/0x1ce
[ 3707.450000] 606f9550:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f9570:  [<60059d64>] __module_text_address+0xd/0x56
[ 3707.450000] 606f9588:  [<6005cf4c>] is_module_text_address+0x9/0x11
[ 3707.450000] 606f9598:  [<60049577>] __kernel_text_address+0x21/0x47
[ 3707.450000] 606f95b8:  [<6001a722>] show_trace+0x8e/0x95
[ 3707.450000] 606f95c0:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f95e8:  [<6002d7e3>] show_regs+0x2b/0x30
[ 3707.450000] 606f9608:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f9618:  [<6001b92b>] segv_handler+0x0/0x6f
[ 3707.450000] 606f9628:  [<6003a114>] __do_softirq+0x12b/0x141
[ 3707.450000] 606f96e8:  [<6001b994>] segv_handler+0x69/0x6f
[ 3707.450000] 606f9708:  [<60018e76>] sigio_handler+0x58/0x5d
[ 3707.450000] 606f9728:  [<60028f15>] sig_handler_common+0x84/0x98
[ 3707.450000] 606f97b0:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f97d0:  [<6006ae91>] pagevec_lru_move_fn+0x2f/0xc5
[ 3707.450000] 606f9828:  [<600164e4>] _einittext+0x2108/0x34b4
[ 3707.450000] 606f9838:  [<600157c8>] _einittext+0x13ec/0x34b4
[ 3707.450000] 606f9918:  [<600164e4>] _einittext+0x2108/0x34b4
[ 3707.450000] 606f9a58:  [<60029005>] sig_handler+0x2d/0x38
[ 3707.450000] 606f9a78:  [<60028c3b>] hard_handler+0x6b/0x9d
[ 3707.450000] 606f9b68:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000]
[ 3707.450000]
[ 3707.450000] Modules linked in:
[ 3707.450000] Pid: 4713, comm: umount Not tainted 3.1.0-09588-ga4d4636-dirty
[ 3707.450000] RIP: 0033:[<000000004071a547>]
[ 3707.450000] RSP: 0000007fbf9a6498  EFLAGS: 00000246
[ 3707.450000] RAX: fffffffffffffda RBX: 000000000060eae0 RCX: ffffffffffffffff
[ 3707.450000] RDX: fffffffffffff58 RSI: 0000000000000000 RDI: 000000000060eae0
[ 3707.450000] RBP: 000000000060eab0 R08: 0000000000000000 R09: 0000000000000000
[ 3707.450000] R10: 0000000000000000 R11: 0000000000000246 R12: 000000000060eb50
[ 3707.450000] R13: 0000007fbf9a6688 R14: 0000000000000000 R15: 0000000000000000
[ 3707.450000] Call Trace:
[ 3707.450000] 606f9498:  [<6001bb84>] panic_exit+0x2f/0x45
[ 3707.450000] 606f94b8:  [<6004f5b7>] notifier_call_chain+0x5f/0x96
[ 3707.450000] 606f94f0:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f9508:  [<6004f62c>] atomic_notifier_call_chain+0x13/0x15
[ 3707.450000] 606f9518:  [<60531513>] panic+0xf6/0x1ce
[ 3707.450000] 606f9550:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f9570:  [<60059d64>] __module_text_address+0xd/0x56
[ 3707.450000] 606f9588:  [<6005cf4c>] is_module_text_address+0x9/0x11
```

```
[ 3707.450000] 606f9598:  [<60049577>] __kernel_text_address+0x21/0x47
[ 3707.450000] 606f95b8:  [<6001a722>] show_trace+0x8e/0x95
[ 3707.450000] 606f95c0:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f95e8:  [<6002d7e3>] show_regs+0x2b/0x30
[ 3707.450000] 606f9608:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f9618:  [<6001b92b>] segv_handler+0x0/0x6f
[ 3707.450000] 606f9628:  [<6003a114>] __do_softirq+0x12b/0x141
[ 3707.450000] 606f96e8:  [<6001b994>] segv_handler+0x69/0x6f
[ 3707.450000] 606f9708:  [<60018e76>] sigio_handler+0x58/0x5d
[ 3707.450000] 606f9728:  [<60028f15>] sig_handler_common+0x84/0x98
[ 3707.450000] 606f97b0:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000] 606f97d0:  [<6006ae91>] pagevec_lru_move_fn+0x2f/0xc5
[ 3707.450000] 606f9828:  [<600164e4>] _einittext+0x2108/0x34b4
[ 3707.450000] 606f9838:  [<600157c8>] _einittext+0x13ec/0x34b4
[ 3707.450000] 606f9918:  [<600164e4>] _einittext+0x2108/0x34b4
[ 3707.450000] 606f9a58:  [<60029005>] sig_handler+0x2d/0x38
[ 3707.450000] 606f9a78:  [<60028c3b>] hard_handler+0x6b/0x9d
[ 3707.450000] 606f9b68:  [<603f9e3c>] ceph_d_prune+0x65/0x6e
[ 3707.450000]
Terminated
```

History

#1 - 01/03/2012 11:48 AM - Sage Weil

- Target version set to v3.3

#2 - 01/09/2012 03:40 PM - Sage Weil

- Status changed from New to Resolved