

Arvados - Bug #12134

[arv-mount] Fix test deadlock by using new llfuse in test suite

08/16/2017 09:57 PM - Tom Clegg

Status:	Resolved	Start date:	08/16/2017
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	FUSE	Estimated time:	0.00 hour
Target version:	2017-09-13 Sprint		
Description			
Peter's patch for the thread-cancel bug has been merged upstream, but there is no new release yet.			
Instead of waiting for a new release, we can update run-tests.sh to build its own version ("1.2.1arvados1"?) from source and install it into \$VENVDIR during the dependencies/setup phase. Then "python setup.py install" in services/fuse should use that custom version instead of going to pypi, and we should see no more test suite deadlocks.			
This is the relevant fix:			
https://bitbucket.org/nikratio/python-llfuse/commits/8aab6579089bb0b07423a41bca84c4654b2f9b81			
Subtasks:			
Task # 12143: Exercise test suite with patched llfuse			Resolved
Task # 12187: Review 12134-llfuse-patch			Resolved
Related issues:			
Related to Arvados - Story #8345: [FUSE] Support llfuse 0.42+		Resolved	02/03/2016
Related to Arvados - Bug #10805: [FUSE] Upgrade llfuse to 1.2, fix deadlock i...		Resolved	01/04/2017

Associated revisions

Revision 84f5a47b - 09/01/2017 03:29 PM - Tom Clegg

Merge branch '12134-llfuse-patch'

refs #12134

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tclegg@veritasgenetics.com>

History

#1 - 08/17/2017 06:12 PM - Tom Clegg

Even with new llfuse, test_replace (tests.test_unmount.UnmountTest) hangs sometimes, but it wakes up if I look in /proc/self/mountinfo for stray fuse mounts and hit them with "echo 1 > /sys/fs/fuse/connections/NNN/abort".

#2 - 08/17/2017 08:50 PM - Tom Clegg

- Status changed from New to In Progress

#3 - 08/18/2017 01:40 PM - Tom Clegg

Tests hung on test_tmp_rewrite (tests.test_tmp_collection.TmpCollectionTest). Destroying fuse mounts had no effect, but SIGKILLing the innermost python program revived the test run.

Again in test_tmp_snapshots (tests.test_tmp_collection.TmpCollectionTest).

```
14694 tom      20    0  11896   3872   2900 S   0.0  0.1   0:00.08      \- run-tests.sh
 9248 tom      20    0 2728408 202196 15776 S   0.3  5.0   2:43.15      \- python
9404 tom      20    0  16956   1044    916 S   0.0  0.0   0:00.01      \- sed
9405 tom      20    0  497612  11512   5820 S   0.0  0.3   0:02.04      \- keepst
ore
9416 tom      20    0  16956   1080    948 S   0.0  0.0   0:00.00      \- sed
9417 tom      20    0  503440   9632   5928 S   0.0  0.2   0:01.86      \- keepst
ore
4799 tom      20    0 2497948 186832   2860 S   0.0  4.6   0:00.00      \- python
```

9396 (one of 9248's threads) is doing waitpid(4799, ...), sleep 100ms, repeat.

4799 is waiting in futex().

Killing 4799 just resulted in a new process in the same place doing the same thing, whether or not I had aborted the fuse mounts. Had to kill 9248 to abort the test suite. Judging by strace, 9396 was the multiprocessing module maintaining its 1-worker pool. I suspect this behavior (multiprocessing keeping an idle worker pool alive) is fine, the problem is just that some other thread is deadlocked. (But it is possible that the way multiprocessing keeps a worker pool alive is helping that other thread get deadlocked.)

#4 - 08/18/2017 02:54 PM - Tom Clegg

```
-----  
Ran 83 tests in 292.678s
```

```
OK  
Exception in thread Thread-83:  
Traceback (most recent call last):  
  File "/usr/lib/python2.7/threading.py", line 801, in __bootstrap_inner  
    self.run()  
  File "/usr/lib/python2.7/threading.py", line 754, in run  
    self.__target(*self.__args, **self.__kwargs)  
  File "/usr/lib/python2.7/multiprocessing/pool.py", line 366, in _handle_tasks  
    debug('task handler got sentinel')  
TypeError: 'NoneType' object is not callable
```

```
Exception in thread Thread-82:  
Traceback (most recent call last):  
  File "/usr/lib/python2.7/threading.py", line 801, in __bootstrap_inner  
    self.run()  
  File "/usr/lib/python2.7/threading.py", line 754, in run  
    self.__target(*self.__args, **self.__kwargs)  
  File "/usr/lib/python2.7/multiprocessing/pool.py", line 330, in _handle_workers  
    debug('worker handler exiting')  
TypeError: 'NoneType' object is not callable
```

```
/home/tom/src/arvados/build/run-tests.sh: line 574: 26482 Segmentation fault      "${3}python" setup.py ${short  
t:++-short-tests-only} test ${testargs[$1]}
```

#5 - 08/29/2017 01:26 PM - Tom Clegg

at commit:7c92bc77b, ran fuse tests 50 times in a row with no deadlocks, interrupted by this (non-deadlocking) failure:

```
=====  
FAIL: runTest (tests.test_mount.FuseModifyFileTest)  
-----  
Traceback (most recent call last):  
  File "/home/tom/src/arvados/services/fuse/tests/test_mount.py", line 348, in runTest  
    self.pool.apply(fuseModifyFileTestHelperReadEndContents, (self.mounttmp,))  
  File "/usr/lib/python2.7/multiprocessing/pool.py", line 244, in apply  
    return self.apply_async(func, args, kwds).get()  
  File "/usr/lib/python2.7/multiprocessing/pool.py", line 567, in get  
    raise self._value  
AssertionError: 'plnp' != 'blub'
```

```
-----  
Ran 83 tests in 317.388s
```

12134-llfuse-patch @ [99c19e1539aabb8053ee9221f62744bf76d63737](https://github.com/99c19e1539aabb8053ee9221f62744bf76d63737)

#6 - 08/30/2017 07:03 PM - Tom Clegg

- Target version changed from 2017-08-30 Sprint to 2017-09-13 Sprint

#7 - 08/31/2017 02:59 PM - Lucas Di Pentima

Tried to run the tests repeatedly on my VM, but llfuse requires cython 0.24+ and Debian Jessie doesn't have it. The updates LGTM, but I don't know if this would create problems with the rest of the team, for example I know Peter runs tests using arvbox.

#8 - 09/01/2017 03:39 PM - Tom Clegg

- Status changed from In Progress to Feedback

#9 - 09/01/2017 03:40 PM - Tom Clegg

We'll have to update arvbox to stretch to make it pass tests.

#10 - 09/08/2017 02:20 AM - Tom Clegg

- Status changed from Feedback to Resolved

No test deadlocks since merge.