

Arvados - Bug #15910

[crunch-run] Crash while writing output files

12/04/2019 03:58 PM - Tom Clegg

Status: Resolved	Start date: 12/04/2019
Priority: Normal	Due date:
Assigned To: Tom Clegg	% Done: 100%
Category:	Estimated time: 0.00 hour
Target version: 2020-01-02 Sprint	
Description crunch-run process disappears while writing outputs (~1GB in ~1700 files in 1 dir). Suspect this is caused by a data race: there's a concurrent map write that sometimes causes panics and complaints from the Go race detector in a test case.	
Subtasks: Task # 15914: Review 15910-crunch-run-crash Resolved	
Related issues: Related to Arvados - Bug #15946: [crunch-run] [collectionfs] Deadlock while w... Resolved 12/23/2019	

Associated revisions

Revision 9c2fb29f - 12/05/2019 03:45 PM - Tom Clegg

Merge branch '15910-crunch-run-crash'

refs #15910

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

History

#3 - 12/04/2019 04:31 PM - Tom Clegg

15910-crunch-run-crash @ [5f9e11ae95dc2b768ccc7e5f165be06b6aabdb2](#) -- <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1683/>

#4 - 12/04/2019 07:06 PM - Lucas Di Pentima

Is the new test intended to expose the bug? I've locally tried to run it with an older version of fs_collection.go and it doesn't fail.

#5 - 12/04/2019 08:22 PM - Lucas Di Pentima

From chat:

lucas: you might need to run it a few times before it crashes -- but "test services/crunch-run -race" (from run-tests interactive mode) seems to find it reliably

Running sdk/go/arvados tests with --repeat 50 on a older collection_fs.go file didn't failed for me.

I've also tried locally running the services/crunch-run -race tests using the interactive mode and they still fail with the fix:

```
[...]  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000002Z Goodbye  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000001Z Hello world!  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000002Z Goodbye  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000001Z Hello world!  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000002Z Goodbye  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000001Z Hello world!  
zzzzz-zzzzzzzzzzzzzzzzzzzzz 2015-12-29T15:51:45.000000003Z Goodbye  
OK: 71 passed  
--- FAIL: TestCrunchExec (52.93s)  
    testing.go:809: race detected during execution of test  
FAIL  
exit status 1  
FAIL    git.curoverse.com/arvados.git/services/crunch-run    53.010s  
===== services/crunch-run tests -- FAILED  
===== test services/crunch-run -- 55s  
Failures (1):
```

Fail: services/crunch-run tests (55s)

Are these race tests being run on our Jenkins test pipelines? If not, can we make Jenkins run them or is there any limitation on how they work?

#6 - 12/04/2019 08:56 PM - Tom Clegg

Jenkins doesn't enable the race detector. I see crunch-run test races here too (although the first few are just races between tested code and the tests themselves, so they might not indicate real bugs).

Both should be addressed. I don't think we should block this bugfix behind either of them. It would be good to verify that the race detector confirms this bugfix for you like it does for me, though.

I should have said:

"test sdk/go/arvados -race" (from run-tests interactive mode) seems to find it reliably

(The new test case in sdk/go/arvados is the one that sets off the race detector without the fix, but not with the fix.)

#7 - 12/04/2019 09:32 PM - Lucas Di Pentima

I've tried sdk/go/arvados -race and I'm getting a "signal: killed" message previous to the failed test report. Checking the RAM usage, it goes up to 100% utilization (even when I upped my VM to 9.4 GiB) just before it crashes. I also tried going from 2 to 8 vCPUs and got the same result.

So I'm not able to confirm on my end but if you think it shouldn't block the fix, please go ahead.

#8 - 12/09/2019 03:31 PM - Tom Clegg

- Status changed from *In Progress* to *Resolved*

#9 - 12/22/2019 12:34 AM - Tom Clegg

- Related to Bug #15946: [crunch-run] [collections] Deadlock while writing output collection added

#10 - 01/22/2020 02:55 PM - Peter Amstutz

- Release set to 22