

Arvados - Bug #7870

[Crunch] crunch-dispatch fails to mark a job failed when its retry loop aborts

11/26/2015 06:45 PM - Ward Vandewege

Status:	Resolved	Start date:	12/01/2015
Priority:	Normal	Due date:	
Assigned To:	Brett Smith	% Done:	100%
Category:	Crunch	Estimated time:	0.00 hour
Target version:	2015-12-02 sprint		
Description			
See https://workbench.c97qk.arvadosapi.com/pipeline_instances/c97qk-d1hrv-stxx1wltwlzssp5			
Node manager log:			
<pre>2015-11-26_18:44:34.74860 2015-11-26 18:44:34 arvnodeman.arvados_nodes[64971] DEBUG: ArvadosNodeListMonitorActor (at 140615574139536) sending poll 2015-11-26_18:44:34.75008 2015-11-26 18:44:34 arvnodeman.jobqueue[64971] DEBUG: JobQueueMonitorActor (at 140615574836752) sending poll 2015-11-26_18:44:34.75613 2015-11-26 18:44:34 arvnodeman.cloud_nodes[64971] DEBUG: CloudNodeListMonitorActor (at 140615594757200) sending poll 2015-11-26_18:44:36.31902 2015-11-26 18:44:36 arvnodeman.arvados_nodes[64971] DEBUG: ArvadosNodeListMonitorActor (at 140615574139536) got response with 233 items 2015-11-26_18:44:37.12386 2015-11-26 18:44:37 arvnodeman.jobqueue[64971] DEBUG: Sending server wishlist: (empty) 2015-11-26_18:44:37.12394 2015-11-26 18:44:37 arvnodeman.jobqueue[64971] DEBUG: JobQueueMonitorActor (at 140615574836752) got response with 0 items 2015-11-26_18:44:37.12487 2015-11-26 18:44:37 arvnodeman.daemon[64971] DEBUG: Standard_D4: idle nodes 0, wishlist size 0 2015-11-26_18:44:37.12501 2015-11-26 18:44:37 arvnodeman.daemon[64971] DEBUG: Standard_D3: idle nodes 0, wishlist size 0 2015-11-26_18:44:37.12517 2015-11-26 18:44:37 arvnodeman.daemon[64971] DEBUG: Standard_D2: idle nodes 0, wishlist size 0 2015-11-26_18:44:37.12535 2015-11-26 18:44:37 arvnodeman.daemon[64971] DEBUG: Standard_D1: idle nodes 0, wishlist size 0 2015-11-26_18:44:44.12915 2015-11-26 18:44:44 arvnodeman.cloud_nodes[64971] DEBUG: CloudNodeListMonitorActor (at 140615594757200) got response with 0 items</pre>			
Subtasks:			
Task # 7877: Review 7870-crunch-dispatch-retry-fail-lock-wip			Resolved

Associated revisions

Revision fc5257c1 - 12/01/2015 09:22 PM - Brett Smith

Merge branch '7870-crunch-dispatch-retry-fail-lock-wip'

Closes #7870, #7877.

History

#1 - 11/26/2015 06:46 PM - Ward Vandewege

- Description updated

#2 - 11/26/2015 06:54 PM - Tom Clegg

According to the log message (and c97qk workbench dashboard) the queue is empty, in which case wishlist=0 is correct. Was there a job queued that has since been cancelled?

#3 - 11/28/2015 10:12 PM - Ward Vandewege

Tom Clegg wrote:

According to the log message (and c97qk workbench dashboard) the queue is empty, in which case wishlist=0 is correct. Was there a job queued that has since been cancelled?

There is a job queued. Also listed on the workbench dashboard. How come you didn't see it?

https://workbench.c97qk.arvadosapi.com/pipeline_instances/c97qk-d1hrv-stxx1wltwzssp5

#4 - 11/30/2015 02:47 PM - Ward Vandewege

Looks like this is yet another instance of "you can't run two job dispatchers at the same time":

/etc/service/crunch-dispatch-jobs-1/log/main

```
@400000005657da2437923c04.s:2015-11-25_16:06:38.86565 dispatch: job c97qk-8i9sb-3k4o5augc494rm0 end
@400000005657da2437923c04.s:2015-11-25_16:25:43.88271 dispatch: waiting for nodes for c97qk-8i9sb-3k4o5augc494rm0
@400000005657da2437923c04.s:2015-11-25_17:18:32.54096 dispatch: c97qk-8i9sb-3k4o5augc494rm0
: crunch-dispatch was stopped during job's tempfail retry loop
@400000005657da2437923c04.s:2015-11-25_17:18:32.92287 dispatch: tried to mark job c97qk-8i9sb-3k4o5augc494rm0
as failed but it was already locked by someone else
```

vs

/etc/service/crunch-dispatch-jobs-0/log/main

```
2015-11-25_16:03:17.20886 dispatch: ["git", "--git-dir=/var/lib/arvados/internal.git", "tag", "c97qk-8i9sb-3k4o5augc494rm0", "4b84269f160c4d64cb46b2d3bc2a404a962e793f"]
2015-11-25_16:03:17.21574
2015-11-25_16:03:17.22272 dispatch: sudo -E -u crunch LD_LIBRARY_PATH= PATH=/var/www/arvados-api/shared/vendor_bundle/ruby/2.1.0/bin:/usr/local/rvm/gems/ruby-2.1.5/bin:/usr/local/rvm/gems/ruby-2.1.5@global/bin:/usr/local/rvm/rubies/ruby-2.1.5/bin:/usr/local/rvm/bin:/usr/local/bin:/usr/local/sbin:/bin:/sbin:/usr/bin:/usr/sbin:/usr/X11R6/bin:/usr/local/arvados/src/services/crunch PERLLIB=/usr/local/arvados/src/sdk/perl/lib PYTHONPATH= RUB_YLIB=/usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib GEM_PATH= salloc --chdir= --immediate --exclusive --no-kill --job-name=c97qk-8i9sb-3k4o5augc494rm0 --nodelist=compute2 /usr/local/arvados/src/services/crunch/crunch-job --job-api-token 5oqv77bzqieoondrt63uyjkwvve9ffgnedi2gdev6rl2nylf --job c97qk-8i9sb-3k4o5augc494rm0 --git-dir /var/lib/arvados/internal.git
2015-11-25_16:03:17.22952 dispatch: job c97qk-8i9sb-3k4o5augc494rm0
2015-11-25_16:03:17.22956 dispatch: child 20641 start 2015-11-25_16:03:17
2015-11-25_16:03:17.84845 dispatch: update compute2 state to {:state=>"alloc", :job=>"c97qk-8i9sb-3k4o5augc494rm0" }
2015-11-25_16:03:18.14644 c97qk-8i9sb-3k4o5augc494rm0
! salloc: error: Unable to allocate resources: Requested nodes are busy
2015-11-25_16:03:18.17740 dispatch: child 20641 exit 75
2015-11-25_16:03:18.17747 dispatch: job c97qk-8i9sb-3k4o5augc494rm0 end
2015-11-25_16:06:34.85513 dispatch: update compute2 state to {:state=>"alloc", :job=>nil}
2015-11-25_16:06:36.77112 dispatch: update compute2 state to {:state=>"down", :job=>nil}
2015-11-25_17:18:31.61005 Received TERM signal
2015-11-25_17:18:33.00783 Stopping crunch-dispatch-jobs-0
2015-11-25_17:18:34.15911 /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/spec_set.rb:92:in `block in materialize': Could not find curb-0.8.8 in any of the sources (Bundler::GemNotFound)
2015-11-25_17:18:34.15922 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/spec_set.rb:85:in `map!'
2015-11-25_17:18:34.15924 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/spec_set.rb:85:in `materialize'
2015-11-25_17:18:34.15926 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/definition.rb:140:in `specs'
2015-11-25_17:18:34.15928 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/definition.rb:185:in `specs_for'
2015-11-25_17:18:34.15930 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/definition.rb:174:in `requested_specs'
2015-11-25_17:18:34.15932 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/environment.rb:18:in `requested_specs'
2015-11-25_17:18:34.15934 from /usr/local/rvm/gems/ruby-2.1.5/gems/bundler-1.10.6/lib/bundler/runtime.rb:13:in `setup'
...
```

#5 - 11/30/2015 06:49 PM - Brett Smith

- Subject changed from [Node manager] node manager thinks 'server wishlist' is empty when there is a job waiting to run to [Crunch] crunch-dispatch fails to mark a job failed when its retry loop aborts

- Category set to Crunch

This is the critical section of logs:

```
dispatch: c97qk-8i9sb-3k4o5augc494rm0: crunch-dispatch was stopped during job's tempfail retry loop
dispatch: tried to mark job c97qk-8i9sb-3k4o5augc494rm0 as failed but it was already locked by someone else
```

When a job fails because of node failures, the same crunch-dispatch that originally launched it must retry it as many times as desired, because the API token that holds the lock only exists in its memory.

If that crunch-dispatch gets a TERM process, it will stop retrying the job, and mark that job failed. It turns out this last step, itself, is failing. When it tries to mark the job failed, it tries to lock the job, and that fails because the job's first crunch-job process already got the lock.

crunch-dispatch needs to be taught not to acquire the lock in this case, because it effectively already has it.

#6 - 11/30/2015 07:03 PM - Brett Smith

Ward Vandewege wrote:

There is a job queued. Also listed on the workbench dashboard. How come you didn't see it?

https://workbench.c97qk.arvadosapi.com/pipeline_instances/c97qk-d1hrv-stxx1wltwzssp5

The job you're referring to ([c97qk-8i9sb-3k4o5augc494rm0](#)) was in the Running state. That's not the same thing as Queued.

#7 - 12/01/2015 02:34 PM - Brett Smith

- Status changed from New to In Progress
- Assigned To set to Brett Smith
- Target version set to 2015-12-02 sprint
- Story points set to 0.5

#8 - 12/01/2015 08:37 PM - Peter Amstutz

I can't really figure out a good way to test this (I tried, I can at least verify there are no syntax errors that prevent execution), but I understand and agree with the logic, so LGTM.

#9 - 12/01/2015 09:25 PM - Brett Smith

- Status changed from In Progress to Resolved
- % Done changed from 0 to 100

Applied in changeset arvados|commit:fc5257c18b24ab0e28b248655dcabf9665bf3.