

## Arvados - Story #17755

### Test singularity support on a cloud cluster by running some real workflows

06/04/2021 01:46 PM - Peter Amstutz

<b>Status:</b> Resolved	<b>Start date:</b> 09/03/2021
<b>Priority:</b> Normal	<b>Due date:</b>
<b>Assigned To:</b> Ward Vandewege	<b>% Done:</b> 100%
<b>Category:</b> Crunch	<b>Estimated time:</b> 0.00 hour
<b>Target version:</b> 2021-09-15 sprint	
<b>Description</b> On a cloud cluster with singularity enabled (requires new compute image + config changes), run the the WGS workflow with small chr19 inputs, eg <a href="#">pirca-xvhdp-m4r8n9qmoh5ggra</a>	
<b>Subtasks:</b> Task # 17796: Review <span style="float: right;"><b>Resolved</b></span>	
<b>Related issues:</b>	
Related to Arvados - Story #17296: Singularity proof of concept	<b>Resolved</b> 05/27/2021
Related to Arvados Epics - Story #16305: Singularity support	<b>In Progress</b> 01/01/2021 09/30/2021
Related to Arvados - Bug #18027: [crunch-run] upload failed (singularity)	<b>Resolved</b> 09/03/2021
Related to Arvados - Bug #18026: [singularity] crunch-run image cache race co...	<b>Resolved</b> 08/24/2021
Blocked by Arvados - Bug #17816: singularity not setting working directory	<b>Resolved</b> 07/14/2021

#### Associated revisions

##### Revision 6ef187b8 - 08/06/2021 01:29 PM - Ward Vandewege

Singularity requires squashfs-tools to build images.

refs #17755

Arvados-DCO-1.1-Signed-off-by: Ward Vandewege <[ward@curii.com](mailto:ward@curii.com)>

##### Revision d59079ca - 08/25/2021 07:34 PM - Peter Amstutz

Merge branch '17755-singularity-race' into main refs #17755

Arvados-DCO-1.1-Signed-off-by: Peter Amstutz <[peter.amstutz@curii.com](mailto:peter.amstutz@curii.com)>

##### Revision 9695271c - 09/03/2021 06:43 PM - Ward Vandewege

Merge branch '17755-add-singularity-to-compute-image'

closes #17755

Arvados-DCO-1.1-Signed-off-by: Ward Vandewege <[ward@curii.com](mailto:ward@curii.com)>

#### History

##### #1 - 06/04/2021 01:46 PM - Peter Amstutz

- Status changed from New to In Progress

##### #2 - 06/04/2021 01:46 PM - Peter Amstutz

- Status changed from In Progress to New

- Tracker changed from Bug to Story

##### #3 - 06/04/2021 01:47 PM - Peter Amstutz

- Category set to Crunch

##### #4 - 06/04/2021 01:49 PM - Peter Amstutz

- Description updated

##### #5 - 06/08/2021 05:46 PM - Peter Amstutz

- Description updated

- Subject changed from Test singularity support on 9tee4 by running some real workflows to Test singularity support on a cloud cluster by running some real workflows

**#6 - 06/09/2021 04:09 PM - Peter Amstutz**

- Assigned To set to Ward Vandewege

**#7 - 06/11/2021 02:25 PM - Ward Vandewege**

- Status changed from New to In Progress

**#8 - 06/11/2021 03:28 PM - Ward Vandewege**

- Related to Story #17296: Singularity proof of concept added

**#9 - 06/11/2021 09:44 PM - Ward Vandewege**

I have built a new compute node image for ce8i5 that has the singularity binary 3.5.2 (that's an **old** version) in it, cf. commit: 784a3f24d37819186a52ea2c67e15e5bd8639076 on branch 17755-add-singularity-to-compute-image. That commit builds singularity from source while creating the packer image which seems suboptimal (it requires a lot of build-time dependencies, and it takes a while).

I've run the diagnostics workflow with it at [ce8i5-xvhdp-el23s4fjsrp1mjb](#).

Observations:

- the conversion of a ~240MiB docker tar file took over 5 minutes (ouch!!), about 4 of which seem to be in the invocation of `mksquashfs`
- this conversion is invisible to the user, it's part of the "loading the image" stage so it's not considered when (e.g.) workbench shows workflow step durations.
- the `hasher3` step ran on the same node as the `hasher1` step, but it repeated the singularity import step. Ouch.

[ce8i5-xvhdp-ia0zn3njin6lwbj](#): docker: 0m9s(6m queued)

[ce8i5-xvhdp-el23s4fjsrp1mjb](#): singularity: 0m7s(11m queued)

**#10 - 06/16/2021 08:12 PM - Ward Vandewege**

I upgraded the ce8i5 compute image to use singularity 3.7.4 and switched ce8i5 over to singularity again. The amount of time it takes to convert the docker image to a singularity image seems unchanged (i.e. it's very slow).

The CWL tests were running, and all the jobs started failing because they are trying to write to a read-only filesystem. E.g. [ce8i5-xvhdp-6aj1q0l31rae2y4](#), [ce8i5-xvhdp-x95gwpw2qdf3113](#), etc. See <https://ci.arvados.org/job/run-cwl-test-ce8i5/387/>. I filed [#17816](#) for this problem.

I then switched ce8i5 back to use docker.

Further testing on this ticket is probably pointless until we fix [#17816](#).

**#11 - 06/16/2021 08:26 PM - Ward Vandewege**

- Related to Bug #17816: singularity not setting working directory added

**#12 - 06/16/2021 08:32 PM - Ward Vandewege**

- Related to deleted (Bug #17816: singularity not setting working directory)

**#13 - 06/16/2021 08:32 PM - Ward Vandewege**

- Blocks Bug #17816: singularity not setting working directory added

**#14 - 06/16/2021 08:32 PM - Ward Vandewege**

- Blocks deleted (Bug #17816: singularity not setting working directory)

**#15 - 06/16/2021 08:33 PM - Ward Vandewege**

- Blocked by Bug #17816: singularity not setting working directory added

**#16 - 06/23/2021 03:43 PM - Peter Amstutz**

- Target version changed from 2021-06-23 sprint to 2021-07-07 sprint

**#17 - 07/07/2021 03:17 PM - Tom Clegg**

- Target version changed from 2021-07-07 sprint to 2021-07-21 sprint

**#18 - 07/21/2021 03:11 PM - Peter Amstutz**

- Target version changed from 2021-07-21 sprint to 2021-08-04 sprint

**#19 - 07/23/2021 07:11 PM - Ward Vandewege**

Revisiting this with latest main (which has the fix from [#17816](#)).

I built a new compute image for ce8i5 with main + my 17755-add-singularity-to-compute-image branch at [c2aceca339ec3a6f3d853865cebd0efe348ff518](#), image id: ce8i5-compute-v1627063559

The diagnostics workflow worked fine, as before: [https://workbench.ce8i5.arvadosapi.com/container\\_requests/ce8i5-xvhdp-6f6xwvqjihusey0](https://workbench.ce8i5.arvadosapi.com/container_requests/ce8i5-xvhdp-6f6xwvqjihusey0)

Re-running one of the cwl tests that failed last time succeeded:

[https://workbench.ce8i5.arvadosapi.com/container\\_requests/ce8i5-xvhdp-wqwktdy4ijpvrgx#Log](https://workbench.ce8i5.arvadosapi.com/container_requests/ce8i5-xvhdp-wqwktdy4ijpvrgx#Log)

I kicked off the cwl tests workflow at <https://ci.arvados.org/job/run-cwl-test-ce8i5/420/>

It failed (errors to be checked), and it took 14 hours (!) as opposed to the normal 3+ hours.

The next batch of singularity fixes ([#17813](#)) has been merged, which includes the fix to do the image conversion just once. The cwl tests are running against it at <https://ci.arvados.org/job/run-cwl-test-ce8i5/423/>. That run completed at 2 hours and 31 minutes, about 35 minutes faster than with Docker (!). One test failure remains, because we do not support the Docker ENTRYPOINT feature yet.

**#20 - 08/03/2021 01:53 PM - Peter Amstutz**

- Related to Story #16305: Singularity support added

**#21 - 08/03/2021 06:33 PM - Ward Vandewege**

Preparing to run [pirca-xvhdp-m4r8n9qmoh5ggra](#). The costanalyzer says these node types were used in that run on pirca:

m5.2xlarge  
m5.4xlarge  
m5.large  
m5.xlarge  
t3.small

Our quota on ce8i5 (azure) is tiny, so maybe we'll move this testing over to tordo (aws) where we have a bit more room. Building compute image with singularity for tordo at [https://ci.arvados.org/view/Packer%20Images/job/packer-build-compute-image/138/...](https://ci.arvados.org/view/Packer%20Images/job/packer-build-compute-image/138/)

That image built, I've reconfigured tordo to use Singularity with it. I'm gathering the files needed to run this workflow at <https://workbench.tordo.arvadosapi.com/projects/tordo-j7d0g-q56odvtnbo37ajt>.

**#22 - 08/04/2021 03:30 PM - Peter Amstutz**

- Target version changed from 2021-08-04 sprint to 2021-08-18 sprint

**#23 - 08/06/2021 01:32 PM - Ward Vandewege**

Ward Vandewege wrote:

Preparing to run [pirca-xvhdp-m4r8n9qmoh5ggra](#). The costanalyzer says these node types were used in that run on pirca:

m5.2xlarge  
m5.4xlarge  
m5.large  
m5.xlarge  
t3.small

Our quota on ce8i5 (azure) is tiny, so maybe we'll move this testing over to tordo (aws) where we have a bit more room. Building compute image with singularity for tordo at [https://ci.arvados.org/view/Packer%20Images/job/packer-build-compute-image/138/...](https://ci.arvados.org/view/Packer%20Images/job/packer-build-compute-image/138/)

That image built, I've reconfigured tordo to use Singularity with it. I'm gathering the files needed to run this workflow at <https://workbench.tordo.arvadosapi.com/projects/tordo-j7d0g-q56odvtnbo37ajt>.

Hmm, tordo seems to be having problems running its usual diag jobs because `singularity build` requires mksquashfs, from the squashfs-tools package (cf. <https://ci.arvados.org/view/Arvados%20build%20pipeline/job/diagnostics-tordo/259/console>). I've added that dependency to the 17755-add-singularity-to-compute-image branch and am building a new image with it for Tordo in <https://ci.arvados.org/view/Packer%20Images/job/packer-build-compute-image/139/console>. Curious I didn't see that on ce8i5. I suppose the azure Debian base image includes that package, and the AWS one does not.

New AMI for tordo is ami-0d9c1d43a4700db01, the config has been updated, and the diagnostics job succeeded at <https://ci.arvados.org/view/Arvados%20build%20pipeline/job/diagnostics-tordo/260/console> with singularity.

Good. Now, the actual chromosome 19 workflow.

The gist of it is to check out <https://github.com/arvados/arvados-tutorial> and then something like arvados-cwl-runner cwl/wgs-processing-wf.yml/wgs-processing-wf-chr19.yml

I copied all the collections referenced in wgs-processing-wf-chr19.yml to my 17755 test project at [tordo-j7d0g-q56odvtnbo37ajt](https://tordo-j7d0g-q56odvtnbo37ajt).

And then:

```
arvados-cwl-runner --project-uuid tordo-j7d0g-q56odvtnbo37ajt
./cwl/wgs-processing-wf.cwl ./yml/wgs-processing-wf-chr19.yml
INFO /usr/bin/arvados-cwl-runner 2.3.0.dev20210729201354, arvados-python-client 2.3.0.dev20210726143205, cwltool 3.0.20210319143721
INFO Resolved './cwl/wgs-processing-wf.cwl' to 'file:///home/ward/arvados-tutorial/WGS-processing/cwl/wgs-processing-wf.cwl'
INFO Using cluster tordo (https://workbench2.tordo.arvadosapi.com/)
INFO Upload local files: "head.html" "generatereport.py" "tail.html"
INFO Uploaded to 4fce132bc4ae9b4cc844115246a6bd41+175 (tordo-4zz18-9p4undcuyh75got)
INFO ['docker', 'pull', 'broadinstitute/gatk:4.1.7.0']
4.1.7.0: Pulling from broadinstitute/gatk
9ff7e2e5f967: Pull complete
59856638ac9f: Pull complete
6f317d6d954b: Pull complete
a9dde5e2a643: Pull complete
f7e0d1ec4184: Pull complete
8ec82ad80803: Pull complete
Digest: sha256:192fedf9b9d65809da4a2954030269e3f311d296e6b5e4c6c7deec12c7fe84b2
Status: Downloaded newer image for broadinstitute/gatk:4.1.7.0
docker.io/broadinstitute/gatk:4.1.7.0
INFO Uploading Docker image broadinstitute/gatk:4.1.7.0
2021-08-06 21:12:40 arvados.arv_put[16368] INFO: Creating new cache file at /home/ward/.cache/arvados/arv-put/6e8880e401836b16215bb738675fcb39
4022M / 4022M 100.0% 2021-08-06 21:13:48 arvados.arv_put[16368] INFO:
2021-08-06 21:13:48 arvados.arv_put[16368] INFO: Collection saved as 'Docker image broadinstitute gatk:4.1.7.0 sha256:e7b5b'
tordo-4zz18-bc2garec0qvov3i
2021-08-06 21:13:49 cwltool[16368] INFO: ['docker', 'pull', 'curii/clinvar-report']
Using default tag: latest
Error response from daemon: pull access denied for curii/clinvar-report, repository does not exist or may require 'docker login': denied: requested access to the resource is denied
2021-08-06 21:13:49 cwltool[16368] ERROR: Workflow error, try again with --debug for more information:
cwl/helper/gvcf-to-vcf.cwl:7:5: Command '['docker', 'pull', 'curii/clinvar-report']' returned
non-zero exit status 1. (X-Request-Id: req-lqjx8gknskluif6s7g0m)
```

Hmmm, looks like I missed that one.

We should probably publish that container on dockerhub...

```
ward@shell:~/arvados-tutorial/WGS-processing$ arv-copy tordo-j7d0g-q56odvtnbo37ajt
pirca-4zz18-rn8581yvkk73lmc: 1067M / 1067M 100.0%
2021-08-06 21:16:19 arvados.arv-copy[16704] INFO:
2021-08-06 21:16:19 arvados.arv-copy[16704] INFO: Success: created copy with uuid tordo-4zz18-4zqcmpzmbikh0w
```

Take 2:

```
ward@shell:~/arvados-tutorial/WGS-processing$ arvados-cwl-runner --project-uuid tordo-j7d0g-q56odvtnbo37ajt
./cwl/wgs-processing-wf.cwl ./yml/wgs-processing-wf-chr19.yml
INFO /usr/bin/arvados-cwl-runner 2.3.0.dev20210729201354, arvados-python-client 2.3.0.dev20210726143205, cwltool 3.0.20210319143721
INFO Resolved './cwl/wgs-processing-wf.cwl' to 'file:///home/ward/arvados-tutorial/WGS-processing/cwl/wgs-processing-wf.cwl'
INFO Using cluster tordo (https://workbench2.tordo.arvadosapi.com/)
INFO Upload local files: "head.html" "generatereport.py" "tail.html"
INFO Using collection 4fce132bc4ae9b4cc844115246a6bd41+175 (tordo-4zz18-9p4undcuyh75got)
INFO ['docker', 'pull', 'curii/bwa-samtools']
Using default tag: latest
Error response from daemon: pull access denied for curii/bwa-samtools, repository does not exist or may require 'docker login': denied: requested access to the resource is denied
ERROR Workflow error, try again with --debug for more information:
cwl/helper/bwamem-samtools-view.cwl:7:5: Command '['docker', 'pull', 'curii/bwa-samtools']' returned
non-zero exit status 1.
```

Jeez. More copying:

```
ward@shell:~/arvados-tutorial/WGS-processing$ arv-copy pirca-4zz18-mjt3nkmmuuo4uos --project-uuid tordo-j7d0g-q56odvtnbo37ajt
pirca-4zz18-mjt3nkmmuuo4uos: 999M / 999M 100.0%
2021-08-06 21:21:46 arvados.arv-copy[17076] INFO:
2021-08-06 21:21:46 arvados.arv-copy[17076] INFO: Success: created copy with uuid tordo-4zz18-8r0gqx1l14jcpfv
```

And one more try:

```
ward@shell:~/arvados-tutorial/WGS-processing$ arvados-cwl-runner --project-uuid tordo-j7d0g-q56odvtnbo37ajt
./cwl/wgs-processing-wf.cwl ./yml/wgs-processing-wf-chr19.yml
INFO /usr/bin/arvados-cwl-runner 2.3.0.dev20210729201354, arvados-python-client 2.3.0.dev20210726143205, cwltool 3.0.20210319143721
INFO Resolved './cwl/wgs-processing-wf.cwl' to 'file:///home/ward/arvados-tutorial/WGS-processing/cwl/wgs-processing-wf.cwl'
INFO Using cluster tordo (https://workbench2.tordo.arvadosapi.com/)
INFO Upload local files: "generatereport.py" "tail.html" "head.html"
INFO Using collection 4fcel32bc4ae9b4cc844115246a6bd41+175 (tordo-4zz18-9p4undcuyh75got)
INFO ['docker', 'pull', 'biocontainers/fastqc:v0.11.9_cv6']
v0.11.9_cv6: Pulling from biocontainers/fastqc
9ff7e2e5f967: Already exists
59856638ac9f: Already exists
6f317d6d954b: Already exists
a9dde5e2a643: Already exists
675cac559d07: Pull complete
0e6dd20ee818: Pull complete
374c558e71da: Pull complete
0df3c64bb19a: Pull complete
e936d7784cf9: Pull complete
4dfd8d164cc5: Pull complete
473490461998: Pull complete
8f5e491552e6: Pull complete
a66ab3a674d9: Pull complete
18f922275a6f: Pull complete
ef8214645dd0: Pull complete
Digest: sha256:c740cf547548689724bf43b3e923fea6a7588c17d59b7cbd76fa9ac7544102c5
Status: Downloaded newer image for biocontainers/fastqc:v0.11.9_cv6
docker.io/biocontainers/fastqc:v0.11.9_cv6
INFO Uploading Docker image biocontainers/fastqc:v0.11.9_cv6
2021-08-06 21:24:52 arvados.arv_put[17135] INFO: Creating new cache file at /home/ward/.cache/arvados/arv-put/28720ca8e76a2892c0ea812e2e0e0ab6
1001M / 1001M 100.0% 2021-08-06 21:25:06 arvados.arv_put[17135] INFO:
2021-08-06 21:25:06 arvados.arv_put[17135] INFO: Collection saved as 'Docker image biocontainers fastqc:v0.11.9_cv6 sha256:33755'
tordo-4zz18-h6aj03stuegf0hu
2021-08-06 21:25:08 cwltool[17135] INFO: Using collection cache size 256 MiB
2021-08-06 21:25:10 arvados.cwl-runner[17135] INFO: [container WGS processing workflow scattered over samples] submitted container_request tordo-xvhdp-4dc271h513rh7ai
2021-08-06 21:25:10 arvados.cwl-runner[17135] INFO: Monitor workflow progress at https://workbench2.tordo.arvadosapi.com/processes/tordo-xvhdp-4dc271h513rh7ai
```

It's running now. I should have used --no-wait, but I'm in screen so it's all right.

#### #24 - 08/06/2021 11:14 PM - Ward Vandewege

That almost worked. One step got stuck in the 'Queued' state:

[https://workbench.tordo.arvadosapi.com/container\\_requests/tordo-xvhdp-iwwdgh2w5wyprix#Log](https://workbench.tordo.arvadosapi.com/container_requests/tordo-xvhdp-iwwdgh2w5wyprix#Log)

Looks like the container was restarted twice. The third copy appears just stuck.

Nothing obviously wrong on the compute node:

```
root@ip-10-253-254-170:/home/crunch# singularity --version
singularity version 3.7.4

root@ip-10-253-254-170:/var/log# df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            970M     0  970M   0% /dev
tmpfs           196M   18M   179M   9% /run
/dev/nvme1n1p1  7.7G   3.0G   4.4G  41% /
tmpfs           980M     0   980M   0% /dev/shm
tmpfs           5.0M   4.0K   5.0M   1% /run/lock
tmpfs           980M     0   980M   0% /sys/fs/cgroup
/dev/nvme1n1p15 124M   278K   124M   1% /boot/efi
```

```
tmpfs          196M    0 196M   0% /run/user/1000
/dev/mapper/tmp  47G   81M  47G   1% /tmp
tmpfs          196M    0 196M   0% /run/user/1001
```

USER	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
root	2	0.0	0.0	0	0	?	S	21:26	0:00	[kthreadd]
root	3	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [rcu_gp]
root	4	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [rcu_par_gp]
root	6	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kworker/0:0H-kblockd]
root	8	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [mm_percpu_wq]
root	9	0.0	0.0	0	0	?	S	21:26	0:00	\_ [ksoftirqd/0]
root	10	0.0	0.0	0	0	?	I	21:26	0:00	\_ [rcu_sched]
root	11	0.0	0.0	0	0	?	I	21:26	0:00	\_ [rcu_bh]
root	12	0.0	0.0	0	0	?	S	21:26	0:00	\_ [migration/0]
root	14	0.0	0.0	0	0	?	S	21:26	0:00	\_ [cpuhp/0]
root	15	0.0	0.0	0	0	?	S	21:26	0:00	\_ [cpuhp/1]
root	16	0.0	0.0	0	0	?	S	21:26	0:00	\_ [migration/1]
root	17	0.0	0.0	0	0	?	S	21:26	0:00	\_ [ksoftirqd/1]
root	19	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kworker/1:0H-kblockd]
root	20	0.0	0.0	0	0	?	S	21:26	0:00	\_ [kdevtmpfs]
root	21	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [netns]
root	22	0.0	0.0	0	0	?	S	21:26	0:00	\_ [kauditd]
root	23	0.0	0.0	0	0	?	S	21:26	0:00	\_ [khungtaskd]
root	24	0.0	0.0	0	0	?	S	21:26	0:00	\_ [oom_reaper]
root	25	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [writeback]
root	26	0.0	0.0	0	0	?	S	21:26	0:00	\_ [kcompactd0]
root	27	0.0	0.0	0	0	?	SN	21:26	0:00	\_ [ksmd]
root	28	0.0	0.0	0	0	?	SN	21:26	0:00	\_ [khugepaged]
root	29	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [crypto]
root	30	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kintegrityd]
root	31	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kblockd]
root	32	0.0	0.0	0	0	?	S	21:26	0:00	\_ [watchdogd]
root	33	0.0	0.0	0	0	?	I	21:26	0:00	\_ [kworker/1:1-events]
root	34	0.0	0.0	0	0	?	I	21:26	0:00	\_ [kworker/u4:1-events_unbound]
root	36	0.0	0.0	0	0	?	S	21:26	0:00	\_ [kswapd0]
root	52	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kthrotld]
root	53	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [ipv6_addrconf]
root	62	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kstrp]
root	67	0.0	0.0	0	0	?	I	21:26	0:00	\_ [kworker/1:2-mm_percpu_wq]
root	104	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [nvme-wq]
root	105	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [nvme-reset-wq]
root	106	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [nvme-delete-wq]
root	107	0.0	0.0	0	0	?	I	21:26	0:00	\_ [kworker/u4:2-events_unbound]
root	115	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [ena]
root	184	0.0	0.0	0	0	?	S	21:26	0:00	\_ [jbd2/nvme1n1p1-]
root	185	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [ext4-rsv-conver]
root	229	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kworker/0:1H-kblockd]
root	257	0.0	0.0	0	0	?	I	21:26	0:00	\_ [kworker/0:2-events]
root	528	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kworker/1:1H-kblockd]
root	807	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kdmflush]
root	1254	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kdmflush]
root	1259	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kcryptd_io]
root	1260	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kcryptd]
root	1261	0.0	0.0	0	0	?	S	21:26	0:00	\_ [dmccrypt_write]
root	1263	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kworker/u5:1-kcryptd]
root	1264	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [kworker/u5:2-kcryptd]
root	1347	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfsalloc]
root	1348	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs_mru_cache]
root	1349	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-buf/dm-1]
root	1350	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-data/dm-1]
root	1351	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-conv/dm-1]
root	1352	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-cil/dm-1]
root	1353	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-reclaim/dm-]
root	1354	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-log/dm-1]
root	1355	0.0	0.0	0	0	?	I<	21:26	0:00	\_ [xfs-eofblocks/d]
root	1356	0.0	0.0	0	0	?	S	21:26	0:00	\_ [xfsaild/dm-1]
root	3687	0.0	0.0	0	0	?	I	22:17	0:00	\_ [kworker/0:0-events_power_efficient]
root	6027	0.0	0.0	0	0	?	I	23:12	0:00	\_ [kworker/0:1-xfs-cil/dm-1]
root	6186	0.0	0.0	0	0	?	I<	23:13	0:00	\_ [kworker/u5:0-kcryptd]
root	1	0.0	0.5	22116	10224	?	Ss	21:26	0:02	/sbin/init
root	247	0.0	0.5	42948	11668	?	Ss	21:26	0:05	/lib/systemd/systemd-journald
root	265	0.0	0.2	20068	4880	?	Ss	21:26	0:00	/lib/systemd/systemd-udev
root	381	0.0	0.2	9488	5720	?	Ss	21:26	0:00	/sbin/dhclient -4 -v -i -pf /run/dhclient.ens5.pid -lf /var/lib/dhcp/dhclient.ens5.leases -I -df /var/lib/dhcp/dhclient6.ens5.leases ens5
root	483	0.0	0.2	9488	5644	?	Ss	21:26	0:00	/sbin/dhclient -6 -v -pf /run/dhclient6.ens5.

```

pid -lf /var/lib/dhcp/dhclient6.ens5.leases -I -df /var/lib/dhcp/dhclient.ens5.leases -nw ens5
message+ 579 0.0 0.2 8808 4132 ? Ss 21:26 0:00 /usr/bin/dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd-activation --syslog-only
root 581 0.0 0.3 19520 7336 ? Ss 21:26 0:00 /lib/systemd/systemd-logind
root 582 0.0 0.2 225824 4280 ? Ssl 21:26 0:01 /usr/sbin/rsyslogd -n -iNONE
root 584 0.0 0.1 10228 2664 ? Ss 21:26 0:00 /usr/sbin/cron -f
_chrony 590 0.0 0.1 4816 2228 ? S 21:26 0:00 /usr/sbin/chronyd -F -1
_chrony 594 0.0 0.0 4816 268 ? S 21:26 0:00 \_ /usr/sbin/chronyd -F -1
root 591 0.0 0.0 5612 1688 tty1 Ss+ 21:26 0:00 /sbin/agetty -o -p -- \u --noclear tty1 linux
root 592 0.0 0.1 7136 2052 ttyS0 Ss+ 21:26 0:00 /sbin/agetty -o -p -- \u --keep-baud 115200,38400,9600 ttyS0 vt220
root 593 0.0 0.3 15852 6708 ? Ss 21:26 0:00 /usr/sbin/sshd -D
root 625 0.0 0.3 16612 7828 ? Ss 21:26 0:00 \_ sshd: admin [priv]
admin 651 0.0 0.2 16612 4848 ? S 21:26 0:00 | \_ sshd: admin@notty
root 1268 0.0 0.3 16612 7956 ? Ss 21:26 0:00 \_ sshd: admin [priv]
admin 1274 0.0 0.2 16784 5348 ? S 21:26 0:01 | \_ sshd: admin@notty
root 6025 0.0 0.4 16608 8076 ? Ss 23:12 0:00 \_ sshd: crunch [priv]
crunch 6043 0.0 0.2 16608 4940 ? S 23:12 0:00 \_ sshd: crunch@pts/0
crunch 6044 0.0 0.2 9648 4736 pts/0 Ss 23:12 0:00 \_ -bash
root 6058 0.0 0.1 11664 3700 pts/0 S 23:12 0:00 \_ sudo -s
root 6059 0.0 0.1 8716 3632 pts/0 S 23:12 0:00 \_ /bin/bash
root 6222 0.0 0.1 12512 3204 pts/0 R+ 23:14 0:00 \_ ps auxwf
admin 628 0.0 0.4 21148 9164 ? Ss 21:26 0:00 /lib/systemd/systemd --user
admin 629 0.0 0.1 22888 2304 ? S 21:26 0:00 \_ (sd-pam)
root 1315 0.2 2.2 1327060 44344 ? Sl 21:26 0:14 /var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df -no-detach --detach --stdin-env --runtime-engine=singularity tordo-dz642-xm3fb15he5o2aul
root 1325 0.1 2.4 1320308 49556 ? Sl 21:26 0:12 \_ /usr/share/python3/dist/python3-arvados-fuse/bin/python /usr/bin/arv-mount --foreground --allow-other --read-write --storage-classes default --crunchst at-interval=10 --file-cache 268435456 --mount-by-pdh by_id --mount-by-id by_uuid /tmp/crunch-run.tordo-dz642-xm3fb15he5o2aul.806337907/keep928487990
root 1417 0.1 4.1 1074032 83872 ? Ssl 21:26 0:09 /usr/sbin/dockerd -H fd:// --default-ulimit nofile=10000:10000 --dns 10.253.0.2
root 1424 0.1 1.6 878344 32680 ? Ssl 21:26 0:12 \_ docker-containerd --config /var/run/docker-containerd/containerd/containerd.toml --log-level info
root 1546 0.0 1.0 32536 21896 ? Ss 21:27 0:00 /usr/share/python3/dist/arvados-docker-cleaner/bin/python /usr/bin/arvados-docker-cleaner
crunch 6029 0.0 0.4 21160 9120 ? Ss 23:12 0:00 /lib/systemd/systemd --user
crunch 6030 0.0 0.1 23080 2556 ? S 23:12 0:00 \_ (sd-pam)

```

From /var/log/auth.log:

```

Aug 6 21:26:56 ip-10-253-254-170 systemd-logind[581]: New session 3 of user admin.
Aug 6 21:26:56 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/usr/bin/cat /var/run/arvados-instance-secret
Aug 6 21:26:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:26:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:26:56 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/usr/bin/docker ps -q
Aug 6 21:26:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:26:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:26:56 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/usr/bin/sh -c set -e; dstdir="/var/lib/arvados/"; dstfile="/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df"; mkdir -p "$dstdir"; touch "$dstfile"; chmod 0755 "$dstdir" "$dstfile"; cat >"$dstfile"
Aug 6 21:26:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:26:57 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:26:57 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:26:57 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:26:57 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:26:57 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --detach --stdin-env --runtime-engine=singularity tordo-dz642-xm3fb15he5o2aul
Aug 6 21:26:57 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:26:57 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:27:06 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:27:06 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:27:06 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:27:16 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:27:16 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:27:16 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:27:26 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:27:26 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)

```

```
Aug 6 21:27:26 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:27:36 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:27:36 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:27:36 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:27:46 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:27:46 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:27:46 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 21:27:56 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 21:27:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 21:27:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
```

....

and on and on and on

....

```
Aug 6 23:25:56 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 23:25:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 23:25:56 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
Aug 6 23:26:06 ip-10-253-254-170 sudo: admin : TTY=unknown ; PWD=/home/admin ; USER=root ; COMMAND=/var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
Aug 6 23:26:06 ip-10-253-254-170 sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Aug 6 23:26:06 ip-10-253-254-170 sudo: pam_unix(sudo:session): session closed for user root
```

Running that command manually shows

```
root@ip-10-253-254-170:/var/log# /var/lib/arvados/crunch-run~59eb599fd48e090ea6b99794587612df --list
tordo-dz642-xm3fbl5he5o2aul
```

So we get the uuid of the container back. That's expected.

I guess something happened to the arv-mount process, it got stuck?

```
root@ip-10-253-254-170:/var/log# v /proc/1325/fd/
total 0
dr-x----- 2 root root 0 Aug 6 21:26 ./
dr-xr-xr-x 9 root root 0 Aug 6 21:26 ../
lr-x----- 1 root root 64 Aug 6 21:26 0 -> /dev/null
l-wx----- 1 root root 64 Aug 6 21:26 1 -> 'pipe:[16552]'
l-wx----- 1 root root 64 Aug 6 21:26 2 -> 'pipe:[16553]'
lrwx----- 1 root root 64 Aug 6 21:26 3 -> /run/lock/crunch-run-tordo-dz642-xm3fbl5he5o2aul.lock*
lrwx----- 1 root root 64 Aug 6 21:26 4 -> 'socket:[17495]'
lrwx----- 1 root root 64 Aug 6 23:34 5 -> /dev/fuse
lrwx----- 1 root root 64 Aug 6 23:34 6 -> 'socket:[17503]'
```

```
root@ip-10-253-254-170:/var/log# v /run/lock/crunch-run-tordo-dz642-xm3fbl5he5o2aul.lock
-rwx----- 1 root root 50 Aug 6 21:26 /run/lock/crunch-run-tordo-dz642-xm3fbl5he5o2aul.lock*
```

## #25 - 08/06/2021 11:42 PM - Ward Vandewege

Trying to cancel this job, the cancel button in wb doesn't work, this is the error I see in the controller logs:

```
Aug 06 23:41:16 tordo.arvadosapi.com arvados-controller[13414]: {"PID":13414,"RequestID":"req-3x51wipebp66b8px
x8hg","level":"info","msg":"response","remoteAddr":"127.0.0.1:41130","reqBytes":1982,"reqForwardedFor":"10.253
.0.41","reqHost":"tordo.arvadosapi.com","reqMethod":"POST","reqPath":"/arvados/v1/container_requests/tordo-xvhd
p-iwwdqh2w5wyprix","reqQuery":"","respBody":{"errors":{"//railsapi.internal/arvados/v1/container_requests/
tordo-xvhd-p-iwwdqh2w5wyprix: 422 Unprocessable Entity: Mounts cannot be modified in state 'Final' ({}"/keep/
35d4f4c5720a1e1c7b5c583cc901c3e0+3765\\"=\u003e{\\"kind\\"=\u003e{\\"collection\\",\\"portable_data_ha
sh\\"=\u003e{\\"35d4f4c5720a1e1c7b5c583cc901c3e0+3765\\"}, \\"/tmp\\"=\u003e{\\"capacity\\"=\u003e107
3741824, \\"kind\\"=\u003e{\\"tmp\\"}, \\"/var/spool/cwl\\"=\u003e{\\"capacity\\"=\u003e1073741824,
\\"kind\\"=\u003e{\\"tmp\\"}}, {\\"/keep/35d4f4c5720a1e1c7b5c583cc901c3e0+3765\\"=\u003e{\\"capacity\\
\"=\u003e0, \\"commit\\"=\u003e{\\"\\\"}, \\"content\\"=\u003e{enil, \\"device_type\\"=\u003e{\\"\\\"},
\\"exclude_from_output\\"=\u003e{efalse, \\"git_url\\"=\u003e{\\"\\\"}, \\"kind\\"=\u003e{\\"collection\\
\"}, \\"path\\"=\u003e{\\"\\\"}, \\"portable_data_hash\\"=\u003e{\\"35d4f4c5720a1e1c7b5c583cc901c3e0+3765
\\"}, \\"repository_name\\"=\u003e{\\"\\\"}, \\"uuid\\"=\u003e{\\"\\\"}, \\"writable\\"=\u003e{efalse}, \
\"/tmp\\"=\u003e{\\"capacity\\"=\u003e1073741824, \\"commit\\"=\u003e{\\"\\\"}, \\"content\\"=\u003e
enil, \\"device_type\\"=\u003e{\\"\\\"}, \\"exclude_\",\"respBytes\":1618,\"respStatus\":\"Unprocessable Entity\",
\"respStatusCode\":422,\"time\":\"2021-08-06T23:41:16.438539926Z\",\"timeToStatus\":0.018596,\"timeTotal\":0.018626,\"tim
eWriteBody\":0.000030}
```



This sounds vaguely familiar. This is a second bug, then.

I'll kill the node with a-d-c.

```
curl -H "Authorization: Bearer $TOKEN" -X GET http://10.253.0.44:9006/arvados/v1/dispatch/instances | jq .
{
  "items": [
    {
      "instance": "i-0321dc62106d21d98",
      "address": "10.253.254.104",
      "price": 0.0208,
      "arvados_instance_type": "t3small",
      "provider_instance_type": "t3.small",
      "last_container_uuid": "tordo-dz642-mitu9vi55h7femh",
      "last_busy": "2021-08-07T00:07:16.063745476Z",
      "worker_state": "running",
      "idle_behavior": "run"
    },
    {
      "instance": "i-0362e2a61c62cfaf2",
      "address": "10.253.254.170",
      "price": 0.0208,
      "arvados_instance_type": "t3small",
      "provider_instance_type": "t3.small",
      "last_container_uuid": "tordo-dz642-xm3fb15he5o2aul",
      "last_busy": "2021-08-07T00:07:16.078001157Z",
      "worker_state": "running",
      "idle_behavior": "run"
    }
  ]
}
```

```
tordo:~# curl -H "Authorization: Bearer $TOKEN" -X POST http://10.253.0.44:9006/arvados/v1/dispatch/instances/kill?instance_id=i-0362e2a61c62cfaf2 | jq .
```

```
tordo:~# curl -H "Authorization: Bearer $TOKEN" -X GET http://10.253.0.44:9006/arvados/v1/dispatch/instances | jq .
{
  "items": [
    {
      "instance": "i-0321dc62106d21d98",
      "address": "10.253.254.104",
      "price": 0.0208,
      "arvados_instance_type": "t3small",
      "provider_instance_type": "t3.small",
      "last_container_uuid": "tordo-dz642-mitu9vi55h7femh",
      "last_busy": "2021-08-07T00:07:56.023842384Z",
      "worker_state": "running",
      "idle_behavior": "run"
    },
    {
      "instance": "i-0362e2a61c62cfaf2",
      "address": "10.253.254.170",
      "price": 0.0208,
      "arvados_instance_type": "t3small",
      "provider_instance_type": "t3.small",
      "last_container_uuid": "tordo-dz642-xm3fb15he5o2aul",
      "last_busy": "2021-08-07T00:07:46.041616399Z",
      "worker_state": "shutdown",
      "idle_behavior": "run"
    }
  ]
}
```

A new node was created:

```
tordo:~# curl -H "Authorization: Bearer $TOKEN" -X GET http://10.253.0.44:9006/arvados/v1/dispatch/instances | jq .
{
  "items": [
    {
      "instance": "i-0321dc62106d21d98",
      "address": "10.253.254.104",
      "price": 0.0208,
      "arvados_instance_type": "t3small",

```

```

    "provider_instance_type": "t3.small",
    "last_container_uuid": "tordo-dz642-mitu9yi55h7femh",
    "last_busy": "2021-08-07T00:08:56.02153052Z",
    "worker_state": "running",
    "idle_behavior": "run"
  },
  {
    "instance": "i-06dd557531f9df013",
    "address": "10.253.254.243",
    "price": 0.0208,
    "arvados_instance_type": "t3small",
    "provider_instance_type": "t3.small",
    "last_container_uuid": "tordo-dz642-xm3fbl5he5o2aul",
    "last_busy": "2021-08-07T00:08:56.047446083Z",
    "worker_state": "running",
    "idle_behavior": "run"
  }
]
}

```

And this node ran the job without trouble: [https://workbench.tordo.arvadosapi.com/container\\_requests/tordo-xvhdp-iwwdqh2w5wyprix#Status](https://workbench.tordo.arvadosapi.com/container_requests/tordo-xvhdp-iwwdqh2w5wyprix#Status).

I think we have a missing error check somewhere early in the singularity docker loading code. There was a (transient?) problem on this particular node, and that made the container get stuck.

#### #26 - 08/09/2021 02:15 PM - Peter Amstutz

The first two failures I think are due to a race condition between two crunch-run processes trying to convert and caching the singularity image.

This attempt found the collection, but apparently before the cached image had been created. (not supposed to happen.)

```

2021-08-06T21:26:16.608212665Z Using Docker image id
"sha256:337550d506a3fc77e30292bba95108f1cd34a33719f0dd997d0de0540522def7"
2021-08-06T21:26:16.608240987Z Loading Docker image from keep
2021-08-06T21:26:17.188261275Z building singularity image
2021-08-06T21:26:17.223285265Z [singularity build /tmp/crunch-run.tordo-dz642-pezh87oegh5fgbs7
.099258167/keep207354666/by_uuid/tordo-4zz18-t0wx4utpwx4ligf/image.sif docker-archive:///tmp/crunch-run-singularity-917581112/image.tar]
2021-08-06T21:26:26.466105517Z INFO: Starting build...
2021-08-06T21:26:26.466105517Z FATAL: While performing build: conveyor failed to get: Error loading tar component
337550d506a3fc77e30292bba95108f1cd34a33719f0dd997d0de0540522def7.json: open /tmp/crunch-run-singularity-917581112/image.tar: no
such file or directory
2021-08-06T21:26:26.466234171Z error in Run: While loading container image: exit status 255
2021-08-06T21:26:26.466268708Z error in CaptureOutput: error scanning files to copy to output: lstat "/var/spool/cwl": lstat /tmp/crunch-run.
tordo-dz642-pezh87oegh5fgbs7.099258167/tmp701045228: no such file or directory
2021-08-06T21:26:26.605452965Z Cancelled

```

On the second attempt, it tried to create a collection with the same temporary name (down to the exact timestamp?) and that failed.

```

2021-08-06T21:26:47.149336231Z Executing container 'tordo-dz642-amjt50vzn4qyn4n'
...
2021-08-06T21:26:47.972965997Z error in Run: While loading container image: error creating 'singularity image for
sha256:337550d506a3fc77e30292bba95108f1cd34a33719f0dd997d0de0540522def7 2021-08-06T21:26:47Z' collection: request failed:
https://workbench.tordo.arvadosapi.com/arvados/v1/collections: 422 Unprocessable Entity: //railsapi.internal/arvados/v1/collections: 422 Unprocessable
Entity: #<ActiveRecord::RecordNotUnique: PG::UniqueViolation: ERROR: duplicate key value violates unique constraint
"index_collections_on_owner_uuid_and_name"
2021-08-06T21:26:47.972965997Z DETAIL: Key (owner_uuid, name)=(tordo-j7d0g-7p82g804nk5l7gx, singularity image for
sha256:337550d506a3fc77e30292bba95108f1cd34a33719f0dd997d0de0540522def7 2021-08-06T21:26:47Z) already exists.
2021-08-06T21:26:47.972965997Z : INSERT INTO "collections" ("owner_uuid", "created_at", "modified_by_user_uuid", "modified_at",
"portable_data_hash", "updated_at", "uuid", "manifest_text", "name", "properties", "delete_at", "file_names", "trash_at", "current_version_uuid")
VALUES ($1, $2, $3, $4, $5, $6, $7, $8, $9, $10, $11, $12, $13, $14) RETURNING "id"> (req-8t57dq95orqsvelydce)
2021-08-06T21:26:48.136224600Z Cancelled

```

#### #27 - 08/09/2021 02:35 PM - Peter Amstutz

The 3rd failure is puzzling, if you check stderr it actually did run to the end, but failed the upload. This happened after the intervention of being stuck in queued state and killing the compute node.

```

2021-08-07T00:08:37.343893355Z Executing container 'tordo-dz642-xm3fbl5he5o2aul'
...
2021-08-07T00:08:39.823076366Z Waiting for container to finish
2021-08-07T00:10:47.551657400Z error in CaptureOutput: error scanning files to copy to output: lstat "/var/spool/cwl": lstat /tmp/crunch-run.
tordo-dz642-xm3fbl5he5o2aul.609138299/tmp536330368: no such file or directory

```

**#28 - 08/09/2021 02:58 PM - Peter Amstutz**

Was the "queued" container actually "locked"? (probably but would be good to confirm).

Because it didn't get far enough, the container record was reused for the 2nd attempt, all the logs in keep are only the second attempt. Fortunately there are logs stored on the API server:

<https://workbench.tordo.arvadosapi.com/containers/tordo-dz642-xm3fbl5he5o2aul#Log>

Although they don't tell us very much. It hung sometime after starting arv-mount.

This might be related to the event handling deadlock we've been seeing.

**#29 - 08/17/2021 07:17 PM - Ward Vandewege**

Peter Amstutz wrote:

Was the "queued" container actually "locked"? (probably but would be good to confirm).

Because it didn't get far enough, the container record was reused for the 2nd attempt, all the logs in keep are only the second attempt. Fortunately there are logs stored on the API server:

<https://workbench.tordo.arvadosapi.com/containers/tordo-dz642-xm3fbl5he5o2aul#Log>

Although they don't tell us very much. It hung sometime after starting arv-mount.

This might be related to the event handling deadlock we've been seeing.

I'll make a new compute image before doing further testing: ami-0ba36796c859d41a2 has the arv-mount fix built in.

**#30 - 08/17/2021 07:25 PM - Ward Vandewege**

- Related to Bug #18027: [crunch-run] upload failed (singularity) added

**#31 - 08/17/2021 07:25 PM - Ward Vandewege**

- Related to Bug #18026: [singularity] crunch-run image cache race condition added

**#32 - 08/17/2021 09:02 PM - Ward Vandewege**

Peter Amstutz wrote:

Was the "queued" container actually "locked"? (probably but would be good to confirm).

Because it didn't get far enough, the container record was reused for the 2nd attempt, all the logs in keep are only the second attempt. Fortunately there are logs stored on the API server:

<https://workbench.tordo.arvadosapi.com/containers/tordo-dz642-xm3fbl5he5o2aul#Log>

Although they don't tell us very much. It hung sometime after starting arv-mount.

This might be related to the event handling deadlock we've been seeing.

Now that the fix for the arv-mount deadlock has been merged, I've built a new compute image with that fix (ami-0ba36796c859d41a2), and kicked off a new test run with

```
ward@shell:~/arvados-tutorial/WGS-processing$ arvados-cwl-runner --project-uuid tordo-j7d0g-g56odvtmbo37ajt
./cwl/wgs-processing-wf.cwl ./yml/wgs-processing-wf-chr19.yml
INFO /usr/bin/arvados-cwl-runner 2.3.0.dev20210813182157, arvados-python-client 2.3.0.dev20210810185721, cwltool 3.0.20210319143721
INFO Resolved './cwl/wgs-processing-wf.cwl' to 'file:///home/ward/arvados-tutorial/WGS-processing/cwl/wgs-processing-wf.cwl'
INFO Using cluster tordo (https://workbench2.tordo.arvadosapi.com/)
INFO Upload local files: "tail.html" "head.html" "generatereport.py"
INFO Using collection 4fce132bc4ae9b4cc844115246a6bd41+175 (tordo-4zz18-9p4undcuyl75got)
INFO Using collection cache size 256 MiB
INFO [container WGS processing workflow scattered over samples] submitted container_request tordo-xvhdp-911ekye8xxwcb7r
INFO Monitor workflow progress at https://workbench2.tordo.arvadosapi.com/processes/tordo-xvhdp-911ekye8xxwcb7
```

```

r
INFO [container WGS processing workflow scattered over samples] tordo-xvhdp-911ekye8xxwcb7r is Final
INFO Overall process status is success
INFO Final output collection 431b75fdad437345c678fe58163ad2ec+1243
INFO Output at https://workbench2.tordo.arvadosapi.com/collections/431b75fdad437345c678fe58163ad2ec+1243
{
  "gvcf": [
    {
      "basename": "ERR2122553.g.vcf.gz",
      "class": "File",
      "format": "http://edamontology.org/format_3016",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122553.g.vcf.gz",
      "secondaryFiles": [
        {
          "basename": "ERR2122553.g.vcf.gz.tbi",
          "class": "File",
          "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122553.g.vcf.gz.tbi"
        }
      ],
      "size": 100428289
    },
    {
      "basename": "ERR2122554.g.vcf.gz",
      "class": "File",
      "format": "http://edamontology.org/format_3016",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122554.g.vcf.gz",
      "secondaryFiles": [
        {
          "basename": "ERR2122554.g.vcf.gz.tbi",
          "class": "File",
          "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122554.g.vcf.gz.tbi"
        }
      ],
      "size": 83396874
    },
    {
      "basename": "ERR2122555.g.vcf.gz",
      "class": "File",
      "format": "http://edamontology.org/format_3016",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122555.g.vcf.gz",
      "secondaryFiles": [
        {
          "basename": "ERR2122555.g.vcf.gz.tbi",
          "class": "File",
          "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122555.g.vcf.gz.tbi"
        }
      ],
      "size": 100428289
    },
    {
      "basename": "ERR2122554.g.vcf.gz",
      "class": "File",
      "format": "http://edamontology.org/format_3016",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122554.g.vcf.gz",
      "secondaryFiles": [
        {
          "basename": "ERR2122554.g.vcf.gz.tbi",
          "class": "File",
          "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122554.g.vcf.gz.tbi"
        }
      ],
      "size": 83396874
    },
    {
      "basename": "ERR2122555.g.vcf.gz",
      "class": "File",
      "format": "http://edamontology.org/format_3016",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122555.g.vcf.gz",
      "secondaryFiles": [
        {
          "basename": "ERR2122555.g.vcf.gz.tbi",
          "class": "File",
          "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122555.g.vcf.gz.tbi"
        }
      ],
      "size": 68070153
    }
  ],
}

```

```

{
  "basename": "ERR2122556.g.vcf.gz",
  "class": "File",
  "format": "http://edamontology.org/format_3016",
  "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122556.g.vcf.gz",
  "secondaryFiles": [
    {
      "basename": "ERR2122556.g.vcf.gz.tbi",
      "class": "File",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122556.g.vcf.gz.tbi"
    }
  ],
  "size": 80162449
},
{
  "basename": "ERR2122557.g.vcf.gz",
  "class": "File",
  "format": "http://edamontology.org/format_3016",
  "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122557.g.vcf.gz",
  "secondaryFiles": [
    {
      "basename": "ERR2122557.g.vcf.gz.tbi",
      "class": "File",
      "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122557.g.vcf.gz.tbi"
    }
  ],
  "size": 92916518
}
],
"report": [
  {
    "basename": "ERR2122553.html",
    "class": "File",
    "format": "http://edamontology.org/format_2331",
    "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122553.html",
    "size": 213477
  },
  {
    "basename": "ERR2122554.html",
    "class": "File",
    "format": "http://edamontology.org/format_2331",
    "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122554.html",
    "size": 203560
  },
  {
    "basename": "ERR2122555.html",
    "class": "File",
    "format": "http://edamontology.org/format_2331",
    "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122555.html",
    "size": 224649
  },
  {
    "basename": "ERR2122556.html",
    "class": "File",
    "format": "http://edamontology.org/format_2331",
    "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122556.html",
    "size": 210935
  },
  {
    "basename": "ERR2122557.html",
    "class": "File",
    "format": "http://edamontology.org/format_2331",
    "location": "keep:431b75fdad437345c678fe58163ad2ec+1243/ERR2122557.html",
    "size": 225496
  }
]
}
INFO Final process status is success

```

Hmm, that reused everything, not quite what I wanted.

Trying again:

```

ward@shell:~/arvados-tutorial/WGS-processing$ arvados-cwl-runner --disable-reuse --project-uuid
tordo-j7d0g-q560dvt-nbo37ajt ./cwl/wgs-processing-wf.cwl ./yml/wgs-processing-wf-chr19.yml

```

```
INFO /usr/bin/arvados-cwl-runner 2.3.0.dev20210813182157, arvados-python-client 2.3.0.dev20210810185721, cwltool 3.0.20210319143721
INFO Resolved './cwl/wgs-processing-wf.cwl' to 'file:///home/ward/arvados-tutorial/WGS-processing/cwl/wgs-processing-wf.cwl'
INFO Using cluster tordo (https://workbench2.tordo.arvadosapi.com/)
INFO Upload local files: "generatereport.py" "tail.html" "head.html"
INFO Using collection 4fcel132bc4ae9b4cc844115246a6bd41+175 (tordo-4zz18-9p4undcuyh75got)
INFO Using collection cache size 256 MiB
INFO [container WGS processing workflow scattered over samples] submitted container_request tordo-xvhdp-bb8pjmzh93hkzmz
INFO Monitor workflow progress at https://workbench2.tordo.arvadosapi.com/processes/tordo-xvhdp-bb8pjmzh93hkzmz
```

...

This run completed successfully without obvious issues: [https://workbench.tordo.arvadosapi.com/container\\_requests/tordo-xvhdp-bb8pjmzh93hkzmz](https://workbench.tordo.arvadosapi.com/container_requests/tordo-xvhdp-bb8pjmzh93hkzmz).

This container started at 5:09 PM 8/17/2021. It completed in 1h18m at 6:27 PM 8/17/2021.

It has runtime of 1h16m(2mqueued) and used 5h18m of node allocation time (4.2x scaling).

### #33 - 08/18/2021 03:00 PM - Peter Amstutz

- Target version changed from 2021-08-18 sprint to 2021-09-01 sprint

### #34 - 09/01/2021 03:13 PM - Peter Amstutz

- Target version changed from 2021-09-01 sprint to 2021-09-15 sprint

### #35 - 09/03/2021 02:06 PM - Ward Vandewege

17755-add-singularity-to-compute-image is ready for review at [72d7d41944006d1f48f570784d4fe56b9812b0c8](#)

It builds singularity from source, it would be nice if there was a better way.

### #36 - 09/03/2021 05:58 PM - Lucas Di Pentima

LGTM, thanks!

### #37 - 09/03/2021 06:44 PM - Ward Vandewege

- Status changed from In Progress to Resolved

Applied in changeset [arvados|9695271cd15811612cb25f53fb38f133f747cd3a](#).