

Arvados - Story #15002

[API] Admin can prevent reuse by cancelling a completed container

03/20/2019 02:57 PM - Tom Clegg

| | | | |
|--|-------------------|------------------------|-----------------|
| Status: | Resolved | Start date: | 04/12/2019 |
| Priority: | Normal | Due date: | |
| Assigned To: | Eric Biagiotti | % Done: | 100% |
| Category: | | Estimated time: | 0.00 hour |
| Target version: | 2019-04-24 Sprint | | |
| Description | | | |
| Background: Sometimes it is desirable to avoid reusing a specific completed container, without disabling reuse for the corresponding step in a workflow. Example: a container exited 0 with bogus output, but it has never™ done that before and will probably never™ do it again, so it's not worth updating the workflow to use a new image/version. | | | |
| Proposed feature: | | | |
| Relax the "frozen in final state" constraints slightly, so an admin can use the API to update a container's state from Complete to Cancelled. | | | |
| Subtasks: | | | |
| Task # 15097: Review 15002-cancel-completed-container | | | Resolved |

Associated revisions

Revision 5756eb03 - 04/17/2019 09:28 PM - Tom Clegg

Merge branch '15002-container-permissions'

refs #15002

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

Revision 348756ad - 04/17/2019 09:43 PM - Eric Biagiotti

Merge remote-tracking branch 'origin/master' into 15002-cancel-completed-container

refs #15002

Arvados-DCO-1.1-Signed-off-by: Eric Biagiotti <ebiagiotti@veritasgenetics.com>

Revision aec31c21 - 04/18/2019 06:41 PM - Eric Biagiotti

Merge branch '15002-cancel-completed-container'

refs #15002

Arvados-DCO-1.1-Signed-off-by: Eric Biagiotti <ebiagiotti@veritasgenetics.com>

History

#2 - 03/20/2019 03:37 PM - Tom Morris

- Story points set to 1.0

#3 - 04/03/2019 02:07 PM - Tom Morris

- Target version changed from To Be Groomed to Arvados Future Sprints

#4 - 04/10/2019 04:21 PM - Tom Morris

- Assigned To set to Eric Biagiotti

- Target version changed from Arvados Future Sprints to 2019-04-24 Sprint

#5 - 04/11/2019 06:31 PM - Eric Biagiotti

- Status changed from New to In Progress

#6 - 04/12/2019 07:53 PM - Eric Biagiotti

Latest at: [e1b8a1b0bdb46d0e162ef7794c06b08d0a5fffa5](https://ci.curoverse.com/view/Developer/job/developer-run-tests/1190/)

- Created an Admin_state_transition hash for Complete => [Cancelled] that gets merged into the State_transition hash if the user is an admin.
- Added a test for setting a completed container's state to cancelled by admin and non-admin
- Created a new admin page and added detail to the containers API page.

Unit tests: <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1190/>

Note: In the container model, we check that the value for container state has actually changed before validating, which means that if a user tries to update to the current state (via arv container update or whatever), it passes validation. Even though nothing changed, this may be confusing to the user as it implies that this type of command is allowed. Let me know if you think this should be changed.

#7 - 04/12/2019 08:51 PM - Tom Clegg

I see how the story description and API docs imply otherwise, but non-admin users are already prohibited from updating containers at all (except the container process itself can update its own progress fields while state==Running). So I don't think we need a special "admin state transitions" here.

In docs:

For example, if a container exited successfully but produced bad output, it may not be worth updating the workflow to use a new image/version. Instead, changing the state of a container to Cancelled will disable reuse of the specific container.

I think this should be presented as an expedient thing to do *in addition to* fixing the bug, rather than as an alternative. We can't force people to fix the bugs, but I don't think we should nudge them in the wrong direction.

Maybe: "... it may not be feasible to update the workflow immediately ... Meanwhile, changing the state ..."

...will change it's state to...

its

...where -u is the UUID of the container:

...where [xxxxx-xxxxx-xxxxxxxxxxxxxxxx](#) is the UUID of the container

all containers that have been cancelled after completion

Either this should say "exited 0 and were then cancelled" or the condition should be ["exit_code", "<>", null]

\ No newline at end of file

(manage-containers.html.textile.liquid) should have one.

Cancelled (admin only)

All state changes are admin only.

The admin page title should probably be something like "Controlling container reuse" -- hint at the user's question, which they know, rather than the answer, which they don't. Same goes for the leading sentence. And I think it needs to be added to `_config.yml` so it shows up in the left nav.

#8 - 04/17/2019 07:46 PM - Tom Clegg

15002-container-permissions @ [5e95c9b723e36cf80e0b9c1bf02206520503d4f1](https://ci.curoverse.com/view/Developer/job/developer-run-tests/1191/)
<https://ci.curoverse.com/view/Developer/job/developer-run-tests/1191/>

#9 - 04/17/2019 08:33 PM - Eric Biagiotti

Latest at [7e0a38f70392822f362bd94f9d9093554c8f351a](https://ci.curoverse.com/view/Developer/job/developer-run-tests/1193/)

- Addressed the documentation issues from Note 7. I already added `manage-containers.html.textile.liquid` to `_config.yml`, so it should be showing up in the left navigation panel. Let me know if this still doesn't work for you.
- Merged in your changes and simplified Complete => Cancelled transition in the container model.
- Updated TestGetLockUnlockCancel test to accept the new state transition.

Unit tests: <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1193/>

#10 - 04/18/2019 02:14 PM - Tom Clegg

at [d97f1f245](#)

- this branch is a good candidate for squashing to reduce noise in the git history (code that never made it in)
- the "Controlling container reuse" page should probably be called `controlling-reuse.html` instead of `manage-containers.html`
- link title in `api/methods/containers.html` should be updated too
- maybe the doc page should allude to the fact that only an admin can do this -- or is it enough that it's in the "admin" section? I'm not sure whether we should rely on people to notice the nav clues when they arrive at this page via web search or the link on the API page.
- maybe "in your workflow" → "in all affected workflows"? (It's probably not "yours" and there can easily be more than one)
- maybe "disable reuse as the workflow continues to run" → "prevent it from being reused in subsequent workflows"?

It looks like this used to test propagating errors from the API.

```
err = cq.Cancel(arvadostest.CompletedContainerUUID)
c.Check(err, check.ErrorMatches, `.*State cannot change from Complete to Cancelled.*`)
```

Now that that's not an error, this part of the test seems superfluous, given the successful `Cancel()` calls above. Might as well remove it. Maybe check the returned error message in one of the non-nil error cases above instead?

#11 - 04/18/2019 06:36 PM - Eric Biagiotti

Addressed the comments from note 10, rebased on master, and squashed into one commit at [30fc42873b2c82e72a95393eb053ab1f7052618](#)

Unit Tests: <https://ci.curoverse.com/view/Developer/job/developer-run-tests/1194/>

#12 - 04/18/2019 06:40 PM - Tom Clegg

LGTM, thanks!

#13 - 04/18/2019 07:02 PM - Eric Biagiotti

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

#14 - 05/21/2019 10:27 PM - Tom Morris

- Release set to 15