

## Arvados - Bug #14977

### [arvados-dispatch-cloud] kill crunch-run procs for containers that are deleted or have state=Cancelled when dispatcher starts up

03/15/2019 06:22 PM - Tom Clegg

<b>Status:</b>	Resolved	<b>Start date:</b>	03/18/2019
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>	Tom Clegg	<b>% Done:</b>	100%
<b>Category:</b>	Crunch	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	2019-03-27 Sprint		
<b>Description</b>			
<p>Currently, a container that has state==Cancelled when arvados-dispatch-cloud starts up will never be added to the container queue, even if its UUID appears on an instance's probe result. Also, a container that has been deleted from the database will never have an entry added/updated in the dispatcher's container queue.</p> <p>The scheduler's sync() func is responsible for killing unneeded crunch-run processes, but it only looks at the container queue, so these crunch-run processes are allowed to run forever.</p> <p>Proposed solution:</p> <p>In (*scheduler.Scheduler)sync(), kill anything returned by sch.pool.Running() that isn't returned by sch.queue.Entries(). This should be safe from "kill crunch-run before seeing its UUID in the queue" races:</p> <ul style="list-style-type: none"><li>• at least one "get entire queue from controller/database" has succeeded before the first call to sync()</li><li>• UUIDs are added to Running() only during (*Scheduler)runQueue(), which does not run concurrently with (*Scheduler)sync().</li></ul> <p>In (*container.Queue)poll(), if a container's UUID is in the local queue but is not returned by the API calls that request that specific UUID, delete it from the local queue. The "get missing containers" loop will need to be more careful to avoid accidentally deleting containers when the API server chooses to return less than a full page of results.</p>			
<b>Subtasks:</b>			
Task # 14982: Review 14977-kill-if-not-in-queue			<b>Resolved</b>
<b>Related issues:</b>			
Blocks Arvados - Story #14807: [arvados-dispatch-cloud] Features/fixes needed...		<b>Resolved</b>	<b>01/29/2019</b>

#### Associated revisions

##### Revision 1c6c860b - 03/18/2019 03:11 PM - Tom Clegg

Merge branch '14977-kill-if-not-in-queue'

fixes #14977

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <[tclegg@veritasgenetics.com](mailto:tclegg@veritasgenetics.com)>

#### History

##### #1 - 03/15/2019 06:23 PM - Tom Clegg

- Status changed from New to In Progress

##### #2 - 03/15/2019 08:41 PM - Tom Clegg

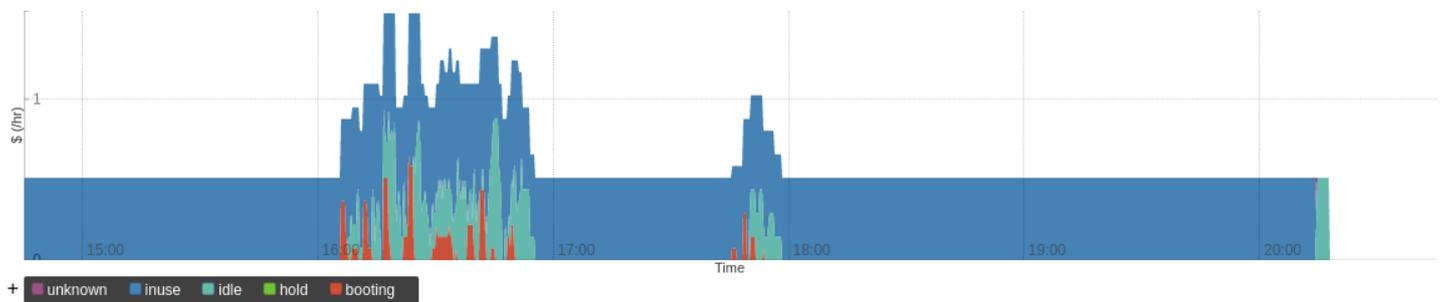
14977-kill-if-not-in-queue @ [2a748e79c3a72454d70e40f39fcad9dabf4943cc](https://github.com/Arvados/arvados-dispatch-cloud/commit/2a748e79c3a72454d70e40f39fcad9dabf4943cc)

This also fixes a different startup bug: "fix stale locks" was not waiting for the pool to load its initial instance list, so it would always return immediately, and the scheduler would create too many new instances at startup ("container is locked but no instances are available to run it").

Successfully tested on c97qk.

##### #3 - 03/15/2019 08:46 PM - Tom Clegg

- File 14977.png added



**#4 - 03/18/2019 01:42 PM - Tom Clegg**

- Blocks Story #14807: [arvados-dispatch-cloud] Features/fixes needed before first production deploy added

**#5 - 03/18/2019 01:46 PM - Tom Clegg**

14977-kill-if-not-in-queue @ [6582c5afa53258fcb36d682fb690203930b7b2f6](https://github.com/Arvados/arvados-dispatch-cloud/pull/6582c5afa53258fcb36d682fb690203930b7b2f6)

(reorder statements slightly for clarity)

**#6 - 03/18/2019 02:14 PM - Ward Vandewege**

Tom Clegg wrote:

14977-kill-if-not-in-queue @ [6582c5afa53258fcb36d682fb690203930b7b2f6](https://github.com/Arvados/arvados-dispatch-cloud/pull/6582c5afa53258fcb36d682fb690203930b7b2f6)

(reorder statements slightly for clarity)

LGTM, thanks!

**#7 - 03/18/2019 02:15 PM - Tom Clegg**

More explanation:

If a container is already cancelled when the dispatcher starts up, it never gets added/updated in the dispatcher's queue. Therefore, the scheduler never finds out it has state=Cancelled. So the fix is to kill any container that is not in the queue -- whether because Cancelled, because deleted, or whatever.

**#8 - 03/18/2019 03:53 PM - Tom Clegg**

- Status changed from In Progress to Resolved

- % Done changed from 0 to 100

Applied in changeset [arvados|1c6c860b90c54f1c257ef2016f725176d9cf50c6](https://github.com/Arvados/arvados/pull/1c6c860b90c54f1c257ef2016f725176d9cf50c6).

**#9 - 05/21/2019 10:27 PM - Tom Morris**

- Release set to 15

**Files**

14977.png	17.7 KB	03/15/2019	Tom Clegg
-----------	---------	------------	-----------