

## Arvados - Bug #14920

### [crunch-dispatch-cloud] New Azure instances always have state=unknown instead of state=booting

03/05/2019 06:37 PM - Tom Clegg

<b>Status:</b>	Resolved	<b>Start date:</b>	03/07/2019
<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>	2019-03-13 Sprint		
<b>Description</b>			
<p>Currently, when a-d-c uses the Azure driver, new instances have state=unknown (instead of the expected state=booting) until the boot/run probes pass.</p> <p>The "unknown" state is intended to cover the case where the "list instances" call returns a previously unseen instance ID. In the Azure case, the "create VM" call does not even return the ID of the newly created instance until the instance has finished booting, so until then, the dispatcher's worker pool doesn't recognize that it corresponds to an outstanding "create" call.</p> <p>Some different ways to address this:</p> <ul style="list-style-type: none"><li>• In the Azure driver, return as soon as the instance ID is known, instead of waiting for it to boot. This is how the driver is expected to work, but the Azure client library might not make it easy.</li><li>• In the worker pool, when an unexpected instance ID appears, check whether its "secret token" tag matches an outstanding Create call. This would also cover the "list returns before create" race, which applies to all drivers.</li></ul> <p>The second option seems better.</p> <p>It would also be worth documenting the expected driver behavior in the driver interface definition: Create() should generally return as soon as the new instance's ID is known, but must not return so early that a subsequent call to Instances() might not include the new instance.</p>			
<b>Subtasks:</b>			
Task # 14928: review 14920-unknown-booting-race			<b>Resolved</b>

#### Associated revisions

##### Revision 64e72e28 - 03/07/2019 07:35 PM - Tom Clegg

Merge branch '14920-unknown-booting-race'

fixes #14920

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

#### History

##### #1 - 03/05/2019 07:19 PM - Tom Clegg

- Category set to Crunch

- Target version set to To Be Groomed

##### #2 - 03/06/2019 02:23 PM - Tom Clegg

- Description updated

##### #3 - 03/07/2019 01:58 AM - Tom Clegg

- Status changed from New to In Progress

- Assigned To set to Tom Clegg

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

14920-unknown-booting-race @ [e49978c5d9bece2a1db646f36cdf346414dd8813](https://github.com/Arvados/arvados/commit/e49978c5d9bece2a1db646f36cdf346414dd8813)

##### #4 - 03/07/2019 05:44 PM - Ward Vandewege

Tom Clegg wrote:

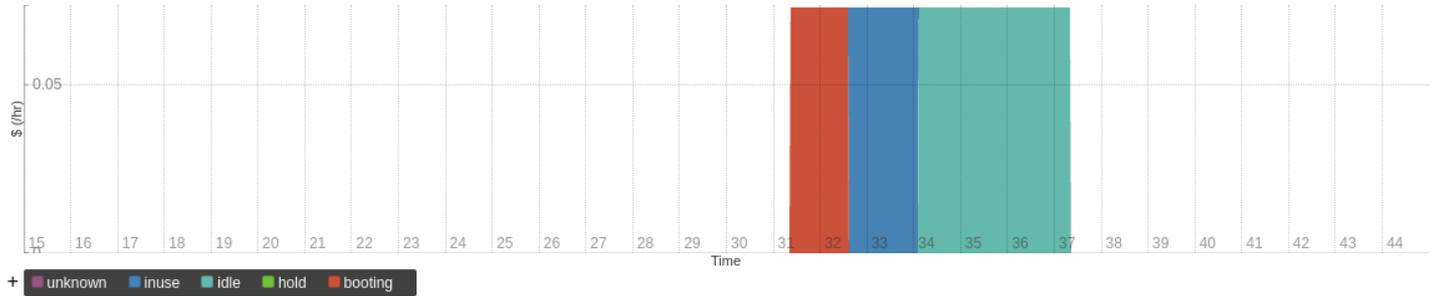
LGTM. I like that the code is a lot more elegant now!

**#5 - 03/07/2019 07:41 PM - Tom Clegg**

- File *14920-fixed.png* added

I noticed while testing that metrics didn't reflect the idle→running change right away when starting a container. With that fixed:

Cost by node state



**#6 - 03/07/2019 08:55 PM - Tom Clegg**

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

Applied in changeset [arvados|64e72e2842da35ef1616a6a499731a6ed0a832b5](#).

**#7 - 03/13/2019 06:54 AM - Tom Morris**

- Target version changed from *Arvados Future Sprints* to *2019-03-13 Sprint*

**#8 - 05/21/2019 10:27 PM - Tom Morris**

- Release set to 15

**Files**

---

14920-fixed.png	20 KB	03/07/2019	Tom Clegg
-----------------	-------	------------	-----------