

Arvados - Bug #15942

[a-d-c] Doesn't fail when resources can't be satisfied?

12/18/2019 08:13 PM - Peter Amstutz

Status:	Resolved	Start date:	12/19/2019
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2020-01-02 Sprint		
Description			
https://workbench.ce8i5.arvadosapi.com/container_requests/ce8i5-xvhdp-zi71xvgt3cbjismq#Status			
This container request asks for 2 vCPUs.			
The cluster config only has one node type, which has 1 vCPU.			
arvados-dispatch-cloud probably should have rejected it, but it scheduled it instead. Unfortunately then we asked Docker for 2 CPUs and it didn't like that:			
2019-12-18T19:53:51.826922218Z Creating Docker container			
2019-12-18T19:53:51.832452845Z error in Run: While creating container: Error response from daemon: Range of CPUs is from 0.01 to 1.00, as there are only 1 CPUs available			
Subtasks:			
Task # 15943: Review 15942-cancel-unsatisfiable			Resolved
Related issues:			
Related to Arvados - Bug #16036: lib/dispatchcloud/container IntegrationSuite...			Duplicate

Associated revisions

Revision 1080a6b8 - 12/21/2019 10:12 PM - Tom Clegg

Merge branch '15942-cancel-unsatisfiable'

refs #15942

Arvados-DCO-1.1-Signed-off-by: Tom Clegg <tom@tomclegg.ca>

History

#1 - 12/18/2019 08:13 PM - Peter Amstutz

- Status changed from New to In Progress

#2 - 12/18/2019 08:17 PM - Peter Amstutz

- Description updated

#3 - 12/19/2019 03:17 AM - Tom Clegg

```
ii arvados-controller 1.4.2.20191212205527-1 amd64 Arvados cluster controller daemon
ii arvados-dispatch-cloud 1.4.2.20191205154502-1 amd64 Arvados cluster cloud dispatch
```

a-d-c logs show the expected error ("cancel container with no suitable instance type") followed by attempts to run the container anyway.

```
{"ContainerUUID": "ce8i5-dz642-reervfgxy9hla0f", "PID": 119132, "error": "constraints not satisfiable by any configured instance type", "level": "warning", "msg": "cancel container with no suitable instance type", "time": "2019-12-18T19:53:30.666716050Z"}
{"ContainerUUID": "ce8i5-dz642-reervfgxy9hla0f", "InstanceType": "alv2", "PID": 119132, "Priority": 562948376722770419, "State": "Locked", "level": "info", "msg": "adding container to queue", "time": "2019-12-18T19:53:30.749278632Z"}
{"Address": "10.47.0.8", "ContainerUUID": "ce8i5-dz642-reervfgxy9hla0f", "Instance": "/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Compute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho", "InstanceType": "alv2", "PID": 119132, "level": "info", "msg": "crunch-run process started", "time": "2019-12-18T19:53:30.782655808Z"}
```

```
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process ended","time":"2019-12-18T19:53:40.605524849Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process started","time":"2019-12-18T19:53:50.778935999Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process ended","time":"2019-12-18T19:54:00.602361070Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process started","time":"2019-12-18T19:54:10.721237795Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process ended","time":"2019-12-18T19:54:20.601629762Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process started","time":"2019-12-18T19:54:30.709696522Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process ended","time":"2019-12-18T19:54:40.605795842Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process started","time":"2019-12-18T19:54:50.734759199Z"}
{"Address":"10.47.0.8","ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"Instance":"/subscriptions/3fa048dc-aa38-4820-85ba-68498da5f26b/resourceGroups/ce8i5/providers/Microsoft.Com
pute/virtualMachines/compute-99cb150b26149780de44b929577e1aed-3w5kue90h391hho","InstanceType":"alv2","PID":119
132,"level":"info","msg":"crunch-run process ended","time":"2019-12-18T19:55:00.594075637Z"}
{"ContainerUUID":"ce8i5-dz642-reervfqxy9hla0f",
,"PID":119132,"State":"Cancelled","level":"info","msg":"dropping container from queue","time":"2019-12-18T19:
55:10.634223282Z"}
```

controller shows multiple cycles of lock (by dispatcher) + unlock (by worker node):

```
{"PID":41701,"RequestID":"req-k00k7pe4rkuhtv3y6o65","level":"info","msg":"request","remoteAddr":"127.0.0.1:455
74","reqBytes":0,"reqForwardedFor":"52.147.168.174","reqHost":"ce8i5.arvadosapi.com","reqMethod":"GET","reqPat
h":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f","reqQuery":"include_trash=true\u0026select=%5B%22uuid%2
2%5D","time":"2019-12-18T19:53:21.922992435Z"}
{"PID":41701,"RequestID":"req-k00k7pe4rkuhtv3y6o65","level":"info","msg":"response","remoteAddr":"127.0.0.1:45
574","reqBytes":0,"reqForwardedFor":"52.147.168.174","reqHost":"ce8i5.arvadosapi.com","reqMethod":"GET","reqPa
th":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f","reqQuery":"include_trash=true\u0026select=%5B%22uuid%
22%5D","respBytes":1564,"respStatus":"OK","respStatusCode":200,"time":"2019-12-18T19:53:21.934605230Z","timeTo
Status":0.011520,"timeTotal":0.011590,"timeWriteBody":0.000070}
{"PID":41701,"RequestID":"req-53z0j6pu0x7dvnuod779","level":"info","msg":"request","remoteAddr":"127.0.0.1:456
64","reqBytes":0,"reqForwardedFor":"52.147.168.174","reqHost":"ce8i5.arvadosapi.com","reqMethod":"POST","reqPa
th":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f/lock","reqQuery":"","time":"2019-12-18T19:53:30.6674096
56Z"}
{"PID":41701,"RequestID":"req-53z0j6pu0x7dvnuod779","level":"info","msg":"response","remoteAddr":"127.0.0.1:45
664","reqBytes":0,"reqForwardedFor":"52.147.168.174","reqHost":"ce8i5.arvadosapi.com","reqMethod":"POST","reqP
ath":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f/lock","reqQuery":"","respBytes":174,"respStatus":"OK",
"respStatusCode":200,"time":"2019-12-18T19:53:30.747416617Z","timeToStatus":0.079848,"timeTotal":0.079990,"tim
eWriteBody":0.000142}
{"PID":41701,"RequestID":"req-12e4zygc41dlwik6t9j2","level":"info","msg":"request","remoteAddr":"127.0.0.1:456
90","reqBytes":42,"reqForwardedFor":"52.177.181.122","reqHost":"ce8i5.arvadosapi.com","reqMethod":"PUT","reqPa
th":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f","reqQuery":"","time":"2019-12-18T19:53:31.858865397Z"}
{"PID":41701,"RequestID":"req-12e4zygc41dlwik6t9j2","level":"info","msg":"response","remoteAddr":"127.0.0.1:45
690","reqBytes":42,"reqForwardedFor":"52.177.181.122","reqHost":"ce8i5.arvadosapi.com","reqMethod":"PUT","reqP
ath":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f","reqQuery":"","respBytes":1564,"respStatus":"OK","res
pStatusCode":200,"time":"2019-12-18T19:53:31.889957654Z","timeToStatus":0.031041,"timeTotal":0.031081,"timeWri
teBody":0.000040}
{"PID":41701,"RequestID":"req-1p3oilkv0upa8l1f45oxq","level":"info","msg":"request","remoteAddr":"127.0.0.1:457
96","reqBytes":0,"reqForwardedFor":"52.147.168.174","reqHost":"ce8i5.arvadosapi.com","reqMethod":"POST","reqPa
th":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f/lock","reqQuery":"","time":"2019-12-18T19:53:50.6462894
03Z"}
```

```
{"PID":41701,"RequestID":"req-1p3oilkv0upa81f45oxq","level":"info","msg":"response","remoteAddr":"127.0.0.1:45796","reqBytes":0,"reqForwardedFor":"52.147.168.174","reqHost":"ce8i5.arvadosapi.com","reqMethod":"POST","reqPath":"/arvados/v1/containers/ce8i5-dz642-reervfqxy9hla0f/lock","reqQuery":"","respBytes":174,"respStatus":"OK","respStatusCode":200,"time":"2019-12-18T19:53:50.741234388Z","timeToStatus":0.094782,"timeTotal":0.094930,"timeWriteBody":0.000148}
```

#4 - 12/19/2019 03:51 AM - Tom Clegg

(*lib/dispatcher/container.Queue)addEnt() detects that the container can't run and (instead of actually adding a queue entry) starts a goroutine to deal with it. That goroutine calls Lock(), which calls /arvados/v1/containers/\$uuid/lock and updates the queue with the response using the generic "addEnt() or update existing entry" code. In this case there's no existing entry, so it calls addEnt() again.

This time, RuntimeConstraints is empty, because the lock and unlock APIs have an implicit "select" param that doesn't include runtime_constraints -- so the container gets dispatched as if it had no runtime constraints, i.e., on the lowest priced node type.

I think the easiest solution is for updateWithResp (called by Cancel and apiUpdate) to just discard the response, instead of adding an entry, if the relevant container is no longer in the queue. I don't see any cases where it makes sense to add a new entry when a just-locked/unlocked/cancelled container has disappeared from the queue and needs to be re-added.

#5 - 12/19/2019 03:52 AM - Tom Clegg

- Assigned To set to Tom Clegg

#6 - 12/19/2019 04:51 AM - Tom Clegg

15942-cancel-unsatisfiable @ [6162a7fc5b213446b808313770d0d7e491fa5723](https://ci.arvados.org/view/Developer/job/developer-run-tests/1696/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/1696/>

#7 - 12/19/2019 02:26 PM - Tom Clegg

15942-cancel-unsatisfiable @ [ef48a49d302ca25136bbb2ed96319a8b56b46250](https://ci.arvados.org/view/Developer/job/developer-run-tests/1696/) with a test that confirms the fix.

#8 - 12/19/2019 09:55 PM - Lucas Di Pentima

Sorry for the delay, I was getting more familiar with the code. Although your explanation makes perfect sense and the fix LGTM, I tried running the test with the pre-fix queue.go code and it succeeded. Maybe there's an issue with the test itself?

#9 - 12/21/2019 10:12 PM - Tom Clegg

Indeed, if I un-fix the bug, the test only fails for me about 80% of the time. I've updated the test to listen to the queue's update channel instead, and (although I still don't think it's airtight) this caught the reintroduced bug 15/15 times.

15942-cancel-unsatisfiable @ [d01a136c64bb5e342542b7de4da45e09e0961dc1](https://ci.arvados.org/view/Developer/job/developer-run-tests/1696/)

#10 - 12/23/2019 01:41 PM - Lucas Di Pentima

Test works great on my side too, please merge. Thanks!

#11 - 12/23/2019 02:06 PM - Tom Clegg

- Status changed from In Progress to Resolved

#12 - 12/30/2019 09:24 PM - Peter Amstutz

- Target version changed from 2020-01-15 Sprint to 2020-01-02 Sprint

#13 - 01/16/2020 07:41 PM - Tom Clegg

- Related to Bug #16036: lib/dispatchcloud/container IntegrationSuite.TestGetLockUnlockCancel is flaky added

#14 - 01/22/2020 02:56 PM - Peter Amstutz

- Release set to 22