

Arvados - Bug #10182

Provide more reasonable error messages for memory issues during container dispatch

10/04/2016 01:54 PM - Tom Morris

Status:	New	Start date:	10/04/2016
Priority:	Normal	Due date:	
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	Arvados Future Sprints		
Description			
A customer received these two errors (logged in RT #136):			
<pre>Error response from daemon: Cannot start container f1082c7aeec7be8b8ad1e45b5cf1273457f262e71a7909e7485a3c280c8c4dd4: [8] System error: open /sys/fs/cgroup/memory/system.slice/docker-f1082c7aeec7be8b8ad1e45b5cf1273457f262e71a7909e7485a3c280c8c4dd4.scope/memory.memsw.limit_in_bytes: no such file or directory</pre>			
<pre>Error response from daemon: Cannot start container 6a0f0f39045bad7032893f3f926f4b1a2caee3d7caa4f061e46d41f52a763965: [8] System error: write /sys/fs/cgroup/memory/system.slice/docker-6a0f0f39045bad7032893f3f926f4b1a2caee3d7caa4f061e46d41f52a763965.scope/memory.memsw.limit_in_bytes: invalid argument</pre>			
Which Ward interpreted as "is missing runtime_contraints, the job ran out of memory" and "has a runtime_constraint for ram but it seems to be set too low, the job ran out of memory."			
We should make the logged error messages be more like Ward's version than the current cryptic versions.			

History

#1 - 11/24/2016 10:17 AM - Joshua Randall

I ran into pretty much exactly these two error messages (from the same job) after upgrading systemd to the latest version (v230 from jessie-backports in my case), which appears to have issues with docker. The underlying problem seems to be that the system.slice directory is no longer present in that version.

The workaround was to switch docker to not use systemd for managing cgroups:
<https://github.com/docker/docker/issues/17653#issuecomment-155609224>

If the fix for this issue obscured the error messages that come from docker, I'd never have figured out the real problem, so whatever the fix is here should probably make sure the errors are (also) logged.

#2 - 11/26/2016 02:18 PM - Ward Vandewege

Joshua Randall wrote:

I ran into pretty much exactly these two error messages (from the same job) after upgrading systemd to the latest version (v230 from jessie-backports in my case), which appears to have issues with docker. The underlying problem seems to be that the system.slice directory is no longer present in that version.

The workaround was to switch docker to not use systemd for managing cgroups:
<https://github.com/docker/docker/issues/17653#issuecomment-155609224>

If the fix for this issue obscured the error messages that come from docker, I'd never have figured out the real problem, so whatever the fix is here should probably make sure the errors are (also) logged.

I agree - note that this ticket was a bit out of date - I also figured out a couple weeks ago that the other failure that can lead to this error is the cgroup thing you identified. Interpreting errors could be useful to give users a hint of what may be going on, but we shouldn't obscure the underlying errors.

#3 - 08/29/2017 01:48 PM - Tom Morris

- Target version set to Arvados Future Sprints