

Arvados - Bug #8828

[Crunch] be more resilient when crunchrunner is not available; also don't test for crunchrunner on api server

03/30/2016 01:15 AM - Ward Vandewege

Status:	Resolved	Start date:	03/31/2016
Priority:	Normal	Due date:	
Assigned To:	Peter Amstutz	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2016-04-13 sprint		
Description			
<p>I noticed two small issues introduced after the changes in 8815 - I should have spotted them in review, sorry. Specifically, we had the diagnostics fail on c97qk with:</p>			
<pre>2016-03-30_01:08:05 c97qk-8i9sb-etuzc7s3j88iym9 1470 0 stderr Running [docker.io run --name=c97qk-ot0gb-n67uuhx6ng9hzb-0 --attach=stdout --attach h=stderr --attach=stdin -i --cidfile=/tmp/crunch-job/c97qk-ot0gb-n67uuhx6ng9hzb-0.cid --sig-proxy --memory=3346971k --memory-swap=3346971k --volume=/tmp/crunch-job/src:/tmp/crunch-job/src:ro --vo lume=/tmp/crunch-job/opt:/tmp/crunch-job/opt:ro --volume=/tmp/crunch-job/task/compute3.1.keep/by_p dh:/keep:ro --volume=/tmp/crunch-job/task/compute3.1.keep/tmp:/keep_tmp --volume=/tmp --volume=:/u sr/local/bin/crunchrunner --volume=/etc/ssl/certs/ca-certificates.crt:/etc/arvados/ca-certificates .crt --env=TASK_KEEPMOUNT_TMP=/keep_tmp --env=CRUNCH_GIT_ARCHIVE_HASH=8e18a89ea517fde50d24eb17b884 bc86 --env=CRUNCH_SRC=/tmp/crunch-job/src --env=JOB_UUID=c97qk-8i9sb-etuzc7s3j88iym9 --env=TASK_QS EQUENCE=0 --env=CRUNCH_REFRESH_TRIGGER=/tmp/crunch_refresh_trigger --env=ARVADOS_API_HOST=c97qk.ar vadosapi.com --env=TASK_TMPDIR=/tmp/crunch-job-task-work/compute3.1 --env=JOB_WORK=/tmp/crunch-j -work --env=CRUNCH_TMP=/tmp/crunch-job --env=TASK_SLOT_NODE=compute3 --env=CRUNCH_SRC_URL=/var/lib /arvados/internal.git --env=JOB_SCRIPT=hash --env=CRUNCH_WORK=/tmp/crunch-job/work --env=CRUNCH_NO DE_SLOTS=1 --env=TASK_SEQUENCE=0 --env=TASK_WORK=/tmp/crunch-job-task-work/compute3.1 --env=JOB_PA RAMETER_INPUT=1724fc6b2145c148b894a8da81132ef8+53 --env=ARVADOS_API_TOKEN=42qrvz14riharlxo9qqdigha lbu1022iuoyrlj859nbfx8bfyk --env=CRUNCH_JOB_BIN=/usr/local/arvados/src/services/crunch/crunch-j --env=TASK_UUID=c97qk-ot0gb-n67uuhx6ng9hzb --env=TASK_SLOT_NUMBER=1 --env=TASK_KEEPMOUNT=/keep -- env=CRUNCH_JOB_UUID=c97qk-8i9sb-etuzc7s3j88iym9 --env=CRUNCH_SRC_COMMIT=4d4c3442e04310d7a88894c105 a7cf351fd9f373 --env=CRUNCH_INSTALL=/tmp/crunch-job/opt --env=HOME=/tmp/crunch-job-task-work/compu te3.1 f30fae7189adac0948eef3b3386e9ef254f69a8187f9eab99004e2d3650605cd /bin/sh -c python -c "from pkg_resources import get_distribution as get; print \"Using Arvados SDK version\", get(\"arvados-p ython-client\").version">&2 2>/dev/null; mkdir -p "/tmp/crunch-job-work" "/tmp/crunch-job-task-wor k/compute3.1" && if which stdbuf >/dev/null ; then exec stdbuf --output=0 --error=0 \tmp/cru nch-job/src/crunch_scripts/hash ; else exec \tmp/crunch-job/src/crunch_scripts/hash ; fi]</pre>			
<pre>2016-03-30_01:08:05 c97qk-8i9sb-etuzc7s3j88iym9 1470 0 stderr invalid value ":/usr/local/bin/crunchrunner" for flag --volume: bad format for volu mes: /usr/local/bin/crunchrunner</pre>			
<pre>2016-03-30_01:08:05 c97qk-8i9sb-etuzc7s3j88iym9 1470 0 stderr See 'docker.io run --help'.</pre>			
<p>Two issues here:</p>			
<p>a) the 'which crunchrunner' lookup is apparently happening on the API server, not on the compute node: the compute node had crunchrunner installed, but the API server did not. I installed it there, and that fixed the problem. Clearly, the test should happen on the compute node.</p>			
<p>b) when 'which crunchrunner' doesn't find the executable, we shouldn't try to append half a volume statement to the docker run command, that breaks the invocation (and thus fails the job).</p>			
Subtasks:			
Task # 8861: Review 8828-which-crunchrunner			Resolved
Related issues:			
Related to Arvados - Story #8815: [Crunch] crunch-job bind mounts crunchrunner...		Resolved	03/29/2016

Associated revisions

Revision 49743c08 - 03/31/2016 10:29 PM - Peter Amstutz

Merge branch '8828-which-crunchrunner' closes #8828

History

#1 - 03/30/2016 01:18 AM - Ward Vandewege

- Description updated

#2 - 03/31/2016 08:42 PM - Peter Amstutz

Pushed branch 8828-which-crunchrunner

#3 - 03/31/2016 09:01 PM - Brett Smith

- Status changed from New to In Progress

- Assigned To set to Peter Amstutz

- Target version set to 2016-04-13 sprint

#4 - 03/31/2016 09:14 PM - Brett Smith

Both the VOLUME_CERTS declarations need to specify that the certs are being mounted at /etc/arvados/ca-certificates.crt. Otherwise this will break the behavior specified in [#8815](#).

With that fix, this is good to merge, thanks.

(I wonder when we're going to hit limit on the maximum size of a single command line...)

#5 - 03/31/2016 10:30 PM - Peter Amstutz

Brett Smith wrote:

Both the VOLUME_CERTS declarations need to specify that the certs are being mounted at /etc/arvados/ca-certificates.crt. Otherwise this will break the behavior specified in [#8815](#).

With that fix, this is good to merge, thanks.

Whoops, thanks for catching that, that's what I get for rushing. Fixed, tested with arvbox, merged, pushd.

(I wonder when we're going to hit limit on the maximum size of a single command line...)

#6 - 03/31/2016 10:35 PM - Peter Amstutz

- Status changed from In Progress to Resolved

- % Done changed from 0 to 100

Applied in changeset arvados|commit:49743c080265b270693154d7a327d0433b0a7dbe.