

Arvados - Bug #16834

[a-d-c] containers going from running -> non-running state repeatedly

09/11/2020 05:42 PM - Ward Vandewege

Status:	Resolved	Start date:	09/14/2020
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	Crunch	Estimated time:	0.00 hour
Target version:	2020-09-23 Sprint		

Description

This is **arvados-dispatch-cloud 2.0.3-1**.

I noticed something odd on our *lugli* grafana graphs:

20200911-container-yoyo.png

There are 2 containers that are going from running to non-running state repeatedly:

```
Sep 11 17:22:24 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.95", "ContainerUUID":"lugli-dz642-lrg004e66zorxld", "Instance":"i-0fb85bb32ca92f0c5", "InstanceType":"c5large", "PID":582, "level":"info", "msg":"crunch-run process detected", "time":"2020-09-11T17:22:24.370207112Z"}
Sep 11 17:22:24 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.74", "ContainerUUID":"lugli-dz642-n2gnjy6o3prkcam", "Instance":"i-02cb2a6d6f6e33895", "InstanceType":"c5large", "PID":582, "level":"info", "msg":"crunch-run process detected", "time":"2020-09-11T17:22:24.370286849Z"}
Sep 11 17:22:24 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.95", "ContainerUUID":"lugli-dz642-lrg004e66zorxld", "Instance":"i-0fb85bb32ca92f0c5", "InstanceType":"c5large", "PID":582, "Reason":"not in queue", "level":"info", "msg":"killing crunch-run process", "time":"2020-09-11T17:22:24.370368549Z"}
Sep 11 17:22:24 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.74", "ContainerUUID":"lugli-dz642-n2gnjy6o3prkcam", "Instance":"i-02cb2a6d6f6e33895", "InstanceType":"c5large", "PID":582, "Reason":"not in queue", "level":"info", "msg":"killing crunch-run process", "time":"2020-09-11T17:22:24.370395251Z"}
Sep 11 17:22:29 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.95", "ContainerUUID":"lugli-dz642-lrg004e66zorxld", "Instance":"i-0fb85bb32ca92f0c5", "InstanceType":"c5large", "PID":582, "Signal":15, "level":"info", "msg":"sending signal", "time":"2020-09-11T17:22:29.370501187Z"}
Sep 11 17:22:29 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.74", "ContainerUUID":"lugli-dz642-n2gnjy6o3prkcam", "Instance":"i-02cb2a6d6f6e33895", "InstanceType":"c5large", "PID":582, "Signal":15, "level":"info", "msg":"sending signal", "time":"2020-09-11T17:22:29.370587727Z"}
Sep 11 17:22:29 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.95", "ContainerUUID":"lugli-dz642-lrg004e66zorxld", "Instance":"i-0fb85bb32ca92f0c5", "InstanceType":"c5large", "PID":582, "level":"info", "msg":"crunch-run process ended", "time":"2020-09-11T17:22:29.387498437Z"}
Sep 11 17:22:29 api.lugli.arvadosapi.com arvados-dispatch-cloud[582]: {"Address":"10.255.254.74", "ContainerUUID":"lugli-dz642-n2gnjy6o3prkcam", "Instance":"i-02cb2a6d6f6e33895", "InstanceType":"c5large", "PID":582, "level":"info", "msg":"crunch-run process ended", "time":"2020-09-11T17:22:29.387750208Z"}
```

They are in *cancelled* state according to Workbench.

There are indeed two AWS VMs online, but docker is running nothing:

```
root@ip-10-255-254-74:/home/admin# docker ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS
PORTS	NAMES			

and

```
root@ip-10-255-254-95:/home/admin# docker ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS
--------------	-------	---------	---------	--------

PORTS NAMES

A full ps on the latter shows that arv-mount is still there, but the rest of our 'onion' is not:

```
root@ip-10-255-254-95:/home/admin# ps auxwwwf
USER      PID  %CPU  %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root         2  0.0  0.0      0     0 ?        S      Sep10   0:00 [kthreadd]
root         3  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [rcu_gp]
root         4  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [rcu_par_gp]
root         6  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kworker/0:0H-kblockd]
root         8  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [mm_percpu_wq]
root         9  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [ksoftirqd/0]
root        10  0.0  0.0      0     0 ?        I      Sep10   0:00 \_ [rcu_sched]
root        11  0.0  0.0      0     0 ?        I      Sep10   0:00 \_ [rcu_bh]
root        12  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [migration/0]
root        14  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [cpuhp/0]
root        15  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [cpuhp/1]
root        16  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [migration/1]
root        17  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [ksoftirqd/1]
root        19  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kworker/1:0H-kblockd]
root        20  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [kdevtmpfs]
root        21  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [netns]
root        22  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [kauditd]
root        23  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [khungtaskd]
root        24  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [oom_reaper]
root        25  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [writeback]
root        26  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [kcompactd0]
root        27  0.0  0.0      0     0 ?        SN     Sep10   0:00 \_ [ksmd]
root        28  0.0  0.0      0     0 ?        SN     Sep10   0:00 \_ [khugepaged]
root        29  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [crypto]
root        30  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kintegrityd]
root        31  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kblockd]
root        32  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [watchdogd]
root        36  0.0  0.0      0     0 ?        S      Sep10   0:38 \_ [kswapd0]
root        52  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kthrotld]
root        53  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [ipv6_addrconf]
root        63  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kstrp]
root       110  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [nvme-wq]
root       111  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [nvme-reset-wq]
root       113  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [ena]
root       114  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [nvme-delete-wq]
root       138  0.0  0.0      0     0 ?        I      Sep10   0:00 \_ [kworker/u4:4-events_unbound]
root       183  0.0  0.0      0     0 ?        I<     Sep10   0:01 \_ [kworker/u5:0-kcryptd]
root       185  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [jbd2/nvme1n1p1-]
root       186  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [ext4-rsv-conver]
root       230  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kworker/1:1H-kblockd]
root       351  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kworker/0:1H-kblockd]
root        804  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kdmflush]
root       1155  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kdmflush]
root       1159  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kcryptd_io]
root       1160  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [kcryptd]
root       1161  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [dmccrypt_write]
root       1164  0.0  0.0      0     0 ?        I<     Sep10   0:01 \_ [kworker/u5:2-kcryptd]
root       1196  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfsalloc]
root       1199  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs_mru_cache]
root       1202  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-buf/dm-1]
root       1203  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-data/dm-1]
root       1204  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-conv/dm-1]
root       1205  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-cil/dm-1]
root       1206  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-reclaim/dm-]
root       1207  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-log/dm-1]
root       1208  0.0  0.0      0     0 ?        I<     Sep10   0:00 \_ [xfs-eofblocks/d]
root       1209  0.0  0.0      0     0 ?        S      Sep10   0:00 \_ [xfsaild/dm-1]
root       3436  0.0  0.0      0     0 ?        I      Sep10   0:00 \_ [kworker/0:0-events]
root      11630  0.0  0.0      0     0 ?        I      15:54   0:00 \_ [kworker/0:1-cgroup_destroy]
root      16205  0.0  0.0      0     0 ?        I      17:17   0:00 \_ [kworker/1:2-xfs-cil/dm-1]
root      17289  0.0  0.0      0     0 ?        I      17:28   0:00 \_ [kworker/u4:1-events_unbound]
```

```

root      17890  0.0  0.0      0      0 ?          I    17:29   0:00  \_ [kworker/1:0-events_power_efficient]
root      1      0.0  0.1 104040  7100 ?          Ss   Sep10   0:02  /sbin/init
root      254      0.0  0.2  47656 10752 ?          Ss   Sep10   0:13  /lib/systemd/systemd-journald
root      264      0.0  0.0  20120  972 ?          Ss   Sep10   0:00  /lib/systemd/systemd-udev
root      384      0.0  0.0   9488  3056 ?          Ss   Sep10   0:00  /sbin/dhclient -4 -v -i -pf /run/dhclient.ens5.pid -lf /var/lib/dhcp/dhclient.ens5.leases -I -df /var/lib/dhcp/dhclient6.ens5.leases -nw ens5
root      492      0.0  0.0   9488  2100 ?          Ss   Sep10   0:00  /sbin/dhclient -6 -v -pf /run/dhclient6.ens5.pid -lf /var/lib/dhcp/dhclient6.ens5.leases -I -df /var/lib/dhcp/dhclient.ens5.leases -nw ens5
message+  584      0.0  0.0   8800  556 ?          Ss   Sep10   0:00  /usr/bin/dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd-activation --syslog-only
root      586      0.0  0.0 225824  1312 ?          Ssl  Sep10   0:01  /usr/sbin/rsyslogd -n -iNONE
root      588      0.0  0.0  19520  960 ?          Ss   Sep10   0:00  /lib/systemd/systemd-logind
root      591      0.0  0.0  10228  1200 ?          Ss   Sep10   0:00  /usr/sbin/cron -f
_chnony   598      0.0  0.0   4816  368 ?          S    Sep10   0:00  /usr/sbin/chronyd -F -1
_chnony   599      0.0  0.0   4816  268 ?          S    Sep10   0:00  \_ /usr/sbin/chronyd -F -1
root      610      0.0  0.0   5612  112 tty1       Ss+  Sep10   0:00  /sbin/agetty -o -p -- \u --noclear tty1 linux
root      611      0.0  0.0   7136  128 ttyS0     Ss+  Sep10   0:00  /sbin/agetty -o -p -- \u --keep-baud 115200,38400,9600 ttyS0 vt220
root      614      0.0  0.0  15852  2036 ?          Ss   Sep10   0:00  /usr/sbin/sshd -D
root      1043     0.0  0.0  16612  1088 ?          Ss   Sep10   0:00  \_ sshd: admin [priv]
admin     1062     0.0  0.0  16848  3000 ?          S    Sep10   0:06  | \_ sshd: admin@notty
root      17232    0.0  0.2  16616  7872 ?          Ss   17:27   0:00  \_ sshd: admin [priv]
admin     17238    0.0  0.1  16616  4692 ?          S    17:27   0:00  \_ sshd: admin@pts/0
admin     17239    0.0  0.1   9644  4672 pts/0     Ss   17:27   0:00  \_ -bash
root      17252    0.0  0.0  11660  3720 pts/0     S    17:27   0:00  \_ sudo -s
root      17253    0.0  0.0   8716  3668 pts/0     S    17:27   0:00  \_ /bin/bash
root      19266    0.0  0.0  12512  3164 pts/0     R+   17:41   0:00  \_ ps auxwww
wf
admin     1046     0.0  0.1  21264  4464 ?          Ss   Sep10   0:00  /lib/systemd/systemd --user
admin     1047     0.0  0.0 104880  2312 ?          S    Sep10   0:00  \_ (sd-pam)
root      1237     0.0  1.4 1074276 54512 ?          Ssl  Sep10   1:26  /usr/sbin/dockerd -H fd:// --default-ulimit nofile=10000:10000 --dns 10.255.0.2
root      1244     0.1  0.5  952160 20456 ?          Ssl  Sep10   2:18  \_ docker-containerd --config /var/run/docker/containerd/containerd.toml --log-level info
root      1477     0.0  0.4  33716 16772 ?          Ss   Sep10   0:00  /usr/share/python3/dist/arvados-docker-cleaner/bin/python /usr/bin/arvados-docker-cleaner
root      4570     0.0 10.4 1608588 396860 ?          Sl   Sep10   0:04  /usr/share/python2.7/dist/python-arvados-fuse/bin/python /usr/bin/arv-mount --foreground --allow-other --read-write --crunchstat-interval=10 --file-cache 268435456 --mount-by-pdh by_id /tmp/crunch-run.lugli-dz642-lrg004e66zorxld.326060118/keep689459133

```

And /var/lock:

```

root@ip-10-255-254-95:/home/admin# ls -laF /var/lock/
total 4
drwxrwxrwt  4 root root 120 Sep 10 15:47 ./
drwxr-xr-x 25 root root 820 Sep 11 17:30 ../
-rwx----- 1 root root  0 Sep 10 15:39 crunch-run-all.lock*
-rwx----- 1 root root  50 Sep 10 15:47 crunch-run-lugli-dz642-lrg004e66zorxld.lock*
drwx----- 2 root root  40 Sep 10 15:39 lvm/
drwxr-xr-x  2 root root  40 Sep 10 15:39 subsys/

root@ip-10-255-254-95:/home/admin# cat /var/lock/crunch-run-lugli-dz642-lrg004e66zorxld.lock
{"UUID": "lugli-dz642-lrg004e66zorxld", "PID": 4560}

```

Subtasks:

Task # 16836: Review 16834-stale-run-lock

Resolved

Associated revisions

Revision [d6598fd6](#) - 09/15/2020 02:49 PM - Tom Clegg

Merge branch '16834-stale-run-lock'

fixes #16834

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

History

#1 - 09/11/2020 05:43 PM - Ward Vandewege

- Description updated

#2 - 09/11/2020 05:52 PM - Ward Vandewege

- Description updated

- File 20200911-container-yoyo.png added

#3 - 09/11/2020 05:56 PM - Ward Vandewege

- Description updated

#4 - 09/11/2020 08:31 PM - Ward Vandewege

- Description updated

#5 - 09/11/2020 09:14 PM - Tom Clegg

This seems to be happening:

- arv-mount inherits the lock on /var/lock/crunch-run-\$uid from crunch-run
- crunch-run calls "arv-mount --unmount" during shutdown, but doesn't wait for arv-mount to exit, because sometimes arv-mount deadlocks
- arv-mount never exits
- crunch-run --list fails to flock(), so it concludes the container is still running, and reports that to a-d-c
- crunch-run --kill fails to kill the PID listed in the lockfile, so it concludes the process has exited, and reports that to a-d-c
- a-d-c cycles endlessly: "list reports crunch-run is running" → "kill crunch-run because container is cancelled" → "kill reports crunch-run has exited"

Proposed solution:

- crunch-run --list should send signal 0 to the PID listed in each lockfile
- if that fails, that means crunch-run has exited and some other process has the lock, so return that information to a-d-c (e.g., "\$uid orphan" instead of just "\$uid")
- when a-d-c sees an "orphan" tag beside a container UUID in a probe result, the instance should be listed as busy but a-d-c should **not** consider that container to be "running", and instead should log a warning and start a timer. If an instance keeps reporting an orphan for longer than a configured interval (default 30 seconds?) a-d-c should put that instance in "drain" state.

#6 - 09/13/2020 04:52 AM - Tom Clegg

16834-stale-run-lock @ [224f4793b7c147cb3ef9a261e3a60cddb557e1f1](https://ci.arvados.org/view/Developer/job/developer-run-tests/2091/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2091/>

#7 - 09/14/2020 01:23 PM - Tom Clegg

- Assigned To set to Tom Clegg

- Status changed from New to In Progress

- Category set to Crunch

#8 - 09/15/2020 02:00 PM - Ward Vandewege

Tom Clegg wrote:

16834-stale-run-lock @ [224f4793b7c147cb3ef9a261e3a60cddb557e1f1](https://ci.arvados.org/view/Developer/job/developer-run-tests/2091/) -- <https://ci.arvados.org/view/Developer/job/developer-run-tests/2091/>

lib/dispatchcloud/worker/worker.go

in the

```
for _, s := range strings.Split(string(stdout), "\n") {
```

loop, it would be nice to have some comments on the various values that stdout can have and perhaps what they mean. Or just put a comment that refers to the relevant place in `lib/crunchrun/background.go` for the possible values, I suppose.

Otherwise, LGTM, thanks.

#9 - 09/15/2020 02:53 PM - Anonymous

- % Done changed from 0 to 100

- Status changed from In Progress to Resolved

Applied in changeset [arvados|d6598fd6339e6219a7103781433356dfde546527](#).

#10 - 10/07/2020 02:12 AM - Peter Amstutz

- Release set to 25

Files

20200911-container-yoyo.png	14.7 KB	09/11/2020	Ward Vandewege
-----------------------------	---------	------------	----------------