

Arvados - Bug #3052

[Crunch] stdout from non-task programs invoked by crunch-job should propagate to job log.

06/18/2014 03:54 PM - Tom Clegg

Status:	Resolved	Start date:	10/02/2014
Priority:	Normal	Due date:	
Assigned To:	Peter Amstutz	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	2014-10-08 sprint		
Description			
See for example the log output for qr1hi-8i9sb-u6iqw0e7tgr9rc1			
Subtasks:			
Task # 4039: Review 3052-crunch-log-stdout			Resolved

Associated revisions

Revision 62790d76 - 10/02/2014 02:51 PM - Peter Amstutz

Merge branch '3052-crunch-log-stdout' closes #3052

History

#1 - 07/04/2014 03:01 PM - Tom Clegg

- Target version set to Arvados Future Sprints

#2 - 09/29/2014 03:15 PM - Ward Vandewege

- Description updated

#3 - 09/29/2014 03:15 PM - Ward Vandewege

- Subject changed from stdout from crunch tasks should propagate to job log. to [Crunch] stdout from crunch tasks should propagate to job log.

#4 - 09/29/2014 07:06 PM - Ward Vandewege

- Target version changed from Arvados Future Sprints to 2014-10-08 sprint

#5 - 09/29/2014 07:07 PM - Peter Amstutz

- Assigned To set to Peter Amstutz

#6 - 09/29/2014 07:13 PM - Peter Amstutz

- Status changed from New to In Progress

#7 - 09/29/2014 08:30 PM - Peter Amstutz

- The culprit is this code in #read_pipes of crunch-dispatch.rb:

```
# Throw away child stdout
begin
  j[:stdout].read_nonblock(2**20)
rescue Errno::EAGAIN, EOFError
end
```

- Related to this, in crunch-job, tasks currently have their stdout and stderr fused into a single stream that goes to stderr. We want to restore the separation.

#8 - 09/29/2014 08:58 PM - Tom Clegg

- Subject changed from [Crunch] stdout from crunch tasks should propagate to job log. to [Crunch] stdout from non-task programs invoked by crunch-job should propagate to job log.

Peter Amstutz wrote:

- Related to this, in crunch-job, tasks currently have their stdout and stderr fused into a single stream that goes to stderr. We want to restore the separation.

(but not here... added [#4028](#))

#9 - 10/01/2014 07:34 PM - Tom Clegg

(Recording question from IRC)

Why is this ([a4679fb](#)) better than combining the crunch-job child's stderr and stdout into a single unix pipe and leaving the existing buffer/propagation code alone?

This way, if a line arrives on stderr when stdout has printed a partial line, we get the intact stderr line followed by an intact stdout line, rather than a stdout line interrupted by a stderr line.

#10 - 10/02/2014 02:38 PM - Tim Pierce

Review at [a4679fb](#). Just one style nit:

crunch_dispatch.rb:337: can we use JSON dictionary style here? buf: {stdout: ", stderr: "} If only because it always makes me dizzy when we combine dictionary styles in the same piece of code.

Otherwise LGTM!

#11 - 10/02/2014 02:55 PM - Anonymous

- *Status changed from In Progress to Resolved*

Applied in changeset arados|commit:62790d76daf181a457d8f47741721967cd9cb7d9.