

Arvados - Bug #11070

[arvados-ws] job logs are not being updated in real time

02/07/2017 09:59 PM - Ward Vandewege

Status:	Resolved	Start date:	
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	API	Estimated time:	0.00 hour
Target version:	2017-02-15 sprint		
Description			
In workbench, job logs are updated in real time via websockets. On our test installations with arvados-ws, that is no longer working. According to the arvados-ws logs, workbench is subscribing to the updates, so the problem appears to be somewhere in arvados-ws.			

Associated revisions

Revision 08c06bc8 - 02/10/2017 10:08 PM - Tom Clegg

Merge branch '11070-ws-listener-problem'

closes #11070

History

#1 - 02/07/2017 10:00 PM - Ward Vandewege

- Description updated

#2 - 02/08/2017 05:55 AM - Tom Clegg

Lots of connections from arvados-ws to postgresql sitting in CLOSE_WAIT, which doesn't seem right:

```
9tee4:~# ls -lart --full-time /proc/2322/fd/
total 0
dr-xr-xr-x 9 root root 0 2016-12-28 19:27:58.596931467 +0000 ..
dr-x----- 2 root root 0 2017-01-06 12:22:12.549920285 +0000 .
lrwx----- 1 root root 64 2017-01-31 18:05:35.954714735 +0000 508 -> socket:[661990687]
lrwx----- 1 root root 64 2017-02-01 23:23:12.710549212 +0000 99 -> socket:[129111377]
lrwx----- 1 root root 64 2017-02-01 23:23:12.710549212 +0000 98 -> socket:[129110706]
...
lrwx----- 1 root root 64 2017-02-01 23:23:12.710549212 +0000 1 -> socket:[18678]
lr-x----- 1 root root 64 2017-02-01 23:23:12.710549212 +0000 0 -> /dev/null
lrwx----- 1 root root 64 2017-02-08 05:32:54.873000662 +0000 513 -> socket:[664409635]
9tee4:~# netstat -nape | grep arvados-ws | grep tcp6 | wc -l
509
9tee4:~# netstat -nape | grep arvados-ws | grep CLOSE_WAIT | wc -l
500
```

#3 - 02/08/2017 06:30 PM - Tom Clegg

I ran tcpdump on the active connection, and I could see the once-per-minute "ping" between arvados-ws and postgresql. But if I generate a log entry (e.g., by modifying my user record) there is no evidence of a "notify" event on the wire.

#4 - 02/08/2017 09:52 PM - Tom Clegg

I think I have it: when we hit a connection problem, we stopped our own event loop but

- didn't tell the pq library's listener to stop listening (which is why there was still 1 ping every minute on the wire), and
- didn't exit the program (so clients could still connect to the webserver goroutine, and wait for events that will never come).

#5 - 02/09/2017 09:24 AM - Tom Clegg

11070-ws-listener-problem @ [0fc6eaead0bf7691e99d19e74ec33636909001a7](https://github.com/Arvados/arvados-ws/commit/0fc6eaead0bf7691e99d19e74ec33636909001a7)

testing on 9tee4 now.

to test behavior after postgresql connection loss (as postgres superuser):

- select * from pg_stat_activity;
- select pg_terminate_backend(pid) from pg_stat_activity where pid=XXXXX;

#6 - 02/09/2017 09:25 AM - Tom Clegg

- *Category set to API*
- *Status changed from New to In Progress*
- *Assigned To set to Tom Clegg*

#7 - 02/10/2017 10:15 PM - Tom Clegg

- *Status changed from In Progress to Resolved*
- *% Done changed from 0 to 100*

Applied in changeset arvados|commit:08c06bc813d96634b362e8c5736341fb2d874f59.

#8 - 02/15/2017 05:07 PM - Tom Clegg

- *Target version set to 2017-02-15 sprint*