

## openQA Project - action #66664

circleci jobs do not retry, are aborted with "Too long with no output (exceeded 10m0s): context deadline exceeded"

2020-05-09 15:14 - okurz

<b>Status:</b>	Resolved	<b>Start date:</b>	2020-05-09
<b>Priority:</b>	High	<b>Due date:</b>	
<b>Assignee:</b>	cdywan	<b>% Done:</b>	0%
<b>Category:</b>	Concrete Bugs	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Ready		
<b>Difficulty:</b>	hard		
<b>Description</b>			
<b>Observation</b>			
See <a href="https://app.circleci.com/pipelines/github/os-autoinst/openQA/2816/workflows/c00e8812-3cc0-4d07-ab9b-f2d5092a6704/jobs/26595/steps">https://app.circleci.com/pipelines/github/os-autoinst/openQA/2816/workflows/c00e8812-3cc0-4d07-ab9b-f2d5092a6704/jobs/26595/steps</a> and <a href="https://app.circleci.com/pipelines/github/os-autoinst/openQA/2817/workflows/1553c384-dbef-48a9-9b3a-91dc2a6477d8/jobs/26605/steps">https://app.circleci.com/pipelines/github/os-autoinst/openQA/2817/workflows/1553c384-dbef-48a9-9b3a-91dc2a6477d8/jobs/26605/steps</a> . Tests get stuck within the 1st of 3 retries and are aborted with "Too long with no output (exceeded 10m0s): context deadline exceeded". Maybe processes left over in the background?			
<b>Acceptance criteria</b>			
<ul style="list-style-type: none"><li>• <b>AC1:</b> The test output within circle CI reveals the real problem, not just "context deadline exceeded"</li></ul>			
<b>Suggestions</b>			
<ul style="list-style-type: none"><li>• Research about circleci meaning of "context deadline exceeded"</li><li>• Create artificial test run that can trigger this problem</li><li>• Try with applying timeout values within individual test modules that are in sum smaller than the global time limit of 10m, e.g. with use <code>OpenQA::Test::TimeLimit '20'</code>;</li><li>• Or extend that limit. Maybe there is a circleCI configuration variable for that?</li></ul>			
<b>Related issues:</b>			
Related to openQA Project - action #71464: Show output of failed tests within...		<b>Resolved</b>	<b>2020-09-17</b>
Copied to openQA Project - action #71467: add specific timeout values to all ...		<b>Resolved</b>	

## History

#1 - 2020-07-03 20:38 - okurz

- Target version set to Ready

#2 - 2020-07-22 13:36 - okurz

- Description updated

- Status changed from New to Workable

#3 - 2020-07-30 08:17 - okurz

- Description updated

- Difficulty set to hard

#4 - 2020-08-20 12:03 - cdywan

- Try with applying timeout values within individual test modules that are in sum smaller than the global time limit of 10m, e.g. with use `OpenQA::Test::TimeLimit '20'`;
- Or extend that limit. Maybe there is a circleCI configuration variable for that?

We can try and set the timeout for no output by overriding [no\\_output\\_timeout](#) although 10 minutes with no output is already a long time a test that takes a third of the time.

#### #5 - 2020-08-20 12:08 - cdywan

- Status changed from *Workable* to *In Progress*

- Assignee set to *cdywan*

<https://github.com/os-autoinst/openQA/pull/3322>

#### #6 - 2020-08-20 13:01 - mkittler

I've seen your PR but it doesn't look like any of the CI runs referenced in the ticket should have taken so long - especially not without producing output at all.

This one looks like the test did not terminate after a failure:

```
[13:46:45] t/05-scheduler-full.t ..... 3/?
# Failed test 'Allocated maximum number of jobs that could have been allocated'
# at t/05-scheduler-full.t line 206.
# got: '9'
# expected: '10'
# Looks like you failed 1 test of 1.
[13:46:45] t/05-scheduler-full.t ..... 4/?
# Failed test 'Simulation of heavy unstable load'
# at t/05-scheduler-full.t line 243.
Died at t/05-scheduler-full.t line 206.
# Tests were run but no plan was declared and done_testing() was not seen.
make[2]: *** [Makefile:168: test-unit-and-integration] Terminated
make[1]: *** [Makefile:163: test-with-database] Terminated
make: *** [Makefile:154: test-scheduler] Terminated
```

Too long with no output (exceeded 10m0s): context deadline exceeded

And the same counts for:

```
[13:21:21] t/full-stack.t .. 19/? # javascript console output: [
# {
# level => "SEVERE",
# message => "http://localhost:48265/asset/e27d1d8f48/ws_console.js 8 WebSocket connection to 'ws://
127.0.0.1:20013/YVXT621ivYM9EOMj/ws' failed: Error in connection establishment: net::ERR_CONNECTION_REFUSED",
# source => "javascript",
# timestamp => 1589030505999,
# },
# ]

# Failed test 'no errors or warnings on javascript console, waiting for connection opened'
# at /home/squamata/project/t/lib/OpenQA/SeleniumTest.pm line 260.
# Looks like you failed 1 test of 7.
[13:21:21] t/full-stack.t .. 21/?
# Failed test 'wait until developer console becomes available'
# at t/full-stack.t line 146.
[13:21:21] t/full-stack.t .. 22/? [13:21:21] t/full-stack.t .. 24/? [13:21:21] t/full-stack.t .. 26/? [13:21:2
1] t/full-stack.t .. 28/? [13:21:21] t/full-stack.t .. 30/? [13:21:21] t/full-stack.t .. 32/? [13:21:21] t/fu
l-stack.t .. 34/? [13:21:21] t/full-stack.t .. 35/? [13:21:21] t/full-stack.t .. 37/? [13:21:21] t/full-stack.
t .. 39/? [13:21:21] t/full-stack.t .. 41/? [13:21:21] t/full-stack.t .. 48/? [13:21:21] t/full-stack.t .. 49/
? [13:21:21] t/full-stack.t .. 50/? [13:21:21] t/full-stack.t .. 51/? [13:21:21] t/full-stack.t .. 54/? [13:21
:21] t/full-stack.t .. 56/? [13:21:21] t/full-stack.t .. 58/? [13:21:21] t/full-stack.t .. 60/? [13:21:21] t/f
ull-stack.t .. 62/? [13:21:21] t/full-stack.t .. 64/? [13:21:21] t/full-stack.t .. 66/? [13:21:21] t/full-stac
k.t .. 68/? [13:21:21] t/full-stack.t .. 70/? [13:21:21] t/full-stack.t .. 72/? [13:21:21] t/full-stack.t .. 7
5/? [13:21:21] t/full-stack.t .. 78/? [13:21:21] t/full-stack.t .. 83/? [13:21:21] t/full-stack.t .. 85/? [13:
21:21] t/full-stack.t .. 86/? [13:21:21] t/full-stack.t .. 93/?
# Failed test 'job 5 running'
# at t/full-stack.t line 289.
make[2]: *** [Makefile:168: test-unit-and-integration] Terminated
make[1]: *** [Makefile:163: test-with-database] Terminated
make: *** [Makefile:150: test-fullstack] Terminated
```

Too long with no output (exceeded 10m0s): context deadline exceeded

So I'm wondering why the tests didn't just fail here. Maybe one of the child processes was still running?

#### #7 - 2020-09-03 14:45 - cdywan

- I tried out `no_output_timeout` to verify that a long-running command with no output is not the issue and I think we can conclude that it's not useful here.
- I added `TimeLimit` to scheduler and fullstack tests, and it seems useful to have because this way we get usable test plan records when the tests exceed the limit (if CircleCI kills the process we don't).
- I added a grep 'not ok' to tools/retry to record test failures.
  - We can see the failures between re-runs in the log (artifacts don't account for that).
  - We can see failures even if the test plan is broken (because of CircleCI or exceptions in test cases).

I still have no satisfactory explanation for the never-ending test process. My recommendation would be to proceed with adding `TimeLimit` and eventually we'll get proper failures if these are indeed failures.

#### #8 - 2020-09-03 15:15 - cdywan

cdywan wrote:

I still have no satisfactory explanation for the never-ending test process. [...]

Might be worth mentioning, I have another idea, and that's to investigate why make will continue to run even if it receives `C` or gets killed. Not sure if anyone investigated before, but it seems superficially similar in that the process gets stuck - killing chromedriver on my system reproduces the never-ending make process.

#### #9 - 2020-09-07 13:06 - cdywan

cdywan wrote:

- I added `TimeLimit` to scheduler and fullstack tests, and it seems useful to have because this way we get usable test plan records when the tests exceed the limit (if CircleCI kills the process we don't).

<https://github.com/os-autoinst/openQA/pull/3367>

#### #10 - 2020-09-08 07:00 - okurz

<https://app.circleci.com/pipelines/github/os-autoinst/openQA/4109/workflows/ef3cacca-7dee-408f-8e83-3af26002f2d3/jobs/39239/steps> is another good example of the problem. `t/25-cache-service.t` failed in the first try, the second try is started but then is aborted by circleCI as it's stuck. So clearly there is a problem in `t/25-cache-service.t` but I guess a test module specific timeout smaller than 10m can help.

#### #11 - 2020-09-09 15:36 - cdywan

- I added a grep 'not ok' to tools/retry to record test failures.
  - We can see the failures between re-runs in the log (artifacts don't account for that).
  - We can see failures even if the test plan is broken (because of CircleCI or exceptions in test cases).

<https://github.com/os-autoinst/openQA/pull/3378>

#### #12 - 2020-09-10 15:04 - cdywan

okurz wrote:

<https://app.circleci.com/pipelines/github/os-autoinst/openQA/4109/workflows/ef3cacca-7dee-408f-8e83-3af26002f2d3/jobs/39239/steps> is another good example of the problem. `t/25-cache-service.t` failed in the first try, the second try is started but then is aborted by circleCI as it's stuck. So clearly there is a problem in `t/25-cache-service.t` but I guess a test module specific timeout smaller than 10m can help.

My latest hypothesis is that the cache service never comes available in the "stuck" case which has a 600s timeout, and which would never produce any output because it uses `note` in each iteration (and I reproduced this after many re-runs):

<https://github.com/os-autoinst/openQA/pull/3382>

#### #13 - 2020-09-15 08:14 - cdywan

There's the semi-related item of `25-cache-service.t` I was also investigating, where All subtests passed but `prove` concludes with an exit code of 255, and a message claiming Looks like your test exited with -1 which can be observed locally. Since the issue of this ticket is sorted now, I filed [#71368](#) to track the cache service issue.

#### #14 - 2020-09-15 08:14 - cdywan

- Status changed from *In Progress* to *Feedback*

#### #15 - 2020-09-16 08:22 - okurz

we still have cases, e.g.

<https://app.circleci.com/pipelines/github/os-autoinst/openQA/4208/workflows/a42f68e0-60c7-46e0-88f3-2ace61b3c2b1/jobs/40235> . IMHO we should still try to set a timeout within circleci that is higher than the timeout we apply on make level. This is independant of setting a timeout within each test module.

[cdywan](#) did you follow up on [#66664#note-4](#) ?

**#16 - 2020-09-17 10:07 - okurz**

- Related to action #71464: Show output of failed tests within circleci log view, like on local prove calls, not only "Tests" tab added

**#17 - 2020-09-17 10:10 - okurz**

- Copied to action #71467: add specific timeout values to all applicable openQA tests in t/ added

**#18 - 2020-09-18 21:14 - okurz**

<https://github.com/os-autoinst/openQA/pull/3407> for the actual bump of the "no output" timeout.

**#19 - 2020-09-19 13:21 - okurz**

- Status changed from Feedback to Resolved

With <https://github.com/os-autoinst/openQA/pull/3407> and <https://github.com/os-autoinst/openQA/pull/3398> now both merged I think we can call this done.