

## openQA Project - action #71554

### unstable/flaky/sporadic t/full-stack.t test failing in script waits on CircleCI

2020-09-19 13:14 - okurz

<b>Status:</b>	Resolved	<b>Start date:</b>	2020-09-19
<b>Priority:</b>	Normal	<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>			
<b>Description</b>			
<b>Observation</b>			
Since recently t/full-stack.t is more unstable.			
<b>Steps to reproduce</b>			
Probably reproducible locally with			
<pre>make test FULLSTACK=1 TESTS=t/full-stack.t</pre>			
<b>Suggestions</b>			
Bisect where the regression comes from and fix that to be stable locally and within CI.			
<b>Related issues:</b>			
Related to openQA Project - action #37638: Flaky fullstack test: 'Test 3 is s...		<b>Resolved</b>	<b>2018-06-21</b>
Related to openQA Project - action #59043: Fix unstable/flaky full-stack test...		<b>Resolved</b>	<b>2019-11-04</b>
Copied from openQA Project - action #71551: unstable/flaky/sporadic t/04-sche...		<b>Resolved</b>	<b>2020-09-19</b>

## History

### #1 - 2020-09-19 13:14 - okurz

- Copied from action #71551: unstable/flaky/sporadic t/04-scheduler.t test failing added

### #2 - 2020-09-23 14:34 - mkittler

I've created <https://github.com/os-autoinst/openQA/pull/3405> to better track down the problem. I'll have a look at the full stack test when checking CI failures of my PRs but so far I'm not quite sure what the problem is.

### #3 - 2020-09-24 05:13 - okurz

- Related to action #37638: Flaky fullstack test: 'Test 3 is scheduled' at t/full-stack.t added

### #4 - 2020-09-24 05:13 - okurz

- Related to action #59043: Fix unstable/flaky full-stack test, i.e. remove sleep, and ui tests added

### #5 - 2020-09-24 05:21 - okurz

It seems to become worse now, e.g.

<https://app.circleci.com/pipelines/github/os-autoinst/openQA/4317/workflows/66701e42-dd43-4159-824e-d8ec08883956/jobs/41463> shows

```
timeout -s SIGINT -k 5 -v  $\$( (20 * (3 + 1) ) )$ m tools/retry prove -l --harness TAP::Harness::JUnit --timer --mer
ge t/full-stack.t
Retry 1 of 3 ...
[17:59:23] t/full-stack.t .. 92/? make[2]: *** [Makefile:174: test-unit-and-integration] Terminated
make[1]: *** [Makefile:169: test-with-database] Terminated
make: *** [Makefile:154: test-fullstack] Terminated
```

Too long with no output (exceeded 30m0s): context deadline exceeded

so the test job is not even finishing within 30m but the logfile in

[https://circle-production-customer-artifacts.s3.amazonaws.com/picard/forke/58f7029dc9e77c000129905e/46416941/5f6b8c9e29253478672eb817-0-build/artifacts/artifacts/full-stack.t?X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Date=20200924T051432Z&X-Amz-SignedHeaders=host&X-Amz-Expires=60&X-Amz-Credential=AKIAJR3Q6CR467H7Z55A%2F20200924%2Fus-east-1%2Fs3%2Faws4\\_request&X-Amz-Signature=6562cafe3e6873cd02d008bc1837d7dad6d8924b65ccffc262ec9f1780401456](https://circle-production-customer-artifacts.s3.amazonaws.com/picard/forke/58f7029dc9e77c000129905e/46416941/5f6b8c9e29253478672eb817-0-build/artifacts/artifacts/full-stack.t?X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Date=20200924T051432Z&X-Amz-SignedHeaders=host&X-Amz-Expires=60&X-Amz-Credential=AKIAJR3Q6CR467H7Z55A%2F20200924%2Fus-east-1%2Fs3%2Faws4_request&X-Amz-Signature=6562cafe3e6873cd02d008bc1837d7dad6d8924b65ccffc262ec9f1780401456) shows what looks like the test running just fine (albeit probably super slow) until it is aborted by circleci. Unfortunately neither the test module timeout nor the timeout on make level trigger. Normal runs take about 5m, e.g. see <https://app.circleci.com/pipelines/github/os-autoinst/openQA/4320/workflows/6e1fdc89-9482-4a3d-9a3e-b78135abbe6e/jobs/41464> so I guess we can at least tweak some timeouts: <https://github.com/os-autoinst/openQA/pull/3415>  
This is only fighting the symptoms, not addressing the root cause for his problem so not assigning the ticket to myself yet.

#### #6 - 2020-09-28 16:32 - cdywan

okurz wrote:

It seems to become worse now, e.g. <https://app.circleci.com/pipelines/github/os-autoinst/openQA/4317/workflows/66701e42-dd43-4159-824e-d8ec08883956/jobs/41463> shows

```
timeout -s SIGINT -k 5 -v $((20 * (3 + 1) ))m tools/retry prove -l --harness TAP::Harness::JUnit --timer -
-merge t/full-stack.t
Retry 1 of 3 ...
[17:59:23] t/full-stack.t .. 92/? make[2]: *** [Makefile:174: test-unit-and-integration] Terminated
make[1]: *** [Makefile:169: test-with-database] Terminated
make: *** [Makefile:154: test-fullstack] Terminated
```

Too long with no output (exceeded 30m0s): context deadline exceeded

so the test job is not even finishing within 30m but the logfile in [https://circle-production-customer-artifacts.s3.amazonaws.com/picard/forke/58f7029dc9e77c000129905e/46416941/5f6b8c9e29253478672eb817-0-build/artifacts/artifacts/full-stack.t?X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Date=20200924T051432Z&X-Amz-SignedHeaders=host&X-Amz-Expires=60&X-Amz-Credential=AKIAJR3Q6CR467H7Z55A%2F20200924%2Fus-east-1%2Fs3%2Faws4\\_request&X-Amz-Signature=6562cafe3e6873cd02d008bc1837d7dad6d8924b65ccffc262ec9f1780401456](https://circle-production-customer-artifacts.s3.amazonaws.com/picard/forke/58f7029dc9e77c000129905e/46416941/5f6b8c9e29253478672eb817-0-build/artifacts/artifacts/full-stack.t?X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-Date=20200924T051432Z&X-Amz-SignedHeaders=host&X-Amz-Expires=60&X-Amz-Credential=AKIAJR3Q6CR467H7Z55A%2F20200924%2Fus-east-1%2Fs3%2Faws4_request&X-Amz-Signature=6562cafe3e6873cd02d008bc1837d7dad6d8924b65ccffc262ec9f1780401456) shows what looks like the test running just fine (albeit probably super slow) until it is aborted by circleci. Unfortunately neither the test module timeout nor the timeout on make level trigger. Normal runs take about 5m, e.g. see <https://app.circleci.com/pipelines/github/os-autoinst/openQA/4320/workflows/6e1fdc89-9482-4a3d-9a3e-b78135abbe6e/jobs/41464> so I guess we can at least tweak some timeouts: <https://github.com/os-autoinst/openQA/pull/3415>  
This is only fighting the symptoms, not addressing the root cause for his problem so not assigning the ticket to myself yet.

Another piece we've not considered in this puzzle is Javascript. I suspect the code gets stuck waiting for the result panel of job 8. The loop never times out. And it's a loop based on number of iteration, not a timeout, which means it's as slow as the javascript and sleep calls make it in practice.

<https://github.com/os-autoinst/openQA/pull/3430>

#### #7 - 2020-09-29 07:33 - cdywan

- Status changed from Workable to In Progress

- Assignee set to cdywan

#### #8 - 2020-09-29 08:37 - cdywan

Note that I also evaluated past jobs on CI and the only failures I could find were due to Javascript getting stuck, and it's not failing a lot actually. I will keep an eye on it anyway, though (and that's part of Feedback).

#### #9 - 2020-09-29 08:38 - cdywan

- Subject changed from unstable/flaky/sporadic t/full-stack.t test failing to flaky t/full-stack.t test failing in script waits on CircleCI

- Description updated

#### #10 - 2020-09-29 13:35 - okurz

- Subject changed from flaky t/full-stack.t test failing in script waits on CircleCI to unstable/flaky/sporadic t/full-stack.t test failing in script waits on CircleCI

- Status changed from In Progress to Feedback

- Priority changed from Urgent to Normal

I included "unstable/flaky/sporadic" in the subject line to have a higher chance to find this ticket again when searching subjects :)

You created <https://github.com/os-autoinst/openQA/pull/3430>, I merged that now. As you stated the problem seems to have been again less severe lately so we can track in Feedback with lower priority now. thx

#### #11 - 2020-10-08 19:13 - okurz

I created <https://github.com/os-autoinst/openQA/pull/3455> to mark t/full-stack.t as stable and faster (reduced timeout). Do you plan any further work

here or what feedback you are waiting for?

**#12 - 2020-10-15 15:37 - cdywan**

- *Status changed from Feedback to Resolved*

I think it's fine now