

openQA Project - action #65082

flaky/sporadic unstable failure in t/43-scheduling-and-worker-scalability.t

2020-03-31 11:48 - okurz

Status:	Workable	Start date:	2020-03-31
Priority:	Low	Due date:	
Assignee:		% Done:	0%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:	future		
Difficulty:			

Description

Observation

Seen in

<https://circlci.com/api/v1.1/project/github/os-autoinst/openQA/19175/output/111/0?file=true&allocation-id=5e83023bf2a75b19962e1646-0-build%2F6148E966> but also in at least one other case:

```
[08:55:42] t/43-scheduling-and-worker-scalability.t .. 1/?
# Failed test 'all jobs done'
# at t/43-scheduling-and-worker-scalability.t line 196.
# got: '4'
# expected: '5'

# Failed test 'all jobs passed'
# at t/43-scheduling-and-worker-scalability.t line 197.
# got: '4'
# expected: '5'
# Looks like you failed 2 tests of 3.

[08:55:42] t/43-scheduling-and-worker-scalability.t .. 2/?
# Failed test 'assign and run jobs'
# at t/43-scheduling-and-worker-scalability.t line 199.

[08:55:42] t/43-scheduling-and-worker-scalability.t .. 3/? # Looks like you failed 1 test of 4.

[08:55:42] t/43-scheduling-and-worker-scalability.t .. Dubious, test returned 1 (wstat 256, 0x100)
Failed 1/4 subtests
[08:56:13]

Test Summary Report
-----
t/43-scheduling-and-worker-scalability.t (Wstat: 256 Tests: 4 Failed: 1)
  Failed test: 2
  Non-zero exit status: 1
Files=2, Tests=9, 88.8589 wallclock secs ( 0.59 usr 0.04 sys + 183.69 cusr 28.88 csys = 213.20 CPU)
Result: FAIL
Retrying...
[08:56:13] t/05-scheduler-full.t .....
[08:56:13] t/05-scheduler-full.t ..... 1/?
[08:56:13] t/05-scheduler-full.t ..... 2/?
[08:56:13] t/05-scheduler-full.t ..... 3/?
[08:56:13] t/05-scheduler-full.t ..... 4/?
[08:56:13] t/05-scheduler-full.t ..... 5/?

[08:56:13] t/05-scheduler-full.t ..... ok 53057 ms ( 0.43 usr 0.03 sys + 146.9
0 cusr 12.32 csys = 159.68 CPU)
[08:57:06] t/43-scheduling-and-worker-scalability.t ..
[08:57:06] t/43-scheduling-and-worker-scalability.t .. 1/?
# Failed test 'all jobs done'
```

```

# at t/43-scheduling-and-worker-scalability.t line 196.
# got: '4'
# expected: '5'

# Failed test 'all jobs passed'
# at t/43-scheduling-and-worker-scalability.t line 197.
# got: '4'
# expected: '5'
# Looks like you failed 2 tests of 3.

[08:57:06] t/43-scheduling-and-worker-scalability.t .. 2/?
# Failed test 'assign and run jobs'
# at t/43-scheduling-and-worker-scalability.t line 199.

[08:57:06] t/43-scheduling-and-worker-scalability.t .. 3/? # Looks like you failed 1 test of 4.

[08:57:06] t/43-scheduling-and-worker-scalability.t .. Dubious, test returned 1 (wstat 256, 0x100)
Failed 1/4 subtests
[08:57:39]

Test Summary Report
-----
t/43-scheduling-and-worker-scalability.t (Wstat: 256 Tests: 4 Failed: 1)
  Failed test: 2
  Non-zero exit status: 1
Files=2, Tests=9, 85.2037 wallclock secs ( 0.57 usr  0.04 sys + 177.56 cusr 16.58 csys = 194.75 CPU)
Result: FAIL
Retrying...
[08:57:39] t/05-scheduler-full.t .....
[08:57:39] t/05-scheduler-full.t ..... 1/?
[08:57:39] t/05-scheduler-full.t ..... 2/?
[08:57:39] t/05-scheduler-full.t ..... 3/?
[08:57:39] t/05-scheduler-full.t ..... 4/?
[08:57:39] t/05-scheduler-full.t ..... 5/?

[08:57:39] t/05-scheduler-full.t ..... ok 55655 ms ( 0.45 usr  0.03 sys + 119.1
2 cusr 12.33 csys = 131.93 CPU)
[08:58:35] t/43-scheduling-and-worker-scalability.t ..
[08:58:35] t/43-scheduling-and-worker-scalability.t .. 1/?
# Failed test 'all jobs done'
# at t/43-scheduling-and-worker-scalability.t line 196.
# got: '4'
# expected: '5'

# Failed test 'all jobs passed'
# at t/43-scheduling-and-worker-scalability.t line 197.
# got: '4'
# expected: '5'
# Looks like you failed 2 tests of 3.

[08:58:35] t/43-scheduling-and-worker-scalability.t .. 2/?

```

that is failing after 3 retries with always the same failure.

Acceptance criteria

- **AC1:** t/43-scheduling-and-worker-scalability.t is stable, e.g. 100/100 runs on circleCI as well as local

Suggestions

- Run `make test TESTS=t/43-scheduling-and-worker-scalability.t SCALABILITY_TEST=1 RETRY=100`

History

#1 - 2020-03-31 11:49 - okurz

- Status changed from In Progress to Feedback

<https://github.com/os-autoinst/openQA/pull/2885>

#2 - 2020-04-07 09:51 - okurz

<https://github.com/os-autoinst/openQA/pull/2882> is the first part from Martchus bumping the internal timeout waiting for jobs to finish. <https://github.com/os-autoinst/openQA/pull/2885#issuecomment-610291020> describes a new, unexpected problem in the scheduler "full" test.

#3 - 2020-05-15 14:26 - cdywan

- Description updated

#4 - 2020-06-26 10:15 - okurz

- Status changed from Feedback to Workable

- Assignee deleted (okurz)

- Priority changed from High to Low

I have not seen the test module t/43-scheduling-and-worker-scalability.t failing on the top level since we have a RETRY=3 and also my recent changes to the test setup, e.g. using dynamic-but-ensured-to-be-free ports. To complete this ticket one would still need to ensure stability without RETRY=3

#5 - 2020-07-23 09:26 - okurz

- Target version set to Ready

#6 - 2021-02-23 15:00 - okurz

<https://github.com/os-autoinst/openQA/pull/3714> is a related change that tried to stabilize t/43-scheduling-and-worker-scalability.t . Maybe that's already enough and the test can be run without retries now.

#7 - 2021-02-23 15:35 - mkittler

The fix will not help with this one because it only prevents us from running into BAIL_OUT('Unable to assign jobs to (idling) workers');. Previously it was possible that not a single worker has already been *fully* registered. However, in the case of this ticket the test was already further but the jobs haven't been processed fast enough.

In fact, I could just reproduce this case locally (with the previous fix present) via `time env runs=400 SCALABILITY_TEST=1 "$OPENQA_BASEDIR/repos/okurz-github-scripts/count_fail_ratio" openqa-test t/43-scheduling-and-worker-scalability.t`:

```
not ok 2 - all jobs done

# Failed test 'all jobs done'
# at t/43-scheduling-and-worker-scalability.t line 211.
# got: '4'
# expected: '5'
not ok 3 - all jobs passed

# Failed test 'all jobs passed'
# at t/43-scheduling-and-worker-scalability.t line 212.
# got: '4'
# expected: '5'
# All jobs:
# - id: 1, state: done, result: passed, reason: none
# - id: 2, state: done, result: passed, reason: none
# - id: 3, state: assigned, result: none, reason: none
# - id: 4, state: done, result: passed, reason: none
# - id: 5, state: done, result: passed, reason: none
1..3
# Looks like you failed 2 tests of 3.
not ok 2 - assign and run jobs
```

But that was also the only failing test run out of 400. Normally it takes 6 polling attempts (# Waiting until all jobs are done, try 6) on my machine. In the failed case there were 125 attempts before (# Waiting until all jobs are done, try 125).

#8 - 2021-04-29 07:42 - okurz

- Target version changed from Ready to future