

openQA Project - action #34504

[tools][sporadic] Job's auto_duplicate fails to duplicate job dependencies

2018-04-09 09:50 - dasantiago

Status:	Resolved	Start date:	2018-04-09
Priority:	Normal	Due date:	
Assignee:	dasantiago	% Done:	20%
Category:	Feature requests	Estimated time:	0.00 hour
Target version:	Done		
Difficulty:			
Description			
I deleted by mistake the poo#32854 instead of my comment.			
This is happens sporadically and needs investigation, we can observe most of the times 'flaky' travis tests that fails on it [1].			
Apparently auto_duplicate fails to copy part of the job dependencies, it doesn't happen only on CI tests, we have other reports mentioning this issue already.			
AC:			
<ul style="list-style-type: none">• Identify and fix the issue• Stress tests inside our unit-test suite that shows that auto_duplicate is idempotent			
1: Test failure points to https://github.com/os-autoinst/openQA/blob/dda08666d6f21473a990e9afcbda6be8a8280b2c/t/05-scheduler-dependencies.t#L785 auto_duplicate()			
Related issues:			
Related to openQA Project - action #35914: Changes to Job::duplicate		Resolved	2018-05-04

History

#1 - 2018-04-09 09:52 - dasantiago

This can be reproduced by the test:

```
time while perl t/05-scheduler-dependencies.t; do echo next; done
```

#2 - 2018-04-09 09:54 - dasantiago

There's two problems with this test:

1-

This query <https://github.com/os-autoinst/openQA/blob/master/lib/OpenQA/Schema/Result/Jobs.pm#L735> doesn't enforce any type of order and the is_deeply from Test::More contains some limitations, it requires the whole data structure to match exactly. There is no place for any flexibility. Although the following chained dependencies are equivalent the test will fail:

[\$jobB2_h->{id}, \$jobC2_h->{id}, \$jobD2_h->{id}] contains:

```
[ 100032,  
 100030,  
 100031  
]
```

\$jobA2_h->{children}->{Chained} contains:

```
[ 100030,  
 100031,  
 100032  
]
```

The dependencies are the same, but it will fail.

2- Some dependencies are not returned.
Still debugging

#3 - 2018-04-11 08:02 - dasantiago

I'm still analyzing the point 2 mentioned above and via analysis of the SQL queries i can confirm that in the cases where it fails, there some insert jobs in the DB that aren't being done causing the tests to fail.

#4 - 2018-04-12 08:14 - szarate

- File scheduler-dependencies-failure.patch added

After running while prove --verbose --color t/05-scheduler-dependencies.t; do echo "Not failed"; done a test (different than auto_duplicate()) fails too. Which means that sorting just hides part of the problem.

```
# [
# 100030,
# 100031,
# 100032
# ]
# [
# 100030,
# 100031,
# 100032
# ]
ok 176 - jobA2 has jobB2, jobC2 and jobD2 as children
[debug] new job 100042
ok 177 - job cloned
ok 178 - job has jobA2 as parent
ok 179 - job cloned
ok 180 - job has jobA2 as parent
ok 181 - job cloned
ok 182 - job has jobA2 as parent
ok 183 - jobA2 is indeed jobA clone
[debug] new job 100043
ok 184 - job correctly not cloned
ok 185 - job has jobA3 as parent
ok 186 - job correctly not cloned
ok 187 - job has jobA3 as parent
ok 188 - job correctly not cloned
ok 189 - job has jobA3 as parent

[debug] new job 100051

ok 190 - job cloned

not ok 191 - job has jobA2 as parent

# Failed test 'job has jobA2 as parent'

# at t/05-scheduler-dependencies.t line 793.
# Structures begin differing at:

# $got->[0] = Does not exist

# $expected->[0] = '100051'
```

#5 - 2018-04-12 08:48 - dasantiago

szarate wrote:

After running while prove --verbose --color t/05-scheduler-dependencies.t; do echo "Not failed"; done a test (different than auto_duplicate()) fails too. Which means that sorting just hides part of the problem.

```
# [
# 100030,
```

```

# 100031,
# 100032
# ]
# [
# 100030,
# 100031,
# 100032
# ]
ok 176 - jobA2 has jobB2, jobC2 and jobD2 as children
[debug] new job 100042
ok 177 - job cloned
ok 178 - job has jobA2 as parent
ok 179 - job cloned
ok 180 - job has jobA2 as parent
ok 181 - job cloned
ok 182 - job has jobA2 as parent
ok 183 - jobA2 is indeed jobA clone
[debug] new job 100043
ok 184 - job correctly not cloned
ok 185 - job has jobA3 as parent
ok 186 - job correctly not cloned
ok 187 - job has jobA3 as parent
ok 188 - job correctly not cloned
ok 189 - job has jobA3 as parent

[debug] new job 100051

ok 190 - job cloned

not ok 191 - job has jobA2 as parent

# Failed test 'job has jobA2 as parent'

# at t/05-scheduler-dependencies.t line 793.
# Structures begin differing at:

# $got->[0] = Does not exist

# $expected->[0] = '100051'

```

It doesn't hide. It's a different problem that i mentioned on point 2.

The sort just fixes the test. The problem on test 191, is that the clone job isn't being created. This is the real problem.

#6 - 2018-04-12 10:01 - szarate

- File error.log added

Run after \$schema->storage->debug(1);

#7 - 2018-04-12 10:05 - szarate

- File passing.log added

#8 - 2018-04-12 11:27 - szarate

- File expected_error.log added

#9 - 2018-04-13 14:53 - szarate

Saving comments from Mudler from poo#32858:

As promised moving back into Ready if takes more than 1 day.

Didn't had much luck with that - it is hard to reproduce locally (just 1 out of ~30 execution run of scheduler_dependencies test fails for me) so it's expensive in terms of time, and difficult to identify the real issue as well (and be sure it's fixed, as we have already a lot of false positives).

At least i can exclude that when it does happen - my hunch in the comment before is not reached, so must be so

mething else.

Tried wrapping everything in a transaction as well (since it is searching and cloning, recursively, multiple i
nvocations could create race conditions) but made tests fails more horribly, and that road takes for sure more
than one day - i'm afraid we will have to refactor this if starts to become even more problematic.

#10 - 2018-04-18 09:36 - dasantiago

- % Done changed from 0 to 80

Issue found and fixed, but it broke some other tests. I'm fixing the remaining tests.

#11 - 2018-04-19 14:41 - dasantiago

Only UI tests (not related) are failing in travis

#12 - 2018-04-19 14:41 - dasantiago

- % Done changed from 80 to 100

#13 - 2018-04-23 08:59 - dasantiago

- % Done changed from 100 to 90

Need to improve the tests before it can be closed.

#14 - 2018-04-24 14:16 - dasantiago

- % Done changed from 90 to 100

The changes were merged yesterday.

More tests are implemented and the PR already done.

#15 - 2018-04-25 10:53 - dasantiago

- Status changed from In Progress to Resolved

#16 - 2018-05-02 13:06 - EDiGiacinto

- Status changed from Resolved to In Progress

Reopening since now duplication is not working properly

<https://openqa.suse.de/tests/1659294#settings>

<https://openqa.suse.de/tests/1652849>

<https://openqa.suse.de/tests/1644104#settings>

<https://openqa.suse.de/tests/overview?distri=caasp&version=3.0&build=0073&groupid=134>

#17 - 2018-05-02 17:23 - dasantiago

EDiGiacinto wrote:

Reopening since now duplication is not working properly

<https://openqa.suse.de/tests/1659294#settings>

<https://openqa.suse.de/tests/1652849>

<https://openqa.suse.de/tests/1644104#settings>

<https://openqa.suse.de/tests/overview?distri=caasp&version=3.0&build=0073&groupid=134>

You have to be more specific in what's wrong with the duplication.

Is the skipped job in <https://openqa.suse.de/tests/1644104#settings> ?

And ahow about the other jobs? It's only on caasp jobs that this problem is happening?

#18 - 2018-05-03 07:10 - EDiGiacinto

dasantiago wrote:

EDiGiacinto wrote:

Reopening since now duplication is not working properly

<https://openqa.suse.de/tests/1659294#settings>
<https://openqa.suse.de/tests/1652849>
<https://openqa.suse.de/tests/1644104#settings>
<https://openqa.suse.de/tests/overview?distri=caasp&version=3.0&build=0073&groupid=134>

You have to be more specific in what's wrong with the duplication.

Adding in CC then who might explain you better, you have QA Engineers responsible for those tests in your room ;)

Basically - there are jobs that now are even decoupled from the cluster at all when you hit the restart button.

Is the skipped job in <https://openqa.suse.de/tests/1644104#settings> ?

It's not a matter of the job results, it's that they are not tied anymore in the same cluster after restarting them. This is more important when automatic restarts comes in place.

And ahow about the other jobs? It's only on caasp jobs that this problem is happening?

No, also ses, and potentially all MM jobs.

#19 - 2018-05-03 09:42 - mkravec

There is ~10% chance that CaaSP cluster will have incomplete job. If incomplete happens, then:

- before this change it failed in "organized" way - jobs started cloning and losing dependencies, so they never recovered and eventually it all died
- now it all goes "kaboom" and weird things happen

For example:

- <https://openqa.suse.de/tests/1661011#settings> - how does this job have 2 children QAM-CaaSP-admin
- <https://openqa.suse.de/tests/1661007> was cloned but lost 1 dependency during that
- <https://openqa.suse.de/tests/1652849> how can this job find 8 new dependencies after being cloned (it recovered & passed fine at at end)
- <https://openqa.suse.de/tests/1652962> how can incomplete job have higher ID than one that passed (1652905) - causing incomplete result being displayed in result overview

I reschedule ISO when this happens.

#20 - 2018-05-03 10:11 - dasantiago

My changes only adds a dependency that was missing, so because of that change it shouldn't lose anything.

For example:

<https://openqa.suse.de/tests/1661011#settings>

This probably was introduced with my change, but the losing deps? :-(

Is there any multimachine test environment that i can use for my needs? Or is there any way to simulate a multi machine environment?

#21 - 2018-05-03 10:12 - dasantiago

- % Done changed from 100 to 20

#22 - 2018-05-04 07:53 - EDiGiacinto

dasantiago wrote:

My changes only adds a dependency that was missing, so because of that change it shouldn't lose anything.

Yes and no, as it's not only adding, before it was just skipping - if you look closely at <https://github.com/os-autoinst/openQA/pull/1623/files#diff-85ae48e70a5c110c9e439c3a5ea28d5fR759> you can also skip other child deps (see next()) and ignore the other conditions down while cycling, that could call duplicate (recursively, again) and i suspect that can bring also to loose duplicated jobs dependencies; even if unpleasant, duplicate() looks buggy, to fix properly this we would need to rewrite it from scratch

For example:

<https://openqa.suse.de/tests/1661011#settings>

This probably was introduced with my change, but the losing deps? :-(

Is there any multimachine test environment that i can use for my needs? Or is there any way to simulate a multi machine environment?

#23 - 2018-05-04 14:10 - szarate

- Status changed from In Progress to Resolved

Setting to resolved, work will continue in: poo#35914 as the initial problem here was solved.

#24 - 2018-05-04 14:10 - szarate

- Related to action #35914: Changes to Job::duplicate added

#25 - 2018-05-04 14:13 - szarate

- Target version changed from Current Sprint to Done

Files

scheduler-dependencies-failure.patch	1.11 KB	2018-04-12	szarate
error.log	29.1 KB	2018-04-12	szarate
passing.log	28.4 KB	2018-04-12	szarate
expected_error.log	28.7 KB	2018-04-12	szarate