

openQA Project - action #66619

OpenQA jobs roll back to the wrong snapshot on hard test failure

2020-05-07 11:54 - MDoucha

Status:	Rejected	Start date:	2020-05-07
Priority:	Low	Due date:	
Assignee:		% Done:	0%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:	future		
Difficulty:			
Description			
<p>When a job includes multiple modules that create a snapshot, VM rollback appears to always use the very first snapshot instead of the last one.</p> <p>Example: https://openqa.suse.de/tests/4203253#step/AD044/6 Module AD043 failed and triggered VM rollback. The remaining modules then fail with the following error:</p> <pre>/tmp/aiodio/junkfile: No such file or directory</pre> <p>This means that the VM was rolled back all the way to boot_ltp. But it was supposed to use the snapshot created by create_junkfile_ltp.</p> <p>This does not appear to be a new issue. The same error appears in all LTP aiodio jobs which failed since VM rollback was enabled for them by https://github.com/os-autoinst/os-autoinst-dstri-opensuse/pull/9264 Oldest known example: https://openqa.suse.de/tests/3987350#step/AD037/6</p>			

History

#1 - 2020-05-07 12:11 - asmorodskyi

please be aware of <https://progress.opensuse.org/issues/60443> there might be some connection between them

#2 - 2020-05-07 13:28 - MDoucha

asmorodskyi wrote:

please be aware of <https://progress.opensuse.org/issues/60443> there might be some connection between them

I don't see how that ticket is related. This one is specifically about rolling back to the wrong snapshot. Aside from that the job ran all the way to the end and isotovideo exited gracefully.

My best guess is that QEMU just creates a new snapshot named something like lastgood.001, lastgood.002 and so on instead of overwriting lastgood as we want it to.

#3 - 2020-05-09 12:08 - okurz

I also do not see the relation to [#60443](#). As you created this ticket in "openQA Infrastructure", why do you think this is specific to our infrastructure? Your description and comments makes it sound like a generic problem within os-autoinst. So should we move this to [openQA Project](#) instead?

#4 - 2020-05-11 08:12 - MDoucha

- Project changed from openQA Infrastructure to openQA Project

#5 - 2020-07-03 19:43 - okurz

- Category set to Concrete Bugs
- Priority changed from Normal to Low
- Target version set to future

In the current form I am not sure how to help. Could you please fill the ticket description based on the defect template in <https://progress.opensuse.org/projects/openqav3/wiki#Defects>? E.g. especially how to reproduce, etc.

#6 - 2020-07-07 08:43 - MDoucha

Test suite that will reproduce the bug (4 test modules):

1. boot to console and save a VM snapshot (test_flags: milestone => 1)
2. create a test file and save another VM snapshot (test_flags: milestone => 1)
3. die()
4. check whether the test file created in step 2 exists

#7 - 2021-03-15 09:27 - MDoucha

OpenQA is still rolling back to the wrong snapshot on test module timeout:

<https://openqa.suse.de/tests/5653440#step/AD074/8>

#8 - 2021-03-17 13:25 - okurz

Not sure what you expected. No one worked on this ticket and we do not currently plan to do, so this should be expected, right?

#9 - 2021-03-17 14:33 - cdywan

- Category changed from Concrete Bugs to Feature requests

This does not appear to be a new issue. The same error appears in all LTP aiodio jobs which failed since VM rollback was enabled for them by

<https://github.com/os-autoinst/os-autoinst-distri-opensuse/pull/9264>

Oldest known example: <https://openqa.suse.de/tests/3987350#step/AD037/6>

As per the description, this should actually be considered a **Feature Request**, if this never worked ☐☐

#10 - 2021-03-17 16:46 - MDoucha

okurz wrote:

Not sure what you expected. No one worked on this ticket and we do not currently plan to do, so this should be expected, right?

I expected someone to raise the priority of this ticket because it's breaking tests.

cdywan wrote:

As per the description, this should actually be considered a **Feature Request**, if this never worked ☐☐

Fixing a broken existing feature to behave according to official documentation does not count as a "feature request" in my book.

#11 - 2021-03-17 17:25 - cdywan

MDoucha wrote:

okurz wrote:

Not sure what you expected. No one worked on this ticket and we do not currently plan to do, so this should be expected, right?

I expected someone to raise the priority of this ticket because it's breaking tests.

I think it'd be good to have a quick call to clarify what we're looking at. Because I see a known issue that rarely pops up but would be nice to look into, and you seem to see a bug that needs to be fixed.

Not sure if you get notifications, so also poking on RC.

#12 - 2021-03-18 10:34 - MDoucha

- Category changed from Feature requests to Concrete Bugs

cdywan wrote:

I think it'd be good to have a quick call to clarify what we're looking at. Because I see a known issue that rarely pops up but would be nice to look into, and you seem to see a bug that needs to be fixed.

Not sure if you get notifications, so also poking on RC.

I see a bug in fairly important OpenQA feature that is used by dozens of testsuites to guarantee certain test flow. And the test flow guarantee is broken by this bug. It is only a minor annoyance to me personally because 1) the LTP tests in question do not fail every time so I can simply restart them to get the missing test results and 2) when the bug happens, the LTP tests will crash in an obvious manner. But this bug may break other tests in more subtle ways that are not readily apparent which will result in regressions slipping through QA. Especially anything following a test module with `always_rollback` flag needs to be considered broken if it depends on some cross-module setup done between the first and any later VM snapshot.

#13 - 2021-03-18 14:52 - cdywan

- Priority changed from Low to High

- Target version changed from future to Ready

- The case of <https://openqa.suse.de/tests/3987350#step/AD036/7> fails due to a btrfs bug in 12sp4 which triggers the snapshot issue because of rollback on failure
- We also have cases like <https://openqa.suse.de/tests/5670834> where everything "passes" despite the rollback being broken. This could be one out of 162 tests affected. From the result and from the logs you can't tell, unless you inspect the files by hand.

Messages like Loading a VM snapshot lastgood don't reveal if the correct snapshot was loaded. When the bug is hit, this actually goes back to the first one, not the previous one. I'm wondering if we could extend logging to be able to see e.g. what snapshot was produced [] []

Note: Due to git/needle logs the autoinst-log.txt is very big, but downloading and locally grepping for VM snapshot works.

I'd like to tentatively suggest this be *High* because this can cause silent regressions in tests relying on rollback and testing the wrong snapshots - and it doesn't seem like that's a work-around to avoid it

Note I left it at *New* because I'd still like to confirm

- is this a bug in backend code or in isotovideo
- is there a work-around - I can't find a way to detect this from logs so far
- can we identify more affected jobs - for now I assume LTP and pam tests are affected

#14 - 2021-03-18 15:33 - mkittler

So we also have cases like <https://openqa.suse.de/tests/5670834> where everything "passes" despite the rollback being broken.

Right. Normally the rollback only happens on failures but this test uses the `always_rollback` flag so it relies on the rollback working.

Module AD043 failed and triggered VM rollback.

But this module is actually `tests/kernel/run_ltp.pm`, right? Maybe what makes your test special is that "AD043" is not a real test module? So autotest it actually still at "create_junkfile_ltp" which is the last module corresponding to a real Perl module (`tests/kernel/create_junkfile_ltp.pm`).

#15 - 2021-03-29 14:46 - MDoucha

I've tried writing a set of reproducer modules and I can't reproduce the issue.

Reproducer code: <https://github.com/mdoucha/os-autoinst-distrib-opensuse/commit/d72eece4eda301a2c42c9348fd147d97d13d9267>

12SP2@x86_64: <https://openqa.suse.de/tests/5730930>

12SP2@ppc64le: <https://openqa.suse.de/tests/5730960>

The failure in module `bang` is intentional. Successful reproduction would cause another failure either in `check_post_rollback` or `check_post_crash`. It doesn't matter whether the 3 check modules have `always_rollback` or not.

mkittler wrote:

Module AD043 failed and triggered VM rollback.

But this module is actually `tests/kernel/run_ltp.pm`, right? Maybe what makes your test special is that "AD043" is not a real test module? So autotest it actually still at "create_junkfile_ltp" which is the last module corresponding to a real Perl module (`tests/kernel/create_junkfile_ltp.pm`).

All modules after setup in the two jobs above are renamed but I still can't reproduce the issue. I guess that the Btrfs bug that causes those LTP failures is involved somehow, as if the disk got corrupted during `create_junkfile_ltp` and the corruption was saved into the snapshot. But in that case, I don't understand how so many tests pass on a corrupted filesystem. And if they pass only because the corruption is hidden by Btrfs structures cached in RAM, why restoring the snapshot doesn't hide the corruption again?

#16 - 2021-03-29 15:00 - mkittler

Thanks for the investigation. Too bad that the issue is not reproducible. Unfortunately I can not tell you whether your suspicion about btrfs is correct or not. It seems generally plausible but raises the question you've mentioned.

#17 - 2021-04-15 14:27 - cdywan

So unfortunately it seems like we still have no reproducer, although the issue still occurs in production, with and without `always_rollback`. To re-iterate what I think we can realistically do:

1. Re-evaluate workarounds i.e. not using snapshots here
2. Use unique identifiers and make logs more explicit
3. Validate snapshots explicitly

The first one would be up to [MDoucha](#), the other ones I'm thinking of straightforward approaches that don't change behavior for now but would help with identifying whatever the real problem is.

#18 - 2021-04-21 07:24 - cdywan

- *Priority changed from High to Low*
- *Target version changed from Ready to future*

Agreed with [MDoucha](#) to collect more details to find out if this is btrfs-specific. Hence putting this in low/future now

#19 - 2021-06-02 15:20 - MDoucha

- *Status changed from New to Rejected*

Ugh... After trying to reproduce the bug for more than a month and finally succeeding (sort of), it turns out that there are two different junkfiles: `/tmp/aiodio.$$junkfile` and `/tmp/aiodio/junkfile`. The latter is disappearing because it's created *after* the last snapshot. The names are similar enough that I thought all this time it was the same file...

Never mind, this ticket is invalid and I'll go fix out broken setup.