

openQA Project - action #88053

Improve logging in scripts (was: openQA in openQA test fails with cryptic error enumerating jobs)

2021-01-20 10:53 - cdywan

Status:	Workable	Start date:	2021-01-20
Priority:	Low	Due date:	
Assignee:		% Done:	0%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:	future		
Difficulty:			
Description			
<p>start_test appears to fail like so:</p> <pre>Test died: command 'job_id=\$(openqa-client --host http://openqa.opensuse.org --json-output jobs get version=Tumbleweed scope=relevant arch=x86_64 build=\$last_tw_build flavor=NET latest=1 jq '.jobs .[] select(.test == "minimalx") .id')' failed at /var/lib/openqa/cache/openqa1-opensuse/tests/openqa//tests/osautoinst/start_test.pm line 22</pre> <p>So wondering what this means, I attempt to reproduce quickly and get a long <i>list</i> of job id's after a <i>considerable</i> delay. So <code>\$job_id</code> is not a single id, but why would that fail ☐☐</p> <p>Upon inspection of the screenshots there's a more specific error message:</p> <pre>parse error: Invalid numeric literal at line 1, column 6</pre> <p>I still don't know what that means since the list I saw is of the form 12345 23456 45678. But, this should be visible w/o a screenshot.</p>			
Related issues:			
Related to openQA Project - action #82067: Unable to use openqa-clone-custom-...		Resolved	2021-01-08

History

#1 - 2021-01-20 13:56 - cdywan

So I tried to reproduce running the snippet from [start_test.pm in os-autoinst-distri-opensuse](#) thinking it's a quoting issue, because the test output won't seemingly expand `$last_tw_build` unlike `ttest`.

It's worth noting the almost identical code this test fails in is [start_test.pm in os-autoinst-distri-openQA](#) (the zypper command has an extra command guard).

So, I'm guessing there must be something in the environment, and not with the quoting of the embedded shellsript?

#2 - 2021-01-20 14:07 - okurz

- Category set to Concrete Bugs

- Status changed from New to Workable

- Target version set to Ready

First, in theory it should be expected that the openQA API route returns a list of ids. And the `jq` command is already written in a way that it would walk over multiple ids. The assignment to the bash variable would have list instead of a single job then. But the "parse error" sounds like `openqa-client --json-output` would not actually output json output.

Multiple things we could make out of this:

- Find out why `openqa-client --json-output` should output a plain id list (let's start with that)
- Fix the actual issue in <https://github.com/os-autoinst/os-autoinst-distri-openQA/>
- Suggest to use serial terminal by default for terminal tests
- Make sure serial terminal based tests still have a nice video (or at least exclude from the video)

#3 - 2021-01-20 14:55 - tinita

I think I found the bug:

https://openqa.opensuse.org/tests/1594919#step/start_test/15

As you can see, the output from openqa-client is:

```
Output: [INFO] Reading config from /etc/openqa/client.conf
{
  "jobs": [
  ...
```

#4 - 2021-01-20 15:00 - tinita

- Status changed from Workable to In Progress
- Assignee set to tinita

#5 - 2021-01-20 15:16 - cdywan

[tinita](#) Yikes. Good catch! That's an *info* level log, but somehow it's still printed there.

So that explains the "invalid numeral"

#6 - 2021-01-20 16:12 - tinita

Maybe it would be a good idea to make that output depend on if the command was called interactively. If it's piped or redirected to a file, like in this case, the output should not appear. There's a handy module for that: IO::Interactive.

#7 - 2021-01-20 16:46 - tinita

- Status changed from In Progress to Workable
- Assignee deleted (tinita)

#8 - 2021-01-21 16:36 - okurz

- Priority changed from Normal to Urgent

Geez, now I see the relation. It's a recent regression. Please fix urgently

#9 - 2021-01-21 17:00 - tinita

I created a PR to revert <https://github.com/os-autoinst/openQA/pull/3679> for now: <https://github.com/os-autoinst/openQA/pull/3688>

#10 - 2021-01-22 12:36 - okurz

- Priority changed from Urgent to Low
- Target version changed from Ready to future

revert was merged. That is making the issue less severe. The ideas we discussed were that we could improve the logging. Underlying libraries should not log directly but be introspectable, e.g. provide the current configuration to be read by applications. Applications should be quiet by default, only output error messages. Log messages should be written to stderr rather than stdout. Higher levels of verbosity could be enabled with a `-v/--verbose` command line parameter or allowing multiple levels like is done in <https://github.com/os-autoinst/openQA/blob/master/script/openqa-label-all#L27>

#11 - 2021-01-22 12:36 - okurz

- Subject changed from *openQA in openQA test fails with cryptic error enumerating jobs* to *Improve logging in scripts (was: openQA in openQA test fails with cryptic error enumerating jobs)*

#12 - 2021-01-22 13:09 - tinita

- Related to action #82067: *Unable to use openqa-clone-custom-git-refspec - Returns 403 Not authorized added*

#13 - 2021-01-22 13:29 - cdywan

<https://github.com/os-autoinst/openQA/pull/3684> In this PR I proposed the "obvious" fix which led to a more general discussion about logging, or not logging as it were. Closed it for now.