

openQA Project - action #16430

[tools] qcow2 image upload fails with no apparent reason and job ends up incomplete

2017-02-02 22:24 - okurz

Status:	Resolved	Start date:	2017-02-02
Priority:	Normal	Due date:	
Assignee:	coolo	% Done:	0%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:			
Difficulty:			
Description			
observation			
https://openqa.opensuse.org/tests/346076/file/autoinst-log.txt shows			
<pre>21:52:08.5284 2867 running nice ionice qemu-img convert -O qcow2 raid/l1 assets_public/opensuse-42 .3-x86_64-0063-gnome@64bit.qcow2 -c p11-kit: couldn't list directory: /etc/pkcs11/modules: Permission denied Failed to open module: /usr/lib64/qemu/block-iscsi.so: failed to map segment from shared object Failed to open module: /usr/lib64/qemu/block-ssh.so: failed to map segment from shared object 2867: EXIT 0 +++ worker notes +++ end time: 2017-02-02 21:56:48 result: done uploading consoletest_setup-loadavg_consoletest_setup.txt uploading consoletest_setup-psaxf.log uploading install_and_reboot-y2logs.tar.bz2 uploading autoinst.xml Checksum :1888125274 Sizes::58699 uploading video.ogv uploading vars.json uploading serial0.txt uploading autoinst-log.txt</pre>			
and then nothing.			
steps to reproduce			
A lot of job with same symptoms, e.g. take a look for all incompletes on https://openqa.opensuse.org/tests?hoursfresh=24&match=gnome-image but apparently only on openqaworker1			
problem			
H1. openqaworker1 has problems uploading the images H2. The other error messages in the log about "permission denied" and such are related to the problem			
suggestion			
proper log messages about this.			
workaround			
maybe retrigger jobs until any other worker than openqaworker1 is hit?			
Related issues:			
Related to openQA Project - action #17218: Can't publish HDD1: cksum: /var/li...		Resolved	2017-02-20
Related to openQA Project - action #19546: QCOW uploaded as asset won't appea...		Resolved	2017-06-02
Has duplicate openQA Tests - action #19670: [tools][sle][functional][aarch64]...		Rejected	2017-06-08

Has duplicate openQA Tests - action #19986: [sled] checksum failed when publi...
Copied to openQA Project - action #20786: [tools] image upload problems are n...

Rejected 2017-06-22
Closed 2017-07-26

History

#1 - 2017-02-05 22:04 - okurz

another example: <https://openqa.opensuse.org/tests/348306>

#2 - 2017-02-20 16:07 - szarate

- Related to action #17218: Can't publish HDD1: cksum:
/var/lib/openqa/share/factory/hdd/sle-12-SP3-x86_64-0248-textmode@64bit.qcow2.TEMP-IMZWihMn: Permission denied added

#3 - 2017-02-20 16:07 - szarate

poo#17218 might be related here

#4 - 2017-05-09 10:27 - coolo

- Status changed from New to Rejected

17218 is not related - and as this is not an often seen problem I close it as random issue in the infrastructure

#5 - 2017-06-02 12:53 - michalnowak

- Related to action #19546: QCOW uploaded as asset won't appear on NFS pool and Assets tab added

#6 - 2017-06-06 21:30 - okurz

- Status changed from Rejected to In Progress

- Priority changed from High to Low

- Target version set to future

<https://openqa.suse.de/tests/984049/file/autoinst-log.txt>

<https://openqa.suse.de/tests/983968/file/autoinst-log.txt>

two times in the same build, different arch, both times same test suite "cryptlvm". maybe the disk image is just very big in these cases:

```
uploading sle-12-SP3-Server-DVD-aarch64-gnome-encrypted.qcow2
Checksum comparison (actual:expected) 4294967295:4149226778 with size (actual:expected) 0:5163515904
Checksum/size comparison of sle-12-SP3-Server-DVD-aarch64-gnome-encrypted.qcow2 FAILED
```

```
uploading sle-12-SP3-Server-DVD-ppc64le-gnome-encrypted.qcow2
Checksum comparison (actual:expected) 4294967295:194332624 with size (actual:expected) 0:10527047680
Checksum/size comparison of sle-12-SP3-Server-DVD-ppc64le-gnome-encrypted.qcow2 FAILED
```

#7 - 2017-06-08 12:36 - vsvecova

Another example:

<https://openqa.suse.de/tests/987817/file/autoinst-log.txt>

```
uploading SLED-12-SP2-20170607-2-x86_64_for_regression.qcow2
Checksum comparison (actual:expected) 4294967295:1762304692 with size (actual:expected) 0:3093506560
Checksum/size comparison of SLED-12-SP2-20170607-2-x86_64_for_regression.qcow2 FAILED
```

Restarting does not help, yields the same result:

<https://openqa.suse.de/tests/987993>

Yet another:

<https://openqa.suse.de/tests/988271/file/autoinst-log.txt>

```
uploading SLED-12-SP2-20170608-1-x86_64_for_regression.qcow2
Checksum comparison (actual:expected) 4294967295:4195181595 with size (actual:expected) 0:3076052992
Checksum/size comparison of SLED-12-SP2-20170608-1-x86_64_for_regression.qcow2 FAILED
```

#8 - 2017-06-08 15:30 - okurz

- Subject changed from qcow2 image upload fails with no apparent reason and job ends up incomplete to [tools]qcow2 image upload fails with no apparent reason and job ends up incomplete

- Priority changed from Low to Normal

- Target version deleted (future)

ok, so it's a thing. How about we further improve our logging? Someone has suggestions? E.g. when size is zero, what should we output? Or check logfiles on the server?

#9 - 2017-06-08 15:38 - okurz

- Has duplicate action #19670: [tools][sle][functional][aarch64] test failed to upload disk image added

#10 - 2017-06-08 15:38 - okurz

Happened in <https://openqa.suse.de/tests/987904> as well (see [#19670](#))

#11 - 2017-06-20 13:48 - vsvecova

Today's example:

<https://openqa.suse.de/tests/1011148>

```
uploading SLED-12-SP2-20170620-6-x86_64_for_regression.qcow2
Checksum comparison (actual:expected) 4294967295:3593870896 with size (actual:expected) 0:3106909696
Checksum/size comparison of SLED-12-SP2-20170620-6-x86_64_for_regression.qcow2 FAILED
```

#12 - 2017-06-22 06:53 - okurz

- Has duplicate action #19986: [sled] checksum failed when publishing cryptlvm hdd images added

#13 - 2017-06-22 10:57 - dzedro

Build 0437 failure <https://openqa.suse.de/tests/1016254>

#14 - 2017-06-23 10:59 - riafarov

Another example here: <https://openqa.suse.de/tests/1018503#>

#15 - 2017-06-27 08:51 - coolo

- Subject changed from [tools]qcow2 image upload fails with no apparent reason and job ends up incomplete to [tools] qcow2 image upload fails with no apparent reason and job ends up incomplete

the last examples (also as of today) are all about rewriting existant files. For some reason all the cryptlvm hdds do not have the build number in the filename. The examples in [#7](#) have though

#16 - 2017-06-27 09:10 - coolo

- Assignee set to kraih

Debugging <https://openqa.suse.de/tests/1024024> ...

From autostart log:

```
uploading textinfo-logs.tar.bz2
uploading sle-12-SP3-Desktop-DVD-x86_64-gnome-encrypted.qcow2
Checksum comparison (actual:expected) 4294967295:4170063198 with size (actual:expected) 0:7906116096
Checksum/size comparison of sle-12-SP3-Desktop-DVD-x86_64-gnome-encrypted.qcow2 FAILED
uploading video.ogv
```

The upload happened around 05:09 - last successful log:

```
[Tue Jun 27 05:09:49 2017] [17003:debug] moved to
/var/lib/openqa/testresults/01024/01024024-sle-12-SP3-Desktop-DVD-x86_64-Build0274-cryptlvm@64bit/ulogs textinfo-logs.tar.bz2
[Tue Jun 27 05:09:49 2017] [17003:debug] 200 OK (0.009422s, 106.135/s)
```

Next successful log:

```
[Tue Jun 27 05:11:37 2017] [17469:debug] moved to
/var/lib/openqa/testresults/01024/01024024-sle-12-SP3-Desktop-DVD-x86_64-Build0274-cryptlvm@64bit video.ogv
[Tue Jun 27 05:11:37 2017] [17469:debug] 200 OK (0.01602s, 62.422/s)
```

In these 2 minutes we saw 1248 POSTs - so it was a quite busy time. But what is worrying me most: we saw 112 worker restarts due to rss limit. So I'm afraid these restarts happen too early - or too late.

This *might* also be related to the random problems we see on cavil. For reference: this is the restart code: <https://github.com/os-autoinst/openQA/pull/1339>

#17 - 2017-06-27 09:16 - coolo

I raised the rss limit to 580.000 and restarted the job. Let's see

#18 - 2017-06-27 10:28 - coolo

the upload worked now

#19 - 2017-06-27 10:29 - coolo

So now going to insane low values

#20 - 2017-06-27 10:30 - coolo

- Assignee changed from kraith to coolo

#21 - 2017-06-27 11:11 - coolo

this is from apache log about the above:

..

```
[Tue Jun 27 05:00:02.042426 2017] [proxy_http:error] pid_14598:Connection reset by peer: [client 2620:113:80c0:8080:10:160:1:100:47638] AH01102: error reading status line from remote server localhost:9526  
[Tue Jun 27 05:00:02.042465 2017] [proxy:error] [pid 14598] [client 2620:113:80c0:8080:10:160:1:100:47638] AH00898: Error reading from remote server returned by /api/v1/jobs/1024024/artefact
```

#22 - 2017-07-03 15:28 - coolo

My latest findings for job 1039448 support my theory it's a mojo bug. I added debug for upload progress and what we see is a working having a running transaction and restarting due to RSS limit:

```
17:25 <coolo> [Mon Jul 3 16:59:28 2017] [22634:debug] Worker 22634 has accepted the request for /api/v1/jobs/1039448/artefact  
17:25 <coolo> [Mon Jul 3 16:59:28 2017] [22634:debug] Worker exceeded RSS limit "219232 > 200000", restarting  
17:25 <coolo> [Mon Jul 3 16:59:28 2017] [24509:debug] Stopping worker 22634  
17:25 <coolo> [Mon Jul 3 16:59:28 2017] [24509:debug] Worker 22634 stopped
```

The graceful restart restarted by the RSS limit is supposed to finish running transactions first

#23 - 2017-07-04 02:11 - qkzhu

failed at: sle-12-SP3-Desktop-DVD-x86_64-Build0281-regression-installation@64bit <https://openqa.suse.de/tests/1039589>

#24 - 2017-07-04 08:55 - coolo

ok, we know what it is. The mojo worker is restarted when it hits certain limits. But this restart is 'graceful', so it first finishes the current requests.

But the timeout for that is by default 20 seconds - and uploading a 8G file can take longer. So if the restart trigger is hit right after accepting the upload, chances are you don't finish in these 20 seconds and we fail.

I will set this limit to 800, but we will need to face workers that are in long timeouts and do not accept requests anymore. So possibly increase their number as well.

#25 - 2017-07-05 08:26 - kraith

To help with the problem of too many workers not accepting connections because they are shutting down, i've added a [new feature](#) to Mojolicious. The web server can now temporarily spawn a few additional workers as soon as an existing worker begins the graceful shutdown process. The default is for up to two spare workers to be spawned, but that can be tweaked with the -s flag.

#26 - 2017-07-06 09:23 - coolo

- Status changed from In Progress to Resolved

<https://github.com/os-autoinst/openQA/pull/1385> is hotpatched on osd

#27 - 2017-07-26 08:40 - okurz

- Status changed from Resolved to In Progress

<https://openqa.suse.de/tests/1076847/file/autoinst-log.txt> looks exactly like that problem. Could it be a different problem even though the symptoms

show up to be the same? How can we be sure? Some worker logs to check here? Or the logs from the web ui host at the same time?

#28 - 2017-07-26 09:35 - okurz

Also <https://openqa.suse.de/tests/1077066> and <https://openqa.suse.de/tests/1077199> failed, reproducible problem, different workers.

#29 - 2017-07-26 09:46 - coolo

It's a different problem though. The upload is done, but the rename fails.

```
cooloo@openqa:/var/lib/openqa/factory/hdd> ls -l SLES-12-SP3-aarch64-0472-minimal_with_sdk0262_installed.qcow2*
-rw-r--r-- 1 geekotest nogroup 916783104  8. Jul 17:16 SLES-12-SP3-aarch64-0472-minimal_with_sdk0262_installed
.qcow2
-rw-r--r-- 1 geekotest nogroup 921108480 26. Jul 09:59 SLES-12-SP3-aarch64-0472-minimal_with_sdk0262_installed
.qcow2.TEMP-0YQHK5WS
-rw-r--r-- 1 geekotest nogroup 921567232 26. Jul 11:11 SLES-12-SP3-aarch64-0472-minimal_with_sdk0262_installed
.qcow2.TEMP-tKieeJTx
-rw-r--r-- 1 geekotest nogroup 914882560 26. Jul 11:20 SLES-12-SP3-aarch64-0472-minimal_with_sdk0262_installed
.qcow2.TEMP-tzKPM_MJ
```

```
[Wed Jul 26 11:20:31 2017] [4779:debug] Routing to controller "OpenQA::WebAPI::Controller::API::V1::Job" and a
ction "create_artefact"
[Wed Jul 26 11:20:31 2017] [4779:debug] moved to /var/lib/openqa/share/factory/hdd/SLES-12-SP3-aarch64-0472-mi
nimal_with_sdk0262_installed.qcow2.TEMP-tzKPM_MJ
[Wed Jul 26 11:20:31 2017] [4779:error] DBIx::Class::Storage::DBI::_dbh_execute(): DBI Exception: DBD::Pg::st
execute failed: ERROR: duplicate key value violates unique constraint "jobs_
assets_job_id_asset_id"
DETAIL:  Key (job_id, asset_id)=(1077199, 58485) already exists. [for Statement "INSERT INTO jobs_assets ( ass
et_id, created_by, job_id, t_created, t_updated) VALUES ( ?, ?, ?, ?, ? )" wi
th ParamValues: 1='58485', 2='1', 3='1077199', 4='2017-07-26 09:20:31', 5='2017-07-26 09:20:31'] at /usr/share
/openqa/script/./lib/OpenQA/Schema/Result/Jobs.pm line 1084

[Wed Jul 26 11:20:31 2017] [4779:debug] 500 Internal Server Error (0.026118s, 38.288/s)
```

The underlying problem is still that the upload failures aren't reported properly

#30 - 2017-07-26 09:47 - coolo

what kind of test is this with HDD_1 and PUBLISH_HDD_1 being the same value?

#31 - 2017-07-26 09:47 - coolo

- Status changed from In Progress to Resolved

this is something else and your reopen is invalid

#32 - 2017-07-26 09:54 - okurz

- Copied to action #20786: [tools] image upload problems are not reported properly added