

openQA Project - action #88603

Two identical jobs are created from one abandoned test

2021-02-15 12:15 - dzedro

Status:	Resolved	Start date:	2021-02-15
Priority:	High	Due date:	
Assignee:	mkittler	% Done:	0%
Category:	Concrete Bugs	Estimated time:	0.00 hour
Target version:	Ready		
Difficulty:			
Description			
When job is abandoned Reason: abandoned: associated worker openqaworker* re-connected but abandoned the job, then sometimes two new identical jobs are created. This is problem when affected job is publishing qcow2 image [1]. Two jobs write same image and the qcow2 image is corrupted and all following images are failing [2].			
[1] https://openqa.suse.de/tests/5464337 https://openqa.suse.de/tests/5465190			
[2] https://openqa.suse.de/tests/5465200#dependencies			

History

#1 - 2021-02-15 12:17 - dzedro

- Subject changed from *Two parallel jobs are created from one abandoned test* to *Two identical jobs are created from one abandoned test*

#2 - 2021-02-15 12:29 - okurz

- Target version set to *Ready*

#3 - 2021-02-15 12:30 - okurz

- Project changed from *openQA Infrastructure* to *openQA Project*

- Category set to *Concrete Bugs*

Actually changing the project to "openQA Project" which raises awareness and also due to the reason that I would not know how any "infrastructure specifics" could cause this.

#4 - 2021-02-15 12:37 - dzedro

Thanks, I don't know what is causing this. But one of the new jobs does look like restart of the abandoned job with all dependencies, the second one does look like clone with --skip-deps.

#5 - 2021-02-15 14:50 - mkittler

- Assignee set to *mkittler*

1. <https://openqa.suse.de/tests/5464337> is the original job which was responsible to create the qcow image

SLES-15-SP2-x86_64-mru-install-desktop-with-addons-Build20210214-1.qcow2 but incompleting with Reason: abandoned: associated worker openqaworker9:15 re-connected but abandoned the job.

1. Hence it has been cloned as <https://openqa.suse.de/tests/5465089> but this job got cancelled with Reason: quit: worker has been stopped or restarted.

2. Hence it has been cloned once more as <https://openqa.suse.de/tests/5465870> and this job finally passed and has not been further cloned.

2. Everything which happened in the previous point works as intended. However, there's also the "additional" job

<https://openqa.suse.de/tests/5465904> which uploaded the same qcow image and caused the problem. Now the question is where that job comes from.

1. It is a clone of <https://openqa.suse.de/tests/5465094> which is not itself a clone and has also not been scheduled as part of an ISO. This raises the question where this job comes from. I see 2 options:

1. 5465094 was actually cloned or scheduled as part of an ISO but due to some bug the information about that is missing within the database.

2. 5465094 was created via the jobs post API. In this case it was likely the mistake of a user (or script) using that API.

#6 - 2021-02-15 14:55 - mkittler

Looks like 5465094 was cancelled by [dzedro](#) and then "auto duplicated" by the worker: <https://openqa.suse.de/admin/auditlog?eventid=6111495>, <https://openqa.suse.de/admin/auditlog?eventid=6111541>

So the offending job was manually cancelled at 2021-02-14T22:15:02. According to the worker log the worker was still preparing the job execution at this time:

```
Feb 14 05:35:33 openqaworker9 systemd[1]: Starting openQA Worker #15...
Feb 14 05:35:33 openqaworker9 systemd[1]: Started openQA Worker #15.
Feb 14 05:35:34 openqaworker9 worker[12131]: [info] [pid:12131] worker 15:
Feb 14 05:35:34 openqaworker9 worker[12131]: - config file: /etc/openqa/workers.ini
Feb 14 05:35:34 openqaworker9 worker[12131]: - worker hostname: openqaworker9
Feb 14 05:35:34 openqaworker9 worker[12131]: - isotovideo version: 20
Feb 14 05:35:34 openqaworker9 worker[12131]: - websocket API version: 1
Feb 14 05:35:34 openqaworker9 worker[12131]: - web UI hosts: openqa.suse.de
Feb 14 05:35:34 openqaworker9 worker[12131]: - class: caasp_x86_64,tap,qemu_x86_64,openqaworker9
Feb 14 05:35:34 openqaworker9 worker[12131]: - no cleanup: no
Feb 14 05:35:34 openqaworker9 worker[12131]: - pool directory: /var/lib/openqa/pool/15
Feb 14 05:35:34 openqaworker9 worker[12131]: [error] [pid:12131] Unable to lock pool directory: /var/lib/openqa/pool/15 already locked
Feb 14 05:35:34 openqaworker9 worker[12131]: [info] [pid:12131] CACHE: caching is enabled, setting up /var/lib/openqa/cache/openqa.suse.de
Feb 14 05:35:34 openqaworker9 worker[12131]: [info] [pid:12131] Project dir for host openqa.suse.de is /var/lib/openqa/share
Feb 14 05:35:34 openqaworker9 worker[12131]: [info] [pid:12131] Registering with openQA openqa.suse.de
Feb 14 05:35:34 openqaworker9 worker[12131]: [info] [pid:12131] Establishing ws connection via ws://openqa.suse.de/api/v1/ws/599
Feb 14 05:35:34 openqaworker9 worker[12131]: [info] [pid:12131] Registered and connected via websockets with openQA host openqa.suse.de and worker ID 599
Feb 14 05:35:53 openqaworker9 worker[12131]: [debug] [pid:12131] Updating status so job 5465094 is not considered dead.
Feb 14 05:35:53 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
...
Feb 14 23:14:58 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
Feb 14 23:15:03 openqaworker9 worker[12131]: [debug] [pid:12131] Updating status so job 5465094 is not considered dead.
Feb 14 23:15:03 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
Feb 14 23:15:08 openqaworker9 worker[12131]: [debug] [pid:12131] Updating status so job 5465094 is not considered dead.
Feb 14 23:15:08 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
Feb 14 23:15:13 openqaworker9 worker[12131]: [debug] [pid:12131] Updating status so job 5465094 is not considered dead.
Feb 14 23:15:13 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
...
Feb 14 23:23:37 openqaworker9 worker[12131]: [debug] [pid:12131] Updating status so job 5465094 is not considered dead.
Feb 14 23:23:37 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
Feb 14 23:23:42 openqaworker9 worker[12131]: [debug] [pid:12131] Updating status so job 5465094 is not considered dead.
Feb 14 23:23:42 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
Feb 14 23:23:44 openqaworker9 systemd[1]: Stopping openQA Worker #15...
Feb 14 23:23:44 openqaworker9 worker[12131]: [info] [pid:12131] Received signal TERM
Feb 14 23:23:44 openqaworker9 worker[12131]: [debug] [pid:12131] Stopping job 5465094 from openqa.suse.de: 05465094-sle-15-SP2-Server-DVD-Updates-x86_64-Build20210214-1-mru-install-desktop-with-addons@64bit - reason: quit
Feb 14 23:23:44 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/status
Feb 14 23:23:45 openqaworker9 worker[12131]: [info] [pid:12131] +++ worker notes +++
Feb 14 23:23:45 openqaworker9 worker[12131]: [info] [pid:12131] End time: 2021-02-14 22:23:45
Feb 14 23:23:45 openqaworker9 worker[12131]: [info] [pid:12131] Result: quit
Feb 14 23:23:45 openqaworker9 worker[12131]: [debug] [pid:12131] Job 5465094 stopped as incomplete
Feb 14 23:23:45 openqaworker9 worker[12131]: [debug] [pid:12131] Duplicating job 5465094
Feb 14 23:23:45 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/duplicate?dup_type_auto=1
Feb 14 23:23:45 openqaworker9 worker[12131]: [warn] [pid:12131] Unable to upload results of the job because no command server URL or worker ID have been set.
Feb 14 23:23:45 openqaworker9 worker[12131]: [debug] [pid:12131] REST-API call: POST http://openqa.suse.de/api/v1/jobs/5465094/set_done?reason=quit%3A+worker+has+been+stopped+or+restarted&result=incomplete&worker_id=599
Feb 14 23:23:45 openqaworker9 worker[12131]: [debug] [pid:12131] Job 5465094 from openqa.suse.de finished - reason: quit
Feb 14 23:23:45 openqaworker9 worker[12131]: [debug] [pid:12131] Informing openqa.suse.de that we are going of
```

```
fline
Feb 14 23:23:45 openqaworker9 systemd[1]: Stopped openQA Worker #15.
Feb 14 23:35:32 openqaworker9 systemd[1]: Starting openQA Worker #15...
```

Judging by the timestamps the cancellation attempt was done because the job looked stuck. Unfortunately it didn't help because the worker can not react to cancellation though the API in that phase (which is something we could improve). However, the worker service itself has been restarted a few minutes later anyways (maybe manually because the cancellation attempt over the API didn't help). That's the reason why the job ended with the result `user_cancelled` (remarked by the web UI due to the cancellation attempt through the API) and Reason: quit: worker has been stopped or restarted (remarked by the worker after cancelling the job after receiving SIGTERM). That the worker duplicated the job is how it is supposed to work. (When shutting down a worker we want another worker to take over which means starting the job from scratch and marking the previous job as incomplete.)

But I guess that 5465094 was cloned resulting in 5465904 is not really the cause of the problem here. If the original job 5465094 could have been executed as expected it would have uploaded the qcow image as well. So this leaves still the question where the offending job 5465094 comes from. By the way, the other "clone chain" starting with <https://openqa.suse.de/tests/5464337> has been scheduled as part of "sle-15-SP2-Server-DVD-Updates-x86_64-20210214-1" so here it is clear where it comes from.

[dzedro](#) Do you have any idea where 5465094 comes from? I've checked the openQA database and there are no dependencies of that job. (Even though the dependencies tab is not displayed for a cloned job the database would still have the dependencies.) That let's me think option 2.1.2 from my last comment is the more likely one. That means we simply have two jobs scheduled (for whatever reason) which uploaded the same qcow image. openQA's restarting/cloning logic doesn't seem at fault here.

#7 - 2021-02-15 15:27 - mkittler

- Status changed from New to Feedback

#8 - 2021-02-15 15:44 - mkittler

Note that there are really no other traces of 5465094 besides the mentioned audit events (cancellation and restart). So it is hard to tell from my side where it comes from. The fact that the job does *not* occur in any other restart audit events makes it look like there was no interference with the other clone chain.

Besides the problem mentioned in the ticket description I found two other problems:

1. It is bad that the worker keeps waiting for the cache service forever. The job 5465094 ran for almost 20 hours and was only stopped because the worker received SIGTERM.
2. It is bad that the worker doesn't react to REST-API calls, e.g. for cancelling, while waiting for the cache service.

#9 - 2021-02-16 09:47 - dzedro

I have no idea where 5465094 comes from. It's hard to track as there is nothing in audit log, it's fresh new job, who knows what triggered it. I have been restarting cache and worker services on openqaworker9 because there were jobs hanging for approx. 18 hours. I will update all new failures of this kind when found, there were tens of such failures on arm e.g. <https://openqa.suse.de/tests/5447958> or <https://openqa.suse.de/tests/5367195>

#10 - 2021-02-16 11:04 - mkittler

Regarding the cache service problem: It looks like openqaworker-cacheservice-minion didn't execute any jobs in the timeframe the job was stuck but there are also no errors:

```
Feb 14 05:38:38 openqaworker9 openqa-worker-cacheservice-minion[3760]: [12351] [i] [#164] Downloading: "SLE-15-SP2-Installer-DVD-x86_64-GM-DVD1.iso"
Feb 14 05:38:38 openqaworker9 openqa-worker-cacheservice-minion[3760]: [12351] [i] [#164] Cache size of "/var/lib/openqa/cache" is 34GiB, with limit 50GiB
Feb 14 05:38:38 openqaworker9 openqa-worker-cacheservice-minion[3760]: [12351] [i] [#164] Downloading "SLE-15-SP2-Installer-DVD-x86_64-GM-DVD1.iso" from "http://openqa.suse.de/tests/5465121/asset/iso/SLE-15-SP2-Installer-DVD-x86_64-GM-DVD1.iso"
Feb 14 05:38:38 openqaworker9 openqa-worker-cacheservice-minion[3760]: [12351] [i] [#164] Content of "/var/lib/openqa/cache/openqa.suse.de/SLE-15-SP2-Installer-DVD-x86_64-GM-DVD1.iso" has not changed, updating last use
Feb 14 05:38:38 openqaworker9 openqa-worker-cacheservice-minion[3760]: [12351] [i] [#164] Finished download
Feb 14 23:23:07 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4095] [i] [#166] Downloading: "SLE-15-SP1-x86_64-Build20210214-1-qam-sles4sap-gnome.qcow2"
Feb 14 23:23:07 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4095] [i] [#166] Cache size of "/var/lib/openqa/cache" is 34GiB, with limit 50GiB
Feb 14 23:23:07 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4095] [i] [#166] Downloading "SLE-15-SP1-x86_64-Build20210214-1-qam-sles4sap-gnome.qcow2" from "http://openqa.suse.de/tests/5465203/asset/hdd/SLE-15-SP1-x86_64-Build20210214-1-qam-sles4sap-gnome.qcow2"
Feb 14 23:23:43 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4165] [i] [#168] Downloading: "SLES-15-SP1-x86_64-mru-install-minimal-with-addons-Build20210214-1-Server-DVD-Updates-64bit.qcow2"
Feb 14 23:23:43 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4165] [i] [#168] Cache size of "/var/lib/openqa/cache" is 34GiB, with limit 50GiB
Feb 14 23:23:43 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4165] [i] [#168] Downloading "SLES-15-SP1-x86_64-mru-install-minimal-with-addons-Build20210214-1-Server-DVD-Updates-64bit.qcow2" from "http://openqa.suse.de/tests/5465903/asset/hdd/SLES-15-SP1-x86_64-mru-install-minimal-with-addons-Build20210214-1-Server-DVD-Updates-64bit.qcow2"
```

```
Feb 14 23:23:43 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4165] [i] [#168] Content of "/var/lib/openqa/cache/openqa.suse.de/SLES-15-SP1-x86_64-mru-install-minimal-with-addons-Build20210214-1-Server-DVD-Updates-64bit.qcow2" has not changed, updating last use
Feb 14 23:23:43 openqaworker9 openqa-worker-cacheservice-minion[3760]: [4165] [i] [#168] Finished download
```

It looks like this failed SQL migration was the problem:

```
Feb 14 03:35:03 openqaworker9 systemd[1]: Started OpenQA Worker Cache Service.
Feb 14 03:35:04 openqaworker9 openqa-workercache-daemon[3115]: [3115] [i] [NaC23h9F] Creating cache directory tree for "/var/lib/openqa/cache"
Feb 14 03:35:04 openqaworker9 openqa-workercache-daemon[3115]: DBD::SQLite::st execute failed: UNIQUE constraint failed: mojo_migrations.name at /usr/share/openqa/script/./lib/OpenQA/CacheService/Model/Cache.pm line 87.
Feb 14 03:35:04 openqaworker9 openqa-workercache-daemon[3115]: at /usr/share/openqa/script/./lib/OpenQA/CacheService/Model/Cache.pm line 87.
Feb 14 03:35:04 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Main process exited, code=exited, status=2/INVALIDARGUMENT
Feb 14 03:35:04 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Unit entered failed state.
Feb 14 03:35:04 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Failed with result 'exit-code'.
Feb 14 03:35:04 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Service RestartSec=100ms expired, scheduling restart.
Feb 14 03:35:04 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Service RestartSec=100ms expired, scheduling restart.
Feb 14 03:35:05 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Service RestartSec=100ms expired, scheduling restart.
Feb 14 03:35:05 openqaworker9 systemd[1]: openqa-worker-cacheservice.service: Service RestartSec=100ms expired, scheduling restart.
Feb 14 03:35:09 openqaworker9 systemd[1]: Stopped OpenQA Worker Cache Service.
Feb 14 03:35:09 openqaworker9 systemd[1]: Started OpenQA Worker Cache Service.
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] [zc6Q80gh] Cache size of "/var/lib/openqa/cache" is 0 Byte, with limit 50GiB
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Listening at "http://127.0.0.1:9530"
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Listening at "http://[::1]:9530"
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Manager 3758 started
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: Web application available at http://127.0.0.1:9530
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: Web application available at http://[::1]:9530
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3767] [i] Worker 3767 started
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3768] [i] Worker 3768 started
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3769] [i] Worker 3769 started
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Creating process id file "/tmp/prefork.pid"
Feb 14 03:35:09 openqaworker9 openqa-workercache-daemon[3758]: [3770] [i] Worker 3770 started
Feb 14 05:05:16 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Stopping worker 3769 gracefully (80 seconds)
Feb 14 05:05:16 openqaworker9 openqa-workercache-daemon[3758]: [10598] [i] Worker 10598 started
Feb 14 05:05:19 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Worker 3769 stopped
Feb 14 05:48:31 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Stopping worker 3767 gracefully (80 seconds)
Feb 14 05:48:31 openqaworker9 openqa-workercache-daemon[3758]: [13126] [i] Worker 13126 started
Feb 14 05:48:32 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Worker 3767 stopped
Feb 14 05:54:46 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Stopping worker 3770 gracefully (80 seconds)
Feb 14 05:54:46 openqaworker9 openqa-workercache-daemon[3758]: [13396] [i] Worker 13396 started
Feb 14 05:54:47 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Worker 3770 stopped
Feb 14 06:05:01 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Stopping worker 3768 gracefully (80 seconds)
Feb 14 06:05:01 openqaworker9 openqa-workercache-daemon[3758]: [13948] [i] Worker 13948 started
Feb 14 06:05:04 openqaworker9 openqa-workercache-daemon[3758]: [3758] [i] Worker 3768 stopped
```

Considering the Cache size of "/var/lib/openqa/cache" is 0 Byte, with limit 50GiB it looks like the cache service started with a fresh database after the SQL error. Since the openqaworker-cacheservice-minion service has not been restarted as well it became likely decoupled from openqaworker-cacheservice as it still had a handle to the old database. Hence no jobs were executed anymore.

Judging by the log the code we've added when trying to fix SQLite issues the last time did not kick in (for the better or worse). ~~Apparently the migration within repair_database worked but the migration done later within init worked not~~ That's because the migration error happened when there was no existing database before and hence the repair_database function returned early (talking about functions within openQA/lib/OpenQA/CacheService/Model/Cache.pm). So this is quite different from the previous SQLite problems (where an existing database was corrupted). [krah](#) noted that it might be a Mojo::SQLite::Migrations bug.

#11 - 2021-03-12 13:01 - mkittler

I'm not sure how to proceed with that. Have you encountered further occurrences in the meantime?

#12 - 2021-03-12 13:20 - dzedro

I have not seen this issue since last comment.

#13 - 2021-03-17 13:29 - mkittler

- *Status changed from Feedback to Resolved*

Then I'd close the ticket for now because I don't know how to proceed/reproduce. Feel free to re-open if you encounter the problem again.

Files

Screenshot_2021-02-15_13-03-39.png	13.1 KB	2021-02-15	dzedro
Screenshot_2021-02-15_13-11-15.png	32.9 KB	2021-02-15	dzedro