

openQA Infrastructure - action #93138

openqa-reload-worker-auto-restart and openqa-worker-auto-restart fail if numofworkers is reduced

2021-05-26 12:35 - nicksinger

Status:	New	Start date:	2021-05-26
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	future		

Description

With

https://gitlab.suse.de/openqa/salt-pillars-openqa/-/merge_requests/320/diffs#0a0ec388f05cea3ca9a77dd911a33555ec192597_38_38mgiessmeier reduced the numofworkers for openqaworker2. Soon after we received an alert that several services on this machine fail. After closer investigation I see the following services failing:

- openqa-reload-worker-auto-restart@36.service loaded failed failed Restarts openqa-worker-auto-restart@36.service as soon as possible without interrupting jobs
- openqa-reload-worker-auto-restart@37.service loaded failed failed Restarts openqa-worker-auto-restart@37.service as soon as possible without interrupting jobs
- openqa-reload-worker-auto-restart@38.service loaded failed failed Restarts openqa-worker-auto-restart@38.service as soon as possible without interrupting jobs
- openqa-worker-auto-restart@37.service loaded failed failed openQA Worker #37

(Note: I don't understand why there is no openqa-worker-auto-restart@36.service nor no openqa-worker-auto-restart@38.service - is this a setup fail?)

These seem to be the related services which numofworkers got reduced by.

Example output for one of these services:

```
openqaworker2:~ # systemctl status openqa-reload-worker-auto-restart@36.service
```

- openqa-reload-worker-auto-restart@36.service - Restarts openqa-worker-auto-restart@36.service as soon as possible without interrupting jobs
Loaded: loaded (/usr/lib/systemd/system/openqa-reload-worker-auto-restart@.service; static; vendor preset: disabled)
Active: failed (Result: exit-code) since Wed 2021-05-26 13:44:13 CEST; 49min ago
Process: 21839 ExecStart=/usr/bin/systemctl reload openqa-worker-auto-restart@36.service (code=exited, status=1/FAILURE)
Main PID: 21839 (code=exited, status=1/FAILURE)

```
May 26 13:44:12 openqaworker2 systemd[1]: Starting Restarts openqa-worker-auto-restart@36.service as soon as possible without interrupting jobs...
May 26 13:44:13 openqaworker2 systemctl[21839]: Job for openqa-worker-auto-restart@36.service canceled.
May 26 13:44:13 openqaworker2 systemd[1]: openqa-reload-worker-auto-restart@36.service: Main process exited, code=exited, status=1/FAILURE
May 26 13:44:13 openqaworker2 systemd[1]: Failed to start Restarts openqa-worker-auto-restart@36.service as soon as possible without interrupting jobs.
May 26 13:44:13 openqaworker2 systemd[1]: openqa-reload-worker-auto-restart@36.service: Unit entered failed state.
May 26 13:44:13 openqaworker2 systemd[1]: openqa-reload-worker-auto-restart@36.service: Failed with result 'exit-code'.
```

and

```
openqaworker2:~ # systemctl status openqa-worker-auto-restart@37.service
```

- openqa-worker-auto-restart@37.service - openQA Worker #37
Loaded: loaded (/usr/lib/systemd/system/openqa-worker-auto-restart@.service; disabled; vendor preset: disabled)
Drop-In: /etc/systemd/system/openqa-worker-auto-restart@.service.d
└─20-nvme-autoformat.conf
Active: failed (Result: timeout) since Wed 2021-05-26 13:45:43 CEST; 48min ago
Process: 18003 ExecStart=/usr/share/openqa/script/worker --instance 37 (code=killed, signal=KILL)

```

)
Main PID: 18003 (code=killed, signal=KILL)

May 26 13:45:43 openqaworker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/ImageMagick-116.png
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: State 'stop-sigterm' timed out. Killing.
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 18003 (worker) with signal SIGKILL.
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 22105 (worker) with signal SIGKILL.
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 23223 (optipng) with signal SIGKILL.
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Main process exited, code=killed, status=9/KILL
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 22105 (worker) with signal SIGKILL.
May 26 13:45:43 openqaworker2 systemd[1]: Stopped openQA Worker #37.
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Unit entered failed state.
May 26 13:45:43 openqaworker2 systemd[1]: openqa-worker-auto-restart@37.service: Failed with result 'timeout'.

```

These services should handle a reduce of numofworkers more gracefully.

Related issues:

Related to openQA Project - action #62441: openqa-worker systemd service can ...	New	2020-01-21
--	-----	------------

History

#1 - 2021-05-26 12:39 - nicksinger

Same for openqaworker5:

● openqa-reload-worker-auto-restart@42.service	loaded failed failed	Restarts openqa-worker-auto-restart@42.service as soon as possible without interrupting jobs
● openqa-reload-worker-auto-restart@43.service	loaded failed failed	Restarts openqa-worker-auto-restart@43.service as soon as possible without interrupting jobs
● openqa-reload-worker-auto-restart@44.service	loaded failed failed	Restarts openqa-worker-auto-restart@44.service as soon as possible without interrupting jobs
● openqa-reload-worker-auto-restart@45.service	loaded failed failed	Restarts openqa-worker-auto-restart@45.service as soon as possible without interrupting jobs

openqa-worker-auto-restart@{42..45}.service seem to be missing all together on that machine. So is it already covered that numofworkers gets reduced and we see a race of some kind here?

#2 - 2021-05-26 12:41 - nicksinger

I've reseted the failed units on openqaworker2 and openqaworker5, alert is good again.

#3 - 2021-05-26 12:42 - nicksinger

- Related to action #62441: openqa-worker systemd service can timeout when stopping added

#4 - 2021-05-26 12:54 - cdywan

- Status changed from New to In Progress

I assume this should be *In Progress*?

nicksinger wrote:

I've reseted the failed units on openqaworker2 and openqaworker5, alert is good again.

Yeah, I just saw the OK

#5 - 2021-05-26 12:56 - nicksinger

- Status changed from In Progress to Workable

No, at max in "Workable". I do not plan to work on the service file fix. I just used the same ticket here to resolve the immediate alert

#6 - 2021-05-27 14:55 - okurz

- Status changed from Workable to New

- Target version set to Ready

#7 - 2021-05-31 14:41 - mkittler

Looks like the worker was really just uploading remaining results, e.g.

```
Mai 26 13:45:40 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/ImageMagick-130.png
Mai 26 13:45:41 openqa-worker2 worker[18003]: [debug] [pid:22105] Uploading artefact ImageMagick-130.png as 6d75de6c7c96e3d8202993f5d49a0c0c
Mai 26 13:45:41 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/.thumbs/ImageMagick-130.png
Mai 26 13:45:41 openqa-worker2 worker[18003]: [debug] [pid:22105] Uploading artefact ImageMagick-130.png as 6d75de6c7c96e3d8202993f5d49a0c0c
Mai 26 13:45:41 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/ImageMagick-85.png
Mai 26 13:45:42 openqa-worker2 worker[18003]: [debug] [pid:22105] Uploading artefact ImageMagick-85.png as 5033d3574a8ee53a78d6e108c20c48
Mai 26 13:45:42 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/.thumbs/ImageMagick-85.png
Mai 26 13:45:42 openqa-worker2 worker[18003]: [debug] [pid:22105] Uploading artefact ImageMagick-85.png as 5033d3574a8ee53a78d6e108c20c48
Mai 26 13:45:42 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/ImageMagick-149.png
Mai 26 13:45:43 openqa-worker2 worker[18003]: [debug] [pid:22105] Uploading artefact ImageMagick-149.png as baf8eb2475cfe7c2e63a76fec41162ae
Mai 26 13:45:43 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/.thumbs/ImageMagick-149.png
Mai 26 13:45:43 openqa-worker2 worker[18003]: [debug] [pid:22105] Uploading artefact ImageMagick-149.png as baf8eb2475cfe7c2e63a76fec41162ae
Mai 26 13:45:43 openqa-worker2 worker[18003]: [debug] [pid:22105] Optimizing /var/lib/openqa/pool/37/testresults/ImageMagick-116.png
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: State 'stop-sigterm' timed out. Killing.
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 18003 (worker) with signal SIGKILL.
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 22105 (worker) with signal SIGKILL.
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 23223 (optipng) with signal SIGKILL.
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Main process exited, code=killed, status=9/KILL
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Killing process 22105 (worker) with signal SIGKILL.
Mai 26 13:45:43 openqa-worker2 systemd[1]: Stopped openQA Worker #37.
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Unit entered failed state.
Mai 26 13:45:43 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Failed with result 'timeout'.
Mai 28 21:02:59 openqa-worker2 systemd[1]: openqa-worker-auto-restart@37.service: Unit cannot be reloaded because it is inactive.
```

This is expected. If the worker is terminated it stops the job but still uploads results (which also involves optimizing PNGs). I guess it is actually useful that logs are uploaded and the job's reason is set accordingly. However, it is indeed questionable to upload videos, images and other results of a job which is abandoned and restarted anyways. So I'd say the worker should still upload logs, mark the job as done (with reason quit) and restart it. That should not take too long and if systemd is still too impatient we can adjust the timeout within the service file. If you agree with that change we can add it as ACs and call the ticket workable.

#8 - 2021-07-09 08:17 - okurz

- Target version changed from Ready to future