

openQA Project - action #80986

coordination # 80142 (Blocked): [saga][epic] Scale out: Redundant/load-balancing deployments of openQA, easy containers, containers on kubernetes

action # 80908 (New): [epic] Continuous deployment (package upgrade or config update) without interrupting currently running openQA jobs

terminate worker process after executing all currently assigned jobs based on config/env variable

2020-12-11 10:51 - okurz

Status: Resolved	Start date: 2020-12-11
Priority: Normal	Due date:
Assignee: mkittler	% Done: 0%
Category: Feature requests	Estimated time: 0.00 hour
Target version: Ready	
Difficulty:	

Description

Motivation

See [#80908](#)

Acceptance criteria

- **DONE AC1:** the worker process terminates after all currently assigned jobs are done based on config/env variable and the systemd service restarts the process
- **DONE AC2:** the worker can still be stopped on shutdown or manual systemctl stop openqa-worker@\$id
- **DONE AC3:** the worker service by default is still kept running if no additional config variable is set
- **DONE AC4:** the worker service by default is still restarted on package upgrades if no additional config variable is set
- **DONE AC5:** deploy updated system configuration on OSD and o3

Suggestions

- Add something like stop() if \$ENV{OPENQA_WORKER_TERMINATE_AFTER_JOBS_DONE};
- Then in <https://github.com/os-autoinst/openQA/blob/master/systemd/openqa-worker%40.service#L15> set the env variable OPENQA_WORKER_TERMINATE_AFTER_JOBS_DONE and configure the service to restart not only on failure (or change worker code to actually stop with failure to trigger a restart)
- Maybe the forced restart of the systemd service on package upgrade can actually be prevented in that mode as well, e.g. override ExecStop or something. Maybe needs to be extended to block normal terminate requests if the above variable is set

Related issues:

Blocks openQA Project - action #80910: openQA workers read updated configurat...	Resolved	2020-12-09
--	-----------------	-------------------

History

#1 - 2020-12-11 10:52 - okurz

- Description updated

#2 - 2020-12-11 14:37 - mkittler

- Assignee set to mkittler

#3 - 2020-12-11 16:55 - mkittler

- Status changed from Workable to In Progress

PR: <https://github.com/os-autoinst/openQA/pull/3636>

#4 - 2020-12-12 04:07 - openqa_review

- Due date set to 2020-12-26

Setting due date based on mean cycle time of SUSE QE Tools

#5 - 2020-12-17 12:50 - mkittler

The PR has been merged. However, the new service file hasn't been enabled in any production instance yet. The o3 workers are rebooted every night anyways so they are likely not of interest here. So I suppose I'll have to come up with a salt change for OSD and maybe I'll test the new service on one worker first.

#6 - 2020-12-17 12:55 - cdywan

mkittler wrote:

The PR has been merged. However, the new service file hasn't been enabled in any production instance yet. The o3 workers are rebooted every night anyways so they are likely not of interest here. So I suppose I'll have to come up with a salt change for OSD and maybe I'll test the new service on one worker first.

Technically that should be an infra ticket, but I think we assumed we want this on production here anyway... although I could offer looking into the salt part, if it helps, I'm getting used to adding stuff there anyways.

#7 - 2020-12-17 16:18 - okurz

keep in mind that this change so far would restart the worker service to re-read the config. What it does not do is prevent jobs to be aborted on package upgrade. That's to be left for the parent epic.

#8 - 2020-12-18 14:29 - okurz

okurz wrote:

keep in mind that this change so far would restart the worker service to re-read the config. What it does not do is prevent jobs to be aborted on package upgrade. That's to be left for the parent epic.

correcting myself. The last part is part of this ticket, see AC4

#9 - 2020-12-18 16:28 - mkittler

About AC4: The restarting is apparently achieved on package-level via the %service_del_postun macro. I don't intend to change this and will rely on the DISABLE_RESTART_ON_UPDATE setting in /etc/sysconfig/services on systems where we want to avoid interrupting running jobs when deploying updates.

It looks like none of the ACs are about enabling this in production. Should I still do this as part of this ticket?

By the way: It makes no sense to do this before the Christmas break so I was focusing on further upstream changes anyways, see <https://github.com/os-autoinst/openQA/pull/3641#issuecomment-748147996>. This further PR is about [#80910#note-7](#) but it takes a different approach than the ticket description suggests. However, in my opinion the "soft-restart" approach I took in that PR fits better with this ticket as it also terminates the worker and relies on the already introduced systemd service file to automatically start the service again. It means that in any case not only the configuration but also the worker code is reloaded.

#10 - 2020-12-18 17:37 - okurz

mkittler wrote:

About AC4: The restarting is apparently achieved on package-level via the %service_del_postun macro. I don't intend to change this and will rely on the DISABLE_RESTART_ON_UPDATE setting in /etc/sysconfig/services on systems where we want to avoid interrupting running jobs when deploying updates.

I consider /etc/sysconfig/services a bit too broad and would offer a solution to our users based on either configuration or selection of appropriate systemd service files.

It looks like none of the ACs are about enabling this in production. Should I still do this as part of this ticket?

I suggest you either do that within this ticket or in another subticket you create as part of the epic or write down in the epic what would need to be done as next step if you don't plan to do that

#11 - 2020-12-21 14:34 - mkittler

Draft for Salt changes: https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/423/diffs

#12 - 2020-12-28 15:08 - cdywan

- Due date changed from 2020-12-26 to 2021-01-08

Updating the due date to account for holidays

#13 - 2021-01-12 12:16 - mkittler

It looks like the automatic restarting on package updates mentioned in [#80986#note-9](#) only works when `openqa-worker.target` is started. So if that target is not started (and only `openqa-worker@.service` or `openqa-worker-auto-restart@.service`) there should be no interference.

Here some further findings regarding `systemd`'s behavior with the `openqa-worker.target` which we might need to take into account when switching to a different service file:

1. When starting `openqa-worker.target`
 1. it always starts `openqa-worker@.service` (also if `openqa-worker-auto-restart@.service` is enabled and `openqa-worker@.service` not).
 2. `openqa-worker-auto-restart@.service` is stopped if already running.
 3. only the first worker instance is started when starting `openqa-worker.target`; also when more worker instances were previously running.
2. When stopping `openqa-worker.target`
 1. `openqa-worker@.service` and `openqa-worker-auto-restart@.service` are stopped and all instances are affected.
3. When starting `openqa-worker-auto-restart@1` while `openqa-worker@1.service` is running the latter is stopped automatically as expected. `openqa-worker.target` remains active if it was active (but it remains active even after manually stopping all worker instances anyways).

PR to document this: <https://github.com/os-autoinst/openQA/pull/3667>

#14 - 2021-01-21 19:13 - cdywan

- Description updated

- Due date deleted (2021-01-08)

We agreed we want the ticket to cover deployment on production, right? Please correct me if I'm wrong. Just trying to make it clearer, what's still missing.

#15 - 2021-01-21 19:17 - cdywan

- Blocks action [#80910](#): `openQA` workers read updated configuration, e.g. `WORKER_CLASS`, whenever they are ready to pick up new jobs added

#16 - 2021-01-23 04:33 - openqa_review

- Due date set to 2021-02-06

Setting due date based on mean cycle time of SUSE QE Tools

#17 - 2021-01-25 17:39 - cdywan

cdywan wrote:

We agreed we want the ticket to cover deployment on production, right? Please correct me if I'm wrong. Just trying to make it clearer, what's still missing.

https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/429

#18 - 2021-01-27 12:27 - mkittler

AC5: The auto-restarting service is now used on OSD workers.

#19 - 2021-01-28 14:09 - mkittler

- Description updated

1. I'd wait to see how well it works on OSD before (manually) reconfiguration the o3 workers (AC5).
2. Note that since the solution involves using a different service AC2 is not fulfilled. However, I'd refrain from updating the regular worker service for now (to add `Restart=always`).
3. As discussed in the epic it would be very desirable to cover idling workers as well to avoid running one more job (per worker slot) with the old configuration. See the epic for this (as it is not covered by any of this tickets ACs).

#20 - 2021-02-11 09:46 - mkittler

- Status changed from *In Progress* to *Resolved*

Configured on o3 yesterday via:

```
for i in aarch64 openqaworker1 openqaworker4 openqaworker7 power8 rebel imagetester; do echo $i && sshpass -p ... ssh root@$i 'for worker_slot in $(systemctl list-units '\''openqa-worker@*.service'\'' | sed -e '\''/.openqa-worker@.*\.service.*!d'\'' -e '\''s|.openqa-worker@(.*)\.service.*|\1|'\'''); do systemctl disable openqa-worker@$worker_slot && systemctl enable openqa-worker-auto-restart@$worker_slot; done'; done
```

Run the same command with swapped service names to revert (if necessary). So far it looks like the change survived the reboot and didn't break anything.

#21 - 2021-02-11 09:46 - mkittler

- *Description updated*

#22 - 2021-02-26 10:33 - okurz

- *Due date deleted (2021-02-06)*