openQA Project - coordination #64881

coordination # 103941 (New): [saga][epic] Scale up: Efficient, event-based handling of storage on new, clean instances

[epic] Reconsider triggering cleanup jobs

2020-03-26 16:12 - mkittler

Status:	New	**Start date:**	2021-08-31
Priority:	Low	**Due date:**	
Assignee:		**% Done:**	38%
Category:	Feature requests	**Estimated time:**	0.00 hour
Target version:	future		
Difficulty:			

Description

Motivation

Currently the cleanup jobs are triggered time-based using a systemd timer. So far this ticket is just a collection of ideas how we can improve.

The time-based trigger might not be frequent enough, especially if the cleanup is aborted in the middle for some reason. E.g. on o3 we see that results/screenshots might pile up a lot which makes the error rate even higher because in consequence the cleanup jobs can take quite long (see https://progress.opensuse.org/issues/55922) and are possible interrupted with the result "worker went away".

Triggering the cleanup blindly when new jobs are scheduled as we did before is not nice either. It means creating tons of Minion jobs which just terminate because a cleanup job is already running and clutter the dashboard.

Acceptance criteria

- **AC1:** Idle instances of openQA, e.g. personal single-user developer instances, only trigger cleanup jobs when quota usage is likely to change, e.g. when new builds or jobs are scheduled or jobs complete
- **AC2:** Cleanup jobs are only triggered when a useful effect is expected, e.g. not 100 times in a row shortly after each other

Suggestions

One possibility to solve this would be that the jobs delete themselves if they can't acquire the lock. Another possibility would be acquiring the lock before creating the job and if that's not possible there will simply be no job (and if it is possible the job needs to adapt the lock).

Note that triggering the cleanup more frequently will not magically solve problems without adjusting quotas for the result storage duration. Now that we keep track of the result size we could additionally add a size-based threshold for results (maybe specify a max. total size and a percentage for each job group).

Subtasks:

action # 97763: Event-based cleanup jobs triggered based on quota size:M	**Resolved**
action # 99258: openQA enables event-based cleanup out of the box	**New**
action # 101376: Use cleanup triggers on finished jobs by default	**New**
action # 101602: Research how to properly communicate these changes based on systemd fi...	**Workable**

Related issues:

Related to openQA Project - coordination #76984: [epic] Automatically remove ...	**New**	**2021-01-21**
Related to openQA Infrastructure - coordination #96974: [epic] Improve/recons...	**Resolved**	**2021-09-20**

History

#1 - 2020-03-28 12:37 - okurz

- Category set to Feature requests

#2 - 2020-07-14 13:37 - okurz

- Target version set to Ready

#3 - 2020-07-22 13:41 - okurz

- Description updated

- Status changed from New to Workable

#4 - 2020-07-26 15:56 - okurz

new mojo minion expiring jobs as in https://github.com/mojolicious/minion/compare/v10.11...v10.12#diff-c112bb3542e98308d12d5ecb10a67abcR2 might help

#5 - 2020-07-28 09:35 - okurz

We discussed this topic in the QA tools weekly meeting 2020-07-28. "Expiring jobs" would expire based on just time and as we want to look into event-based triggers they are not that helpful. But as kraih explained we can go another way: We can create minion jobs that can create locks which expire after a time which can be longer than the time the actual minion job runs. So what we should do:

- Add back trigger for cleanup job
- Use a configurable "dead-time" for locks
- Optional, after that: Based on configuration, call "df", compare free space with a configurable limit and only trigger cleanup job if threshold is exceeded
  - can use https://software.opensuse.org/package/perl-Filesys-Df?search_term=perl-FileSys-Df
  - can mock "df" in tests to simply give back what we want, e.g. "enough free space available" or "free space exceeded"

#6 - 2020-10-18 12:55 - okurz

- Priority changed from Normal to Low

#7 - 2020-11-04 20:27 - okurz

- Related to coordination #76984: [epic] Automatically remove assets+results based on available free space added

#8 - 2020-11-04 20:28 - okurz

- Status changed from Workable to Blocked

- Assignee set to okurz

I think we can look into #76984 first

#9 - 2021-05-05 18:16 - okurz

- Status changed from Blocked to Feedback

We have not yet done #76984 but I think that #88121 brought us further.

mkittler I would very much appreciate your feedback on how you see this ticket here after #88121 . We can discuss here or also have a video chat.

#10 - 2021-05-18 09:49 - okurz

- Status changed from Feedback to Blocked

By default the cleanup systemd timers are started as dependency in systemd/openqa-webui.service . Currently the space aware cleanup is triggered also when the timers trigger. We can consider the cases when "jobs post" or "isos post" is called to trigger the space aware cleanup and not have the systemd timer pulled in by default anymore.

waiting for #91782 which we see as related

We also have the feature of keeping a minimum amount of space but have not enabled it in production yet

#11 - 2021-05-18 09:49 - okurz

- Parent task set to #64746

#12 - 2021-06-28 11:36 - okurz

- Status changed from Blocked to New

- Assignee deleted (okurz)

blocker #91782 resolved

#13 - 2021-07-14 10:16 - cdywan

- Due date set to 2021-08-31

- Status changed from New to Feedback

- Assignee set to cdywan

Some thoughts from the planning poker:

- We may want backups in place for tackling this (#94555)
- Maybe we want event-based triggers instead of (systemd) timers?
- We could have a workshop about this topic? From the user story, personal setups, experience with production setups e.g. assets being cleaned up too soon - see workshop topics

I'll take the ticket and look into a workshop slot with mkittler as a resident expert

#14 - 2021-08-17 10:15 - cdywan

- Status changed from Feedback to Resolved

> We could have a workshop about this topic? From the user story, personal setups, experience with production setups e.g. assets being cleaned up too soon - see workshop topics
> I'll take the ticket and look into a workshop slot with mkittler as a resident expert

Workshop slot taken for this Friday.

#15 - 2021-08-17 10:15 - cdywan

- Status changed from Resolved to Feedback

Meh. Why does Redmine keep flipping around states.

#16 - 2021-08-26 10:54 - cdywan

cdywan wrote:

> > We could have a workshop about this topic? From the user story, personal setups, experience with production setups e.g. assets being cleaned up too soon - see workshop topics
> > I'll take the ticket and look into a workshop slot with mkittler as a resident expert
>
> Workshop slot taken for this Friday.

- https://youtu.be/g331EIPd_jQ
- https://gitlab.suse.de/openqa/salt-states-openqa/-/merge_requests/561
- #97304

#17 - 2021-08-26 10:54 - cdywan

- Related to action #97304: Assets deleted even if there are still pending jobs size:M added

#18 - 2021-08-27 15:06 - mkittler

- Related to coordination #96974: [epic] Improve/reconsider thresholds for skipping cleanup added

#19 - 2021-08-27 15:08 - mkittler

Note that #97304 is not really related except for the tact that it is about cleanup. It is a problem which is independent from *triggering* the cleanup which this ticket is about. It was also not a direct outcome of the workshop or discussed there. I just looked at incompletes with the cleanup topic in mind.

Actually, #96974 is more related. It is basically about the very same problem as this ticket. I've created it to note down a few ideas I had.

#20 - 2021-08-27 15:08 - mkittler

- Related to deleted (action #97304: Assets deleted even if there are still pending jobs size:M)

#21 - 2021-08-31 10:36 - cdywan

- Subject changed from Reconsider triggering cleanup jobs to [epic] Reconsider triggering cleanup jobs

#22 - 2021-08-31 10:42 - okurz

- Tracker changed from action to coordination

#23 - 2021-09-23 18:23 - okurz

cdywan With the single subtask resolved I think the next step could be a subtask about using your new feature switch as default to cover AC1

#24 - 2021-10-22 13:39 - cdywan

- Copied to action #101376: Use cleanup triggers on finished jobs by default added

#25 - 2021-10-22 13:41 - cdywan

- Status changed from Feedback to Blocked

#26 - 2021-12-14 09:59 - okurz

- Parent task changed from #64746 to #103941

#27 - 2021-12-14 10:00 - okurz

- Status changed from Blocked to New

- Assignee deleted (cdywan)

- Target version changed from Ready to future