

openQA Project - action #2240

track job duplications

2014-04-04 09:31 - Inussel

Status:	Closed	Start date:	2014-04-04
Priority:	High	Due date:	
Assignee:	ancorgs	% Done:	100%
Category:		Estimated time:	10.00 hours
Target version:	Sprint 06		
Difficulty:			
Description			
openqa should track duplications of a job so one can see if someone restarted a job or if openqa already restarted the job itself because it crashed.			

History

#1 - 2014-04-04 14:12 - ancorgs

- Assignee set to ancorgs

- Target version set to Sprint 06

#2 - 2014-04-04 14:34 - ancorgs

Duplicating job could be ok if we are able to easily determine which jobs are still relevant and which ones are just simply "old stuff". It's easy if you have a "line" of duplications. I mean, something like A duplicated into B and B duplicated into C and so on. In this case, the only relevant one is the one that have never been cloned or, in some cases, also those which a running or scheduled duplicated. The problem with the current model is that it allows more than one duplication for a given job. So instead of having a line, you can have something like:

- A duplicated into B.
- A duplicated again into C.
- C duplicated into D.

No way to know what is still relevant in a simple way.

So, I think we should never allow more than one duplicated per job, so the sequence is clear. At database level, there are several alternative implementations, but before proceeding I want to be sure that we all agree in that point.

#3 - 2014-04-07 06:44 - Inussel

ancorgs wrote:

So, I think we should never allow more than one duplicated per job, so the sequence is clear. At database level, there are several alternative implementations, but before proceeding I want to be sure that we all agree in that point.

Fair enough. I can't think of a reason why we'd need to duplicate jobs more than one. I guess we'll have a prominent link to the new job when viewing a duplicated one anyways.

#4 - 2014-04-07 15:45 - ancorgs

- Status changed from New to In Progress

- % Done changed from 0 to 50

Correct and safe duplication using database transactions and optimistic locking implemented with tests. Just one hour of work for somebody knowing DBIx::Class and Perl. Almost a whole day to me :-(. Time to do something useful with the new clone_id.

#5 - 2014-04-08 10:37 - ancorgs

- % Done changed from 50 to 70

Clones are displayed in the interface and read by the scheduler. Schema version updated and upgrade scripts committed.

<https://github.com/os-autoinst/openQA/pull/19>

#6 - 2014-04-08 16:09 - ancorgs

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

- % Done changed from 70 to 100

Finally, the first really useful use case for clone_id: hiding old stuff in the jobs list. <https://github.com/os-autoinst/openQA/pull/22> Closing the issue.

#7 - 2014-04-16 07:14 - ancorgs

- Estimated time set to 10.00 h