

Redmine - Defect #1931

Timeout risks: need asynchronous operations

2008-09-22 09:25 - Thomas Lecavelier

Status: New	Start date: 2008-09-22
Priority: High	Due date:
Assignee:	% Done: 0%
Category: SCM	Estimated time: 8.00 hours
Target version:	Affected version:
Resolution:	
Description As notified by Calvin in the forums(http://www.redmine.org/boards/2/topics/2327), big SCM repositories cause web servers to timeout access to repository views. It's not a good idea to encourage people to set long timeout their web server. A solution is to wrap potential long call in asynchronous process: <ol style="list-style-type: none">1. Find out every potential long call2. Wrap it in an asynchronous process3. Add a "waiting hook" on that page, with a js observer to check when the call take end For my part, I think it's an important defect (so the priority) and not a little yard, but a doable one. I'm open to every suggestion about a different way to fix it.	
Related issues:	
Related to Redmine - Defect #1435: slow display of 'repository' tab for huge ...	Closed 2008-06-12
Related to Redmine - Defect #7002: mysql timeout when first browsing a large ...	Closed 2010-11-30
Related to Redmine - Defect #2824: MySQL Deadlock Error when showing a big re...	New 2009-02-24

History

#1 - 2008-09-25 21:04 - Eric Davis

+1 I ended up turning off my SCM's autofetch commits and have a cronjob as a workaround. Problem is, the cronjob only runs every hour so my Redmine is always behind.

#2 - 2008-09-25 21:32 - Nicolas Chucho

Eric Davis wrote:

+1 I ended up turning off my SCM's autofetch commits and have a cronjob as a workaround. Problem is, the cronjob only runs every hour so my Redmine is always behind.

Using cron is the recommended way to do. I've worked on local cache for subversion and git (in branches/nbc for the moment) to speed things up but even than, creating a svnmirror of a repository can take a loooooong time.

We've talked about using something like backgroundrb instead of cron but it's not a priority for the moment. I would like to do other things before that.

I'm very bad in js so if someone can drop me a patch implementing the observer part Thomas speak about, I could try to implement the rails part.

#3 - 2008-10-02 02:43 - Eric Davis

Nicolas Chucho wrote:

I'm very bad in js so if someone can drop me a patch implementing the observer part Thomas speak about, I could try to implement the rails part.

Don't think the observer would be that difficult. You can use `periodically_call_remote` from Rails with it pointed to an action that checks the status of the process. The action would return a busy if the process is still working or the response if it's complete (e.g. the updated svn browser)

We've talked about using something like backgroundrb instead of cron but it's not a priority for the moment. I would like to do other things before that.

I'd agree with you there. Ideally I would like to get some other improvements in Redmine before we optimize things.

#4 - 2011-02-24 10:11 - Etienne Massip

- *Category set to SCM*

#5 - 2013-02-25 11:08 - Gabriel Mazetto

This sounds like Delayedjobs or better Sidekiq could be optionally used here... This is a good time to start thinking about it, because Rails 4 will arrive with a Queue API, so more and more applications will start using async jobs.

#6 - 2015-06-20 12:26 - Toshi MARUYAMA

- *Related to Defect #2824: MySQL Deadlock Error when showing a big repository added*

#7 - 2016-05-08 06:00 - Go MAEDA

- *Description updated*