

Redmine - Feature #6982

Change in behavior of Repository auto-fetch feature

2010-11-25 19:41 - Dayson Pais

Status:	New	Start date:	2010-11-25
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	SCM	Estimated time:	0.00 hour
Target version:			
Resolution:			
Description			
Present Scenario of Repository Fetch			
<p>Presently, the changeset/other information about a repository is fetched each time I click on the Repository tab for a project. This is annoyingly slow if the repository is located on a remote server (like assembla, github, etc). I've been using redmine on a dedicated server with a 100mbps connection and I still don't like the delay. The browsing of the repository should be smooth. Especially when I click the repository tab just a second after the previous click. (The reconnection to the svn server again is redundant and would 'rarely' produce a change set.</p>			
Recommended change in behavior			
<p>Instead of using a cron job to do this at a fixed interval, the easiest solution would be for redmine to implement it within itself. Redmine can store the 'last-fetched-time' for every repository and let the user change this. So every time an http request is made for the repository page, redmine can refresh with the remote svn server only if the time since last refresh is greater than set time. This way, if you are navigating the same page within seconds, it will be seamless.</p>			
Advantages			
<ol style="list-style-type: none">1. Server Independent implementation. (for users who don't have access to cron tab)2. More efficient than cron since if redmine is not used for an entire day, the server doesn't unnecessarily run a changeset fetch more than 100 times in 24 hours. (if the cron is set to 10 minutes). We run 3 copies of redmine and having a change set being fetched regularly (when not even required) is quite not acceptable. Every bit of RAM/Processor matters on a server at times.3. Easy to code (as far as logic is concerned)			
<p>I've had the chance to interact with a few others who use redmine and I think this feature would be acceptable & preferred by a majority part of the community.</p>			
<p>Looking forward to this. I love redmine!</p>			
Related issues:			
Related to Redmine - Defect #7984: Confusing revisions links on repository root		New	2011-03-24

History

#1 - 2010-11-25 20:49 - Ivan Cenov

So every time an http request is made for the repository page, redmine can refresh with the remote svn server **only if the time since last refresh is greater than set time**. This way, if you are navigating the same page within seconds, it will be seamless.

Hi,

I think that there should be a link like **Refresh now!** in this case.

Ivan

#2 - 2010-11-25 21:23 - Dayson Pais

Even that would be a great add-on. But a refresh now button + a refresh interval is the most ideal solution here.

#3 - 2010-11-25 21:41 - Muntek Singh

I can't think of a situation where you would not be better served by simply setting up a post-commit-hook in your repository to run the fetch changesets?

This way the repo is only updated when needed, every time :)

#4 - 2010-11-25 22:50 - Felix Schäfer

- Priority changed from High to Normal

Muntek Singh wrote:

I can't think of a situation where you would not be better served by simply setting up a post-commit-hook in your repository to run the fetch changesets?

Agreed, I think there now are enough options for post-commit hooks (either the runner or the "ping-redmine-over-http-API" solution) that the auto-fetch should only be a very last resort.

Other than that, you're (the OP) writing about fetching from github: redmine only queries local copies of git repositories, the sync to any other repo is at the charge of the admin, so your point is moot in that case.

Last but not least: there are IIRC some SCM adapters (the git one at least, not sure about the svn one, haven't delved in the others) that still have to query information from the repo *each time* the redmine repository page is loaded because it's not cached, so even if we enforce a "only do a complete update every x minutes" policy, there would still be calls to the SCM repo each time the page is loaded.

#5 - 2010-11-26 08:58 - Dayson Pais

[Felix Schäfer](#)

Ignore my github example then. We don't really use git. Most of our projects are still under svn.

Why the need for a post-commit-hook when again, the solution is easily implementable within redmine? It's one less thing to worry about or setup. Besides, telling redmine to fetch the latest changeset is not required until someone is about to click that repository tab.

there would still be calls to the SCM repo each time the page is loaded.

Are you referring to browsing through the repo?

#6 - 2010-11-26 09:12 - Felix Schäfer

Dayson Pais wrote:

Why the need for a post-commit-hook when again, the solution is easily implementable within redmine? It's one less thing to worry about or setup. Besides, telling redmine to fetch the latest changeset is not required until someone is about to click that repository tab.

You seemed to be looking for an efficient solution, the post-commit hook is as efficient as you will get: the info redmine gets and stores about repositories is incrementally updated and only as needed.

With your scenario, a mid-traffic redmine/project will still get enough visits to the repository tab to trigger redmine to fetch changesets every 5 minutes (oder 10 or whatever you set the interval to), which would be as much as you'd get with the cron solution, so I still don't see the big advantage. You either have the choice of simplicity and utter uptodateness with the fetch changeset on each repository page view at the cost of higher load, you can setup the post-commit solution and have uptodateness without the load at the price of a little bit of setup, or half-way between both with the cron solution.

That said, I didn't say I was against implementing something like that, just that it doesn't make (much) sense to me.

there would still be calls to the SCM repo each time the page is loaded.

Are you referring to browsing through the repo?

I'm referring for example to [source:/trunk/lib/redmine/scm/adapters/git_adapter.rb#L87](#) which gets called once for each and every file or directory shown in the repository browser within redmine each time you visit the corresponding repository page. Mind you, that's the one I remember, it might and probably isn't the only one.

#7 - 2010-11-26 13:26 - Dayson Pais

I still don't prefer post commit hooks. Especially on remote server's like assembla, etc creating these hooks are complicated. Why all this when redmine can implement this feature within itself?

It's a substitution or rather a redesign of the way the present feature works. The perceived performance will definitely be more & at the same time having the 'refresh' button is a good tool. Sometimes, the svn list fetch can fail due to any reason, so I can manually refresh in that case.

I don't find this feature to be an overhead or unnecessary. It looks quite part of the system to me.

#8 - 2010-12-04 12:07 - John Yani

I don't understand the need of refreshing repositories.

Could you please explain, is Redmine caching repository structure? Why's that? I could not find that info in redmine docs.

#9 - 2010-12-04 13:10 - Jean-Philippe Lang

Redmine stores some metadata for each changeset (id, commit log and the list of affected files) in the database. Basically, it allows searching the commit logs, displaying changesets in the activity view and displaying commit links (eg. [r4000](#)).

That said, the repository browser will always hit the repository to list files.