

Redmine - Defect #60

(some?) SVN operations are *\*very\** slow

2007-05-12 07:37 - moe moe

Status:	Closed	Start date:	
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Affected version:	
Resolution:			
<b>Description</b> i've been playing with redmine trunk and added my svn repo that currently holds 124 revisions and about 380000 files.  the first click on "SVN Repository" took about 80minutes to complete. that's okay, i understand that some initial crunching is probably not avoidable.  after that first click most operations seemed to be fast, until i clicked on the "revisions" link. redmine is crunching again for another 20minutes now (and still going)...  i guess there must either be a general problem with svn caching or it's a particular problem with the "revisions" link.  whatever it is, it should to be fixed to make redmine suitable for medium sized projects.  (timings are from my athlon64 3500+ with 2G of ram)			

History

#1 - 2007-05-13 12:43 - Jean-Philippe Lang

380,000 files! That's what you call a medium sized project ? :-)

The first 80 min correspond to fetching all existing revisions.  
You could have done this offline for the first import with this command:  
ruby script/runner "Repository.fetch\_changesets"  
RAILS\_ENV="production"

Concerning the "revisions" link, could you tell me how many lines you have in the "changesets" and "changes" table ?  
The revisions link you're speaking about is the link that appears at the bottom of the main repository screen (the one you get when you click on "SVN repo" in the left menu), isn't it ?

Once you get the response for the "revisions" link, could you post the corresponding log. I'd like to view the render/database times.  
This action doesn't perform anything on the repository. It's just a paginated list of the revisions (that are stored in database).  
It shouldn't be that long with only 124 revisions.

#2 - 2007-05-13 14:32 - moe moe

About the number of files: well, in reality its only 1/8 that size but we have some branches that blow it up considerably - and i'd like redmine to see all branches.

Thx for the script/runner hint, I'll use that next time.  
Regardless the initial import time seems okay to me considering the size of the repository.

Yes, I'm talking about the "View revisions" link that

shows  
up under the "Latest revisions" table.

Here's the log for my last attempt (I hope rubyforge doesn't wrap it too bad):

```
Processing RepositoriesController#revisions (for
192.168.1.10 at 2007-05-12 16:57:02) [GET]
Session ID: 65ca3e9ec1bb180b82b588a293aac761
Parameters: {"action"=>"revisions",
"id"=>"1",
"controller"=>"repositories"}
Rendering within layouts/base
Rendering repositories/revisions
Completed in 659.44962 (0 reqs/sec) | Rendering: 0.07530
(0%) | DB: 0.60253 (0%) | 200 OK
[http://svn/repositories/revisions/1]
```

The bottleneck seems to be somewhere in the code, the ruby process does very little i/o but takes up 99% CPU until it eventually finishes.

btw. I could never finish loading page 2 of the the revisions list. It ran for 4h or so until I killed the process.

More figures:

```
select count(*) from changesets;
===== | count(*) |
===== | 40 |
=====
1 row in set (0.00 sec)
```

```
select count(*) from changes;
===== | count(*) |
===== | 105603 |
=====
1 row in set (0.03 sec)
```

This is a experimental repo, so it hasn't really seen many commits since it was imported. Most of the changes are probably due to the branching. (btw this test was done on yet another new import of the same repo, that's why the revision is down to 40 again).

Just for the record: The "view revisions" link seems to be the only way to trigger this problem for me. Everything else is zippy as should be.

**#3 - 2007-05-14 16:31 - Jean-Philippe Lang**

I thought the time was spent on the database. But it seems that it's not the case: DB: 0.60253 (0%)

I loaded my database with 1,000 changesets and 100,000 changes on the same repository, and it works fine (total response time ~0.3s for the "View revisions" link).

I looked at the queries. There may be a problem with a useless association loading.  
Does one of your changeset has most of the 105,603 changes ?  
Don't know if my question is clear... So could you just give me the first result lines for this query:  
SELECT changeset\_id, count(id) AS t FROM changes GROUP BY changeset\_id ORDER BY t DESC

**#4 - 2007-05-14 16:36 - moe moe**

in short: yes.

```
mysql> SELECT changeset_id, count(id) AS t FROM changes
GROUP BY changeset_id ORDER BY t DESC;
=====+ | changeset_id | t |
=====+ | 33 | 84613 | | 60 | 13764 | | 59 | 2907 | | 62 | 2684 | | 49 | 814 | | 58 | 383 | | 65 |
```

163			40		53				47		40				55		37				63		35				72		33				64		13				53		12				61						
12			39		9				34		3				54		3				70		3				73		3				42		2				48		1				50		1				
		66		1					35		1				51		1				67		1				36		1				52		1				68		1				37		1				
69		1				38		1						71		1				56		1				41		1				57		1				43		1				44		1				45	
1				46				1																																											
=====+-----																																																			
41 rows in set (0.03 sec)																																																			

#5 - 2007-05-14 17:25 - Jean-Philippe Lang

OK, problem found.  
I'll try to fix it asap.

#6 - 2007-05-15 17:37 - Jean-Philippe Lang

I fixed this problem in rev 535. I also added pagination for changes when displaying a changeset.

I tested with mysql and a changesets containing 100,000 changes.  
Works now without a hitch :-)

Please, let me know if it's ok for you.

#7 - 2007-05-16 08:15 - moe moe

thank you very much, works perfectly now! :)

#8 - 2007-05-16 13:52 - Jean-Philippe Lang

Thanks ;-)