

Redmine - Feature #6092

Truncate Git revision labels in Activity page/feed and allow configurable length

2010-08-09 17:54 - Phil Moorhouse

Status:	Closed	Start date:	2010-08-09
Priority:	Normal	Due date:	
Assignee:	Toshi MARUYAMA	% Done:	100%
Category:	SCM	Estimated time:	0.00 hour
Target version:	1.1.0		
Resolution:	Fixed		
Description			
<p>Git uses hashes for the revision labels rather than an incremental counter, so a lot of space is used up just displaying the label in the activity page / feed that would otherwise show more useful info (such as a longer part of the commit message).</p> <p>Redmine already truncates the Git revision labels to the first 8 chars on the Repository page, which is more than enough for most projects, so it would be good to see this on the Activity page. If we could also configure the length to be even shorter if required, that would be great. The href for any links would obviously still contain the full revision hash.</p>			
Related issues:			
Related to Redmine - Defect # 3724: Mercurial repositories display revision I...		Closed	2009-08-10
Related to Redmine - Defect # 7047: Git adapter very slow when a commit modif...		New	2010-12-04
Related to Redmine - Defect # 7146: Git adapter lost commits before 7 days fr...		Closed	2010-12-21
Related to Redmine - Defect # 6013: git tab,browsing, very slow -- even after...		Closed	2010-08-02
Related to Redmine - Feature # 2610: Git repository comments should use git c...		Closed	2009-01-28
Duplicated by Redmine - Feature # 3244: Shorten git commit ids in Atom feeds ...		Closed	2009-04-27
Duplicated by Redmine - Feature # 3945: Shorten revision ID for git events		Closed	2009-09-30

Associated revisions

Revision 4613 - 2011-01-02 10:45 - Toshi MARUYAMA

Changing revision label and identifier at SCM adapter level (#3724, #6092)

Contributed by Yuya Nishihara.

History

#1 - 2010-08-12 07:30 - Adam Soltys

- File patch.diff added
- Status changed from New to Resolved
- % Done changed from 0 to 100

Attached patch applies the fix. Also available at <http://github.com/asoltys/redmine>

#2 - 2010-08-12 07:32 - Adam Soltys

Oh, I should mention I didn't implement this as a configurable setting as requested. Maybe later...

#3 - 2010-08-12 09:35 - Felix Schäfer

- Category set to UI

- Status changed from Resolved to New

Please leave the status changes to the devs, thank you. (resolved is for patches in trunk but not in stable yet).

The patch looks good to me, I'm not sure we'd need it configurable either. If we take the git etiquette of 72 chars per line, add the "Revision " and ": " around the revision hash (11 chars total, might be somewhat different in other locales) and 8 chars for the revision hash, we end up with 91 chars, which leaves some space for the other locales to have longer strings.

JB: I like the idea as it is, but I'd rather make a @short_revision attribute each adapter can override at will rather than a hard default like that. I especially think people using SCMs with sequential numbering would be quite surprised to see commit 10000000 followed by commit 10000000 :-)

#4 - 2010-08-12 17:02 - Adam Soltys

Agreed Felix. I'll work on moving it at least into a variable if not into a setting. Also, sorry about changing the ticket status, I'll know now not to do that anymore.

#5 - 2010-08-15 00:06 - Adam Soltys

- File patch.diff added

Ok, I've implemented this as a setting now. I figure this way it's not necessary to have separate values in each adapter. I've defaulted the value to 8. If someone with an svn repository with over 10 million revisions needs to increase this, they can!

#6 - 2010-08-15 03:22 - Toshi MARUYAMA

Adam Soltys wrote:

Ok, I've implemented this as a setting now. I figure this way it's not necessary to have separate values in each adapter. I've defaulted the value to 8. If someone with an svn repository with over 10 million revisions needs to increase this, they can!

Your patch is great!

Can you implement setting of "7.days" at source:tags/1.0.0/app/models/repository/git.rb#L50 and <http://www.redmine.org/attachments/3272/git-fast-browse.patch> in #4773 ?

#7 - 2010-08-15 06:10 - Adam Soltys

- File patch.diff added

Toshi, ok, I did these two things as settings for you, but I don't actually think they should be included in trunk and presented to the average user. They're specific to git, for one thing, and they deal with some pretty obscure internal performance mechanisms.

A better solution in my opinion would be to revisit the git adapter and see if we can get some more performance out of it by default. If github can do it, so can we!

#8 - 2010-08-15 06:32 - Adam Soltys

I just realized I was incorrectly using the `log_display_limit` setting instead of the truncation setting in a couple places of code. Fixed that. Here's the patch again with just the truncation setting addressing the original issue, this time done properly.

I've also placed this change in a "for-eric" branch on my github fork and sent him a pull request. <http://github.com/asoltys/redmine/tree/for-eric>

#9 - 2010-08-15 06:34 - Adam Soltys

- *File revision_truncation_length_as_setting.diff added*

#10 - 2010-08-16 02:04 - Toshi MARUYAMA

Adam Soltys wrote:

Toshi, ok, I did these two things as settings for you, but I don't actually think they should be included in trunk and presented to the average user. They're specific to git, for one thing, and they deal with some pretty obscure internal performance mechanisms.

7.days problem is reported at #6013 and "git log -n 1 FILE_or_DIR" problem is reported at #5096.

#11 - 2010-08-17 05:14 - Adam Soltys

Thanks for pointing to those issues Toshi. I hadn't been following redmine development for the last year or so but I'm glad to see that other people have taken up the charge on these matters. I was exposed first-hand to some of the general performance problems that redmine has with distributed scm's back when I worked on #1406, so I've gained an appreciation and desire to fix these things once and for all.

I still don't agree with exposing these two settings at the moment because I still see them as temporary hacks to solve performance problems that I think we could do better on. In particular, I have some ideas about possibly pre-fetching and caching some repository history or implementing some more sophisticated on-demand lookups and partial history scans that might give us adequate performance without requiring users to sacrifice the "last revision" field or manually specify how far back to look for dirty changesets.

I need to do a bit more thinking and design-work before I tackle a full-on reimplementation of the git and mercurial adapters though. So in the meantime I'm just trying to get my feet wet again with contributing to redmine. I do have plans to become a regular contributor through the coming fall and winter though and repository management and performance will be a central area of interest for me.

#12 - 2010-08-25 22:40 - Jean-Baptiste Barth

Didn't have the time to read everything but thank you guys for your work on these topics.

About the original issue, I agree we should truncate git's sha1 revision names but I have 2 worries about it :

- I don't like the idea of introducing a Setting for that. There are already bunches of useless settings in redmine, and consulting administration panel always give me the feeling that we have too many things here which leads to confusion and bloat software ; I'd rather turn it into a constant in the class in a first approach. I think nobody cares Github truncates revision names to 7 or 8 chars because nobody reads these IDs, they're just here to point to the commit. We should have the same approach, decide a number (6? 7? 8? 10?) and begin with this.
- after some thoughts about it, I think this modification must be applied at adapter level, only to SCM with non sequential revision numbers.

Adam: can you confirm the last patch is the only one to keep ? We can delete other patches so that it's clearer which one to work on.

#13 - 2010-08-26 02:04 - Toshi MARUYAMA

Yuya's implementation of Mercurial overhaul (#4455) shows activity with mercurial style such as [12:ff70262d476d](https://github.com/ff70262d476d) .

You can see at <http://dev.openttdcoop.org/projects/ttrs/activity> .

I pushed yuya's [mq](#) to my github repository [branch](#) .

#14 - 2010-08-26 09:42 - Felix Schäfer

Jean-Baptiste Barth wrote:

About the original issue, I agree we should truncate git's sha1 revision names but I have 2 worries about it : - I don't like the idea of introducing a Setting for that. There are already bunches of useless settings in redmine, and consulting administration panel always give me the feeling that we have too many things here which leads to confusion and bloat software ; I'd rather turn it into a constant in the class in a first approach. I think nobody cares Github truncates revision names to 7 or 8 chars because nobody reads these IDs, they're just here to point to the commit. We should have the same approach, decide a number (6? 7? 8? 10?) and begin with this.

- after some thoughts about it, I think this modification must be applied at adapter level, only to SCM with non sequential revision numbers.

Agreed, I'd go with 8 chars and an extra short_revision virtual attribute in the adapter.

#15 - 2010-08-27 03:09 - Toshi MARUYAMA

Adam Soltys wrote:

In particular, I have some ideas about possibly pre-fetching and caching some repository history or implementing some more sophisticated on-demand lookups and partial history scans that might give us adequate performance without requiring users to sacrifice the "last revision" field or manually specify how far back to look for dirty changesets.

7.days is only git adapter problem.

There are some ideas of git at after [note-13](#) of #4773 .

Because mercurial has revision number, Redmine can know mercurial "last committed revision"(tip) easily.

But mercurial tip **is not latest revision** (see [note-23](#) of #4455).

This is the reason of Redmine 1.0.x mercurial adapter fetching performance problem.

Yuya fixed this problem at [note-144](#) of #4455 and I pushed to my github [hg-patches-svntrunk branch](#) .

#16 - 2010-08-27 03:27 - Adam Soltys

Thanks for looking at this Jean-Baptiste. I agree with you about the importance of avoiding software bloat and I don't like exposing obscure technical settings to the users either.

I'll have to look at the code though to see if it's possible to move this setting into the adapters. I think in some places where the revisions are being shortened we don't have access to the adapter objects, which is why I went with a global setting, but maybe I can figure something out. So yes, the latest patch supersedes the other ones, but hold off on looking at it because I'll probably be coming up with another one soon in light of this latest discussion.

Toshi, thanks again for the info. I looked at Yuya's branch quickly and it seems like he's done a lot of good work that I need to take a closer look at.

#17 - 2010-08-29 19:00 - Yuya Nishihara

Adam Soltys wrote:

Toshi, thanks again for the info. I looked at Yuya's branch quickly and it seems like he's done a lot of good work that I need to take a closer look

| at.

Hi, in short, my patches try to switch revision formats by repository classes (Mercurial or not).
Implemented as follows: - implement Changeset#format_identifier and Revision#format_identifier
- pass changeset or revision object to format_revision and link_to_revision helpers,
so that it can utilize #format_identifier

#18 - 2010-08-31 15:03 - Toshi MARUYAMA

Adam Soltys wrote:

Toshi, ok, I did these two things as settings for you, but I don't actually think they should be included in trunk and presented to the average user. They're specific to git, for one thing, and they deal with some pretty obscure internal performance mechanisms.

A better solution in my opinion would be to revisit the git adapter and see if we can get some more performance out of it by default. If github can do it, so can we!

Github seems to use Ajax for calling "git log -n 1 FILE_or_DIR".

Redmine calls "git log -n 1 FILE_or_DIR" not only in browsing directory tree but also in cat/diff/annotate.

Because entry() of abstract_adapter.rb calls entries() at source:tags/1.0.1/lib/redmine/scm/adapters/abstract_adapter.rb#L84 .

#19 - 2010-09-19 16:17 - Toshi MARUYAMA

- File yuya-issue-4455-mq-ab997af9e-exclude-hg.diff added
- File git-act-before.png added
- File git-act-after.png added

I imported Yuya's [MQ](#) (Mercurial Queues) exclude Mercurial adapter to [github](#) and done [minor change](#).

Now my github branch is [issue-4455-yuya-mq-20100825-exclude-hg](#)

And I attach the patch and git activity page images of source:tags/1.0.1/test/fixtures/repositories/git_repository.tar.gz.

#20 - 2010-09-28 16:17 - Toshi MARUYAMA

Adam Soltys wrote:

I need to do a bit more thinking and design-work before I tackle a full-on reimplementation of the git and mercurial adapters though. So in the meantime I'm just trying to get my feet wet again with contributing to redmine. I do have plans to become a regular contributor through the coming fall and winter though and repository management and performance will be a central area of interest for me.

Adam, please consider #5357.

Mercurial has same problem (#3567).

Yuya fixed this problem at [note-144](#) of #4455.

#21 - 2010-09-30 00:53 - Toshi MARUYAMA

FYI:

[Bazaar merged revisions don't show up](#)

#22 - 2010-10-15 09:13 - Toshi MARUYAMA

Toshi MARUYAMA wrote:

```
| "git log -n 1 FILE_or_DIR" problem is reported at #5096.
```

I found JPL posted a patch for this problem at [note-13](#) of #1435 .
But current redmine support git branch, this patch is incompatible.

#23 - 2010-12-13 15:52 - Toshi MARUYAMA

- File *revision-truncate-20101213.diff* added
- File *git-12-digits-activity.png* added
- File *git-12-digits-annotate.png* added
- File *git-12-digits-browse.png* added

I update the patch to set truncating character number at adapter level.
This patch contains unit tests.

And I pushed to my github.

- <https://github.com/marutosi/redmine/commits/git-activity>
- <https://github.com/marutosi/redmine/commit/710d3f355de44d87ccb3ec0caf49ec77917ca16b>

This patch truncates 8 characters.

You can set truncating number.

- blame/annotate
https://github.com/marutosi/redmine/blob/710d3f355de44d87/lib/redmine/scm/adapters/git_adapter.rb#L271
- Others
<https://github.com/marutosi/redmine/blob/710d3f355de44d87/app/models/repository/git.rb#L39>

I attach truncating 12 characters images.

#24 - 2010-12-15 05:10 - Toshi MARUYAMA

#3244, #3945 and #6092 are duplicate.

#25 - 2010-12-20 12:06 - Toshi MARUYAMA

- File *revision-truncate-20101213-fix-fixtures.diff* added

attachment:revision-truncate-20101213.diff has a lack of fixtures at unit tests.
I fixed it.

revision-truncate-20101213-fix-fixtures.diff is additional patch for attachment:revision-truncate-20101213.diff .

And I pushed to my github.

- <https://github.com/marutosi/redmine/commits/git-activity>
- <https://github.com/marutosi/redmine/commit/325ee1e73bafae5567d3e3a7ba09d1d2181841a>

#26 - 2011-01-02 07:46 - Toshi MARUYAMA

- *Category changed from UI to SCM*
- *Assignee set to Toshi MARUYAMA*

#27 - 2011-01-02 10:55 - Toshi MARUYAMA

- *Status changed from New to Closed*
- *Target version set to 1.1.0*
- *Resolution set to Fixed*

Implemented in svn trunk by r4613 and 1.1 by r4614.

Files

patch.diff	825 Bytes	2010-08-12	Adam Soltys
patch.diff	19.3 KB	2010-08-14	Adam Soltys
patch.diff	29 KB	2010-08-15	Adam Soltys
revision_truncation_length_as_setting.diff	19.3 KB	2010-08-15	Adam Soltys
yuya-issue-4455-mq-ab997af9e-exclude-hg.diff	14.9 KB	2010-09-19	Toshi MARUYAMA
git-act-before.png	31.5 KB	2010-09-19	Toshi MARUYAMA
git-act-after.png	24.3 KB	2010-09-19	Toshi MARUYAMA
revision-truncate-20101213.diff	17.8 KB	2010-12-13	Toshi MARUYAMA
git-12-digits-activity.png	62.7 KB	2010-12-13	Toshi MARUYAMA
git-12-digits-annotate.png	98.3 KB	2010-12-13	Toshi MARUYAMA
git-12-digits-browse.png	181 KB	2010-12-13	Toshi MARUYAMA
revision-truncate-20101213-fix-fixtures.diff	1.09 KB	2010-12-20	Toshi MARUYAMA