

Redmine - Feature #2852

Clean Incoming Email of quoted text "----- Reply above this line -----"

2009-02-27 00:36 - Dan Cameron

Status:	Closed	Start date:	2009-02-27
Priority:	Normal	Due date:	
Assignee:	Eric Davis	% Done:	100%
Category:	Email receiving	Estimated time:	0.00 hour
Target version:	1.1.0		
Resolution:	Fixed		
Description I'm requesting redmine incorporates method of stripping extraneous quoted text from replies. A few other trackers and services handle it simply by adding a hash at the top of the emails sent from the system. Then when someone replies the content below the hash is removed, since it was added by the email client when replying. This is an example of a possible solution: --- Please reply ABOVE THIS LINE to respond to this issue --- Issue #73 has been updated by ... [content] You have received this notification because you have either subscribed to it, or are involved in it. To change your notification preferences, please click here: https://redmine.sproutventure.com/my/account			
Related issues:			
Related to Redmine - Feature #6628: Improvements in truncate email		Closed	2010-10-11
Related to Redmine - Defect #8758: Ignore email keywords after delimiter		Closed	2011-07-07

Associated revisions

Revision 4361 - 2010-11-02 20:00 - Eric Davis

Remove email quotes (>) when searching for incoming email delimiters. #2852 #6628

Changes the email truncation so a delimiter can be prefixed by email quotes or spaces ("> Eric said..."). This will let it be combined with a "---- Reply above ----" delimiter to match "> --- Reply above ---"

Revision 4362 - 2010-11-02 20:00 - Eric Davis

Add a Email Header setting. Useful for adding delimiters to every email.

#2852 #6628

History

#1 - 2009-02-27 00:39 - Dan Cameron

I digress, my example is not a "hash" tag nor do I think a hash would work but something similar could.

Thanks.

#2 - 2009-03-02 09:57 - Kari Jääskeläinen

I agree!

#3 - 2009-03-02 12:13 - Paul Rivier

- Category set to Email notifications

Could we rather set opening and ending tags, and remove region between both, keeping text above and below.

The reason for that is that a lot of people bottom-post, and would probably be hurt by a strict top-posting convention.

#4 - 2009-04-07 09:01 - Marlin Forbes

Please, for the love of Google, don't make it top-posting. You'll start a flame war... :)

#5 - 2009-04-07 10:06 - Paul Rivier

Also, such a full text marker would be hard to internationalize. Maybe we should rather use markers like follows, where (a) | (b) means "either a or b, depending on the user language"

~~bor~~ (Beginning of region) | (Début de la région)

~~eor~~ (End of region) | (Fin de la région)

In this case, *-bor-* and *-eor-* match would be sufficient, and the following text would be meaningless when parsing the incoming mail, so it could be in any language. The same could be done with property fields in emails, like

~~prt~~ (Priority) | (Priorité) :

I'm wondering how we could use the same principle for fields when creating issues by email, from scratch. Maybe the right process would be :

1. send a simple mail to the system, using subject, body, attachments, and any correct field in the form above.
2. the system can send back a pre-formated mail with all available fields that the user can use
3. the user fill them and send back the mail to the system

thoughts ?

#6 - 2009-04-07 16:19 - James Blanding

What about making the text portion of the string configurable? Wouldn't that sort of weasel around the language issue? Perhaps the leading non-alphanumeric bits could be hard-coded, for a bit of self-preservation:

--==-- %s --==--

I suppose the disadvantage would then be that the input from the setting would have to be sanitized. Hmm...

#7 - 2009-04-19 18:08 - James Blanding

- *File redmine-0.8.3-strip_quoted.patch added*

Well, I now have a very basic patch against 0.8.3 that seems to work. I'm submitting it here in the hopes that it will spurn someone who knows what they're doing to implement it better. The patch adds the (hardcoded) string "----- REPLY ABOVE THIS LINE" to the beginning of the mailer templates for issue_add and issue_edit. It also adds a small chunk of code to MailHandler#receive_issue_update to strip off quoted text including and following that string. I'm guessing it would be fairly trivial to adapt it for things other than Issues.

A couple things to note about this patch:

1. I'm just starting to learn Ruby and haven't started learning Rails yet at all...please be gentle. :X
2. The bit added to MailHandler#receive_issue_update is rather ugly (see point 1).
3. Obviously this string shouldn't be hardcoded, at least not in more than one place; ideally it should be a system setting I think.
4. Did I mention this patch is ugly?
5. I have some test code for removing arbitrary chunks of body text between BEGIN and END tags if anyone's interested, but reading it may cause nausea.

#8 - 2009-04-23 10:11 - Kiall Mac Innes

+1 for the idea, but top and bottom posting happens in real life ;) it needs to allow for both!

#9 - 2009-04-26 06:45 - James Blanding

- *File redmine-0.8.3-strip_quoted_top_and_bottom.patch added*

New patch supporting both top and bottom posting. Same disclaimers as previous patch apply.

I found I couldn't just add the REPLY BELOW tag to the issue mailer templates, because the text from the email footer setting would appear below the tag. I located the top-level template ("layout"), and added both tags to that. So I believe this patch should work for all mail notifications.

This approach is not perfect, however, because I suspect the tags will now be present in mails that can't be replied to, such as those related to account activation and password recovery. I think there needs to be some method that gets called from the layout templates to decide whether or not to output the tags based on what type of message is being generated. Not really sure how that would work...it'd be great if someone could suggest something...possibly a method for the Mailer class...? ^^;;

#10 - 2009-08-30 06:49 - Isaac Lo

James Blanding wrote:

New patch supporting both top and bottom posting.

Problem I have with that is that my email system has inbound spam scanning. If emails are determined clean, it adds a footer to the end of the emails. Hence, you will end up with the spam footer at the bottom of replies.

#11 - 2009-08-30 06:50 - Isaac Lo

1

#12 - 2009-11-28 13:19 - Jean-Philippe Lang

- Category changed from *Email notifications* to *Email receiving*

#13 - 2009-12-07 22:53 - Ian Gustafson

+1. This would really go a long way toward making Redmine a viable help desk solution.

#14 - 2010-02-02 01:53 - Owen Barton

I think this may be achievable with configuration now that [#4409](#) is in.

#15 - 2010-02-02 02:14 - Dan Cameron

- Assignee set to *Eric Davis*

Yeah, that new feature works great but I still need to be able to add something this in order to set the option within the administrator. Does anyone know where I need to edit in order for this to happen?

--- Please reply ABOVE THIS LINE to respond to this issue ---

#16 - 2010-02-02 14:19 - Felix Schäfer

Won't be possible even with --- Respond above here ---, because even if you put it right at the beginning of the email, it will still have something like "On someday, somebody wrote:" before it, and it will be quoted.

#17 - 2010-02-02 16:47 - James Blanding

Felix Schäfer wrote:

Won't be possible even with --- Respond above here ---, because even if you put it right at the beginning of the email, it will still have something like "On someday, somebody wrote:" before it, and it will be quoted.

Most quoting I've seen is done with one of a handful of characters: > | :

Alternatively, if you're using a long string as a delimiter, you could approach it from the other angle and just look for a single non-alphanumeric character, followed by a space, followed by the delimiter string.

Removing the text inserted by the MUA before the quoted text would be more difficult I think. Though it's been 9 months since I last used it, when Basecamp introduced handling of replies to Todo item comments, they weren't doing anything to filter this text. It was a bit annoying to have to remember to manually remove it, but if you forget, it's only one extra line...definitely preferable IMO versus having the entire body of the original message.

#18 - 2010-02-03 00:40 - Felix Schäfer

James: If you want to put the line at the top of the mails you generate for issues, edit `@app/views/mailer/issue*` and add the line prefixed with > to the list added in [#4409](#).

#19 - 2010-04-16 18:51 - Jeff Wallace

Has there been any progress with this feature request? I would be happy to help out once we decide what is the best way to strip out quoted replies in emails.

#20 - 2010-04-16 21:39 - Felix Schäfer

Jeff Wallace wrote:

Has there been any progress with this feature request? I would be happy to help out once we decide what is the best way to strip out quoted replies in emails.

I'd say [#4409](#) should resolve this, no?

#21 - 2010-04-16 21:45 - Jeff Wallace

Felix Schäfer wrote:

I'd say [#4409](#) should resolve this, no?

I was hoping for a more *automatic* solution that didn't require users writing something like **CUT**.

My thoughts were to regex for the reply headers generated by Gmail / Outlook and wipe out them out. Here was a sample regex for gmail replies:

```
On (\w|,| )*<redmine@foo.com> wrote:\s((^$|^>.*?$)\s)*
```

It would catch the start of the reply and any consecutive line underneath that was either a blank line or started with >.

#22 - 2010-04-16 21:46 - Jeff Wallace

Whoops, the regex should be:

```
On (\w|,| |:)*<redmine@foo.com> wrote:\s((^$|^>.*?$)\s)*
```

#23 - 2010-06-28 19:52 - Wayne Walter

I have created a patch that solves this problem definitively for notes history when dealing with the public who may use any kind of email client for posting notes which often include full history with each post.

For solution, it imitates the elegant gmail algorithm for handling quoted email. But in Redmine it works whether notes originate from email or the web form.

The benefits of this email quoting solution:

1. Users can write notes above the quoted text.
2. They can also write noted below the quoted text!
3. You can also intersperse comments within the quoted text.
4. It's far simpler than "cutting" off the messages and fiddling with regular expressions plus those disallow benefits 2 and 3 above.
5. Furthermore, it applies Ruby/Rails philosophy of convention over configuration.
6. So zero configuration is necessary unless we add an enable/disable option to this feature. Although it's hard to imagine a user even wanting to see the plethora of repeated text.

How does it work?

1. It simple compares each line of notes to the previous journal entry.
2. If the line matches any line in the previous it's treated as "quoted".
3. When a fresh line is found, it gets displayed plus 3 prior lines for context.
4. All quoted lines are removed from display.
5. Someone in the future might want to make the quoted lines appear in a collapsible div like gmail so you can open and close the quotes.

So the effect is that each note on Redmine shows only the fresh lines of context plus 3 lines above each for context.

Trust me, it's absolutely lovely. This makes it a cinch to read notes history since there is literally zero useless information to wade through.

About the implementation. This is not a "hack". It's a quality patch. My redmine environment is setup to build and run all the unit tests except those that depend on repositories. Setting up to run those tests will be required to totally finish this patch but those tests should pass without a problem.

This code is mostly is in a new gem that I created for email diff which actually contains the algorithm for comparing the notes.

The rest of this modification is mostly some small changes to the `issues_controller.rb`, `journal_helper.rb`, and `journal.rb`

The `issues_controller` gives each journal object a reference to the previous journal object in the history with a nil for the first one.

The journal_helper.ch actually applies this diff by comparing the current notes to the notes from the previous journal. And this happens prior to textilable method.

If you want the end-all elegant solution to viewing notes history in Redmine, this is it!

If anyone else is interested in getting this patch completely finished or including it in Redmine please post back here on this ticket or you can find contact info on my business site <http://www.tickzoom.com>

#24 - 2010-06-28 19:58 - Wayne Walter

There's only one small glitch in this the prevents it from being totally perfect. Somewhere lines are getting split which get past a certain length. That causes the diff algorithm to see them as fresh lines for the next message even though after they the get filtered out.

It would be nice to find this and remove the line length splitter.

If anyone has any idea why they were getting split, that will be most useful. Hopefully, there was never any good reason so we can simply remove that so you don't get the extra 5 to 6 junk lines with each note entry.

Thoughts?

Sincerely,
Wayne

#25 - 2010-06-28 21:17 - Jeff Wallace

I'd like to take a look at the patch.

#26 - 2010-06-29 00:43 - Wayne Walter

I found the cause of the glitch was not Redmine. It's gmail and other clients which automatically wrap lines on word boundaries for outgoing email after 78 or 80 characters.

Gmail obviously accounts for that in their algorithm. I just figured out a way to do that also.

Jeff Wallace. FYI, I have a public project setup for my company that you can look at.

Just go to <http://support.tickzoom.com>. If you browse through the most recent 2 tickets that I'm using for testing, you'll see this thing in action. If you ever find it done, give it 10 minutes and try again until I totally finish fixing the "glitch" above.

Frankly, this has taken some 30 hours of painstaking work to reverse engineer how Gmail works its magic on quoted email text. Not to mention another dozen hours figuring out Redmine.

If anyone or group of people cares to contribute financially then I will send the patch to Redmine and follow through to jump through any hoops for Jean-Phillippe to get it included permanently.

I hope that doesn't sound like a shameless capitalist. Thing is, I have my own commercial project that I must to spend my time on to pay the bills.

I'll post back once the glitch is fixed.

Sincerely,
Wayne

#27 - 2010-06-29 00:48 - Wayne Walter

FYI, at the <http://support.tickzoom.com> site, you can click the quote icon beside any of the message to reply and that will show you the full content of the note.

That way you can see all the repetitive junk from the original email that gets filtered by the new history view patch.

Wayne

#28 - 2010-06-29 02:38 - Wayne Walter

Okay, it's done. Now perfectly handles matching wrapped lines. There's a demo ticket to see it at [\[\[http://support.tickzoom.com/issues/42\]\]](http://support.tickzoom.com/issues/42)

The rest of the testing tickets are getting deleted and we're not populating the site with real data from our old Trac system.

Again, contact me if you want this patch and we can work something out.

Sincerely,
Wayne

#29 - 2010-06-29 02:40 - Wayne Walter

Oops, you can't edit your posts here? I typed. I meant to say we are now populating that site with live data. So only ticket 42 is setup as an example ticket to play with.

Sincerely,
Wayne

#30 - 2010-06-29 07:05 - Felix Schäfer

Wayne Walter wrote:

I hope that doesn't sound like a shameless capitalist. Thing is, I have my own commercial project that I must to spend my time on to pay the bills.

Oh, so you did pay to use redmine? Well, if you did, more karma to you, but we have a relatively strict policy regarding commercial offerings: there's a forum and a wiki page for that, I'll let you figure out where they are, and seeing how much you've already pasted your link around here, I'm not even sure I'd want a link from here to the forum, but as the warning comes late, be my guest. Anyway, make sure your code and its distribution comply with the redmine license.

#31 - 2010-06-29 17:04 - Eric Davis

Wayne Walter wrote:

If anyone or group of people cares to contribute financially then I will send the patch to Redmine and follow through to jump through any hoops for Jean-Phillippe to get it included permanently.

If you don't have a patch you can contribute then we can't consider this code to be real. Like Felix said, this is bordering on a commercial advertisement and is against the site's [Commercial Offerings policy](#).

#32 - 2010-06-30 01:54 - Wayne Walter

To clarify, the purpose is to finish this patch to include it on the open source for everyone to use. This is not an effort to profit the work. If anyone has extra money to donate for me to finish the work let me know. Sorry for confusion.

Pour clarifier, le but de cet effort et pour ajouter dans l'open source pour tout le monde à utiliser. Ce n'est pas un effort à profiter du travail. Si qq'n a de l'argent à contribuer pour finir le patch, faites moi savoir. Désolé pour toute confusion.

#33 - 2010-06-30 02:03 - Wayne Walter

Felix Schäfer wrote:

Wayne Walter wrote:

I hope that doesn't sound like a shameless capitalist. Thing is, I have my own commercial project that I must to spend my time on to pay the bills.

Oh, so you did pay to use redmine? Well, if you did, more karma to you, but we have a relatively strict policy regarding commercial offerings: there's a forum and a wiki page for that, I'll let you figure out where they are, and seeing how much you've already pasted your link around here, I'm not even sure I'd want a link from here to the forum, but as the warning comes late, be my guest. Anyway, make sure your code and its distribution comply with the redmine license.

FYI, "my own commercial project" above is unrelated to Redmine. It's a Forex, Stock, and Futures trading platform called TickZoom. My point is that I have to spend my time working on TickZoom for my customers rather than working for free for users of Redmine. So if anyone wants to donate some money I'll spend the time to get the patch for email quote filtering into Redmine. Otherwise, it's not worth my time.

#34 - 2010-10-25 17:04 - Eric Davis

- Assignee deleted (Eric Davis)

#35 - 2010-11-02 20:00 - Eric Davis

- Status changed from New to Closed
- Assignee set to Eric Davis
- Target version set to 1.1.0
- % Done changed from 0 to 100
- Resolution set to Fixed

I've updated the email truncation in [r4361](#) and [r4362](#) to support reply quotes in front of the delimiter string (>) and to allow setting a standard header for all emails. It's different than this patch because it will let you define whatever truncation keywords you want and will also let you add that truncation to the standard email layout (Admin > Settings > Email notifications).

Hopefully this will cut down on all of the extra content.

#36 - 2010-11-05 18:00 - Anonymous

Great news, many thanks!

#37 - 2012-05-21 06:10 - Gibbet G

Can we improve this by adding a 'Footer' option too? This way if people do not see the 'header' and reply AFTER the email text - it is not lost... Redmine should remove all data within the Header and Footer...

I currently do this with a bad patch to core that even ignores headers and cuts mails from the line 'written on' to the line 'targetted' version... This is obviously not good enough...

May I re-open this ticket?

Files

redmine-0.8.3-strip_quoted.patch	2.84 KB	2009-04-19	James Blanding
redmine-0.8.3-strip_quoted_top_and_bottom.patch	1.99 KB	2009-04-26	James Blanding