

Redmine - Patch #9484

Limit SCM annotate to text files under the maximum file size for viewing

2011-10-28 02:07 - Antonio García-Domínguez

Status:	Closed	Start date:	2011-10-28
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	SCM	Estimated time:	0.00 hour
Target version:	1.3.0		
Description			
<p>After recently upgrading to Redmine 1.2.1 + patch #4905 and moving from Mongrel 1.1.5 to Passenger 3.0.9 (using Ruby 1.8.7 as packaged in Ubuntu's ruby1.8 1.8.7.249-2 package, for amd64), I started getting rogue processes which used up 100% CPU every now and then. Killing these rogue processes using "kill -SIGABRT pid" produced entries like these:</p>			
<pre>Processing RepositoriesController#annotate (for 10.182.221.1 at 2011-10-28 01:51:08) [GET] Parameters: {"action"=>"annotate", "id"=>"sources-fm", "path"=>["trunk", "xmleye-1.2.3_svn473-bp elunit_results.tar.gz"], "controller"=>"repositories"} Rendering template within layouts/base Rendering repositories/annotate ActionView::TemplateError (SIGTERM) on line #17 of app/views/repositories/annotate.rhtml: 14: <table class="filecontent annotate syntaxhl"> 15: <tbody> 16: <% line_num = 1 %> 17: <% syntax_highlight(@path, to_utf8(@annotate.content)).each_line do line %> 18: <% revision = @annotate.revisions[line_num-1] %> 19: <tr class="bloc-<%= revision.nil? ? 0 : colors[revision.identifier revision.revision] % >"> 20: <th class="line-num" id="L<%= line_num %>"><a href="#L<%= line_num %>"><%= line_num %></ a></th> lib/redmine/codeset_util.rb:21:in `replace_invalid_utf8' app/helpers/repositories_helper.rb:146:in `to_utf8_internal' app/helpers/repositories_helper.rb:121:in `to_utf8' app/views/repositories/annotate.rhtml:17:in `_run_rhtml_app47views47repositories47annotate46rh tml' passenger (3.0.9) lib/phusion_passenger/rack/request_handler.rb:96:in `process_request' passenger (3.0.9) lib/phusion_passenger/abstract_request_handler.rb:513:in `accept_and_process _next_request' passenger (3.0.9) lib/phusion_passenger/abstract_request_handler.rb:274:in `main_loop' passenger (3.0.9) lib/phusion_passenger/classic_rails/application_spawner.rb:321:in `start_req uest_handler' passenger (3.0.9) lib/phusion_passenger/classic_rails/application_spawner.rb:275:in `send' passenger (3.0.9) lib/phusion_passenger/classic_rails/application_spawner.rb:275:in `handle_sp awn_application' passenger (3.0.9) lib/phusion_passenger/utils.rb:479:in `safe_fork' passenger (3.0.9) lib/phusion_passenger/classic_rails/application_spawner.rb:270:in `handle_sp awn_application' passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:357:in `__send__' passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:357:in `server_main_loop' passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:206:in `start_synchronously' passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:180:in `start' passenger (3.0.9) lib/phusion_passenger/classic_rails/application_spawner.rb:149:in `start' passenger (3.0.9) lib/phusion_passenger/spawn_manager.rb:219:in `spawn_rails_application' passenger (3.0.9) lib/phusion_passenger/abstract_server_collection.rb:132:in `lookup_or_add' passenger (3.0.9) lib/phusion_passenger/spawn_manager.rb:214:in `spawn_rails_application' passenger (3.0.9) lib/phusion_passenger/abstract_server_collection.rb:82:in `synchronize' passenger (3.0.9) lib/phusion_passenger/abstract_server_collection.rb:79:in `synchronize' passenger (3.0.9) lib/phusion_passenger/spawn_manager.rb:213:in `spawn_rails_application' passenger (3.0.9) lib/phusion_passenger/spawn_manager.rb:132:in `spawn_application' passenger (3.0.9) lib/phusion_passenger/spawn_manager.rb:275:in `handle_spawn_application' passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:357:in `__send__'</pre>			

```
passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:357:in `server_main_loop'
passenger (3.0.9) lib/phusion_passenger/abstract_server.rb:206:in `start_synchronously'
passenger (3.0.9) helper-scripts/passenger-spawn-server:99
```

As you can see, it was trying to show the output of calling 'svn annotate' on a medium (~18MB) tar.gz file, while trying to interpret its contents as UTF-8 text.

We already have a setting for limiting the maximum size of the text files which can be viewed (instead of forcing the download) using 'View'. I think we should use the same limit for annotating the files, and also prevent Redmine from annotating binary files (using the same heuristics as in 'View').

I have attached a patch which modifies 'annotate' so it checks that the file is really a text file under the specified size limit before trying to annotate it. Otherwise, it displays a new error message explaining why it could not be annotated.

Associated revisions

Revision 7728 - 2011-11-05 04:47 - Toshi MARUYAMA

do not annotate text files which exceed the size limit for viewing (#9484)

Contributed by Antonio García-Domínguez.

Revision 7729 - 2011-11-05 04:47 - Toshi MARUYAMA

update i18n locale yamls (#9484)

History

#1 - 2011-10-28 03:18 - Toshi MARUYAMA

Great work!

Could you add a test such as [source:tags/1.2.1/test/functional/repositories_filesystem_controller_test.rb#L115](#) ?

#2 - 2011-10-30 00:13 - Antonio García-Domínguez

Toshi MARUYAMA wrote:

Great work!

Could you add a test such as [source:tags/1.2.1/test/functional/repositories_filesystem_controller_test.rb#L115](#) ?

I'm glad you liked the patch :-). Sorry, I forgot to watch the issue and didn't see your message until now. I'll write it tomorrow: it's midnight right now in Spain. Basically, I just need to test that the proper error is shown when trying to annotate a file that is too big, right?

#3 - 2011-10-30 01:22 - Toshi MARUYAMA

Yes. Please see [source:tags/1.2.1/test/functional/repositories_git_controller_test.rb#L306](#)

#4 - 2011-10-30 16:47 - Antonio García-Domínguez

- *File revised-9484.patch added*

I have revised the patch with a new test case for the Subversion test suite. I had to change the new code a bit as well, as it was breaking one of the functional tests for Git repositories. It seems that with the old patch, a different error message would have been shown when trying to annotate binary files.

I have tried this one with 'rake test RAILS_ENV=test' using a SQLite database: all the tests seem to pass. I haven't run the tests that require an LDAP server, though.

#5 - 2011-11-05 06:27 - Toshi MARUYAMA

- *Status changed from New to Closed*

- *Target version set to 1.3.0*

Committed in trunk [r7728](#) and [r7729](#), thanks.

Files

0001-Fix-for-expensive-annotations.patch	4.1 KB	2011-10-28	Antonio García-Domínguez
revised-9484.patch	4.79 KB	2011-10-30	Antonio García-Domínguez