

RhodeCode CE/EE - Task #4251

[customer] Pull request with subrepos

29.09.2016 13:41 - Marcin Kuzminski [CTO]

Status:	Feedback	Start date:	29.09.2016
Priority:	High	Due date:	
Assignee:	Martin Bornhold	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v4.5	Spent time:	0.00 hour
Sorting:		Commit Number:	
Description			
<p>While a lot of time passed since we've exchanged emails last time, let me share some details about the way we work at R and what kind of Rhodecode improvements would be useful for us.</p> <ul style="list-style-type: none">• Most of our iOS/Mac product repositories are Mercurial• We extensively use Mercurial's subrepositories for internal dependencies: our internal libraries are Mercurial or git repositories, which are added to product repository as a subrepo.• I like github flow: creating feature branches (bookmarks in Mercurial), opening pull request to master (default), code-reviewing this pull request and merging it after review is finished <p>Unfortunately Rhodecode does not allow to create pull request for Mercurial repository in case it contains any subrepo, even if there were no changes in that sub-repositories.</p> <p>So that's the main reason why we can't use github flow with Rhodecode and that's the reason why we have to Crucible.</p> <p>I've shared these details with team about a year ago : https://rhodecode.tenderapp.com/help/discussions/support-tickets/1234-unable-to-create-pull-request</p> <p>Currently Rhodecode is more like a repository storage for us and while post-commit reviews are possible they are rarely used. That's why most people at R almost never see Rhodecode web interface.</p> <p>Let me know if you need any more details or have any workaround that would still allow us to use pull requests.</p>			

History

#1 - 11.10.2016 14:25 - Martin Bornhold

- Status changed from New to In Progress

#2 - 11.10.2016 14:25 - Martin Bornhold

- Assignee set to Martin Bornhold

#3 - 12.10.2016 10:12 - Marcin Kuzminski [CTO]

- Description updated

- Private changed from Yes to No

#4 - 12.10.2016 11:47 - Martin Bornhold

While investigating on this i noticed that we already support pull requests if sub-repositories are used. I tested it with a root mercurial repository with mercurial and git subrepositories. I was able to create pull request from a fork back to the root repo and also to create pull requests from one bookmark to another bookmark inside of the root repo. For me it looks like everything works fine. At least the pull requests got successfully merged, the commits are available in the target repo of the pull request and the subrepositories are available after a clone of the root repo.

Maybe we have a different workflow how to use pull requests or sub repositories or we are missing something here. It would be great if you can provide detailed step by step instructions how to reproduce your case. Also every error or warning message from RhodeCode will be useful.

These are the steps i did to test it:

1. Created a root mercurial repo in RhodeCode and cloned it locally
2. Add mercurial sub repo to the clone and push to the root repo
3. Created a fork of root repo in RhodeCode, cloned it locally, made changes and pushed them
4. Create PR from fork to root repo and merged it.
5. Add git sub repo in fork and pushed it.

6. Create PR from fork to root and merged it.
7. Cloned the root repo a second time and got all sub repos.

#5 - 19.10.2016 11:11 - Martin Bornhold

- Status changed from In Progress to Feedback

#6 - 20.10.2016 10:20 - Jasper van Zuijlen

We, too, run into this problem. Here is how we use it.

Our Rhodocode installation does not allow for anonymous access and the default user rights are 'none'. There are a lot of small separate teams and we don't want to bother the users with repos they do not need. We chose a forking workflow meaning that every developer has his own fork. To get changes back into the central repository a user has to create a pull request. This process of forking, reviewing and merging works very comfortably because of the full rhodocode integration. Superb.

The problem arises with repos that use subrepo(s); and we are using that more and more. When we create a pull request on a repo with subrepo(s) the creation of the pull request fails because of this:

```
2016-10-19 17:22:30.156 INFO [rhodocode.lib.auth] user <AuthUser('id:9[USERNAME] ip:X.X.X.X auth:True')> auth
enticating with:RHODECODE_AUTH IS authenticated on func CompareController:compare
2016-10-19 17:22:30.490 INFO [rhodocode.lib.middleware.request_wrapper] IP: X.X.X.X Request to /main_repo/com
pare/rev@4787f2c66120169c3e35ff3aea0a5722837e1da3...rev@f35ccfb5bc231e0172fd71e1065088e937b93f66 time: 0.397s
2016-10-19 17:22:36.813 INFO [rhodocode.lib.base] IP: X.X.X.X User: <AuthUser('id:9[USERNAME] ip:X.X.X.X auth
:True')> accessed /forks/main_repo-USERNAME/pull-request/new [pullrequests.create]
2016-10-19 17:22:36.839 INFO [rhodocode.lib.auth] user <AuthUser('id:9[USERNAME] ip:X.X.X.X auth:True')> auth
enticating with:RHODECODE_AUTH IS authenticated on func PullrequestsController:create
/opt/rhodocode/store/rvbnwlhks4ixlpnlfmpfwh9sqmm287d-python2.7-SQLAlchemy-0.9.9/lib/python2.7/site-packages/s
qlalchemy/engine/default.py:579: SAWarning: Unicode type received non-unicode bind param value
param.append(processors[key] (compiled_params[key]))
2016-10-19 17:22:37.030 INFO [rhodocode.lib.utils] Logging action: `user_commented_pull_request:23` on repo: `<
Repository('525:main_repo')>` by user: <User('id:9:USERNAME')> ip:X.X.X.X
2016-10-19 17:22:42.604 INFO [root] MAIL SEND TO: [u'admin@email.com']
2016-10-19 17:22:43.215 INFO [rhodocode.lib.middleware.request_wrapper] IP: X.X.X.X Request to /sub_repo time
: 0.029s
2016-10-19 17:22:43.230 ERROR [rhodocode.controllers.pullrequests] Error occurred during sending pull request
Traceback (most recent call last):
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/controllers/pullrequests.py", line 452, in create
    description
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/model/pull_request.py", line 353, in create
    pull_request, created_by_user, 'create')
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/model/pull_request.py", line 377, in _trigger_pull_request_hook
    pull_request=pull_request)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/lib/hooks_utils.py", line 81, in trigger_log_create_pull_request_hook
    extras.update(pull_request.get_api_data())
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/model/db.py", line 3109, in get_api_data
    merge_status = PullRequestModel().merge_status(pull_request)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/model/pull_request.py", line 883, in merge_status
    resp = self._try_merge(pull_request)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/model/pull_request.py", line 937, in _try_merge
    pull_request, target_vcs, target_ref)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/model/pull_request.py", line 968, in _refresh_merge_state
    workspace_id, dry_run=True, use_rebase=use_rebase)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/lib/vcs/backends/base.py", line 408, in merge
    workspace_id, target_ref)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/lib/vcs/backends/hg/repository.py", line 658, in _maybe_prepare_merge_workspace
    self._local_clone(shadow_repository_path)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/lib/vcs/backends/hg/repository.py", line 543, in _local_clone
    hooks=False)
  File "/opt/rhodocode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodocode-enterprise-ce-4.4.1/lib/pyth
on2.7/site-packages/rhodocode/lib/vcs/exceptions.py", line 194, in wrapper
    raise _EXCEPTION_MAP[kind](*e.args)
RepositoryError: http authorization required for http://rhodocodeserver.xxx/rhodocode/sub_repo (in subrepo ext
ernals/sub_repo)
```

i.e. after the local clone of the central repository an update is done on the default branch causing the subrepos to be cloned and updated as well (which I don't think is necessary, really). Since no authentication is forwarded this fails, as described in the exception.

I tried doing the local clone using /opt/rhodecode/store/87myn79yI0h9yb0c524s5yibkkpp9y9w-python2.7-mercurial-3.8.4/bin/hg and managed to get it to work with a hgrc:

```
$ cat /etc/mercurial/hgrc
# system-wide mercurial configuration file
# See hgrc(5) for more information
[auth]
x.prefix = *
x.username = rhodecode
x.password = [PASSWORD]
x.schemes = http https
```

This user was created in rhodecode with superadmin rights.

Unfortunately this doesn't fix the problem because, from what I understood, the hg binary is not used, but instead the API is called directly. This apparently does not load the hgrc. I also understood that the hgrc is generated from the database table 'rhodecode_ui' but adding:

```
[auth] x.prefix * 1
[auth] x.username rhodecode 1
[auth] x.password [PASSWORD] 1
[auth] x.schemes http https 1
```

Or:

```
auth x.prefix * 1
auth x.username rhodecode 1
auth x.password [PASSWORD] 1
auth x.schemes http https 1
```

Doesn't seem to make a difference.

What does seem to work if the following patch:

```
$ diff /opt/rhodecode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodecode-enterprise-ce-4.4.1/lib/python2.7/site-packages/rhodecode/lib/vcs/backends/hg/repository_ori.py /opt/rhodecode/store/zdz5m20x9vhlp3ibbq14195mcshhja31-python2.7-rhodecode-enterprise-ce-4.4.1/lib/python2.7/site-packages/rhodecode/lib/vcs/backends/hg/repository.py
542c542
< self._remote.clone(self.path, clone_path, update_after_clone=True,
---
> self._remote.clone(self.path, clone_path, update_after_clone=False,
```

Any chance that'll make it into the next release?

#7 - 20.10.2016 16:11 - Marcin Kuzminski [CTO]

Update for initial problem:

Hi Marcin,

I've just checked this without positive result on Rhodecode 4.4.2 :

Steps:

- hg repo with one subrepo
- created two bookmarks
- tried to open pull request from one bookmark from other
- got Error occurred during sending pull request
- tried again with branches
- same result
- Logs from vcserver:
(multiple GET / FETCH)
2016-10-20 13:59:00.171 DEBUG [vcserver.base] GET GitFactory@/home/rhodecode/repos/appversioner with cache:True. Context: e9b68891-3d5f-40bd-a00a-805bfacc25e3
2016-10-20 13:59:00.172 DEBUG [vcserver.base] FETCH GitFactory@/home/rhodecode/repos/appversioner repo object from cache. Context: e9b68891-3d5f-40bd-a00a-805bfacc25e3
2016-10-20 13:59:00.172 DEBUG [Pyro4.core] Exception occurred while handling request: Exception("NotGitRepository(u'No git repository was found at /home/rhodecode/repos/appversioner',)",)
2016-10-20 13:59:00.173 DEBUG [vcserver.svn] Invalid Subversion path /home/rhodecode/repos/appversioner
2016-10-20 13:59:00.191 DEBUG [Pyro4.core] Exception occurred while handling request: Exception('http authorization required for https://hg.readable.com/appversioner (in subrepo appversioner)',)
2016-10-20 13:59:00.308 DEBUG [vcserver.base] GET MercurialFactory@/home/rhodecode/repos/experiments/andrian-hg-test with cache:True. Context: a860800a-e10d-43de-91f8-7cb4daa88520

2016-10-20 13:59:00.308 DEBUG [vcserver.base] FETCH MercurialFactory@/home/rhodecode/repos/experiments/andria n-hg-test repo object from cache. Context: a860800a-e10d-43de-91f8-7cb4daa88520 (multiple GET / FETCH)

Let me know if you need any additional information.

Best regards,

#8 - 25.10.2016 13:24 - Martin Bornhold

- Status changed from Feedback to In Progress

#9 - 26.10.2016 13:27 - Jasper van Zuijlen

Hi,

After some further investigation I too wasn't able to get it to work. The cause is that, when a pull-request is created or when a merge is made, Rhodecode does an update as well. The first case I disabled locally but the second case that is not possible (I think?). The problem in that is that the authentication info is not flowed to the sub-repos and since the global hgrc is not used (only a simulation via database settings which appears not to work the same). When anonymous login is disabled and/or the rights for the default user are set to none, the rhodecode local update is denied access, as it shows up as an anonymous request.

Once could argue that when rhodecode tries to access itself, I should allow itself to do so.

Regards,
Jasper

#10 - 26.10.2016 13:31 - Marcin Kuzminski [CTO]

We're currently investigating:

- passing in credentials to the subrepo call. However this produces several security problems.
- using subpath of Mercurial to map URL into a local path on filesystem
- patching mercurial to disable pulling in subrepos

As for the question on RhodeCode doing a self-auth. I think worth thinking about. Maybe some kind of combination of localhost access + authentication token would work but again top priority is not to create any vector of attack, or credentials leak here.

#11 - 30.10.2016 21:07 - Redmine Integration

pullrequest created by marcink (status: under_review). <https://internal-code.rhodecode.com/rhodecode-enterprise-ce/pull-request/2823>

#12 - 31.10.2016 00:12 - Redmine Integration

- Status changed from In Progress to Resolved

Commit e2506d632a44 by Marcin Kuzminski marcin@rhodecode.com on default branch changed this issue.

<https://internal-code.rhodecode.com/rhodecode-enterprise-ce/changeset/e2506d632a4475c80ceec60f9fe2abcece283296>

#13 - 31.10.2016 00:12 - Redmine Integration

pullrequest merged by marcink (status: approved). <https://internal-code.rhodecode.com/rhodecode-enterprise-ce/pull-request/2823>

#14 - 31.10.2016 00:15 - Marcin Kuzminski [CTO]

- Status changed from Resolved to In Progress

accidentally linked wrong commit

#15 - 10.11.2016 17:52 - Marcin Kuzminski [CTO]

- Status changed from In Progress to Feedback

We managed to make the basic example with authentication work by disabling the subrepos. In This way you can use pull requests with subrepos even without having access to subrepos. We hope this will enable support with your use cases.

We tested a basic example with adding a subrepo path to a non-public rhodecode instance. A full open/merge scenario is now working.

Relevant commits:

- <https://code.rhodecode.com/rhodecode-enterprise-ce/compare/rev@e31abe6080b9c2d08eb0f9617d5613f7914f9abf...rev@ebe0247cd154889e46e1b1638efa15728fc577d0>
- <https://code.rhodecode.com/rhodecode-vcserver/compare/rev@5682a2ff63e2d0d10d3e1e4aa81010bff74b5011...rev@2582dee7dd4c0f1028fd5a55cdf1913c32cb520a>