

RhodeCode CE/EE - Support #5269

Upgrade from RC EE 3.7.1 to RC EE 4.x

05.04.2017 10:24 - Anton Nikolaev

Status: New	Start date: 05.04.2017
Priority: Normal	Due date:
Assignee:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version:	Spent time: 0.00 hour
Sorting:	Affected Version:
Commit Number:	

Description

Hi all.

Sorry for my english

I used this instruction when upgrade from RC EE 3.7.1 to RC EE 4.x
<https://docs.rhodecode.com/RhodeCode-Control/tasks/upgrade-to-latest.html>

After that I have two problems

- Russian charset was broken For example

CYB -> Ъ;ЂЃ'

- Some push/pull (similar to 3.7.1) have stopped working or started to run very slowly.

History

#1 - 05.04.2017 10:37 - Marcin Kuzminski [CTO]

Hi,

Please check if default_encoding is set to proper value inside rhodecode.ini, if you don't use UTF8

as for slow push/pull, maybe you need to increase the number of workers=N inside rhodecode.ini/vcserver.ini ? We didn't have any reports on performance decrease after 3.X -> 4.X migration.

#2 - 24.04.2017 00:14 - Anton Nikolaev

- File before_after.png added

Information about versions for testing

```
[user@rhodecode ~]$ rccontrol status
```

```
- NAME: enterprise-2
- STATUS: RUNNING
- TYPE: Enterprise
- VERSION: 3.7.1
- URL: http://192.168.122.51:8080
```

```
- NAME: vcserver-1
- STATUS: RUNNING
- TYPE: VCSServer
- VERSION: 3.7.1
- URL: http://127.0.0.1:10001
```

```
[user@rhodecode ~]$ rccontrol status
RHODECODE CONTROL VERSION: 1.13.1
```

```
- NAME: enterprise-2
- STATUS: RUNNING
  logs:/home/user/.rccontrol/enterprise-2/enterprise.log
- VERSION: 4.2.1 Enterprise
```

```
- VCS: vcserver-1
- URL: http://192.168.122.118:8080
- CONFIG: /home/user/.rccontrol/enterprise-2/rhocode.ini
```

```
- NAME: vcserver-1
- STATUS: RUNNING
  logs:/home/user/.rccontrol/vcserver-1/vcserver.log
- VERSION: 4.2.1 VCSServer
- URL: http://127.0.0.1:10001
- CONFIG: /home/user/.rccontrol/vcserver-1/vcserver.ini
```

Marcin Kuzminski

Please check if default_encoding is set to proper value inside rhocode.ini, if you don't use UTF8

The same values for default_encoding

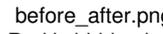
RC Enterprise 3.7.1

```
[user@rhocode ~]$ cat .rccontrol/enterprise-2/rhocode.ini | grep default_encoding
default_encoding = utf8,cp866
```

RC Enterprise 4.2.1

```
[user@rhocode ~]$ cat .rccontrol/enterprise-2/rhocode.ini | grep default_encoding
default_encoding = utf8,cp866
```

But

 before_after.png
Red is hidden information

Interface localization to russian - OK

Latin characters - OK

Russian characters in files - OK

Russian characters in commit messages, users first name and second name, groups descriptions and etc was **broken**. Examples in the green rectangle on screenshot.

My solution for every column in DB which contain informations (commit messages, groups descriptions and etc)

```
ALTER TABLE tablename MODIFY column TEXT CHARSET latin1;
ALTER TABLE tablename MODIFY column BLOB;
ALTER TABLE tablename MODIFY column TEXT CHARSET utf8;
```

In RC DB it's

Table	Columns
changeset_comments	text
groups	group_description
users	firstname,lastname
etc	etc

Your scripts changing DB encoding from latin1 to utf8 during update from 3.x to 4.x. Why is this necessary?

I think that doing it manually is bad.

I propose two solutions for the next RC version

- 1) Don't change DB encoding during update
- 2) Do the same as I indicated, but in automatic mode.

Marcin Kuzminski

as for slow push/pull, maybe you need to increase the number of workers=N inside rhocode.ini/vcserver.ini ?

The same values for parameter "workers" in both versions.

RC Enterprise 3.7.1

```
[user@rhocode ~]$ cat .rccontrol/vcserver-1/vcserver.ini | grep workers
```

```
# where N is number of RhodeCode Enterprise workers, eg. running 2 instances
# 8 gunicorn workers each would be 2 * 8 * 6 = 96, threadpool_size = 96
[user@rhodecode ~]$ cat .rccontrol/enterprise-2/rhodecode.ini | grep workers
workers = 3
## Sets the number of process workers. You must set `instance_id = *` when this option
#workers = 1
```

RC Enterprise 4.2.1

```
[user@rhodecode ~]$ cat .rccontrol/vcserver-1/vcserver.ini | grep workers
# where N is number of RhodeCode Enterprise workers, eg. running 2 instances
# 8 gunicorn workers each would be 2 * 8 * 6 = 96, threadpool_size = 96
[user@rhodecode ~]$ cat .rccontrol/enterprise-2/rhodecode.ini | grep workers
workers = 3
## Sets the number of process workers. You must set `instance_id = *` when this option
#workers = 1
```

I can't talk about real repositories, so I tested it with this repo - <http://hg.hedgewars.org/hedgewars/>, I have had problems with the push / pull that I described earlier.

What information can I still provide to solve this problem? Metrics, parameters and etc.

Marcin Kuzminski

We didn't have any reports on performance decrease after 3.X -> 4.X migration.

So I'll be the first who will report about it :)

#3 - 24.04.2017 10:13 - Marcin Kuzminski [CTO]

Hi,

Thanks for your feedback. As for database encoding, we wasn't aware of this problem sorry, I believe it's our migration scripts that could be doing this, but it wasn't intended. Can i ask if only new tables were created in utf8 or everything was changed to utf8 ? because we generally recommend using utf8 for everything, and our new tables are created in such mode by default.

Can you show the problem with push with hg --debug push ? also provide your HTTP server(Nginx/Apache) logs if it sits in front of RhodeCode server.

#4 - 18.06.2017 20:07 - Anton Nikolaev

- File rh1.png added

Marcin Kuzminski [staff] wrote:

Hi,

Thanks for your feedback. As for database encoding, we wasn't aware of this problem sorry, I believe it's our migration scripts that could be doing this, but it wasn't intended. Can i ask if only new tables were created in utf8 or everything was changed to utf8 ? because we generally recommend using utf8 for everything, and our new tables are created in such mode by default.

Our company use RC from version 1.x with MySQL. Default encoding in MySQL latin1. Which solutions can you recommend for me?

Can you show the problem with push with hg --debug push ? also provide your HTTP server(Nginx/Apache) logs if it sits in front of RhodeCode server.

I did a lot of experiments with different OS. With HTTP servers and without it.

The result of one experiment

OS:

```
user@u16rh4:~$ uname -a
Linux u16rh4 4.4.0-21-generic #37-Ubuntu SMP Mon Apr 18 18:33:37 UTC 2016 x86_64 x86_64 x86_64 GNU/Linux
```

Install clean RC in clean OS

```
user@u16rh4:~$ rccontrol status
RHODECODE CONTROL VERSION: 1.14.0
```

```
- NAME: community-1
```

```
- STATUS: RUNNING
  logs:/home/user/.rccontrol/community-1/community.log
- VERSION: 4.7.2 Community
- VCS: vcserver-1
- URL: http://192.168.122.95:10002
- CONFIG: /home/user/.rccontrol/community-1/rhocodecode.ini
```

```
- NAME: vcserver-1
- STATUS: RUNNING
  logs:/home/user/.rccontrol/vcserver-1/vcserver.log
- VERSION: 4.7.2 VCServer
- URL: http://127.0.0.1:10001
- CONFIG: /home/user/.rccontrol/vcserver-1/vcserver.ini
```

Without HTTP server

I'm create repo

rh1.png

Clone Hedgewars from official site to non-RC pc

```
$ hg clone http://hg.hedgewars.org/hedgewars/
```

And push it to RC by command from screen with debug

```
$ hg --version
Mercurial Distributed SCM (version 2.8.2)
(see http://mercurial.selenic.com for more information)
```

```
Copyright (C) 2005-2013 Matt Mackall and others
This is free software; see the source for copying conditions. There is NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
```

```
$ hg push http://admin@192.168.122.95:10002/hedgewars --debug
```

Part of log

```
$ hg push http://admin@192.168.122.95:10002/hedgewars --debug
pushing to http://admin@192.168.122.95:10002/hedgewars
using http://192.168.122.95:10002/hedgewars
http auth: user admin, password not set
sending capabilities command
query 1; heads
sending batch command
searching for changes
taking quick initial sample
searching: 2 queries
query 2; still undecided: 12404, sample size is: 100
sending known command
sampling from both directions
searching: 3 queries
query 3; still undecided: 1723, sample size is: 200
sending known command
3 total queries
12432 changesets found
list of changesets:
475c0f2f9d17f770440c1ec533a9c48390523ad8
30f2d1037d5da65660b237d1bd2e37ea6b39f7c8
4eeab397c3c6f9e71873a82a61ab77baa2a92655
. . .
sending unbundle command
sending 643033094 bytes
sending: 64/1255924 kb (0.01%)
sending: 128/1255924 kb (0.01%)
sending: 192/1255924 kb (0.02%)
sending: 256/1255924 kb (0.02%)
sending: 320/1255924 kb (0.03%)
sending: 384/1255924 kb (0.03%)
sending: 448/1255924 kb (0.04%)
sending: 512/1255924 kb (0.04%)
sending: 576/1255924 kb (0.05%)
sending: 640/1255924 kb (0.05%)
sending: 704/1255924 kb (0.06%)
abort: error: Connection reset by peer
$
```

Full log in bash_log.txt

RC without ERROR

```
user@u16rh4:~/rccontrol$ pwd
/home/user/rccontrol
user@u16rh4:~/rccontrol$ grep -r "INFO" . | wc -l
203
user@u16rh4:~/rccontrol$ grep -r "ERROR" . | wc -l
0
```

#5 - 19.06.2017 23:38 - Marcin Kuzminski [CTO]

Hi Anton,

Little bit confused here, is the problem still encoding issues, or is it now a push issue ?

Please take a look at /home/user/rccontrol/vcsserver-1/vcsserver.log for any Exceptions ?

#6 - 20.06.2017 17:03 - Anton Nikolaev

- File vcsserver.log added

Marcin Kuzminski [staff] wrote:

Hi Anton,

Little bit confused here, is the problem still encoding issues, or is it now a push issue ?

In issue's description I wrote

After that I have two problems

Please take a look at /home/user/rccontrol/vcsserver-1/vcsserver.log for any Exceptions ?

Add full log - vcsserver.log

#7 - 29.06.2017 19:21 - Anton Nikolaev

Do you have any new information?

#8 - 29.06.2017 19:36 - Marcin Kuzminski [CTO]

Hi Anton,

No sorry we couldn't figure out anything relevant from the logs. I quick question would it be possible to have somehow a remote session so we can check your instance ourself ? This looks like something obvious to fix, but we're missing it. Maybe a remote sesison would be productive and we could find a cool print for your troubles.

Files

before_after.png	34 KB	23.04.2017	Anton Nikolaev
rh1.png	37 KB	18.06.2017	Anton Nikolaev
vcsserver.log	16.1 KB	20.06.2017	Anton Nikolaev