

CUNY Academic Commons - Bug #2479

Double posting, double notification of message posted via RBE

2013-02-25 04:27 PM - Matt Gold

Status:	Resolved	Start date:	2013-02-25
Priority name:	Urgent	Due date:	
Assignee:	Raymond Hoh	% Done:	50%
Category name:	BuddyPress (misc)	Estimated time:	0.00 hour
Target version:	1.4.21		
Description			
Just received a double notification of a message I posted via RBE to a group. Looks like the message posted twice, as well. Saw something similar when another group member posted earlier today. Please see the attached screenshots. This is for the following group:			

History

#1 - 2013-02-25 06:08 PM - Boone Gorges

- Assignee changed from Boone Gorges to Raymond Hoh
- Priority name changed from Normal to High
- Target version set to 1.4.20

Ray, could you have a look please? Thanks.

#2 - 2013-02-25 06:56 PM - Raymond Hoh

Hi,

Can an admin deactivate RBE from the main CAC site and then reactivate it?

It appears there were problems connecting to GMail yesterday and there might have been a race condition where RBE tried to connect twice after failure.

#3 - 2013-02-25 07:00 PM - Boone Gorges

No problem, Ray - I just toggled it.

#4 - 2013-02-28 10:19 AM - Matt Gold

- Priority name changed from High to Urgent

Hi All --

Just saw another instance of this here - <http://commons.gc.cuny.edu/groups/digital-dissertations/forum/topic/preservation-questions/#post-13006>

Upgrading the priority on this, though, truly, I think it falls somewhere between "high" and "urgent"

#5 - 2013-02-28 12:31 PM - Raymond Hoh

I need to do a better job at detecting IMAP connections. Since this is a pretty hard issue to debug and test for, I might not be able to address this for 1.4.20.

Boone: I'm thinking of ditching WP cron and using a transient marker instead. I have an idea of how all this may work in my head, but not sure if this will work as I implement it ;)

A temporary fix is checking if the forum reply already exists in the database. BP 1.6 introduced this functionality, but RBE bypasses this as the reply check is only used on frontend posts. This can definitely be done for 1.4.20, but it doesn't address the underlying issue.

#6 - 2013-02-28 01:28 PM - Boone Gorges

Boone: I'm thinking of ditching WP cron and using a transient marker instead.

That's OK by me. Here's a super-simple version I posted recently: <https://gist.github.com/boonebgorges/4714650#file-gistfile1-php> (not using

transients, but you get the idea) FWIW, this is basically how wp-cron works, except wp-cron is standardized. That said, wp-cron has weird limits on how often you can run an event - I think it's once every 10 minutes. Is that part of the problem with the IMAP stuff?

A temporary fix is checking if the forum reply already exists in the database.

I think that a good feature for RBE would be general dupe detection. Something like:
- Require an abstract method `BP_Reply_By_Email_Extension::detect_existing_item()`.
- Check that `! $extension->detect_existing_item()` before `$extension->post_by_email()`

I know you're not currently using your Extension class for BP-native content types, but the same principle could apply there too. We're only talking about maybe 4 or 5 different content types, and I'm pretty sure there's enough information between the email (timestamp, user name, etc) and the database item to reliably detect dupes in each case.

I know this isn't a trivial thing to implement, but it seems like it would be a good failsafe, and maybe easier to address in the long run than difficult-to-test IMAP issues.

What do you think?

#7 - 2013-03-01 02:50 PM - Boone Gorges

- Target version changed from 1.4.20 to 1.4.21

#8 - 2013-03-04 03:57 PM - Raymond Hoh

- % Done changed from 0 to 50

I decided to add a stop-gap before I go and overhaul the current WP-cron connection method.

In [commit 12761b2](#), I've added a locking mechanism to RBE.

According to the RBE logs, on occasion, RBE would attempt to connect twice under a minute. This would cause the double posts encountered in this ticket.

The commit sets a transient before connecting to the IMAP server, which expires after 60 seconds. So if a second attempt to connect is made within 60 seconds, that attempt would be stopped.

I think 60 seconds is a long-enough time limit to connect to GMail and should hopefully stop these double connection instances.

I'll also add the duplicate reply exists check in my next commit to CAC.

I think that a good feature for RBE would be general dupe detection. Something like:
- Require an abstract method `BP_Reply_By_Email_Extension::detect_existing_item()`.
- Check that `! $extension->detect_existing_item()` before `$extension->post_by_email()`

Good idea, Boone. Will look into it.

#9 - 2013-03-05 09:33 AM - Boone Gorges

Awesome - thanks, Ray!

#10 - 2013-03-08 03:34 PM - Matt Gold

Hi Guys -- just experienced this again - <http://commons.gc.cuny.edu/groups/gc-digital-fellows/forum/topic/first-year-digital-program-idea/#post-13160>

#11 - 2013-03-11 02:05 PM - Boone Gorges

- Status changed from Assigned to Resolved

Ray - because it sounds like your stopgap is probably going to resolve this issue for now, and I want to clear the milestone, I'm going to mark this one resolved. If you decide to go forward with the dupe check, we can open a separate ticket for the enhancement (or just roll it in the next time we upgrade RBE).

#12 - 2013-03-11 03:23 PM - Raymond Hoh

I'm fine with that, Boone.

I've noticed that the RBE log is getting rather large (it's around 8MB right now). We might want to archive the current log and start fresh.

#13 - 2013-03-12 08:59 AM - Boone Gorges

Good idea, Ray. I've archived it to `rbe.log.1` and wiped the content of `rbe.log`.

#14 - 2013-04-08 12:12 AM - Matt Gold

- File Screen Shot 2013-04-08 at 12.11.22 AM.png added

- File Screen Shot 2013-04-08 at 12.11.33 AM.png added

Hi All --

Caught what I think might be another instance of this in the CUNY IRT group. Please see attached screenshots and let me know what you think. Thanks.

#15 - 2013-04-09 01:20 AM - Raymond Hoh

Hi Matt,

I took a look at the RBE log and checked GMail.

Before the double post occurred, RBE got stuck. When RBE tried to reboot itself, RBE somehow ended up connecting to GMail twice at close to the same time, which led to both connections parsing the same email message and thus, the double post.

This looks to be a fringe case, however let me know if this happens again.

I still haven't added the duplicate forum post check that was mentioned earlier in this thread, but I'll do that tomorrow for bbPress 1 and that should alleviate the problem for the time being.

#16 - 2013-04-09 02:04 AM - Matt Gold

okay -- cool. Thanks, Ray.

Files

Screen Shot 2013-02-25 at 11.25.34 AM.png	154 KB	2013-02-25	Matt Gold
Screen Shot 2013-02-25 at 11.27.05 AM.png	42.8 KB	2013-02-25	Matt Gold
Screen Shot 2013-02-25 at 11.27.13 AM.png	20.8 KB	2013-02-25	Matt Gold
Screen Shot 2013-04-08 at 12.11.22 AM.png	124 KB	2013-04-08	Matt Gold
Screen Shot 2013-04-08 at 12.11.33 AM.png	92.9 KB	2013-04-08	Matt Gold