CUNY Academic Commons - Bug #2796

Sluggish Response on the Site - Sunday 9/15 - 3PM - 6PM

2013-09-15 05:56 PM - scott voth

Status:	Resolved	**Start date:**	2013-09-15
Priority name:	Normal	**Due date:**	
Assignee:	Boone Gorges	**% Done:**	0%
Category name:		**Estimated time:**	0.00 hour
Target version:	1.5.4		

Description

Site seems very slow this afternoon and is taking a long time to respond. Are there server issues?

History

#1 - 2013-09-15 06:10 PM - Boone Gorges

- Status changed from New to Assigned

- Assignee set to Matt Gold

- Priority name changed from Normal to Urgent

- Severity set to Critical

I can see from running top that we're serving lots of pages, and mysqld is getting bogged down. It could be that we're experiencing some sort of DOS attack, or that some bug is ratcheting our traffic in some way. But I can't tell without accessing the logs in /var/log/httpd, and they're currently off-limits to the commons user.

Matt, you'll have to open a ticket with IT to see if they can shed some light on the situation. If they aren't interested in doing active debugging, they can send a copy of the access log, and I can start with that.

#2 - 2013-09-15 06:28 PM - Matt Gold

Thanks for the report, Scott, and for the update, Boone. I've opened a ticket with IT and cc'ed Boone on the message. Hopefully, we'll hear a response soon, but it being Sunday night, I'm not sure . . .

#3 - 2013-09-15 07:43 PM - scott voth

Seems to be fixed now.

#4 - 2013-09-19 12:10 PM - Boone Gorges

- Status changed from Assigned to Rejected

Matt has followed up with IT, and we'll continue to examine this off the thread. I'm going to close the ticket for now as it appears that there's no actual bug on the Commons, but that this was a network issue. I'll reopen with details when and if they emerge, and are Commons-specific.

#5 - 2013-09-25 05:23 PM - Boone Gorges

- Status changed from Rejected to Assigned

- Assignee changed from Matt Gold to Boone Gorges

- Priority name changed from Urgent to Normal

- Target version set to 1.5.4

I've had a chance to examine the logs.

From the looks of things, the Commons is pretty regularly attacked by bots. Most of them are brute-force attacks on wp-login.php. This sort of attack is blocked by some .htaccess rules I instituted some time ago, resulting in a 301 redirect to a static html page. It's highly unlikely that this sort of attack (which is ongoing) is related to the slowdown.

I do see a pretty major anomaly with registration. In the logs, I see a series of attempts that look like this:

142.91.111.107 - - [15/Sep/2013:05:17:31 -0400] "GET /register/ HTTP/1.0" 200 44092
142.91.111.107 - - [15/Sep/2013:05:17:32 -0400] "POST /register/ HTTP/1.0" 200 44649

An IP address GETs /register/, and then POSTs to it. I'm guessing this is a spambot, attempting to register for splogs. Apparently, these attempts are

failing - probably because of our *.cuny.edu whitelist. However, on the Sunday in question, we were getting fielding these requests on the order of 20-40 times per minute. Each one of these attempts loads the entirety of WordPress - twice, actually - before finally disallowing the registration.

For wp-login.php, I'm using a trick that denies POST requests that don't originate from one of our domains - this catches most brute-force login attempts. The same technique won't work for the /register/attack, because they appear to be moving through the /register/ page. (At least, some of them do.) Stuff like CAPTCHAs, honeypots, etc also are unlikely to do much, because (a) we're not having a spam problem, and (b) these techniques require loading WP, which is the root of our problem.

One idea might be to move the /register/ page to a different URL, one that BP spammers are less likely to guess. This is a pretty thin veil, but it might help. Then again, it might cause other problems with incoming links if we were to change it.

Ray, Dom, I'm adding you as watchers in case you have any ideas about how this might be tackled at the Apache level.

#6 - 2013-09-25 05:26 PM - Matt Gold

Thanks for your work on this, Boone.

#7 - 2013-09-25 05:42 PM - Dominic Giglio

I agree that htaccess is where the fix for such attacks/requests belong.

I've been following Jeff Starr over at perishablepress.com for a while now and he has a great set of htaccess rules for optimizing and protecting a WordPress site. He calls them "G Series" blacklists. The current version is the 5G Blacklist. You can read all about it and see what he includes and why here:

http://perishablepress.com/5g-blacklist-2013/

I'm not saying we need the whole thing but there are a bunch of great/advanced htaccess best practices that could help us out. One of his hobbies is server monitoring so he's done a fair amount of tweaking and customization when it comes to that frontline wall provided by apache/htaccess rules.

#8 - 2013-09-25 07:10 PM - Raymond Hoh

I am going to second Jeff Starr's 5G .htaccess list. Was about to post the same thing, Dom!

We can also try using the Cookies for Comments plugin and use the .htaccess rules that that plugin recommends, but use the BP registration slug in place of wp-signup.php as well.

#9 - 2013-09-25 07:23 PM - Dominic Giglio

Great minds think alike!! :-)

#10 - 2013-09-25 09:34 PM - Boone Gorges

Thanks for the thoughtful feedback, guys.

A lot of these 5G rules look pretty good, but I don't see at a glance which ones will help with this specific issue. Maybe the user-agent check? Though that won't help either if the bot is spoofing.

Good idea about the cookie check, Ray. I'd considered a similar trick. It'd make registration impossible for users with cookies turned off, but then again you need cookies to use WP as a logged-in user anyway. I'm going to play with the CFC cookie plugin to see how it works - I see it's doing something clever with an auto-generated css file, and I'd like to understand it better.

#11 - 2013-09-26 12:30 AM - Raymond Hoh

Just created a companion plugin that blocks logins and BP registrations if the cookie generated by Cookies for Comments does not exist.

Feel free to give that a try as well.

Note: The plugin is meant to supplement Cookies for Comments. The .htaccess rules will be infinitely better! With my plugin, you can now use the .htaccess rules on wp-login.php as well.

Something like this:

```
# on the wp-login.php page, check if a POST request is made and check if our special cookie exists
# if not, forbid the login attempt from happening
RewriteCond %{REQUEST_METHOD} POST
RewriteCond %{HTTP_COOKIE} !^.*SPECIAL-COOKIE-HASH-HERE.*$
RewriteRule ^wp-login.php - [F,L]
```

#12 - 2013-10-01 08:41 PM - Boone Gorges

- Status changed from Assigned to Resolved

The plugin looks pretty nifty, Ray. Thanks!

In https://github.com/castiron/cac/commit/cbec00c1642c4a4a8ba9ca8923685441e0ad17ee, I added cookies-for-comments along with Ray's companion plugin, and made the .htaccess changes in support of the cookie check. For the time being, I'm not making any more changes to the .htaccess file - some of them look like good protection against hacks, but that's not exactly what we're looking at in this ticket, and I only want to introduce one potentially breaking change at a time into the voodoo of our .htaccess.

Tentatively marking this as resolved.