

CUNY Academic Commons - Bug #10997

Website disappeared

2019-01-23 03:05 PM - Lisa Rhody

Status:	Resolved	Start date:	2019-01-23
Priority name:	High	Due date:	
Assignee:	Boone Gorges	% Done:	0%
Category name:		Estimated time:	0.00 hour
Target version:	1.14.6		
Description			
We (Steve Zweibel and I) built a site at dataforgood.commons.gc.cuny.edu, and checking it today, it has been overwritten by the basic template. We'd like to recuperate the old site if possible, as we seem to have also lost the data. Thank you!			
Related issues:			
Related to CUNY Academic Commons - Support #11017: site didn't save?		Resolved	2019-01-25

History

#1 - 2019-01-23 03:20 PM - Boone Gorges

- File Screenshot_2019-01-23_14-19-13.png added

I'm seeing a site, but it looks like it has content from the **teaching** template. Can you confirm? See screenshot.

#2 - 2019-01-23 04:26 PM - Boone Gorges

- Priority name changed from Normal to High

I'm looking into this and leaving some notes for myself and Ray:

The mysqlbinlog shows that the site was reset (TRUNCATE...) at 20:20:17 on 2019-01-15. There happens to have been a corresponding post on the dataforgood site edited (apparently legitimately) just seconds before that happened, but I don't see the connection. In fact, I only see a small number of hits on the dataforgood website during the time period in question, which makes me suspect that post editing was happening via Jetpack or XML-RPC or something like that.

My assumption is that the cac-site-templates feature kicked in incorrectly, but I can't understand how it would happen. It is triggered by wpmu_new_blog, which passes the ID of the new site to the hook. It's not possible to spoof, as this value is not provided by the user. It could be that another plugin incorrectly **triggered** this hook, but I haven't been able to see how.

I have to step away but I'll take a closer look later. Ray, if you have any ideas, I'd welcome them.

Lisa, it's likely that I'll be able to piece together the missing posts, but it will be more complex to get the missing settings (theme, etc). I'm hopeful that those parts won't be as hard to reproduce manually.

#3 - 2019-01-23 05:04 PM - Boone Gorges

Another thought is that the site creation process sometimes takes so long [#10701](#) that the process hangs, and it could be that there are very isolated instances where the request is left open for hours or for days, with the wpmu_new_blog routine run at the end of this period of time.

I'm going to run some more tests to decrease the query overhead of site creation (I see that something is writing to the database hundreds of times to copy over some user roles) and I'm also going to put in place a failsafe that prevents the site template clone mechanism from running if the site was created more than, say, five minutes ago.

#4 - 2019-01-23 06:01 PM - Boone Gorges

- Assignee set to Boone Gorges

- Target version set to 1.14.6

Here's one bit of fairly low-hanging fruit in WP, which I've patched with a similar workaround for the Commons:

<https://core.trac.wordpress.org/ticket/46087><https://github.com/cuny-academic-commons/cac/commit/3917844a1f26422adb6f6483f527dd62f057cf7b>

#5 - 2019-01-23 09:05 PM - Boone Gorges

Another update. I've been reading through the bin logs from around the time of the event, and it appears that it overlaps directly with some testing that I was doing related to [#10961](#). That bug caused site creation to abort prematurely in some cases. I'm not 100% sure of the mechanism, but it seems

likely that there's a connection. While this isn't exactly good news, it does strongly suggest that this is a one-off occurrence that came up because of a specific server issue, and is not a software-level bug.

I'll spend some time putting together what I can of the site content.

#6 - 2019-01-23 09:24 PM - Boone Gorges

Scratch that previous comment - I was looking at the wrong thing. The template overwrite happened the day after the firewall issues.

#7 - 2019-01-23 09:29 PM - Boone Gorges

The gap between the last legitimate site update and the data loss was over 24 hours, which means that IT should have a full-fidelity database backup. I've just put in a request to have it restored, so that I can replace the affected tables wholesale.

#8 - 2019-01-23 11:56 PM - Lisa Rhody

Thanks for looking into it in detail, Boone!

#9 - 2019-01-24 03:34 PM - Boone Gorges

- *Status changed from New to Resolved*

Hi Lisa - The content has been restored. Thanks for your patience and sorry for the trouble.

#10 - 2019-01-24 06:18 PM - Lisa Rhody

Thanks, Boone!

#11 - 2019-01-25 03:46 PM - Boone Gorges

- *Related to Support #11017: site didn't save? added*

Files

Screenshot_2019-01-23_14-19-13.png	137 KB	2019-01-23	Boone Gorges
------------------------------------	--------	------------	--------------