

D4Science Infrastructure - Task #12827

Configure CKAN to store .cache files (that are session info) in memory

Nov 06, 2018 03:05 PM - Francesco Mangiacrapa

Status:	Closed	Start date:	Nov 06, 2018
Priority:	Normal	Due date:	
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	Application	Estimated time:	0.00 hour
Target version:	Data Catalogue		
Infrastructure:	Development, Production		
Description			
<p>@tommaso.piccioli@isti.cnr.it has noticed under '/tmp/[ckan_name]/sessions' directory that millions of file are created by CKAN. Investigating the issue we have found (from https://lists.okfn.org/pipermail/ckan-dev/2015-February/008648.html) the following one:</p> <p>The cache directory or the session directory is only used for stuff that is not saved in the database.</p> <p>It is always safe to remove. It is only used for very temporary data.</p> <p>CKAN only uses it in fact to store the little popup messages that come up after you are redirected to a new page.</p> <p>... but those files are never deleted.</p> <p>We need and can (in order to avoid filling the disk) to pass from caching on disk to cacking in memory by installing and configuring memcached (see at https://github.com/memcached/memcached) for pylons and beaker.</p> <p>How to install and use it:</p> <p>https://stackoverflow.com/questions/13781023/how-to-configure-ckan-to-store-sessions-somewhere-other-than-disk-e-g-db-or-me</p> <p>https://beaker.readthedocs.io/en/latest/sessions.html#deleting</p>			

History

#1 - Nov 20, 2018 06:59 PM - Andrea Dell'Amico

- Status changed from New to In Progress

Going to reconfigure the dev instance.

#2 - Nov 20, 2018 07:44 PM - Andrea Dell'Amico

- % Done changed from 0 to 70

Functionality active in dev. The variable

```
ckan_memcache_sessions: True
```

must be created under group_vars/ckan_prod/ckan_prod.yml to enable it on all the production instances. Then, the playbook must be run with the following options:

```
./run.sh ckan.yml -i inventory/hosts.dev -l ckan_dev -t ckan_memcache,memcache
```

(change the inventory file and the target host).

@francesco.mangiacrapa@isti.cnr.it can you check if the dev instance is working correctly?

#3 - Nov 23, 2018 11:49 AM - Francesco Mangiacrapa

Hi @andrea.dellamico@isti.cnr.it,

I just performed some tests in dev and the catalogue ckan-d seems to work fine. I also published an item without any problem.

#4 - Nov 23, 2018 12:22 PM - Andrea Dell'Amico

I see that the sessions directory under /tmp is empty and I see CKAN sessions into the memcache data, so I guess the configuration can be extended

to the production nodes. Maybe starting next monday?

#5 - Nov 26, 2018 05:26 PM - Andrea Dell'Amico

Reconfiguration of the sobigdata ckan servers is ongoing. It will take some time, because the old sessions data is being removed.

#6 - Nov 26, 2018 06:01 PM - Andrea Dell'Amico

- Assignee changed from *_InfraScience Systems Engineer* to *Francesco Mangiacrapa*

- % Done changed from 70 to 90

@francesco.mangiacrapa@isti.cnr.it will decide when reconfigure the remaining CKAN instances.

#7 - Nov 30, 2018 06:43 PM - Andrea Dell'Amico

- Status changed from *In Progress* to *Closed*

- Assignee changed from *Francesco Mangiacrapa* to *_InfraScience Systems Engineer*

- % Done changed from 90 to 100

All the servers have been reconfigured.