

QCAT - User Story #1774

Caching

06 Feb 2018 12:23 - Sebastian Manger

Status:	Closed	Start date:	06 Feb 2018
Priority:	Normal	Due date:	
Assignee:		% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	Maintenance	Spent time:	0.00 hour
Resolution:			
Description			
Intro			
<p>The full configuration (= 'empty paper questionnaire') takes up to 10 seconds to initialize. We therefore use 'caching', so it is initialized only once, but 'served' to people accessing qcat from the cache - which is very fast. While caching is optimized to store many 'small' objects, we store few 'big' objects - so we added an additional workaround. This is now causing problems.</p>			
Decision:			
<ul style="list-style-type: none">• Drill down the cache usage. We have to find out what is really stored there and why it does grow that much.			
What happens if we don't implement this?			
<ul style="list-style-type: none">• qcat.wocat.net will take a long time to load (~10 seconds) for more and more users ('performance spikes'). The more configurations / versions / languages there are, the more often these spikes.• The nicer, but more expensive way would be to change the way configurations are cached - as 'many small items'. See issue #1793			
Technical Detail:			
<p>The configuration objects is stored in memcache and the (uwsgi) process. The latter is a workaround, as serializing the object from memcache in python takes approx. 3 seconds. Therefore, the serialized objects from memcached are stored in the uwsgi process as 'last recently used' objects. As the uwsgi processes have an upper limit of consumed RAM and are restarted after the limit is reached, more configs = more trouble. As of now there are about 500 cache misses per day.</p> <p>But there is still an important open question:</p> <ul style="list-style-type: none">• The size one configuration(technology) in one language is: 20MB• This multiplied with the number of configurations (5) and the number of languages (6) would give: $5 * 6 * 20 \text{ MB} = 600 \text{ MB}$ <p>But the cache keeps growing and reaches 2 GB after a few days.</p>			
Failed fix			
<p>Split configuration caching into sections. Try this asap, with A/B testing (cde staff).</p> <p>This is just another 'plaster', but a refactor of the whole configuration object would be too time consuming.</p>			
Subtasks:			
Task # 1775: Drill down cache usage			Closed
Task # 1776: Fix cache usage			Closed

History

#1 - 06 Feb 2018 15:17 - Sebastian Manger

- File performance spikes.png added
- Description updated

#2 - 08 Feb 2018 11:41 - Lukas Vonlanthen

- Description updated

#3 - 22 Feb 2018 10:11 - Kurt Gerber

- Target version changed from Carbon Benefit to QCAT Backlog

#4 - 27 Feb 2018 10:10 - Kurt Gerber

- Status changed from New to Accepted

- Target version changed from QCAT Backlog to Maintenance

#5 - 28 Jan 2022 13:30 - Kurt Gerber

- Status changed from Accepted to Closed

Files

performance spikes.png	29.7 KB	06 Feb 2018	Sebastian Manger
------------------------	---------	-------------	------------------