

Core - Bug # 63629

Bug # 63692 (Accepted): Memory consumption while bulk inserting

Status:	Resolved	Priority:	Must have
Author:	Lukas Krieger	Category:	Performance
Created:	2014-12-06	Assigned To:	
Updated:	2014-12-10	Due date:	
TYPO3 Version:	6.2		
PHP Version:			
Complexity:			
Is Regression:	No		
Sprint Focus:			
Subject:	High memory consumption in BackendUtility::getPagesTSconfig while using the DataHandler		
Description			
<p>There is a memory consumption problem in</p> <pre>TYPO3\CMS\Backend\Utility::getPagesTSconfig</pre> <p>In the Typo3 Forum there is this thread, describing the investigation, a solution for a patch and also testing the patch compared to the original function.</p> <p>http://forum.typo3.org/index.php/t/206887/</p> <p>Executive Summary:</p> <p>If you import pages with the DataHandler into Typo3, the function mentioned above is creating a huge array saving all TSconfigs for every page (http://typo3.org/api/typo3cms/backend_2_classes_2_utility_2_backend_utility_8php_source.html#l01153 line 1197)</p> <pre>if (\$useCacheForCurrentPageId) { > \$pagesTSconfig_cache[\$id] = \$TSconfig; }</pre> <p>In my case, there are a lot of different keys in that array but only one unique page TS.</p> <p>http://forum.typo3.org/index.php?t=tree&#38;th=206887&#38;goto=722134&#38;#msg_722134</p> <p>On a normal Backend call, the function is caching the pageTS config for only one page. What makes sense because the function is called multiple times BUT only for one single page!</p> <p>Therefore the array has everytime a count of 1.</p> <p>1 pageID => 1 TSconfig</p> <p>The main problem is that multiple pages can have the same TSconfig, like in my DataHandler import.</p> <p>Therefore it would be nice to have a n:n relation for caching the TSconfig in this function.</p> <p>The static cache array</p> <pre>\$pagesTSconfig_cache</pre> <p>is in the very first request always empty.</p>			

Here is the description of a possible patch:

http://forum.typo3.org/index.php?t=tree&th=206887&goto=722198&#msg_722198

What the patch currently does is following:

1. add a second static array for caching unique TSconfig
2. use the pagesTSconfig_cache as a reference to the unique TSconfig stored in the other array

Here you can find some tests:

http://forum.typo3.org/index.php?t=tree&th=206887&goto=722232&#msg_722232

http://forum.typo3.org/index.php?t=tree&th=206887&goto=722311&#msg_722311 (comparing the original function with the patched one and adding conditions in the backend)

An improvement would be maybe to use the CacheManager for storing the caches globally.

Because the parsed page TSconfig is already cached by "TYPO3\CMS\Backend\Configuration\TsConfigParser" on line 1186

```
// Parsing the page TS-Config
$pageTS = implode(LF . '[GLOBAL]' . LF, $TSdataArray);
/* @var $parseObj \TYPO3\CMS\Backend\Configuration\TsConfigParser */
$parseObj = GeneralUtility::makeInstance('TYPO3\CMS\Backend\Configuration\TsConfigParser');
$res = $parseObj->parseTSconfig($pageTS, 'PAGES', $id, $rootLine);
```

Some result:

Task: modifying 280 pages

Original Function: The whole process needs 19 seconds and is using 85MB memory

Modified Function: The whole process needs also 19 seconds but is only using 14MB memory

Task: modifying 150 pages

Original Function: 9 seconds, 51.1 MB memory

Modified Function: 8.7 seconds, 11MB memory

Task: modifying 118 pages

Original Function: 7.2 seconds, 41.7 MB memory

Modified Function: 6.8 seconds, 10 MB memory

Task: modifying 54 pages

Original Function: 3.6 seconds , 23MB

Modified Function: 3.3 seconds, 8.5 MB memory

The patch will follow on Gerrit.

ps.: there are also some mass editing functions in the normal typo3 backend available. This patch could improve the performance if the backend mass editing is also calling this function.

Related issues:

related to Core - Bug # 53370: Cache getPagesTSconfig per ID and rootline

Resolved

2013-11-06

Associated revisions

Revision d5c82727 - 2014-12-08 01:36 - Stephan Großberndt

[BUGFIX] Better caching of TSconfig in BackendUtility for bulk inserts

Changes the cache of pageTSconfig from creating a cached entry for each page to creating a cached entry for each different TSconfig (including userTSconfig) and a hash-cache for each page instead.

Resolves: #63629

Releases: master, 6.2

Change-Id: I3e4c70cd46cec26dda8574c058d828ec2e22a385

Reviewed-on: <http://review.typo3.org/35122>

Reviewed-by: Philipp Gampe <philipp.gampe@typo3.org>

Tested-by: Philipp Gampe <philipp.gampe@typo3.org>

Reviewed-by: Stephan Großberndt <stephan@grossberndt.de>

Reviewed-by: Lukas Krieger <lukas.krieger@me.com>

Tested-by: Lukas Krieger <lukas.krieger@me.com>

Reviewed-by: Markus Klein <klein.t3@reelworx.at>

Tested-by: Markus Klein <klein.t3@reelworx.at>

Revision 22fc23e6 - 2014-12-08 01:43 - Stephan Großberndt

[BUGFIX] Better caching of TSconfig in BackendUtility for bulk inserts

Changes the cache of pageTSconfig from creating a cached entry for each page to creating a cached entry for each different TSconfig (including userTSconfig) and a hash-cache for each page instead.

Resolves: #63629

Releases: master, 6.2

Change-Id: I3e4c70cd46cec26dda8574c058d828ec2e22a385

Reviewed-on: <http://review.typo3.org/35145>

Reviewed-by: Markus Klein <klein.t3@reelworx.at>

Tested-by: Markus Klein <klein.t3@reelworx.at>

History

#1 - 2014-12-06 23:50 - Gerrit Code Review

- Status changed from New to Under Review

Patch set 1 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.

It is available at <http://review.typo3.org/35122>

#2 - 2014-12-07 14:34 - Gerrit Code Review

Patch set 2 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.

It is available at <http://review.typo3.org/35122>

#3 - 2014-12-07 14:48 - Gerrit Code Review

Patch set 3 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.

It is available at <http://review.typo3.org/35122>

#4 - 2014-12-08 01:40 - Gerrit Code Review

Patch set 1 for branch **TYPO3_6-2** of project **Packages/TYPO3.CMS** has been pushed to the review server.

It is available at <http://review.typo3.org/35145>

#5 - 2014-12-08 02:00 - Stephan Großberndt

- *Status changed from Under Review to Resolved*

- *% Done changed from 0 to 100*

Applied in changeset commit:d5c827272d016b51ca916abf575c5aff2bd732d.

#6 - 2014-12-09 07:04 - Philipp Gampe

- *Parent task set to #63692*

#7 - 2014-12-09 13:40 - Stephan Großberndt

As Philipp Gampe pointed out in <https://review.typo3.org/#/c/35186/1/typo3/sysexst/core/Classes/DataHandling/DataHandler.php@1615>

"Please let's use a (static?) member here instead of the static local variable. Static local variables are impossible to mock."

That applies to these caches as well. Create another issue since this is merged?

#8 - 2014-12-09 17:35 - Alexander Opitz

Hi Stephan, thanks for all the work.

Yes, please create a new issue, add a patch and add Oliver Klee as Reviewer.

#9 - 2014-12-10 18:10 - Mathias Schreiber

- *Parent task deleted (#63692)*

#10 - 2014-12-10 18:10 - Mathias Schreiber

- *Parent task set to #63692*