

Core - Bug # 52235

Epic # 55070 (Accepted): Workpackages

Epic # 55065 (New): WP: Overall System Performance (Backend and Frontend)

Bug # 52949 (Resolved): Speed decrease since 4.5

Status:	Resolved	Priority:	Must have
Author:	Jan-Erik Revsbech	Category:	
Created:	2013-09-24	Assigned To:	
Updated:	2014-08-28	Due date:	
TYPO3 Version:	6.1		
PHP Version:			
Complexity:	medium		
Is Regression:	No		
Sprint Focus:			
Subject:	Timeout when copying pages recursively due to cache-clearing overload		
Description	<p>We currently experience a problem resulting in timeouts when copying more than a few pages, especially if the site is mounted over NFS. The problem seems to arise due to several issues.</p> <p>The clear_cache command in DataHandler calls <code>\$GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pageId_' . \$pageId)</code>; for each page it calculates should be cleared (including siblings, and possibly parents, grand-parents etc.). Since this will ask all registered caches to clear for the desired tag, this results in many, many checks for flush caching. Combine this with a massively inefficient cache FileBackend (see #52125) this results in timeouts when copying more than 3-4 pages.</p> <p>The result is even worse, since the clear_cache is called for the page being copied AND for each content element on the page, and for each content element the cache is cleared for its PID (if using workspaces, there is yet another problem with this, but that seems to be fixed in another ticket).</p> <p>So copying a page with 4 content elements, will clear cache for the page 5 times! And for each of these, the siblings and possibly parents are cleared as well.</p> <p>I would suggest a few changes and would like some feedback on this:</p> <ol style="list-style-type: none">1: Make sure that only the cache_pages and possibly cache_pagesection is cleared by only calling <code>\$GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pageId_' . \$pageId)</code>;2: Make DataHandler collect a list of all pages it wants to clear, and then defer the actual clearing to later, to make sure we only clear cache one time for each page. We could either defer until the destructor of DataHandler is called, or we could just put everything3: As suggested in #52125 FileBackend should either be fixed, or not used by default. <p>Feedback is welcome. I will push a changeset to Gerrit with the suggested changes when possible feedback is collected.</p>		
Related issues:			
related to Core - Bug # 51116: Massive speed problem from TYPO3 version 6.x a...	Resolved	2013-08-16	
related to Core - Bug # 52125: Saving records takes ages to complete	Resolved	2013-09-19	
related to Core - Bug # 61017: TCEMAIN.clearCacheCmd does not trigger	Under Review	2014-08-18	
related to Core - Task # 52295: Use SimpleFileBackend for t3lib_110n cache	Resolved	2013-09-26	
duplicated by Core - Bug # 57265: Importing large amounts of data issues many...	Rejected	2014-03-25	

Associated revisions

Revision c9738f8c - 2014-03-25 11:59 - Jan-Erik Revsbech

[BUGFIX] Fix problem with cache being cleared multiple time

When pages or records in general are changed using the DataHandler, the pageCache is cleared for each page multiple times. To fix this, the DataHandler will collect which pages need clearing, and wait with the actual clearing until all operations in cmdMap or dataMap has been processed. This way the cache is only cleared once for each page.

Resolves: #52235

Releases: 6.2

Change-Id: If1aa389991b801413eafa38d7d140f932bb321f2

Reviewed-on: <https://review.typo3.org/25423>

Reviewed-by: Markus Klein

Tested-by: Markus Klein

Reviewed-by: Jigal van Hemert

Reviewed-by: Jan-Erik Revsbech

Tested-by: Jan-Erik Revsbech

Tested-by: Philipp Gampe

Reviewed-by: Benjamin Mack

Tested-by: Benjamin Mack

History

#1 - 2013-09-25 11:41 - Markus Klein

I'd suggest to implement some sort of batch cache clearing in the caches.

```
Example: * $GLOBALS['typo3CacheManager']->getCache('cache_pages')->setBatchFlushing(TRUE); *
$GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pagelId_ ' . 1); *
$GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pagelId_ ' . 2); *
$GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pagelId_ ' . 3); * ... *
$GLOBALS['typo3CacheManager']->getCache('cache_pages')->setBatchFlushing(FALSE);
```

Once the batch flushing is disabled, the actual flushing will be executed.

#2 - 2013-09-26 10:47 - Jan-Erik Revsbech

Markus Klein wrote:

I'd suggest to implement some sort of batch cache clearing in the caches.

```
Example: - $GLOBALS['typo3CacheManager']->getCache('cache_pages')->setBatchFlushing(TRUE);
- $GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pagelId_ ' . 1);
- $GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pagelId_ ' . 2);
- $GLOBALS['typo3CacheManager']->getCache('cache_pages')->flushByTag('pagelId_ ' . 3);
- ...
- $GLOBALS['typo3CacheManager']->getCache('cache_pages')->setBatchFlushing(FALSE);
```

Once the batch flushing is disabled, the actual flushing will be executed.

That is not a bad idea, but perhaps the DataHandler should be the one actually "queueing" things for cache_clearing? It could be done in the CachingLayer, but I'm not sure thats a good idea.

#3 - 2013-09-26 14:45 - Christian Kuhn

just and idea on the 'batch' clearing (don't know if it works out): we have this "shutdown" method in bootstrap ... maybe we could execute the tag clearing there?

#4 - 2013-09-26 16:10 - Christian Kuhn

- Parent task set to #52304

#5 - 2014-01-09 16:41 - Markus Klein

- Parent task changed from #52304 to #52949

#6 - 2014-03-25 03:30 - Ernesto Baschny

This currently also makes importing a t3d file (used by the Distribution Management). Since we will see more people using this, we should try to tackle this issue soon.

Work in progress is here: <https://review.typo3.org/#/c/25423/>

#7 - 2014-03-25 03:35 - Philipp Gampe

- Status changed from New to Under Review

#8 - 2014-03-25 03:40 - Philipp Gampe

- Priority changed from Should have to Must have

#9 - 2014-03-25 09:47 - 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 <https://review.typo3.org/25423>

#10 - 2014-03-25 09:48 - 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 <https://review.typo3.org/25423>

#11 - 2014-03-25 09:59 - Gerrit Code Review

Patch set 4 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.

It is available at <https://review.typo3.org/25423>

#12 - 2014-03-25 10:29 - Gerrit Code Review

Patch set 5 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.

It is available at <https://review.typo3.org/25423>

#13 - 2014-03-25 10:34 - Gerrit Code Review

Patch set 6 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.
It is available at <https://review.typo3.org/25423>

#14 - 2014-03-25 11:32 - Gerrit Code Review

Patch set 7 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.
It is available at <https://review.typo3.org/25423>

#15 - 2014-03-25 11:36 - Gerrit Code Review

Patch set 8 for branch **master** of project **Packages/TYPO3.CMS** has been pushed to the review server.
It is available at <https://review.typo3.org/25423>

#16 - 2014-03-25 13:45 - Jan-Erik Revsbech

- *Status changed from Under Review to Resolved*

- *% Done changed from 0 to 100*

Applied in changeset commit:c9738f8c7af749adffb8b0093ddc852e62f4ba1b.

#17 - 2014-08-28 03:01 - Markus Klein

This caused a regression: #61017