

TYPO3 Core - Bug #90518

DB problems with rootline cache: DELETE cf_cache_rootline double-JOIN with cf_cache_rootline_tags take excessively long

2020-02-24 13:27 - Sybille Peters

Status:	Closed	Start date:	2020-02-24
Priority:	Should have	Due date:	
Assignee:		% Done:	0%
Category:	Caching	Estimated time:	0.00 hour
Target version:		Complexity:	
TYPO3 Version:	9	Is Regression:	
PHP Version:		Sprint Focus:	
Tags:			
Description			
Recently we are experiencing increasing number of exceptions and "too many connections" on TYPO3 9.5.13 in combination with DELETE queries on cf_cache_rootline and cf_cache_rootline_tags that take excessively long.			
Obviously, this is causing massive problems.			
Findings:			
<ul style="list-style-type: none">Increasing number of exceptions: An exception occurred while executing 'INSERT INTO `cf_cache_rootline ... Lock wait timeout exceeded; try restarting transaction Doctrine\DBAL\Exception			
This seems to be correlated with excessively long DELETE queries on cf_cache_rootline and cf_cache_rootline_tags with 2 JOINS.			
<pre># Time: 200221 11:44:21 # User@Host: ----- @ [-----] # Thread_id: 13694456 Schema: ----- QC_hit: No # Query_time: *2212.109482* Lock_time: 0.000340 Rows_sent: 0 Rows_examined: *1007307838* # Rows_affected: *94732* SET timestamp=1582281861; DELETE tags2, cache1 FROM cf_cache_rootline_tags AS tags1 JOIN cf_cache_rootline_tags AS tags2 ON tags1.identifier = tags2.identifier JOIN cf_cache_rootline AS cache1 ON tags1.identifier = cache1. identifier WHERE tags1.tag IN ('pageId_3');</pre>			
(I blanked out some IP addresses etc.)			
in Typo3DatabaseBackend::flushByTags()			

History

#1 - 2020-02-24 13:28 - 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 <https://review.typo3.org/c/Packages/TYPO3.CMS/+63401>

#2 - 2020-02-24 14:07 - Sybille Peters

- Description updated

Update: I made a mistake above. Updated the descri

#3 - 2020-02-24 14:08 - Sybille Peters

- Description updated

#4 - 2020-03-03 12:43 - Sybille Peters

- Description updated

- Status changed from Under Review to New

#5 - 2020-03-03 13:41 - Sybille Peters

- File `mysql_processlist.txt` added

- File `stacktrace.txt` added

This is still an ongoing investigation, but I will summarize what I found:

Initial problem:

Error due to exception "Lock wait timeout" in Doctrine shown both in frontend (as "Oops" and backend) and visible in logs. All occurrences were on rootline cache.

System

TYPO3: 9.5.13

DB-Server: MariaDB

pages: 38791 (not deleted, sys_language_uid=0)

typical number of entries (these vary, of course)

- `cf_cache_rootline`: 15347

- `cf_cache_rootline_tags`: 112814

Findings

- Some occurrences in slow query log. All are similar database statement that take several seconds up to 30 minutes (!!!) in one extreme case:

```
DELETE tags2, cache1 FROM cf_cache_rootline_tags AS tags1
JOIN cf_cache_rootline_tags AS tags2 ON tags1.identifier = tags2.identifier
JOIN cf_cache_rootline AS cache1 ON tags1.identifier = cache1.identifier
WHERE tags1.tag IN ('pageId_3');
```

The tag is not always the same. Here, it is the start page (page 3)

- the occurrences in the slow query log more or less correlate with the errors
- there are almost 40000 pages in the system with a deep page depth of more than 8 in some cases
- some of the entries in `cf_cache_rootline.content` are quite large (up to 40KB). Not only is rootline information for all parent pages written into this, including Page TSconfig
- The tags in `cf_cache_rootline_tags` contain tags for page ids, e.g. "pageId_3" for page id 3. All child pages contain the tag for their parents. So if the above DELETE + double-JOIN is executed on tag 'pageId_3' this means that practically all rootline cache and rootline cache tag entries will be affected in one big scoop.
- On trying to track down error: after saving the start page, some more errors occurred. I looked at show processlist on DB server and saw both a DELETE / JOIN statement on rootline cache and several INSERTs simultaneously

Explanation

This is a preliminary explanation as I am not sure of all facts yet.

The long running DELETE `cf_cache_rootline`, `cf_cache_rootline_tags` affects a lot of (or almost all rows) and would lock out simultaneous SQL statements (on rootline cache tables). On a lot of occasions a rootline cache entry is filled (or queried) so this is pretty likely (simultaneous edit and page load operations).

Open questions

So what is the reason the DELETE takes so long?

Is it the problem that such a DELETE statement gets created (affecting so many rows at once) or should the DB server be able to handle this without problems?

Is it a combination of a lot of pages, a lot of rootline and tags entries being deleted simultaneously? Not enough memory available on DB server?

Where should I look next:

- DB server?
- TYPO3 configuration?
- TYPO3 itself?

Generally, deleting practically all

Changes

The problem no longer occurs now. I made these changes. This has (hopefully) defused the situation:

- Assigned more RAM to Innodb buffer cache
- Update MariaDB server to latest version
- Enabled 'compression' for rootline cache
- Set lifetime on rootline cache to 24h (instead of 30 days which is the default).
- removed some clearCacheCmd in page TSconfig that would result in clearing the start page with a DELETE of practically all rootline entries (effectively triggering the situation described above)
- removed page TSconfig from start page (that was only intended temporarily and was moved to sitepackage) to prevent it from being written to cf_cache_rootline.content of all subpages.

I consider this currently a workaround until I am more sure what to do long term.

#6 - 2020-03-24 21:14 - Sybille Peters

- *Status changed from New to Closed*

Update: On large sites, it is recommended to put rootline cache (and other caches) in e.g. Redis and not DB.

I will close this issue now.

Files

mysql_processlist.txt	56.7 KB	2020-03-03	Sybille Peters
stacktrace.txt	10.7 KB	2020-03-03	Sybille Peters