

TYPO3 Core - Bug #21017

Use separate tables for tags in the caching framework

2009-09-08 18:35 - Oliver Hader

| | | | |
|---|--------------|------------------------|-------------------|
| Status: | Closed | Start date: | 2009-09-08 |
| Priority: | Should have | Due date: | |
| Assignee: | Oliver Hader | % Done: | 0% |
| Category: | Caching | Estimated time: | 0.00 hour |
| Target version: | | Complexity: | |
| TYPO3 Version: | 4.3 | Is Regression: | |
| PHP Version: | 5.2 | Sprint Focus: | |
| Tags: | | | |
| Description | | | |
| The database backend of the caching framework stores tags in a comma separated list. Searching for tags results in a LIKE query which will be very slow when there are many cache entries. (issue imported from #M11903) | | | |
| Related issues: | | | |
| Related to TYPO3 Core - Bug #20863: TCEmain clears cache inefficiently | | Closed | 2009-08-11 |
| Related to TYPO3 Core - Bug #21718: Add functionality to work with caching fr... | | Closed | 2009-11-27 |
| Related to TYPO3 Core - Feature #21688: Add support for IN and EXISTS subqueries | | Closed | 2009-11-25 |

History

#1 - 2009-09-08 18:37 - Oliver Hader

The attached patch uses a separate table to store tags. Thus, it's faster for selecting data. However, due to the lack of JOINS in TYPO3_DB, the flush operations and garbage collection are very inefficient since they have to use iterations over the cache identifiers and will create a lot of DELETE queries...

#2 - 2009-09-08 20:27 - Oliver Hader

Extended TYPO3_DB to support something like

```
DELETE tbl1 FROM tbl1, tbl2 WHERE tbl1.identifier = tbl2.identifier;
```

(this has to be in a separate RFC when posting to the Core List).

#3 - 2009-09-15 20:51 - Martin Kutschker

A DELETE with multiple tables is non-standard.

PostgreSQL has the non-standard USING clause.

```
DELETE FROM tbl1 USING tbl2 WHERE tbl1.identifier = tbl2.identifier;
```

<http://www.postgresql.org/docs/8.4/interactive/sql-delete.html>

A more portable way would be to use sub-selects. Maybe the DBAL team is possible to come up with a portable solution.

#4 - 2009-09-15 20:59 - Martin Kutschker

A note in a forum claimed that this will work with Oracle 9i:

```
DELETE FROM (tbl1, tbl2 WHERE tbl1.identifier = tbl2.identifier);
```

Though a reply mentioned an error when executing this command.

#5 - 2009-09-15 21:07 - Martin Kutschker

SQL Server 2000 seems also to support the extended syntax like Mysql.

[http://msdn.microsoft.com/en-us/library/aa258847\(SQL.80\).aspx](http://msdn.microsoft.com/en-us/library/aa258847(SQL.80).aspx)

It offers also an example how to make the DELETE without this extension. The solution requires sub-selects.

```
/* SQL-92-Standard subquery */
USE pubs
DELETE FROM titleauthor
WHERE title_id IN
(SELECT title_id
FROM titles
WHERE title LIKE '%computers%')
```

```
/* Transact-SQL extension */
USE pubs
DELETE titleauthor
FROM titleauthor INNER JOIN titles
ON titleauthor.title_id = titles.title_id
WHERE titles.title LIKE '%computers%'
```

#6 - 2009-09-15 21:13 - Martin Kutschker

Extending the DELETEquery function like this should do the trick:

```
DELETEquery($table,$where,$table2=",$joinCondition=")
```

This will maybe not allow for all possible SQL query variations, but will be sufficient for most cases.

#7 - 2009-09-16 08:13 - Xavier Perseguers

Hi,

Trying your stuff on Oracle 11 leads to this:

```
DELETE FROM (T1, T2 WHERE T1.column1 = T2.column2);
```

Column 13: Missing SELECT keyword

In addition, it seems that this cannot be achieved without using triggers or ON CASCADE DELETE:

<http://www.orafaq.com/forum/t/55408/0/>

#8 - 2009-09-16 09:47 - Martin Kutschker

Xavier, this is the error I mentioned. The thread you link to suggests this syntax:

```
DELETE FROM (SELECT * FROM tbl1, tbl2 WHERE tbl1.identifier = tbl2.identifier);
```

Anyway, as Oracle supports sub-queries, DBAL can provide a variant suitable for it. And as a fallback it still could use a SELECT and multiple DELETEs.

#9 - 2009-09-16 10:55 - Xavier Perseguers

OK. With SELECT, at least one of the table should have a key, which I don't have ATM otherwise I get the same error as in the thread I linked. But as you said, we certainly can find a way for DBAL to work with this.

#10 - 2009-09-21 15:12 - Oliver Hader

Committed to SVN Trunk (rev. 6025)

#11 - 2009-09-27 22:40 - Martin Kutschker

After re-reading the docs for Mysql, PostgreSQL and SQL server I find that these three databases support THIS syntax:

```
DELETE
[FROM] tbl_name
USING|FROM table_references
WHERE where_condition
```

Note: Mysql and Posgresql use USING, whereas SQL server uses a second FROM. The "table_references" can be anything that is found in the FROM clause of a SELECT statement.

Unfortunately this syntax deletes only from one table. Neither PostgreSQL nor SQL server allow multi-tables in the (first) FROM. The USING (second FROM) is only used for a join.

Still looking for a DBAL way to get a portable simple solution for the cache clearing problem.

Files

| | | | |
|------------------|---------|------------|---------------------|
| 0011903_v0.patch | 9.63 KB | 2009-09-08 | Administrator Admin |
| 0011903_v1.patch | 12.2 KB | 2009-09-08 | Administrator Admin |
| 0011903_v2.patch | 12.2 KB | 2009-09-10 | Administrator Admin |
| 0011903_v4.patch | 14.8 KB | 2009-09-18 | Administrator Admin |
| 0011903_v6.patch | 17.3 KB | 2009-09-21 | Administrator Admin |