

## TYPO3 Core - Task #90143

### Redirects: Poor performance of redirect matching for large redirects table

2020-01-17 14:26 - Mario Lubenka

<b>Status:</b>	Under Review	<b>Start date:</b>	2020-01-17
<b>Priority:</b>	Should have	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>	Link Handling, Site Handling & Routing	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Candidate for patchlevel	<b>Complexity:</b>	medium
<b>TYPO3 Version:</b>	9	<b>Sprint Focus:</b>	
<b>PHP Version:</b>	7.2		
<b>Tags:</b>			

#### Description

### Problem description

The current redirect matching retrieves all redirects from the database table and stores it in an internal cache. This is then fetched and filtered for the route that matches the request.

Related files:

- <https://github.com/TYPO3-CMS/redirects/blob/master/Classes/Service/RedirectService.php>
- <https://github.com/TYPO3-CMS/redirects/blob/master/Classes/Service/RedirectCacheService.php>

With a large redirects table this takes a pretty long time.

I created 350.000 unique redirects and it took the page about 5 seconds to redirect to the target page.

Before you say "350.000 will never be reached": On a long term a heavily edited website might reach that number at some point in time (assuming this feature will continue to exist with further TYPO3 versions).

### Suggested solution

I suggest that the redirect matching is mainly done via the database query.

Most of the logic can be moved into a WHERE clause. Operations which are not available via database operations would then be applied to the remaining records.

With that the above 5 seconds of loading time could be reduced to merely 1 second.

This might also imply that the redirect cache via RedirectCacheService is not needed anymore.

Databases already come with caching mechanisms, as to such I fail to understand why a PHP cache is also needed.

This is a performance improvement which should be applied to TYPO3 9.5 as well.

#### Related issues:

Related to TYPO3 Core - Bug #93287: Avoid creation of unnecessary redirects o...	<b>New</b>	<b>2021-01-14</b>
Related to TYPO3 Core - Epic #93547: Collection of problems with large sites	<b>Accepted</b>	<b>2021-02-19</b>
Related to TYPO3 Core - Bug #91792: Automatically created redirects are not c...	<b>New</b>	<b>2020-07-13</b>

#### History

##### #1 - 2020-01-17 16:49 - 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/+62956>

##### #2 - 2020-01-17 16:51 - 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/c/Packages/TYPO3.CMS/+62956>

##### #3 - 2020-01-17 16:53 - 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/c/Packages/TYPO3.CMS/+62956>

#### #4 - 2020-01-20 15:36 - Mario Lubenka

- File *redirect-measure-time.ods* added

Adding some performance numbers. ;)

## Test Setup

I installed a TYPO3 10.3.0-dev with the Introduction package and created multiple redirects by adding the following code after Bootstrap::init in index.php and reloaded the page ONCE. After that I removed the code.

```
/** @var \TYPO3\CMS\Core\Database\Connection $conn */
$conn = GeneralUtility::makeInstance(ConnectionPool::class)
    ->getConnectionForTable('sys_redirect');
$n = 400000; // 400k redirects
for ($i = 1; $i <= $n; $i++) {
    $conn->insert('sys_redirect', [
        'source_host' => '*',
        'source_path' => '/test' . $i,
        'target' => 't3://page?uid=80'
    ]);
}
```

I did 6 sets with a different number of redirect records in the database (1k up to 400k). Before each set, the typo3temp folder was removed to make sure all code caches are regenerated. During each set, the page load time until redirect header was taken 6 times. The first value was discarded because Cache had to regenerate, leaving us with 5 values for the page load time.

The time was measured using the following command, returning the time in seconds.

```
curl -w %{time_total} --url https://t3-contribution.ddev.site/test10
```

## Results

The full data is in the attached *redirect-measure-time.ods* file. The table below shows the average response time until the redirect happens in milliseconds.

# redirects	Current implementation	Proposed implementation
1000	61ms	53ms
10000	130ms	68ms
50000	426ms	105ms
100000	701ms	141ms
250000	1560ms	263ms
400000	2384ms	346ms

As you can see, the proposed implementation is a lot faster when many redirect records in the database.

With large TYPO3 websites, the current implementation may cause page load times up to 5 seconds, which is what we discovered in our real project.

#### #5 - 2020-01-20 15:38 - Mario Lubenka

No additional indexes needed.

#### #6 - 2020-02-15 14:09 - 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/c/Packages/TYPO3.CMS/+62956>

#### #7 - 2020-02-15 14:18 - 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/c/Packages/TYPO3.CMS/+62956>

#### #8 - 2020-02-15 14:37 - 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/c/Packages/TYPO3.CMS/+62956>

**#9 - 2020-04-07 00:38 - 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/c/Packages/TYPO3.CMS/+62956>

**#10 - 2020-04-07 00:43 - 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/c/Packages/TYPO3.CMS/+62956>

**#11 - 2021-01-14 07:51 - Sybille Peters**

- Related to Bug #93287: Avoid creation of unnecessary redirects on Slug changes added

**#12 - 2021-02-19 07:01 - Sybille Peters**

- Related to Epic #93547: Collection of problems with large sites added

**#13 - 2021-03-19 09:23 - Daniel Goerz**

- Related to Bug #91792: Automatically created redirects are not created using the DataHandler added

**#14 - 2021-09-07 10:29 - Gerrit Code Review**

Patch set 9 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/+62956>

**#15 - 2021-09-11 12:57 - Gerrit Code Review**

Patch set 10 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/+62956>

**#16 - 2021-09-11 13:16 - Gerrit Code Review**

Patch set 11 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/+62956>

**#17 - 2021-09-11 21:50 - Gerrit Code Review**

Patch set 12 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/+62956>

**#18 - 2021-09-11 23:55 - Gerrit Code Review**

Patch set 13 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/+62956>

**#19 - 2021-09-12 12:16 - Gerrit Code Review**

Patch set 14 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/+62956>

**#20 - 2021-09-20 21:46 - Gerrit Code Review**

Patch set 15 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/+62956>

**Files**

---

redirect-measure-time.ods	17.5 KB	2020-01-20	Mario Lubenka
---------------------------	---------	------------	---------------