

Core - Task # 55167

Epic # 55070 (Accepted): Workpackages

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

Epic # 55656 (Accepted): Optimize overall Extbase performance

Story # 55168 (Accepted): Optimize Extbase generic persistence

Status:	Resolved	Priority:	Should have
Author:	Felix Oertel	Category:	Extbase
Created:	2014-01-20	Assigned To:	Felix Oertel
Updated:	2014-04-01	Due date:	2014-03-20
TYPO3 Version:	6.2		
PHP Version:			
Complexity:	hard		
Sprint Focus:			
Subject:	Extbase: avoid parseQuery() / buildQuery() to be called multiple times		
Description	<p>At the moment the extbase persistence runs buildQuery() and parseQuery() for every single getObjectDataByQuery() / getObjectCountByQuery(). When fetching child objects, this is called for every single object. This results in a lot of re-run code. In a simple blog_example page, this is called up to 200 times.</p> <p>We will try to cache the built query as in an live environment always the same queries should be run, just with different parameters. Special care is to be taken for which parts of the query are "static" and about to be built into the cache identifier.</p> <p>We will have to give the developer the chance to disable query caching though, because in some cases it might result in real big data piles. (e. g. matrix-like queries, allowing the front-end to influence the number of queried properties (search)).</p>		
Subtasks:			
Task # 55661: Add functional tests for extbase persistence		Resolved	
Task # 57173: Use of cf_extbase_typo3dbbackend_queries results in fatal		Resolved	
Task # 57130: Extbase Queries with NULL comparison don't work anymore		Resolved	
Task # 57104: Regression: Query building fails for LazyObjectStorage		Resolved	
Task # 57263: Typo3DbQueryParser does not handle LowerCaseInterface correctly		Resolved	
Task # 57338: Make Typo3DbQueryParser Singleton		Resolved	
Related issues:			
related to Core - Task # 53514: Extbase: DataMapper causes redundant queries	Resolved		2013-11-11
related to Core - Task # 57275: Add first level cache for the query cache	Resolved		2014-03-25
related to Core - Task # 57280: Introduce setting to be able to disable Extba...	Resolved		2014-03-25
related to Core - Bug # 57053: Regression: Extbase query builder uses invalid...	Resolved		2014-03-18
follows Core - Task # 55360: Extbase: use preparedStatement in generic persis...	Resolved		2014-03-19

Associated revisions

Revision cf42649b - 2014-03-12 14:34 - Felix Oertel

[TASK] Add functional tests for generic persistence

To backup the changes in Extbase's generic persistence, we add some functional tests.

Resolves: #55661

Related: #55360

Related: #55167

Releases: 6.2

Change-Id: I0b845618b1cec9788ca36ac59e25ea6779f01cb6

Reviewed-on: <https://review.typo3.org/28262>

Reviewed-by: Marc Bastian Heinrichs

Tested-by: Marc Bastian Heinrichs

Reviewed-by: Steffen Ritter

Tested-by: Steffen Ritter

Revision 57caf5a4 - 2014-03-12 14:53 - Ernesto Baschny

[TASK] Follow-up to "Add functional tests for generic persistence"

Test directory changed from "Relation" to "Persistence".

Resolves: #55661

Related: #55360

Related: #55167

Releases: 6.2

Change-Id: I422dbfb6bdad4215bdc70dea775917cf3159605f

Reviewed-on: <https://review.typo3.org/28297>

Reviewed-by: Ernesto Baschny

Tested-by: Ernesto Baschny

Revision 2583fd2a - 2014-03-17 23:24 - Felix Oertel

[TASK] Use queryCache in generic persistence

At the moment extbase persistence is quite slow. To some degree, the query building is responsible for that. For each and every query the `parseQuery()` and `buildQuery()` methods are called.

This patch introduces a query cache to store the query structure. This way a query can be reused, even when fired with different parameters.

- Introduce a parameterIdentifier to the comparison
 - In the way the extbase query object model works, it's hard to identify an object in a way which is unique and still comprehensible. To ease this, an identifier is introduced.
- Move parsing to an external QueryParser
 - All the query parsing is now separated from the storageBackend.
 - The query parsing is mostly left the way it was, it should be re-profiled later.
- introduce cache `extbase_typo3dbbackend_queries`
- No more statement handling by reference
 - Make `getStatementParts()` return the query and parameters in an array instead of handing down the parameters by reference.

- Split up parseQuery in preparsing and real parsing with caching
In every call the query has to be prepared to determine the identifier and extract the parameters. The identifier can then be used to look up the query structure in the cache. If it is not (yet) cached, Extbase will do the real parsing which takes some more time.
- Introduce useQueryCache to querySettings
Like preparedStatements, caching the queries does not always make sense and will spam your cache under certain conditions. It's enabled by default and can be disabled on a per-query-basis.
- Use named parameter's placeholders
Instead of using "?" as parameter's placeholder and relying on the right order of parameters in the array, named placeholders are used to ensure the correct association.
- Move unit tests around to fit the new queryParser

Resolves: #55167

Releases: 6.2

Change-Id: I517c0500cad75cae96fd739f2ff7222db80b42a5

Reviewed-on: <https://review.typo3.org/27555>

Reviewed-by: Anja Leichsenring

Tested-by: Anja Leichsenring

Tested-by: Philipp Gampe

Reviewed-by: Wouter Wolters

Tested-by: Wouter Wolters

Revision b8edbf13 - 2014-03-20 09:07 - Alexander Stehlik

[BUGFIX] Named parameters in Extbase comparison queries

Instead of using unnamend parameters (?) in parseComparison() in the Extbase Typo3DbQueryParser the matching parameter identifier will be used for building the comparison SQL.

Additionally the non existing "\$parameters" method parameter is removed from a lot of method calls.

Resolves: #57053

Related: #55167

Releases: 6.2

Change-Id: I1123275af57e8803ecd792d65f91627cef54f501

Reviewed-on: <https://review.typo3.org/28511>

Reviewed-by: Felix Oertel

Tested-by: Felix Oertel

Reviewed-by: Wouter Wolters

Reviewed-by: Markus Klein

Reviewed-by: Xavier Perseguers

Tested-by: Xavier Perseguers

Revision 8326a71f - 2014-03-24 17:57 - Klaas Johan Kooistra

[BUGFIX] Queries with NULL don't work anymore

See change: I517c0500cad75cae96fd739f2ff7222db80b42a5

At the end of the Typo3DbQueryParser::parseComparison() method the updated operator for NULL values was no longer passed to the parseDynamicOperand, generating = NULL instead of IS NULL.

Solved the issue by adding the operator parameter again.

Change-Id: I13c3cecf0bf4c95789032f31ead192b945794484

Resolves: #57130

Related: #55167

Releases: 6.2

Reviewed-on: <https://review.typo3.org/28592>

Reviewed-by: Markus Klein

Reviewed-by: Felix Oertel

Tested-by: Felix Oertel

Revision 449c81e4 - 2014-03-24 22:32 - Felix Oertel

[BUGFIX] FOLLOWUP Queries with NULL don't work

In <https://review.typo3.org/28592> we apparently did not fix the whole problem. This followup will deal with the edge-case of using a similar built statement, which results in the same hash being generated from the QueryParser.

Resolves: #57130

Related: #55167

Releases: 6.2

Change-Id: I02a478fcabb89d11054f8f0a3c800965edfd7844

Reviewed-on: <https://review.typo3.org/28715>

Reviewed-by: Marc Bastian Heinrichs

Tested-by: Marc Bastian Heinrichs

Reviewed-by: Markus Klein

Tested-by: Markus Klein

History

#1 - 2014-01-20 14:46 - Felix Oertel

- Parent task set to #55168

#2 - 2014-01-20 14:46 - Felix Oertel

- Subject changed from avoid buildQuery() to be called multiple times to avoid parseQuery() / buildQuery() to be called multiple times

#3 - 2014-01-20 14:49 - Alexander Schnitzler

Did you already have a look at ticket #53514? Saves a lot of queries, too.

Edit: wrong ticket number.

#4 - 2014-01-20 15:19 - Felix Oertel

yep, got this on my list, thanks.

#5 - 2014-02-11 19:30 - Gerrit Code Review

- Status changed from Accepted to Under Review
- Estimated time set to 0.00

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/27555>

#6 - 2014-02-14 19:25 - Ernesto Baschny

- Subject changed from *avoid parseQuery() / buildQuery() to be called multiple times* to *Extbase: avoid parseQuery() / buildQuery() to be called multiple times*
- Category set to Extbase
- Estimated time set to 0.00
- TYPO3 Version set to 6.2

#7 - 2014-02-19 21:17 - Gerrit Code Review

- Estimated time set to 0.00

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/27555>

#8 - 2014-03-10 10:52 - 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/27555>

#9 - 2014-03-10 13:14 - 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/27555>

#10 - 2014-03-11 15:51 - 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/27555>

#11 - 2014-03-11 21:36 - Gerrit Code Review

- Estimated time set to 0.00

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/27555>

#12 - 2014-03-11 21:37 - 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/27555>

#13 - 2014-03-14 18:58 - Gerrit Code Review

- Estimated time set to 0.00

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/27555>

#14 - 2014-03-15 21:26 - 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/27555>

#15 - 2014-03-16 15:41 - 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/27555>

#16 - 2014-03-16 15:54 - 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/27555>

#17 - 2014-03-17 23:30 - Anonymous

- Status changed from Under Review to Resolved

Applied in changeset commit:2583fd2a12dd3e740a9739f73aa602c2793999ec.

#18 - 2014-03-18 20:51 - Alexander Stehlik

After the patch was applied I get this SQL error with the news extension:

```
SELECT COUNT(*)
FROM tx_news_domain_model_news
WHERE
(
tx_news_domain_model_news.uid IN (
SELECT uid_local FROM tx_news_domain_model_news_category_mm WHERE uid_foreign=?
```

```

)
AND tx_news_domain_model_news.pid IN ('99999999')
)
AND tx_news_domain_model_news.deleted=0 AND tx_news_domain_model_news.t3ver_state<=0
AND tx_news_domain_model_news.pid<>-1 AND tx_news_domain_model_news.hidden=0
AND tx_news_domain_model_news.starttime<=1395170580 AND (tx_news_domain_model_news.endtime=0 OR
tx_news_domain_model_news.endtime>1395170580)
AND (tx_news_domain_model_news.fe_group="
OR tx_news_domain_model_news.fe_group IS NULL
OR tx_news_domain_model_news.fe_group='0'
OR FIND_IN_SET('0',tx_news_domain_model_news.fe_group)
OR FIND_IN_SET('-1',tx_news_domain_model_news.fe_group)
)
)
AND (tx_news_domain_model_news.sys_language_uid IN (0,-1))

```

This is a var_dump of the sql query parts:

```

array (size=9)
'keywords' =>
array (size=0)
empty
'tables' =>
array (size=1)
'tx_news_domain_model_news' => string 'tx_news_domain_model_news' (length=25)
'unions' =>
array (size=0)
empty
'fields' =>
array (size=1)
'tx_news_domain_model_news' => string 'tx_news_domain_model_news.*' (length=27)
'where' =>
array (size=5)
0 => string '(' (length=1)
1 => string 'tx_news_domain_model_news.uid IN (SELECT uid_local FROM tx_news_domain_model_news_category_mm WHERE
uid_foreign=?)' (length=114)
2 => string ' AND ' (length=5)
3 => string 'tx_news_domain_model_news.pid IN (:AND1pid)' (length=43)
4 => string ')' (length=1)
'additionalWhereClause' =>
array (size=0)
empty
'orderings' =>
array (size=1)
0 => string 'tx_news_domain_model_news.datetime DESC' (length=39)
'limit' => null
'offset' => null

```

This is the error:

```

| You have an error in your SQL syntax; check the manual that corresponds to your MySQL server version for the right syntax to use near '?' AND
| tx_news_domain_model_news.pid IN ('99999999')) AND tx_news_domain_model_n' at line 1

```

I think the problem is the unnamed parameter (?) in the uid_foreign query.

I also pointed out some problems in the code in Gerrit.

#19 - 2014-03-18 21:29 - Alexander Stehlik

Think I found it. Regression in <https://forge.typo3.org/issues/57053>

#20 - 2014-03-20 13:54 - Alexander Stehlik

- *Estimated time set to 0.00*

Found another regression in #57104 :(

#21 - 2014-03-21 10:03 - Peter Beernink

Also #57130 seems to be another regression..

#22 - 2014-04-01 21:30 - Alexander Stehlik

- *Estimated time deleted (0.00)*