

Redmine - Defect #15781

Customfields have a noticeable impact on search performance due to slow database COUNT

2013-12-30 00:56 - Anonymous

Status:	Closed	Start date:	
Priority:	Normal	Due date:	
Assignee:	Jean-Philippe Lang	% Done:	0%
Category:	Custom fields	Estimated time:	0.00 hour
Target version:	2.4.3	Affected version:	
Resolution:	Fixed		

Description

Having any custom fields searchable **majorly** impacts search performance.

This would be helped by using something like Elasticsearch - See #10897 and #9180

For example, with just a few customfields being made searchable, it took 30 seconds for me to search for a single phrase. Inspecting the database query the slow part is the COUNT:

(23445.6ms)

```
EXPLAIN SELECT COUNT(DISTINCT "issues"."id") FROM "issues" LEFT OUTER JOIN "projects" ON "projects"."id" = "issues"."project_id" LEFT OUTER JOIN "journals" ON "journals"."journalized_id" = "issues"."id" AND (journals.private_notes = 'f' OR (projects.status <> 9 AND projects.id IN (SELECT em.project_id FROM enabled_modules em WHERE em.name='issue_tracking')) AND "journals"."journalized_type" = 'Issue' WHERE (projects.status <> 9 AND projects.id IN (SELECT em.project_id FROM enabled_modules em WHERE em.name='issue_tracking')) AND (((LOWER(subject) LIKE '%remotetransportexception%') OR (LOWER(issues.description) LIKE '%remotetransportexception%') OR (LOWER(journals.notes) LIKE '%remotetransportexception%') OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 1) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 1) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 1))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 4) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 4) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 4))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 11) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 11) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 11))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 12) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 12) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 12))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 13) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 13) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 13))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 17) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 17) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 17))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 18) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 18) AND issues.project_id IN (SELECT project_id FROM
```

custom_fields_projects WHERE custom_field_id = 18))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 19) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 19) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 19))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 21) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 21) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 21))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 26) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 26) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 26))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 27) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 27) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 27))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 31) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 31) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 31))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 32) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 32) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 32))) OR issues.id IN (SELECT customized_id FROM custom_values WHERE customized_type='Issue' AND customized_id=issues.id AND LOWER(value) LIKE '%remotetransportexception%' AND custom_values.custom_field_id = 25) AND ((1=1) AND (issues.tracker_id IN (SELECT tracker_id FROM custom_fields_trackers WHERE custom_field_id = 25) AND issues.project_id IN (SELECT project_id FROM custom_fields_projects WHERE custom_field_id = 25)))));

Related issues:

Related to Redmine - Feature # 10897: Offer an advanced issue query language ...	New	
Related to Redmine - Feature # 9180: Improve search system for issues - like ...	New	2011-09-04

Associated revisions

Revision 12481 - 2014-01-05 15:23 - Jean-Philippe Lang

Reduces the number of subqueries when searching with many custom fields set as searchable (#15781).

Revision 13004 - 2014-03-25 09:35 - Toshi MARUYAMA

add ChangeLog note that #15781 was forgotten to merge to v2.4.3.

Revision 13005 - 2014-03-25 09:39 - Toshi MARUYAMA

Merged r13004 from trunk to 2.5-stable.

add ChangeLog note that #15781 was forgotten to merge to v2.4.3.

Merged r13004 from trunk to 2.4-stable.

add ChangeLog note that #15781 was forgotten to merge to v2.4.3.

History

#1 - 2013-12-30 05:48 - Toshi MARUYAMA

- Category set to Custom fields

#2 - 2013-12-30 05:49 - Toshi MARUYAMA

- Related to Feature #10897: Offer an advanced issue query language as an alternative to it's current UI added

#3 - 2013-12-30 05:49 - Toshi MARUYAMA

- Related to Feature #9180: Improve search system for issues - like "context specific search" added

#4 - 2014-01-05 15:29 - Jean-Philippe Lang

- Assignee set to Jean-Philippe Lang

- Target version set to 2.4.3

r12481 should significantly improve that by reducing the number of subqueries. Here are the results for up to 7 searchable custom fields and ~10 000 issues on my dev machine, the table shows the execution time of the count query:

searchable fields	without fix	with fix applied
0	1.5	1.5
1	1.9	1.9
2	2.3	1.9
3	2.7	1.9
4	3.0	2.0
5	3.5	2.0
6	3.9	2.0
7	4.3	2.1

The results may depend on your custom fields visibility (different subqueries will still be done if custom field have different visibilities across project for the current user). If you are able to try the patch, I'd be happy to get some feedback.

#5 - 2014-01-06 05:16 - Anonymous

Thanks Jean-Philippe,

I've tested the patch out and my search query went from 10.1s to 7.5s

While this is an improvement, It's still very slow, I'd suggest that redmine would benefit greatly from utilising elasticsearch.

Jean-Philippe Lang wrote:

...

The results may depend on your custom fields visibility (different subqueries will still be done if custom field have different visibilities across project for the current user). If you are able to try the patch, I'd be happy to get some feedback.

#6 - 2014-01-09 08:40 - Jean-Philippe Lang

- Status changed from New to Closed

- Resolution set to Fixed

Thanks for the feedback, I'm closing it as using an alternative search engine is out of the scope of this request. It's true that sequential scan of many text columns is slow but IMHO using something like elasticsearch is overkill and having to install this along with Redmine would be a deal breaker for a few. Plus, I'm not sure it's designed to run queries with complex conditions on other things than text (eg. private projects, private notes, custom fields visibility vs. user's permissions).

We can get pretty good results by using advanced full text search features of the RDBMS. I've made a POC with Postgres, a few indexes and some SQL changes, the COUNT query with the same data now takes a few milliseconds.

#7 - 2014-02-10 18:19 - Ilya S

why i can't find this revision in 2.4.3.stable ? <http://www.redmine.org/releases/redmine-2.4.3.zip>

#8 - 2014-03-04 05:05 - Toshi MARUYAMA

Ilya S wrote:

| why i can't find this revision in 2.4.3.stable ? <http://www.redmine.org/releases/redmine-2.4.3.zip>

Really. r12481 has not merged to 2.4-stable.
Should we fix target version and ChangeLog?

#9 - 2014-03-25 09:41 - Toshi MARUYAMA

Toshi MARUYAMA wrote:

| Ilya S wrote:

| | why i can't find this revision in 2.4.3.stable ? <http://www.redmine.org/releases/redmine-2.4.3.zip>

| | Really. r12481 has not merged to 2.4-stable.

| | Should we fix target version and ChangeLog?

I have added note to ChangeLog and wiki.