

Redmine - Patch #21608

Project#allowed_to_condition performance

2015-12-26 04:19 - Ondřej Ezr

Status:	Closed	Start date:	
Priority:	Normal	Due date:	
Assignee:	Jean-Philippe Lang	% Done:	0%
Category:	Performance	Estimated time:	0.00 hour
Target version:	3.4.0		
Description			
Project allowed to condition is not using index for enabled_modules subquery. It can cause issues with growing count of projects.			
Related issues:			
Related to Redmine - Patch # 23196: Speed up Project.allowed_to_condition Closed			

Associated revisions

Revision 16242 - 2017-01-21 10:37 - Jean-Philippe Lang

Use EXISTS instead of IN subquery (#21608).

Patch by Ondřej Ezr.

History

#1 - 2015-12-26 06:07 - Go MAEDA

Does this patch resolve #19976 and #19102?

#2 - 2015-12-26 13:14 - Ondřej Ezr

It does not solve problem with list of project ids in those queries, but it speeds them up because the list is not the main problem of those queries.

It will resolve #19976 - it is exactly what slows that query (although, subquery instead of list of project ids would help too).

In #19102 it will speedup the queries for sure, but the list of ids again is caused by project_statement in query.rb and it is not what this issue is targeting, in our company we rewritten a query and solve that, but to extract and generalize that solution we currently have no capacities, but I will try to keep it in mind.

#3 - 2015-12-27 19:14 - Ondřej Ezr

I had gave it quick look and it would be solved by #21611
project_statement is really bad right now

#4 - 2015-12-28 01:25 - Go MAEDA

- Target version set to Candidate for next major release

Ondřej Ezr wrote:

|

2021-01-28

1/4

I had gave it quick look and it would be solved by #21611
project_statement is really bad right now

Thanks for investigating. I have set the target version to "Candidate for next major release".

#5 - 2016-01-03 09:10 - Jean-Philippe Lang

- Status changed from New to Needs feedback
- Target version deleted (Candidate for next major release)

The attached patch contains non Redmine code (plugins/easyproject/*). What am I supposed to do with that?

#6 - 2016-01-03 11:32 - Ondřej Ezr

- File project_visible_condition.diff added

I am really sorry, I have selected wrong file.

#7 - 2016-01-16 10:12 - Go MAEDA

- Status changed from Needs feedback to New
- Target version set to 3.3.0

The patch passed all tests.

I compared query plans when getting Latest News in Home screen.

current trunk

SQL:

```
SELECT "news".* FROM "news" INNER JOIN "projects" ON "projects"."id" = "news"."project_id" WHERE (projects.status <> 9 AND projects.id  
IN (SELECT em.project_id FROM enabled_modules em WHERE em.name='news')) ORDER BY news.created_on DESC LIMIT 5;
```

Query plan:

There is a full scan of enabled_modules table.

```
0|0|1|SEARCH TABLE projects USING INTEGER PRIMARY KEY (rowid=?)  
0|0|0|EXECUTE LIST SUBQUERY 1  
1|0|0|SCAN TABLE enabled_modules AS em  
0|1|0|SEARCH TABLE news USING INDEX news_project_id (project_id=?)  
0|0|0|USE TEMP B-TREE FOR ORDER BY
```

with the patch applied

SQL:

```
SELECT "news".* FROM "news" INNER JOIN "projects" ON "projects"."id" = "news"."project_id" WHERE (projects.status <> 9 AND EXISTS (SELECT 1 AS one FROM enabled_modules em WHERE em.project_id = projects.id AND em.name='news')) ORDER BY news.created_on DESC LIMIT 5;
```

Query plan:

No full scan of enabled_modules table.

```
0|0|0|SCAN TABLE news USING INDEX index_news_on_created_on
0|1|1|SEARCH TABLE projects USING INTEGER PRIMARY KEY (rowid=?)
0|0|0|EXECUTE CORRELATED SCALAR SUBQUERY 1
1|0|0|SEARCH TABLE enabled_modules AS em USING INDEX enabled_modules_project_id (project_id=?)
```

#8 - 2016-01-16 11:29 - Jean-Philippe Lang

Which database are you using? I'm using PostgreSQL 9.3 and I get the same query plans before and after the patch. In some case, it uses the enabled_modules_project_id index but not always (eg. for Issue.visible.count). I made these tests against 1k projects and ~10k lines in enabled_modules.

I was able to make Issue.visible.count use the index by doing a JOIN on enabled_modules instead of a subquery but that would require a bit more refactoring to implement. Compared to the above queries, that would be:

```
SELECT "news".* FROM "news"
INNER JOIN "projects" ON "projects"."id" = "news"."project_id"
INNER JOIN "enabled_modules" ON "enabled_modules"."project_id" = "projects"."id" AND "enabled_modules"."name" = 'news'
WHERE projects.status <> 9
ORDER BY news.created_on DESC LIMIT 5;
```

#9 - 2016-01-18 01:15 - Go MAEDA

Thanks for deep inspection.

I tested on SQLite3. I will test again on MySQL.

#10 - 2016-05-08 09:20 - Jean-Philippe Lang

- Target version changed from 3.3.0 to Candidate for next major release

More feedback needed before applying this change. Full scans can be a problem but correlated subqueries can be a problem too.

#11 - 2016-06-29 19:19 - Jean-Philippe Lang

- Related to Patch #23196: Speed up Project.allowed_to_condition added

#12 - 2017-01-21 10:37 - Jean-Philippe Lang

- Status changed from New to Closed
- Assignee set to Jean-Philippe Lang
- Target version changed from Candidate for next major release to 3.4.0

Patch committed.

#13 - 2017-03-20 08:57 - Go MAEDA

- Subject changed from Project allowed to condition performance to Project#allowed_to_condition performance

#14 - 2018-09-07 15:18 - Jérôme BATAILLE

Hi,

I have a question about this patch, isn't it introducing a slight difference ?

- Before projects were rejected if they had not the module enabled.
- Now, no result is returned, if no project has the module enabled.

Am I wrong, does it matter ?

#15 - 2018-09-07 16:54 - Pavel Rosický

Jérôme BATAILLE wrote:

Hi,
I have a question about this patch, isn't it introducing a slight difference ? - Before projects were rejected if they had not the module enabled.
- Now, no result is returned, if no project has the module enabled.
Am I wrong, does it matter ?

Before: project IN (all projects with [enabled module])

After: EXISTS ([enabled_module]) for a project (not all projects)

If you have many projects and enabled modules, the second version should be faster. Results should be equal.

<https://stackoverflow.com/questions/24929/difference-between-exists-and-in-in-sql/24948>

Files

project_query_speedup.patch	4.19 KB	2015-12-26	Ondřej Ezr
project_visible_condition.diff	953 Bytes	2016-01-03	Ondřej Ezr