

redmineorg-copy202205 - Vote #66688

Migrating from trac is not possible, fails to allocate memory

2022/05/09 14:32 - Admin Redmine

ステータス:	Closed	開始日:	2010/02/25
優先度:	通常	期日:	
担当者:		進捗率:	0%
カテゴリ:	Importers_15	予定工数:	0.00時間
対象バージョン:	0.9.4_18	作業時間:	0.00時間
Redmineorg_URL:	https://www.redmine.org/issues/4931	status_id:	5
category_id:	15	tracker_id:	1
version_id:	18	plus1:	0
issue_org_id:	4931	affected_version:	
author_id:	8477	closed_on:	
assigned_to_id:	0	affected_version_id:	
comments:	7		

説明

When trying to import a trac repository, the process will fail due to memory allocation errors (tested on Debian, with 512mb ram and a few GB of swap).

```
Migrating components.....
Migrating milestones.....
Migrating custom fields..
Migrating tickets.....(TRIMMED DOTS).....rake aborted!
failed to allocate memory

(See full trace by running task with --trace)
```

Eventually, this worked when I commented a few lines in the migration script:

```
turing:~/redmine-0.8# svn diff
Index: lib/tasks/migrate_from_trac.rake
=====
--- lib/tasks/migrate_from_trac.rake      (revision 2903)
+++ lib/tasks/migrate_from_trac.rake      (working copy)
@@ -309,12 +309,12 @@
@@ -309,12 +309,12 @@
text = text.gsub(/[(\d+)\.]/, 'r\1')
# Ticket number re-writing
text = text.gsub(/#(\d+)/) do |s|
-   if $1.length < 10
-       TICKET_MAP[$1.to_i] ||= $1
-       "\#{TICKET_MAP[$1.to_i] || $1}"
-   else
-       s
-   end
+   if $1.length < 10
+       TICKET_MAP[$1.to_i] ||= $1
+       "\#{TICKET_MAP[$1.to_i] || $1}"
+   else
+       s
+   end
end
# We would like to convert the Code highlighting too
# This will go into the next line.
```

However, this is "not optimal" and I would like to have a better solution. See: <http://www.redmine.org/boards/2/topics/8923>

journals

How many tickets do you have in your Trac database?

Can you retry without the useless assignment:

Index: lib/tasks/migrate_from_trac.rake

```
-----
--- lib/tasks/migrate_from_trac.rake      (revision 3487)
+++ lib/tasks/migrate_from_trac.rake      (working copy)
@@ -310,7 +310,7 @@
     # Ticket number re-writing
     text = text.gsub(/#(\d+)/) do |s|
       if $1.length < 10
-         TICKET_MAP[$1.to_i] ||= $1
+         TICKET_MAP[$1.to_i] ||= $1
+         "\#{TICKET_MAP[$1.to_i] || $1}"
       else
         s
       end
     end
   end
end
```

Sorry for the late-update, I have been fighting Debian to get ROR running :)

I have about 1300 tickets, without many interactions (some do mention "ticket:1234" or something similar, but most do not). Your "patch" fixed the issue, now the imported does work.

Let me understand, the memory consumption was $O(n^2)$ when n is the amount of tickets?

Let me understand, the memory consumption was $O(n^2)$ when n is the amount of tickets?

No, the ticket count was not the problem (1300 is ok).

I think you add many or large `#\d+` patterns in some texts (wiki, ticket descriptions...) that were not ticket ids. And they were filling the `TICKET_MAP` hash.

Thanks for the feedback.

Yes, each ticket title contains 2-3 numbers... Now lets work on a backport of this to Debian :)

I have contacted the Debian developers, in order to fix the Debian package for the upcoming release (so Redmine is part of the official distribution). They asked if this bug will be part of 0.9.4 (which apparently will happen, as you tagged it properly).

Anyway, if you are interested, the bug is found here: <http://bugs.debian.org/576682> (this includes another bug in Debian + a patch, feel free to comment on this if you want).

Yes, it will be part of 0.9.4. It was merged in 0.9-stable in r3643.

履歴

#1 - 2022/05/10 17:23 - Admin Redmine

- カテゴリ を Importers_15 にセット

- 対象バージョン を 0.9.4_18 にセット