

Foreman - Bug #5568

Possible memory leak in trends

05/05/2014 08:04 AM - Lukas Zapletal

Status: Closed	
Priority: Normal	
Assignee: Jon McKenzie	
Category: Trends	
Target version: 1.8.1	
Difficulty:	Fixed in Releases:
Triaged:	Found in Releases:
Bugzilla link:	Red Hat JIRA:
Pull request: https://github.com/foreman/foreman/pull/2365	
Description	
By Sean Alderman (https://groups.google.com/forum/#!msg/foreman-users/hi7TagIRE8A/5ZRzll8ChxEJ):	
I am collecting a few trends in foreman using host facts. I'm using a cron entry for the foreman user: <code>/usr/sbin/foreman-rake trends:counter</code> every half hour.	
My foreman machine is running puppet, puppetdb, foreman, foreman-proxy, and postgres, 2 vCPU and 6GB of RAM. It seems that the rake commands are running away with all the memory and swap. Is there a way to control this?	
Here's what top shows currently:	
<pre>PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND 30002 foreman 20 0 6602m 3.3g 104 D 13.3 58.3 17:06.22 /opt/rh/ruby193/root/usr/bin/ruby /opt/rh/ruby193/root/usr/bin/rake trends:clean</pre>	
Seems like a lot of RAM and accumulated CPU time. I can add more RAM, but other than rake commands like this one, the server doesn't seem to utilize the memory allocated.	
Related issues:	
Related to Foreman - Bug #4114: Trends don't scale well and become very slow	Duplicate 01/17/2014

Associated revisions

Revision 040abfa3 - 05/13/2015 03:25 AM - Jon McKenzie

Fixes #5568 - Improves performance of trends:clean rake task

Perform trend counter dupe counting within the database rather than in the Ruby code.

Revision 5b7ac025 - 05/20/2015 07:25 AM - Jon McKenzie

Fixes #5568 - Improves performance of trends:clean rake task

Perform trend counter dupe counting within the database rather than in the Ruby code.

(cherry picked from commit 040abfa32f326de7cd488f34f243377f76fd70ae)

Conflicts:

test/lib/tasks/trends_test.rb

History

#1 - 05/06/2014 07:50 AM - Dominic Cleal

- Related to Bug #4114: Trends don't scale well and become very slow added

#2 - 09/16/2014 10:41 AM - Edson Manners

Seem to have a similar issue. Currently running 1.4.5 on RHEL 6.5.

We have 545 hosts and I'm not running and anytrends manually foreman seems to be running these trends on its own. I'm running this in a KVM VM with 6 cores and 16GB RAM at this point but it's obviously never enough. I'll be upgrading through 1.5 and 1.6 in the coming days and looking to see if that helps.

Top show this after 'service foreman stop; service httpd stop'

```
30143 foreman 20 0 8407m 7.9g 6544 D 11.6 50.9 42:44.19 rake
```

I have to 'kill -9' the rake process as it doesn't die on its own.

```
'ps -lef | grep 30143'
```

```
/opt/rh/ruby193/root/usr/bin/ruby /opt/rh/ruby193/root/usr/bin/rake trends:clean
```

#3 - 02/04/2015 10:29 AM - Shimon Shtein

Can it be related to the amount of data in the trend_counters table?

#4 - 05/07/2015 12:05 PM - Jon McKenzie

We're running into this issue too. I haven't done a ton with ActiveRecord, but the issue are these lines in the trends:clean rake task:

```
counts = TrendCounter.group([:trend_id, :created_at]).count
```

```
dupes = counts.select{ |attrs, count| count > 1 }
```

Rather than doing the aggregation inside the database, the task first executes a very expensive GROUP and then filters the results it needs in plain Ruby.

Again, not an ActiveRecord expert, but I think replacing those two lines with the following should significantly speed it up:

```
dupes = TrendCounter.having('count(*) > 1').group([:trend_id, :created_at]).count
```

This translates the SQL from this:

```
SELECT COUNT(*) AS count_all, trend_id AS trend_id, created_at AS created_at FROM "trend_counters" GROUP BY trend_id, created_at ORDER BY created_at
```

...to this:

```
SELECT COUNT(*) AS count_all, trend_id AS trend_id, created_at AS created_at FROM "trend_counters" GROUP BY trend_id, created_at HAVING count(*) > 1 ORDER BY created_at
```

#5 - 05/09/2015 02:47 PM - Jon McKenzie

Huh, I must have misread this issue initially, which is for trends:counter not trends:clean. Oh well, I submitted a PR for the trends:clean issue anyways

#6 - 05/11/2015 04:37 AM - Dominic Cleal

- Status changed from Assigned to Ready For Testing

- Assignee deleted (Lukas Zapletal)

- Pull request <https://github.com/foreman/foreman/pull/2365> added

- Pull request deleted ()

#7 - 05/13/2015 04:01 AM - Jon McKenzie

- Status changed from Ready For Testing to Closed

- % Done changed from 0 to 100

Applied in changeset [040abfa32f326de7cd488f34f243377f76fd70ae](#).

#8 - 05/13/2015 04:36 AM - Dominic Cleal

- Assignee set to Jon McKenzie

- translation missing: en.field_release set to 50