Foreman - Bug #3886

Facts page is slow to load with many facts stored

12/16/2013 11:37 AM - Dominic Cleal

Status:	New		
Priority:	Normal		
Assignee:			
Category:	Performance		
Target version:			
Difficulty:		**Fixed in Releases:**	
Triaged:		**Found in Releases:**	
Bugzilla link:		**Red Hat JIRA:**	
Pull request:			

Description

https://gist.github.com/xorpaul/7232272

```
Cache read: http://foreman/fact_values?
Started GET "/fact_values" for 1.2.3.4 at 2013-10-30 13:58:39 +0100
Cache read: _session_id:e28f9d66d14e0dac7a4bd3978087a14c
Processing by FactValuesController#index as HTML
  User Load (1.0ms)  SELECT "users".* FROM "users" WHERE "users"."id" = $1 LIMIT 1  [["id", 1]]
Setting current user thread-local variable to admin
  FactValue Load (0.9ms)  SELECT "fact_values".* FROM "fact_values" LIMIT 1
Cache read: entries_per_page
  SQL (3455.2ms)  SELECT "fact_values"."id" AS t0_r0, "fact_values"."value" AS t0_r1, "fact_values"."fact_name_id" AS t0_r2, "fact_values"."host_id" AS t0_r3, "fact_values"."updated_at" AS t0_r4, "fact_values"."created_at" AS t0_r5, "fact_names"."id" AS t1_r0, "fact_names"."name" AS t1_r1, "fact_names"."updated_at" AS t1_r2, "fact_names"."created_at" AS t1_r3, "hosts"."id" AS t2_r0, "hosts"."name" AS t2_r1, "hosts"."ip" AS t2_r2, "hosts"."last_compile" AS t2_r3, "hosts"."last_freshcheck" AS t2_r4, "hosts"."last_report" AS t2_r5, "hosts"."updated_at" AS t2_r6, "hosts"."source_file_id" AS t2_r7, "hosts"."created_at" AS t2_r8, "hosts"."mac" AS t2_r9, "hosts"."root_pass" AS t2_r10, "hosts"."serial" AS t2_r11, "hosts"."puppet_status" AS t2_r12, "hosts"."domain_id" AS t2_r13, "hosts"."architecture_id" AS t2_r14, "hosts"."operatingsystem_id" AS t2_r15, "hosts"."environment_id" AS t2_r16, "hosts"."subnet_id" AS t2_r17, "hosts"."ptable_id" AS t2_r18, "hosts"."medium_id" AS t2_r19, "hosts"."build" AS t2_r20, "hosts"."comment" AS t2_r21, "hosts"."disk" AS t2_r22, "hosts"."installed_at" AS t2_r23, "hosts"."model_id" AS t2_r24, "hosts"."hostgroup_id" AS t2_r25, "hosts"."owner_id" AS t2_r26, "hosts"."owner_type" AS t2_r27, "hosts"."enabled" AS t2_r28, "hosts"."puppet_ca_proxy_id" AS t2_r29, "hosts"."managed" AS t2_r30, "hosts"."use_image" AS t2_r31, "hosts"."image_file" AS t2_r32, "hosts"."uuid" AS t2_r33, "hosts"."compute_resource_id" AS t2_r34, "hosts"."puppet_proxy_id" AS t2_r35, "hosts"."certname" AS t2_r36, "hosts"."image_id" AS t2_r37, "hosts"."organization_id" AS t2_r38, "hosts"."location_id" AS t2_r39, "hosts"."type" AS t2_r40 FROM "fact_values" LEFT OUTER JOIN "fact_names" ON "fact_names"."id" = "fact_values"."fact_name_id" LEFT OUTER JOIN "hosts" ON "hosts"."id" = "fact_values"."host_id" AND "hosts"."type" IN ('Host::Managed') WHERE (fact_names.name <> '_timestamp') ORDER BY "fact_values"."value" ASC NULLS FIRST LIMIT 75 OFFSET 0
   (1711.7ms)  SELECT COUNT(DISTINCT "fact_values"."id") FROM "fact_values" LEFT OUTER JOIN "fact_names" ON "fact_names"."id" = "fact_values"."fact_name_id" LEFT OUTER JOIN "hosts" ON "hosts"."id" = "fact_values"."host_id" AND "hosts"."type" IN ('Host::Managed') WHERE (fact_names.name <> '_timestamp')
  Rendered fact_values/index.html.erb within layouts/application (5533.1ms)
Cache read: views/tabs_and_title_records-1
Read fragment views/tabs_and_title_records-1 (1.1ms)
  Rendered home/_topbar.html.erb (2.6ms)
  Bookmark Load (1.0ms)  SELECT "bookmarks".* FROM "bookmarks" WHERE (controller = 'fact_values') ORDER BY name
  Rendered common/_searchbar.html.erb (5.4ms)
Completed 200 OK in 5557ms (Views: 379.0ms | ActiveRecord: 5169.7ms)
Cache write: _session_id:e28f9d66d14e0dac7a4bd3978087a14c ({:expires_in=>nil})
```

And query plan:

```
Limit  (cost=154639.97..154640.15 rows=75 width=2720)
   ->  Sort  (cost=154639.97..158587.07 rows=1578842 width=2720)
         Sort Key: fact_values.value
         ->  Hash Left Join  (cost=697.21..97574.15 rows=1578842 width=2720)
               Hash Cond: (fact_values.host_id = hosts.id)
               ->  Hash Join  (cost=133.46..63460.01 rows=1578842 width=88)
                     Hash Cond: (fact_values.fact_name_id = fact_names.id)
                     ->  Seq Scan on fact_values  (cost=0.00..37666.54 rows=1579454 width=45)
                     ->  Hash  (cost=101.24..101.24 rows=2578 width=43)
                           ->  Seq Scan on fact_names  (cost=0.00..101.24 rows=2578 width=43)
                                 Filter: ((name)::text <> '_timestamp'::text)
               ->  Hash  (cost=471.38..471.38 rows=7390 width=2632)
                     ->  Seq Scan on hosts  (cost=0.00..471.38 rows=7390 width=2632)
                           Filter: ((type)::text = 'Host::Managed'::text)
(14 rows)
```

xorpaul reports removing the sort drops this to 8 milliseconds, but removing the sort could be problematic for pagination and maybe grouping (nested facts?).

https://github.com/theforeman/foreman/pull/992 has more info + discussion.