

Foreman - Bug #35689

Queries take too long to return when you have lots of facts

10/28/2022 03:39 PM - Guilherme M. Schroeder

Status: Ready For Testing	
Priority: Normal	
Assignee: Adam Ruzicka	
Category: Search	
Target version:	
Difficulty:	Fixed in Releases:
Triaged: No	Found in Releases: 3.4.0
Bugzilla link:	Red Hat JIRA:
Pull request: https://github.com/foreman/foreman/pull/9724	
Description	
<p>Foreman 3.4.0 with ~ 6k hosts + 900 in discovery mode. Receiving facts from Puppet. ~1.8 million facts in the fact_values table.</p> <p>When trying to do a query like "facts.architecture=x86_64" in the search box of "Hosts -> Discovered Hosts", it takes more than 4 minutes (check debug.txt attached).</p> <p>The query is constructed as the following:</p> <pre>irb(main):001:0> puts Host::Discovered.search_for("facts.architecture=x86_64").to_sql SELECT "hosts".* FROM "hosts" INNER JOIN "fact_values" fact_values_588605 ON (hosts.id = fact_values_588605.host_id) INNER JOIN "fact_names" fact_names_588605 ON (fact_names_588605.id = fact_values_588605.fact_name_id) WHERE "hosts"."type" = 'Host::Discovered' AND ((fact_names_588605."name" = 'architecture') AND ("hosts"."id" IN (SELECT "host_id" FROM "fact_values" WHERE "fact_values_588605"."value" = 'x86_64'))) ORDER BY "hosts"."created_at" DESC NULLS LAST => nil irb(main):002:0></pre> <p>It can be simplified to:</p> <pre>SELECT "hosts".* FROM "hosts" INNER JOIN "fact_values" fact_values_588605 ON (hosts.id = fact_values_588605.host_id) INNER JOIN "fact_names" fact_names_588605 ON (fact_names_588605.id = fact_values_588605.fact_name_id) WHERE "hosts"."type" = 'Host::Discovered' AND "hosts"."organization_id" = '33' AND "hosts"."organization_id" = '33' AND fact_names_588605."name" = 'architecture' AND "fact_values_588605"."value" = 'x86_64' ORDER BY "hosts"."created_at" DESC NULLS LAST LIMIT '100' OFFSET '0';</pre> <p>And it takes only 32ms to return.</p>	
Related issues:	
Copied to Discovery - Bug #36491: Queries take too long to return when you ha...	Ready For Testing

History

#1 - 10/28/2022 03:50 PM - Ewoud Kohl van Wijngaarden

- Project changed from Foreman to Discovery
- Category set to Discovery plugin

#2 - 06/06/2023 02:28 PM - Adam Ruzicka

- Project changed from Discovery to Foreman
- Category changed from Discovery plugin to Search

#3 - 06/06/2023 02:33 PM - The Foreman Bot

- Status changed from New to Ready For Testing
- Assignee set to Adam Ruzicka
- Pull request <https://github.com/foreman/foreman/pull/9724> added

#4 - 06/08/2023 12:52 PM - Adam Ruzicka

- Copied to Bug #36491: Queries take too long to return when you have lots of facts added

Files

debug.txt	166 KB	10/28/2022	Guilherme M. Schroeder
-----------	--------	------------	------------------------