

Foreman - Bug #17143

Performance degradation in "non-targeted" search through WebUI or API

10/31/2016 01:12 PM - Konstantin Orekhov

Status: New	
Priority: Normal	
Assignee:	
Category: Search	
Target version:	
Difficulty:	Fixed in Releases:
Triaged:	Found in Releases:
Bugzilla link:	Red Hat JIRA:
Pull request:	
Description	
<p>After upgrading to 1.13.1, I immediately noticed the huge delays in search through the hosts in both WebUI and API (the latter is actually worse than WebUI). Here are few examples:</p>	
<pre>1. time curl -kSs -u admin:\$PASS https://localhost/api/hosts jq '.' grep total "total": 1695, "subtotal": 1695, real 0m0.816s user 0m0.040s sys 0m0.076s</pre>	
<p>But the search through my 1695 hosts on this instances takes 2.5 minutes (!):</p>	
<pre>1. time curl -kSs -u admin:\$PASS https://localhost/api/hosts?search=test jq '.' grep total "total": 1695, "subtotal": 1, real 2m25.147s user 0m0.053s sys 0m0.063s</pre>	
<p>Corresponding logs (WebUI search):</p>	
<pre>2016-10-28 18:39:29 a607439e [app] [I] Parameters: {"utf8"=>"", "search"=>"test"} 2016-10-28 18:40:16 a607439e [app] [I] Rendered hosts/_list.html.erb (10.1ms) 2016-10-28 18:40:16 a607439e [app] [I] Rendered hosts/index.html.erb within layouts/application (18.0ms) 2016-10-28 18:40:17 a607439e [app] [I] Rendered common/_searchbar.html.erb (11.6ms) 2016-10-28 18:40:17 a607439e [app] [I] Rendered layouts/_application_content.html.erb (12.2ms) 2016-10-28 18:40:17 a607439e [app] [I] Rendered home/_user_dropdown.html.erb (1.9ms) 2016-10-28 18:40:17 a607439e [app] [I] Read fragment views/tabs_and_title_records-6 (1.1ms) 2016-10-28 18:40:17 a607439e [app] [I] Rendered home/_topbar.html.erb (43.9ms) 2016-10-28 18:40:17 a607439e [app] [I] Rendered layouts/base.html.erb (45.5ms) 2016-10-28 18:40:17 a607439e [app] [I] Completed 200 OK in 47691ms (Views: 73.3ms ActiveRecord: 47585.0ms)</pre>	
<p>API:</p>	
<pre>2016-10-29 13:44:51 8a9d5e4a [app] [I] Started GET "/api/hosts?search=test" for 127.0.0.1 at 2016-10-29 13:44:51 -0700 2016-10-29 13:44:51 8a9d5e4a [app] [I] Processing by Api::V2::HostsController#index as JSON 2016-10-29 13:44:51 8a9d5e4a [app] [I] Parameters: {"search"=>"test", "apiv"=>"v2"} 2016-10-29 13:44:51 8a9d5e4a [app] [I] Authorized user admin(Admin User) 2016-10-29 13:46:10 8a9d5e4a [app] [I] Rendered api/v2/hosts/index.json.rabl within api/v2/layouts/index_layout (58.6ms) 2016-10-29 13:47:28 8a9d5e4a [app] [I] Completed 200 OK in 156710ms (Views: 62.3ms ActiveRecord: 156519.5ms)</pre>	
<p>This was not the case in previous versions. "targeted" search (search=name~something) still works perfectly fine. Is this an intended change in behavior?</p>	
Related issues:	

History**#1 - 11/01/2016 04:08 AM - Dominic Cleal**

- Category set to Search

Please include debug logs with the 'sql' logger enabled: <https://theforeman.org/manuals/1.13/index.html#7.2Debugging>

#2 - 11/01/2016 04:20 PM - Konstantin Orekhov

Here you go - <https://gist.github.com/korekhov/7ad0fddae6e330f1655305c626bb4808>

#3 - 11/01/2016 04:21 PM - Konstantin Orekhov

There's also a thread for this issue - <https://groups.google.com/forum/#!topic/foreman-users/oT68AXrskV0>

#4 - 11/02/2016 04:00 AM - Dominic Cleal

The slowest request is:

```
2016-11-01 13:01:23 e58c75eb [sql] [D] SQL (79753.0ms) SELECT DISTINCT `hosts`.`id` FROM `hosts` LEFT OUTER JOIN `host_status` ON `host_status`.`host_id` = `hosts`.`id` LEFT OUTER JOIN `compute_resources` ON `compute_resources`.`id` = `hosts`.`compute_resource_id` LEFT OUTER JOIN `hostgroups` ON `hostgroups`.`id` = `hosts`.`hostgroup_id` LEFT OUTER JOIN `operatingsystems` ON `operatingsystems`.`id` = `hosts`.`operatingsystem_id` LEFT OUTER JOIN `nics` ON `nics`.`host_id` = `hosts`.`id` LEFT OUTER JOIN `tokens` ON `tokens`.`host_id` = `hosts`.`id` LEFT OUTER JOIN `models` ON `models`.`id` = `hosts`.`model_id` LEFT OUTER JOIN `nics` `primary_interfaces_hosts_join` ON `primary_interfaces_hosts_join`.`host_id` = `hosts`.`id` AND `primary_interfaces_hosts_join`.`primary` = 1 LEFT OUTER JOIN `domains` ON `domains`.`id` = `primary_interfaces_hosts_join`.`domain_id` LEFT OUTER JOIN `realms` ON `realms`.`id` = `hosts`.`realm_id` LEFT OUTER JOIN `environments` ON `environments`.`id` = `hosts`.`environment_id` LEFT OUTER JOIN `architectures` ON `architectures`.`id` = `hosts`.`architecture_id` LEFT OUTER JOIN `images` ON `images`.`id` = `hosts`.`image_id` LEFT OUTER JOIN `nics` `primary_interfaces_hosts` ON `primary_interfaces_hosts`.`host_id` = `hosts`.`id` AND `primary_interfaces_hosts`.`primary` = 1 LEFT OUTER JOIN `nics` `primary_interfaces_hosts_join_2` ON `primary_interfaces_hosts_join_2`.`host_id` = `hosts`.`id` AND `primary_interfaces_hosts_join_2`.`primary` = 1 LEFT OUTER JOIN `subnets` ON `subnets`.`id` = `primary_interfaces_hosts_join_2`.`subnet_id` AND `subnets`.`type` = 'Subnet::ipv4' LEFT OUTER JOIN `nics` `primary_interfaces_hosts_join_3` ON `primary_interfaces_hosts_join_3`.`host_id` = `hosts`.`id` AND `primary_interfaces_hosts_join_3`.`primary` = 1 LEFT OUTER JOIN `subnets` `subnet6s_hosts` ON `subnet6s_hosts`.`id` = `primary_interfaces_hosts_join_3`.`subnet6_id` AND `subnet6s_hosts`.`type` = 'Subnet::ipv6' LEFT OUTER JOIN `nics` `provision_interfaces_hosts` ON `provision_interfaces_hosts`.`host_id` = `hosts`.`id` AND `provision_interfaces_hosts`.`provision` = 1 LEFT OUTER JOIN `discovery_rules` ON `discovery_rules`.`id` = `hosts`.`discovery_rule_id` LEFT OUTER JOIN `host_salt_modules` ON `host_salt_modules`.`host_id` = `hosts`.`id` LEFT OUTER JOIN `salt_modules` ON `salt_modules`.`id` = `host_salt_modules`.`salt_module_id` LEFT OUTER JOIN `salt_environments` ON `salt_environments`.`id` = `hosts`.`salt_environment_id` LEFT OUTER JOIN `smart_proxies` ON `smart_proxies`.`id` = `hosts`.`salt_proxy_id` WHERE `hosts`.`type` IN ('Host::Managed') AND ((`hosts`.`name` LIKE '%test%' OR `hosts`.`comment` LIKE '%test%' OR `models`.`name` LIKE '%test%' OR `hostgroups`.`name` LIKE '%test%' OR `hostgroups`.`title` LIKE '%test%' OR `hostgroups`.`title` LIKE '%test%' OR `domains`.`name` LIKE '%test%' OR `realms`.`name` LIKE '%test%' OR `environments`.`name` LIKE '%test%' OR `architectures`.`name` LIKE '%test%' OR `compute_resources`.`name` LIKE '%test%' OR `images`.`name` LIKE '%test%' OR `operatingsystems`.`name` LIKE '%test%' OR `operatingsystems`.`description` LIKE '%test%' OR `operatingsystems`.`title` LIKE '%test%' OR `operatingsystems`.`major` LIKE '%test%' OR `operatingsystems`.`minor` LIKE '%test%' OR `nics`.`ip` LIKE '%test%' OR `nics`.`ip` LIKE '%test%' OR `nics`.`mac` LIKE '%test%' OR `subnets`.`network` LIKE '%test%' OR `subnets`.`name` LIKE '%test%' OR `subnets`.`network` LIKE '%test%' OR `subnets`.`name` LIKE '%test%' OR `hosts`.`uuid` LIKE '%test%' OR `nics`.`mac` LIKE '%test%' OR `operatingsystems`.`name` LIKE '%test%' OR `operatingsystems`.`description` LIKE '%test%' OR `operatingsystems`.`title` LIKE '%test%' OR `operatingsystems`.`major` LIKE '%test%' OR `operatingsystems`.`minor` LIKE '%test%' OR `discovery_rules`.`name` LIKE '%test%' OR `salt_modules`.`name` LIKE '%test%' OR `salt_environments`.`name` LIKE '%test%' OR `smart_proxies`.`name` LIKE '%test%')) ORDER BY `hosts`.`name` ASC LIMIT 20 OFFSET 0
```

#5 - 11/02/2016 04:22 AM - Ohad Levy

my guess is that now we search on both hosts and nic tables.

Could you run that query with explain?

thanks

#6 - 11/02/2016 04:37 PM - Konstantin Orekhov

One thing I forgot to mention. This is the same DB that was migrated from 1.7.1 to 1.13.1 and this fix applied to it - <https://groups.google.com/forum/#!msg/foreman-users/qkmfeP2R-ss/vUE3FJHsBQAJ>

Maybe I need to reindex? How, if so?

BTW, "explain analyze" fail for me with this kind of error:

You have an error in your SQL syntax; check the manual that corresponds to your MySQL server version for the right syntax to use near 'analyze SELECT DISTINCT `hosts`.`id` FROM `hosts` LEFT OUTER JOIN `host_status` ' at line 1

If I just run "explain". I get this (if that helps at all):

```
1 SIMPLE hosts ALL
PRIMARY,index_hosts_on_name,index_hosts_on_last_report,index_hosts_on_installed_at,host_arch_id_ix,host_os_id_ix,host_env_id_ix,host_mediu
m_id_ix,host_group_id_ix,index_hosts_on_certname,index_hosts_on_type,hosts_compute_resource_id_fk,hosts_image_id_fk,hosts_model_id_fk,h
osts_ptable_id_fk,hosts_puppet_ca_proxy_id_fk,hosts_puppet_proxy_id_fk,hosts_location_id_fk,hosts_organization_id_fk,index_hosts_on_compute_
profile_id,hosts_realms_id_fk NULL NULL NULL 1620 Using where; Using temporary; Using filesort
1 SIMPLE host_status ref index_host_status_on_host_id index_host_status_on_host_id 4 foreman.hosts.id 1 Using index; Distinct
1 SIMPLE compute_resources eq_ref PRIMARY PRIMARY 4 foreman.hosts.compute_resource_id 1 Distinct
1 SIMPLE hostgroups eq_ref PRIMARY PRIMARY 4 foreman.hosts.hostgroup_id 1 Distinct
1 SIMPLE operatingsystems eq_ref PRIMARY PRIMARY 4 foreman.hosts.operatingsystem_id 1 Distinct
1 SIMPLE nics ref index_by_host index_by_host 5 foreman.hosts.id 20 Distinct
1 SIMPLE tokens ref index_tokens_on_host_id index_tokens_on_host_id 5 foreman.hosts.id 1 Using index; Distinct
1 SIMPLE models eq_ref PRIMARY PRIMARY 4 foreman.hosts.model_id 1 Distinct
1 SIMPLE primary_interfaces_hosts_join ref index_by_host index_by_host 5 foreman.hosts.id 20 Using where; Distinct
1 SIMPLE domains eq_ref PRIMARY PRIMARY 4 foreman.primary_interfaces_hosts_join.domain_id 1 Distinct
1 SIMPLE realms eq_ref PRIMARY PRIMARY 4 foreman.hosts.realm_id 1 Distinct
1 SIMPLE environments eq_ref PRIMARY PRIMARY 4 foreman.hosts.environment_id 1 Distinct
1 SIMPLE architectures eq_ref PRIMARY PRIMARY 4 foreman.hosts.architecture_id 1 Distinct
1 SIMPLE images eq_ref PRIMARY PRIMARY 4 foreman.hosts.image_id 1 Distinct
1 SIMPLE primary_interfaces_hosts ref index_by_host index_by_host 5 foreman.hosts.id 20 Using where; Distinct
1 SIMPLE primary_interfaces_hosts_join_2 ref index_by_host index_by_host 5 foreman.hosts.id 20 Using where; Distinct
1 SIMPLE subnets eq_ref PRIMARY,index_subnets_on_type PRIMARY 4 foreman.primary_interfaces_hosts_join_2.subnet_id 1
Using where; Distinct
1 SIMPLE primary_interfaces_hosts_join_3 ref index_by_host index_by_host 5 foreman.hosts.id 20 Using where; Distinct
1 SIMPLE subnet6s_hosts eq_ref PRIMARY,index_subnets_on_type PRIMARY 4 foreman.primary_interfaces_hosts_join_3.subnet6_id
Using where; Distinct
1 SIMPLE provision_interfaces_hosts ref index_by_host index_by_host 5 foreman.hosts.id 20 Using where; Distinct
1 SIMPLE discovery_rules eq_ref PRIMARY PRIMARY 4 foreman.hosts.discovery_rule_id 1 Distinct
1 SIMPLE host_salt_modules ALL NULL NULL NULL NULL 1 Using where; Distinct; Using join buffer (Block Nested Loop)
1 SIMPLE salt_modules eq_ref PRIMARY PRIMARY 4 foreman.host_salt_modules.salt_module_id 1 Distinct
1 SIMPLE salt_environments eq_ref PRIMARY PRIMARY 4 foreman.hosts.salt_environment_id 1 Distinct
1 SIMPLE smart_proxies eq_ref PRIMARY PRIMARY 4 foreman.hosts.salt_proxy_id 1 Using where; Distinct
```

#7 - 11/02/2016 05:44 PM - Konstantin Orekhov

Please see "explain extended" outputs I just posted at the request from Tomer Brisker:

<https://groups.google.com/d/msg/foreman-users/oT68AXrskV0/XeqYmhQZCgAJ>
<https://groups.google.com/d/msg/foreman-users/oT68AXrskV0/dcc7ER4ZCgAJ>

#8 - 11/03/2016 09:33 AM - Tomer Brisker

- Related to Bug #17196: Add indexes to the host_salt_modules and hostgroup_salt_modules tables added