

Foreman - Bug #25788

Redundant taxonomy enumeration causes significant performance degradation.

01/06/2019 01:16 PM - Shimon Shtein

Status:	New	
Priority:	Normal	
Assignee:		
Category:		
Target version:		
Difficulty:		Fixed in Releases:
Triaged:	No	Found in Releases:
Bugzilla link:		Red Hat JIRA:
Pull request:		

Description

In used_organization_ids and used_location_ids a call to get_taxonomy_ids is used with which_organization/which_location sent in as a parameter respectively. It looks like which_location already contains all sub-locations and needs not to be expanded further. Each expansion like this causes another sql query per each sub-location.

How to reproduce:

Create a big locations tree with 3 root nodes:

America
USA
LA
WA
NY
Canada
...
Europe
Asia

Assign around 100 subnets to a single sub-loc (for example LA)

Now add only LA to a non-admin user.
Perform api/v2/subnets query with this user.

Clone the user, but add to it all three root locations (America, Europe, Asia).
Perform api/v2/subnets query with the cloned user.

Observe time difference between requests