

## Foreman - Bug #25788

### Redundant taxonomy enumeration causes significant performance degradation.

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

<b>Status:</b> New	
<b>Priority:</b> Normal	
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b>	
<b>Difficulty:</b>	<b>Fixed in Releases:</b>
<b>Triaged:</b> No	<b>Found in Releases:</b>
<b>Bugzilla link:</b>	<b>Red Hat JIRA:</b>
<b>Pull request:</b>	
<b>Description</b>	
<p>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.</p>	
<p>How to reproduce:</p>	
<p>Create a big locations tree with 3 root nodes:</p>	
<p>America USA LA WA NY Canada ... Europe Asia</p>	
<p>Assign around 100 subnets to a single sub-loc (for example LA)</p>	
<p>Now add only LA to a non-admin user. Perform api/v2/subnets query with this user.</p>	
<p>Clone the user, but add to it all three root locations (America, Europe, Asia). Perform api/v2/subnets query with the cloned user.</p>	
<p>Observe time difference between requests</p>	