

Foreman - Bug #10329

report::expire still running very slowly (at least on PostgreSQL)

04/30/2015 02:15 AM - Dirk Heinrichs

Status: New	
Priority: High	
Assignee:	
Category: Puppet Reports	
Target version:	
Difficulty:	Fixed in Releases:
Triaged:	Found in Releases: 1.7.4
Bugzilla link:	Red Hat JIRA:
Pull request:	
Description	
<p>This is a followup to #8565. With many long reports in the database (large messages and logs tables), expiring reports takes quite a long time on our machine. Sometimes more than a day so that multiple jobs sum up and start hogging the CPU. Here's how the situation looks like:</p>	
<pre>foreman=> select count(*) from reports; count ----- 1951 (1 row)</pre>	
<pre>foreman=> select count(*) from messages; count ----- 704462 (1 row)</pre>	
<pre>foreman=> select count(*) from logs; count ----- 802815 (1 row)</pre>	
<pre>foreman=> explain DELETE FROM "messages" WHERE (id not IN (SELECT DISTINCT message_id FROM "logs"));</pre>	
QUERY PLAN	
<pre>----- Delete on messages (cost=0.42..9806323789.43 rows=352231 width=6) -> Seq Scan on messages (cost=0.42..9806323789.43 rows=352231 width=6) Filter: (NOT (SubPlan 1)) SubPlan 1 -> Materialize (cost=0.42..27142.52 rows=279366 width=4) -> Unique (cost=0.42..24653.69 rows=279366 width=4) -> Index Only Scan using index_logs_on_message_id on logs (cost=0.42..226 46.65 rows=802815 width=4) (7 rows)</pre>	

History

#1 - 05/07/2015 07:48 AM - Dirk Heinrichs

Just wanted to add the link to the [blog post](#) giving a possible solution to the problem.