

Foreman - Feature #31274

Reports generation is memory heavy

11/06/2020 11:55 AM - Lukas Zapletal

Status:	Ready For Testing	
Priority:	Normal	
Assignee:	Lukas Zapletal	
Category:	Templates	
Target version:		
Difficulty:		Fixed in Releases:
Triaged:	No	Found in Releases:
Bugzilla link:		Red Hat JIRA:
Pull request:	https://github.com/theforeman/foreman/pull/8136	
Description		
<p>The way we generate reports is memory unfriendly. All rows are stored in an array as hashes and then dumped into string which is returned to the client. In background job case, string is stored in the db or emailed.</p> <p>Refactoring: Rewrite report_headers, report_row and report_render helpers so they immediately stream data into an IO object. Do not use any intermediate data structure, all formats we currently support (JSON, YAML, CSV, plaintext) do support this. When called via API, data can be streamed directly over TCP. In case of background processing, data can be streamed via SQL LOAD postgresql directly into blob/text.</p> <p>This could be implemented without any user-facing changes - the three helpers will remain the same, but they will immediately stream content instead of storing anything in memory. The last call (report_render) will output optional footer and close the stream.</p>		
Related issues:		
Related to Foreman - Bug #31200: Puma memory leak		Need more information

History

#1 - 11/06/2020 11:56 AM - Lukas Zapletal

- Related to Bug #31200: Puma memory leak added

#2 - 11/12/2020 05:48 PM - The Foreman Bot

- Status changed from New to Ready For Testing

- Assignee set to Lukas Zapletal

- Pull request <https://github.com/theforeman/foreman/pull/8136> added