

Pulp - Story #6024

As a plugin writer, I can iterate over repository versions in a memory efficient manner

01/18/2020 09:55 PM - gmbnomis

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	gmbnomis	% Done:	100%
Category:		Estimated time:	0:00 hour
Sprint/Milestone:	3.1.0	Tags:	
Platform Release:		Sprint:	
Groomed:	No	Quarter:	
Sprint Candidate:	No		

Description

Motivation

Plugins need to iterate over the content in a repository version to create a publication. Pulp should offer a batching iterator over content in order to limit the memory footprint and to allow efficient use of content batches in DB operations.

Solution

We could make `batch_qs()` available to plugin writers. But this generator is hard to use correctly because of the ordering requirement (even in pulpcore it has not been used correctly).

Instead, add a generator to `RepositoryVersion`, which does "the right thing" by default and is customizable by plugin writers if necessary.

For an example use case, see

https://github.com/gmbnomis/pulp_cookbook/blob/batched_publication_upstream/pulp_cookbook/app/tasks/publishing.py#L87-L117

Associated revisions

Revision 2a9b2882 - 01/21/2020 08:55 PM - gmbnomis

Add a batching content iterator `content_batch_qs()` to `RepositoryVersion`

Plugins need to iterate over the content in a repository version to create a publication.

`RepositoryVersion.content_batch_qs()` generates content batches to efficiently iterate over all content.

fixes: #6024 <https://pulp.plan.io/issues/6024>

History

#1 - 01/18/2020 10:28 PM - gmbnomis

- Description updated

#2 - 01/18/2020 11:09 PM - gmbnomis

PR is at: <https://github.com/pulp/pulpcore/pull/509>

#3 - 01/21/2020 06:38 PM - bmbouter

- Status changed from NEW to POST

PR available at: <https://github.com/pulp/pulpcore/pull/509>

#4 - 01/21/2020 11:08 PM - gmbnomis

- Status changed from POST to MODIFIED

- % Done changed from 0 to 100

Applied in changeset [pulpcore|2a9b28821b204329946536d878e0e1f7687ed521.](#)

#5 - 01/31/2020 12:24 PM - bmbouter

- Sprint/Milestone set to 3.1.0

#6 - 01/31/2020 12:25 PM - bmbouter

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE