

RPM Support - Issue #4031

MemoryError during sync of a repo with many errata

09/20/2018 12:43 PM - ttereshc

Status:	CLOSED - NOTABUG	Start date:	
Priority:	Normal	Due date:	
Assignee:		Estimated time:	0:00 hour
Category:		Groomed:	No
Sprint/Milestone:	Pulp 3 RPM MVP	Sprint Candidate:	No
Severity:	2. Medium	Tags:	
Version:		Sprint:	
Platform Release:		Quarter:	
OS:			
Triaged:	No		
Description			
To reproduce, sync EPEL7 repo https://dl.fedoraproject.org/pub/epel/7/x86_64/			
<pre>Sep 20 07:48:51 pulp3.dev rq[6934]: MemoryError Sep 20 07:48:51 pulp3.dev rq[6934]: Traceback (most recent call last): Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6/site-p ackages/rq/worker.py", line 793, in perform_job Sep 20 07:48:51 pulp3.dev rq[6934]: rv = job.perform() Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6/site-p ackages/rq/job.py", line 599, in perform Sep 20 07:48:51 pulp3.dev rq[6934]: self._result = self._execute() Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6/site-p ackages/rq/job.py", line 605, in _execute Sep 20 07:48:51 pulp3.dev rq[6934]: return self.func(*self.args, **self.kwargs) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/devel/pulp_rpm/pulp_rpm/app/tasks/synchr onizing.py", line 69, in synchronize Sep 20 07:48:51 pulp3.dev rq[6934]: loop.run_until_complete(pipeline) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/usr/lib64/python3.6/asyncio/base_events.py", line 468 , in run_until_complete Sep 20 07:48:51 pulp3.dev rq[6934]: return future.result() Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/devel/pulp/plugin/pulpcore/plugin/stages /api.py", line 58, in create_pipeline Sep 20 07:48:51 pulp3.dev rq[6934]: await asyncio.gather(*futures) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/devel/pulp/plugin/pulpcore/plugin/stages /content_unit_stages.py", line 185, in __call__ Sep 20 07:48:51 pulp3.dev rq[6934]: await self._post_save(batch) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/devel/pulp_rpm/pulp_rpm/app/tasks/synchr onizing.py", line 324, in _post_save Sep 20 07:48:51 pulp3.dev rq[6934]: UpdateCollectionPackage.objects.bulk_create(update_collect ion_packages_to_save) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6/site-p ackages/django/db/models/manager.py", line 82, in manager_method Sep 20 07:48:51 pulp3.dev rq[6934]: return getattr(self.get_queryset(), name)(*args, **kwargs) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6/site-p ackages/django/db/models/query.py", line 457, in bulk_create Sep 20 07:48:51 pulp3.dev rq[6934]: objs_with_pk, objs_without_pk = partition(lambda o: o.pk i s None, objs) Sep 20 07:48:51 pulp3.dev rq[6934]: File "/home/vagrant/.virtualenvs/pulp/lib64/python3.6/site-p ackages/django/utils/functional.py", line 387, in partition Sep 20 07:48:51 pulp3.dev rq[6934]: results[predicate(item)].append(item) Sep 20 07:48:51 pulp3.dev rq[6934]: MemoryError</pre>			

History

#1 - 09/20/2018 12:45 PM - ttereshc

Maybe related to [#4014](#) (relations between collections UpdateRecords are incorrect)

#2 - 09/20/2018 03:37 PM - daviddavis

- *Sprint/Milestone set to Pulp 3 RPM MVP*

#3 - 09/21/2018 03:09 PM - ttereshc

- *Status changed from NEW to CLOSED - NOTABUG*

The problem is no longer observed after fixing [#4014](#).
Feel free to reopen, if you hit it.

I believe one can still reproduce the issue if they put 8K+ collections in one erratum but it's unlikely to happen in real-life repositories.

#4 - 04/26/2019 10:33 PM - bmbouter

- *Tags deleted (Pulp 3)*