

Pulp - Issue #7147

AttributeError: 'Package' object has no attribute '_remote_artifact_saver_cas'

07/15/2020 03:21 PM - ttereshc

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	dkliban@redhat.com	Estimated time:	0:00 hour
Category:		Groomed:	No
Sprint/Milestone:	3.8.0	Sprint Candidate:	No
Severity:	2. Medium	Tags:	
Version:		Sprint:	Sprint 83
Platform Release:		Quarter:	
OS:			
Triaged:	Yes		

Description

To reproduce, in pulp 2:

- create repo A with the feed http://mirror.linux.duke.edu/pub/centos/8/BaseOS/x86_64/kickstart/ and sync
- create repo B with the same feed and sync
- migrate both repos

Observe the error below.

```
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: Traceback (most recent call last):
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/usr/local/lib/pulp/lib64/python3.6/site-packages/rq/worker.py", line 886, in perform_job
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: rv = job.perform()
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/usr/local/lib/pulp/lib64/python3.6/site-packages/rq/job.py", line 664, in perform
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: self._result = self._execute()
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/usr/local/lib/pulp/lib64/python3.6/site-packages/rq/job.py", line 670, in _execute
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: return self.func(*self.args, **self.kwargs)
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulp-2to3-migration/pulp_2to3_migration/app/tasks/migrate.py", line 127, in migrate_from_pulp2
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: migrate_content(plan)
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulp-2to3-migration/pulp_2to3_migration/app/migration.py", line 34, in migrate_content
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: plugin.migrator.migrate_content_to_pulp3()
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulp-2to3-migration/pulp_2to3_migration/app/plugin/rpm/migrator.py", line 140, in migrate_content_to_pulp3
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: loop.run_until_complete(dm.create())
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/usr/lib64/python3.6/asyncio/base_events.py", line 484, in run_until_complete
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: return future.result()
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulp-2to3-migration/pulp_2to3_migration/app/plugin/content.py", line 84, in create
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: await pipeline
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagr
```

```
ant/devel/pulpcore/pulpcore/plugin/stages/api.py", line 225, in create_pipeline
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: await asyncio.gather(*futures)
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulpcore/pulpcore/plugin/stages/api.py", line 43, in __call__
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: await self.run()
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulpcore/pulpcore/plugin/stages/artifact_stages.py", line 225, in run
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: RemoteArtifact.objects.bulk_get_or_create(self._needed_remote_artifacts(batch))
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: File "/home/vagrant/devel/pulpcore/pulpcore/plugin/stages/artifact_stages.py", line 264, in _needed_remote_artifacts
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: for content_artifact in d_content.content._remote_artifact_saver_cas:
Jul 15 12:21:58 pulp2-nightly-pulp3-source-centos7.rhgoose.example.com rq[18841]: AttributeError: 'Package' object has no attribute '_remote_artifact_saver_cas'
```

Observations:

- there are almost duplicated declarative content units in a [batch](#) which only differ by remote (that's expected, if one content unit has multiple LCE in pulp2)
- If I ignore the exception and proceed to the next `d_content` [here](#), I see that the even batches are failing and the odd ones are correct. E.g. batch#1 won't raise `AttributeError`, all content in the batch#2 will, batch#3 will work fine, batch#4 will raise `AttributeError`.

Related issues:

Related to Pulp - Backport #7702: Backport Request: 7147 (AttributeError: 'Pa...

CLOSED - CURRENTRELEASE

Associated revisions

Revision 8dcdcd2d - 10/08/2020 04:32 PM - dkliban@redhat.com

Handles duplicate content in the stages API.

If the very first Content in a batch has already been prefetched in a previous call to `prefetch_related_objects`, Django does not set the attribute specified in `to_attr` when prefetching the rest of the items in the batch. As a result, the `'_remote_artifact_saver_cas'` attribute is not set on the Content instances and `AttributeError` is raised when `'_remote_artifact_saver_cas'` is accessed by the `RemoteArtifactSaverStage`.

This patch removes any existing `'_remote_artifact_saver_cas'` attributes on Content before calling into Django's `prefetch_related_objects()`.

fixes: #7147 <https://pulp.plan.io/issues/7147>

History

#1 - 07/15/2020 03:24 PM - ttereshc

- Description updated

#2 - 07/28/2020 03:40 PM - ttereshc

- Project changed from Migration Plugin to Pulp

#3 - 07/28/2020 03:41 PM - ttereshc

- Description updated

#4 - 07/28/2020 04:41 PM - fao89

- Triaged changed from No to Yes

- Sprint set to Sprint 78

#5 - 08/07/2020 04:33 PM - rchan

- Sprint changed from Sprint 78 to Sprint 79

#6 - 08/24/2020 01:10 PM - rchan

- Sprint changed from Sprint 79 to Sprint 80

#7 - 09/04/2020 05:24 PM - rchan

- Sprint changed from Sprint 80 to Sprint 81

#8 - 09/18/2020 09:29 PM - rchan

- Sprint changed from Sprint 81 to Sprint 82

#9 - 09/22/2020 07:34 PM - dkliban@redhat.com

- Status changed from NEW to ASSIGNED

- Assignee set to dkliban@redhat.com

#10 - 09/27/2020 05:36 PM - pulpbot

- Status changed from ASSIGNED to POST

PR: <https://github.com/pulp/pulpcore/pull/937>

#11 - 10/02/2020 06:05 PM - rchan

- Sprint changed from Sprint 82 to Sprint 83

#12 - 10/08/2020 04:32 PM - dkliban@redhat.com

- Status changed from POST to MODIFIED

Applied in changeset [pulpcore|8dcddcd2dfb6d1f217e9725db9bbead1eef86c079](#).

#13 - 10/13/2020 04:55 PM - jsherril@redhat.com

- Related to Backport #7702: Backport Request: 7147 (AttributeError: 'Package' object has no attribute '_remote_artifact_saver_cas') to 3.7 added

#14 - 10/20/2020 05:44 PM - ttereshc

- Sprint/Milestone set to 3.8.0

#15 - 10/20/2020 09:38 PM - pulpbot

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE

#16 - 10/20/2020 09:39 PM - pulpbot

PR: <https://github.com/pulp/pulpcore/pull/975>