

## Docker Support - Story #3497

As a user, I can upload skopeo tar that references existing layers.

03/16/2018 03:13 PM - tomckay@redhat.com

<b>Status:</b>	CLOSED - CURRENTRELEASE	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	jortel@redhat.com	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0:00 hour
<b>Sprint/Milestone:</b>	2.17.0	<b>Tags:</b>	Pulp 2
<b>Platform Release:</b>	2.17.0	<b>Sprint:</b>	Sprint 38
<b>Target Release - Docker:</b>		<b>Quarter:</b>	
<b>Groomed:</b>	Yes		
<b>Sprint Candidate:</b>	Yes		

**Description**

**Motivation:**

To save bandwidth and improve efficiency, layers referred to in a skopeo tar manifest.json that already exist should be allowed to be omitted from the tar.

As I am implementing "docker push" and "skopeo copy docker:", both tools first check if a layer exists before attempting to push it. Since pulp expects all the layers to be part of the upload this leads to having to get the layer from pulp, add it to the tar file, then give tar file to pulp.

**Solution:**

The upload API will support uploading archives that contain only a subset of layers needed to create the full container image. Archive must contain manifest.json file. When such an archive is uploaded to Pulp, it will be added to the repository only if the missing layers are already present in the repository. If the layers are missing from both the archive being uploaded and repository the upload will fail.

### Associated revisions

**Revision 891eff64 - 06/06/2018 09:02 PM - jortel@redhat.com**

Support upload tarball with incomplete images. closes #3497

**Revision 891eff64 - 06/06/2018 09:02 PM - jortel@redhat.com**

Support upload tarball with incomplete images. closes #3497

**Revision 891eff64 - 06/06/2018 09:02 PM - jortel@redhat.com**

Support upload tarball with incomplete images. closes #3497

**Revision 891eff64 - 06/06/2018 09:02 PM - jortel@redhat.com**

Support upload tarball with incomplete images. closes #3497

### History

**#1 - 03/16/2018 03:40 PM - dkliban@redhat.com**

This seems reasonable to me. Our documentation would need to state that this type of upload has a race condition with the orphan cleanup task.

The upload code for pulp\_docker would try to find any blobs missing from the tarball already in Pulp. If the blob is found, an association with the repository is made. If not found, an error is returned.

**#2 - 03/16/2018 03:51 PM - dalley**

- Tracker changed from Issue to Story
- % Done set to 0

**#3 - 03/16/2018 04:43 PM - ipanova@redhat.com**

Update: we'll make a query to the destination repo of the upload and check if the blob mentioned in manifest.json is present. In case it is present in the repo, proceed to the next blob. If the blob is not present, check its existence in the tar file. In case it is not in the tar file - fail upload.

**#4 - 05/28/2018 05:32 PM - ipanova@redhat.com**

- Sprint Candidate changed from No to Yes

**#5 - 05/29/2018 08:32 PM - dkliban@redhat.com**

- Description updated

**#6 - 05/30/2018 02:40 PM - ipanova@redhat.com**

- Description updated

**#7 - 05/30/2018 02:41 PM - ipanova@redhat.com**

- Description updated
- Groomed changed from No to Yes

**#8 - 06/01/2018 04:14 PM - rchan**

- Sprint set to Sprint 38

**#9 - 06/04/2018 09:10 PM - jortel@redhat.com**

- Status changed from NEW to ASSIGNED
- Assignee set to jortel@redhat.com

**#10 - 06/04/2018 10:08 PM - jortel@redhat.com**

[ipanova@redhat.com](mailto:ipanova@redhat.com) wrote:

Update: we'll make a query to the destination repo of the upload and check if the blob mentioned in manifest.json is present. In case it is present in the repo, proceed to the next blob. If the blob is not present, check its existence in the tar file. In case it is not in the tar file - fail upload.

@ipanova, here you suggest that a blob not included in the uploaded archive must already exist in pulp AND must already be associated with the destination repository. If the blob exists in pulp but is not associated with the destination repository, why fail the upload instead of associating it as @dkliban suggested in [#3497-1](#)? Is your distinction intentional?

**#11 - 06/04/2018 10:13 PM - tomckay@redhat.com**

@jortel - It is intentional. Using this paradigm, a user can't upload a manifest that references layers that they don't really have access to. By forcing the layer to exist in the repository already, this forces the initial upload of the layer as proof of ownership.

**#12 - 06/05/2018 05:37 PM - jortel@redhat.com**

[tomckay@redhat.com](mailto:tomckay@redhat.com) wrote:

@jortel - It is intentional. Using this paradigm, a user can't upload a manifest that references layers that they don't really have access to. By forcing the layer to exist in the repository already, this forces the initial upload of the layer as proof of ownership.

Okay. Thank you, @ipanova.

**#13 - 06/06/2018 06:38 PM - jortel@redhat.com**

- Status changed from ASSIGNED to POST

[https://github.com/pulp/pulp\\_docker/pull/236](https://github.com/pulp/pulp_docker/pull/236)

**#14 - 06/12/2018 04:29 PM - jortel@redhat.com**

- Status changed from POST to MODIFIED

- % Done changed from 0 to 100

Applied in changeset [891eff64b28ef1cc37faf4609635fda615cadb3a](#).

**#15 - 06/18/2018 02:56 PM - dkliban@redhat.com**

- Platform Release set to 2.16.2

**#16 - 06/18/2018 06:08 PM - dkliban@redhat.com**

- Platform Release deleted (2.16.2)

**#17 - 06/18/2018 11:57 PM - amacdona@redhat.com**

- Platform Release set to 2.16.2

**#18 - 06/19/2018 12:09 AM - amacdona@redhat.com**

- Platform Release deleted (2.16.2)

**#19 - 07/09/2018 04:09 PM - ipanova@redhat.com**

- Platform Release set to 2.17.0

**#20 - 07/09/2018 05:32 PM - rchan**

- Sprint/Milestone set to 2.17.0

**#21 - 07/09/2018 05:33 PM - rchan**

Adding to 2.17.0 milestone. This is one of the required deliverables.

**#22 - 08/07/2018 06:31 PM - ipanova@redhat.com**

- Platform Release changed from 2.17.0 to master

**#23 - 08/07/2018 06:31 PM - ipanova@redhat.com**

- Platform Release deleted (master)

**#24 - 08/07/2018 06:36 PM - ipanova@redhat.com**

- Platform Release set to 2.17.0

**#25 - 08/14/2018 11:17 PM - ipanova@redhat.com**

- Status changed from MODIFIED to 5

**#26 - 08/31/2018 10:43 PM - ipanova@redhat.com**

- Status changed from 5 to CLOSED - CURRENTRELEASE

**#27 - 04/15/2019 10:12 PM - bmbouter**

- Tags Pulp 2 added