

## Container Support - Story #4172

### Remove duplicate tags from repository during sync

11/19/2018 06:18 PM - amacdona@redhat.com

<b>Status:</b>	CLOSED - CURRENTRELEASE	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	amacdona@redhat.com	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0:00 hour
<b>Sprint/Milestone:</b>	1.0.0	<b>Tags:</b>	
<b>Platform Release:</b>		<b>Sprint:</b>	Sprint 46
<b>Groomed:</b>	Yes	<b>Quarter:</b>	
<b>Sprint Candidate:</b>	Yes		
<b>Description</b>			
<p>The docker plugin has repository version uniqueness constraints for Tags. For example the Tag with name "latest" points to a Manifest List. Another Tag with name "latest" (pointing to another Manifest List) cannot be added to the repository version while the first "latest" tag remains. Note: This repository version could have a "latest" that points to a manifest instead of a manifest list.</p> <p>At this point, we don't think this uniqueness constraint can exist in the db, so we need to enforce the constraint at sync time (and later during add/remove.</p> <p>During sync, before the new tags are added to the repository version, corresponding tags that are already in the repository version should be removed.</p> <p>Heres an example: RepositoryVersion1 contains a tag, "latest" which points to ManifestList1. A sync retrieves a new Tag "latest" that points to ManifestList2. Before "latest"-&gt; ManifestList2 is added to the repository version, "latest"-&gt;ManifestList1 must be removed.</p> <p>For implementation, this probably should be a new Stage, which could be the last item in the pipeline. (Add/remove stages are added outside of the plugin-defined pipeline.)</p>			
<b>Related issues:</b>			
Blocked by Pulp - Story #3934: As a plugin writer, I can have a stage that re...		<b>CLOSED - CURRENTRELEASE</b>	
Copied to Container Support - Test #4299: Remove duplicate tags from reposito...		<b>NEW</b>	

#### Associated revisions

##### Revision 7f79cd73 - 12/18/2018 09:25 PM - amacdona@redhat.com

Remove non-unique tag names during sync

When syncing, if a tag name references a different manifest/manifest list it must be removed from the repo. To do this simply, it was more convinient to split the Tag model into ManifestTag and ManifestListTag.

<https://pulp.plan.io/issues/4172> fixes #4172

##### Revision f871b7d5 - 12/18/2018 09:25 PM - amacdona@redhat.com

Distribute tags according to requested media type

Return the artifacts associated with a tag with type failover. Try manifest list first, then manifest schema 2. If both fail, log the warning and return a 404. This provides a convinient location for an on-the-fly schema converter.

These changes are necessary to use the split-tag models for distribution, so they will be merged as a part of #4172.

<https://pulp.plan.io/issues/4172> re #4172

**Revision 7f79cd73 - 12/18/2018 09:25 PM - amacdona@redhat.com**

Remove non-unique tag names during sync

When syncing, if a tag name references a different manifest/manifest list it must be removed from the repo. To do this simply, it was more convenient to split the Tag model into ManifestTag and ManifestListTag.

<https://pulp.plan.io/issues/4172> fixes #4172

**Revision f871b7d5 - 12/18/2018 09:25 PM - amacdona@redhat.com**

Distribute tags according to requested media type

Return the artifacts associated with a tag with type failover. Try manifest list first, then manifest schema 2. If both fail, log the warning and return a 404. This provides a convenient location for an on-the-fly schema converter.

These changes are necessary to use the split-tag models for distribution, so they will be merged as a part of #4172.

<https://pulp.plan.io/issues/4172> re #4172

**Revision 7f79cd73 - 12/18/2018 09:25 PM - amacdona@redhat.com**

Remove non-unique tag names during sync

When syncing, if a tag name references a different manifest/manifest list it must be removed from the repo. To do this simply, it was more convenient to split the Tag model into ManifestTag and ManifestListTag.

<https://pulp.plan.io/issues/4172> fixes #4172

**Revision f871b7d5 - 12/18/2018 09:25 PM - amacdona@redhat.com**

Distribute tags according to requested media type

Return the artifacts associated with a tag with type failover. Try manifest list first, then manifest schema 2. If both fail, log the warning and return a 404. This provides a convenient location for an on-the-fly schema converter.

These changes are necessary to use the split-tag models for distribution, so they will be merged as a part of #4172.

<https://pulp.plan.io/issues/4172> re #4172

**Revision 7f79cd73 - 12/18/2018 09:25 PM - amacdona@redhat.com**

Remove non-unique tag names during sync

When syncing, if a tag name references a different manifest/manifest list it must be removed from the repo. To do this simply, it was more convenient to split the Tag model into ManifestTag and ManifestListTag.

<https://pulp.plan.io/issues/4172> fixes #4172

Revision f871b7d5 - 12/18/2018 09:25 PM - amacdona@redhat.com

Distribute tags according to requested media type

Return the artifacts associated with a tag with type failover. Try manifest list first, then manifest schema 2. If both fail, log the warning and return a 404. This provides a convenient location for an on-the-fly schema converter.

These changes are necessary to use the split-tag models for distribution, so they will be merged as a part of #4172.

<https://pulp.plan.io/issues/4172> re #4172

## History

---

**#1 - 11/28/2018 04:14 PM - ipanova@redhat.com**

- Groomed changed from No to Yes
- Sprint Candidate changed from No to Yes

**#2 - 11/28/2018 04:27 PM - amacdona@redhat.com**

- Related to Story #3934: As a plugin writer, I can have a stage that removes duplicates added

**#3 - 11/28/2018 05:44 PM - amacdona@redhat.com**

- Related to deleted (Story #3934: As a plugin writer, I can have a stage that removes duplicates)

**#4 - 11/28/2018 05:44 PM - amacdona@redhat.com**

- Blocked by Story #3934: As a plugin writer, I can have a stage that removes duplicates added

**#5 - 11/28/2018 06:00 PM - amacdona@redhat.com**

- Sprint set to Sprint 46

**#6 - 12/07/2018 02:03 PM - amacdona@redhat.com**

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

**#7 - 12/18/2018 10:32 PM - amacdona@redhat.com**

- Status changed from ASSIGNED to MODIFIED
- % Done changed from 0 to 100

Applied in changeset [7f79cd73ee9eee532fa9a7adb816fd7abc6c19dc](https://pulp.plan.io/changesets/details?id=7f79cd73ee9eee532fa9a7adb816fd7abc6c19dc).

**#8 - 12/20/2018 02:25 AM - amacdona@redhat.com**

[https://github.com/pulp/pulp\\_docker/pull/282](https://github.com/pulp/pulp_docker/pull/282)

**#9 - 01/07/2019 05:40 PM - bherring**

- Copied to Test #4299: Remove duplicate tags from repository during sync added

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

- Tags deleted (Pulp 3)

**#11 - 11/20/2019 07:19 PM - ipanova@redhat.com**

- *Project changed from Docker Support to Container Support*

**#12 - 12/13/2019 06:44 PM - bmbouter**

- *Status changed from MODIFIED to CLOSED - CURRENTRELEASE*

**#13 - 01/22/2020 08:00 PM - ipanova@redhat.com**

- *Sprint/Milestone set to 1.0.0*