

Docker Support - Story #3450

As a user I would like to limit the tags we sync for docker repos

03/09/2018 07:19 PM - jmontleo@redhat.com

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

Description

Motivation:

We're currently sync'ing some very large openshift repos from brew. In reality we're only interested in the current 3.9 and newer releases/tags but we're getting everything, which is weighing in at 1.2TB at present and growing with each puddle for each version.

It would be nice to be able to limit the tags we sync.

Solution:

Provide a config on the DockerImporter config called "tags". This config will accept a list of tags to sync.

Provide a CLI option called "--tags". The user will be able to specify "--tags" when creating a docker repository. The user will also be able to override the "--tags" importer config at sync time, by specifying it as part of the override_configs.

When the "tags" option is provided:

Pulp will sync just those tags + manifests associated with the tags and all the corresponding layers. Sync will only be performed using the immediate download policy. There will be no support for the on_demand download policy.

Note: If the list of tags was changed, and tag X is not wanted anymore, a manual removal of that tag X should occur in the local repo.

Note: If some invalid tags provided are invalid, Pulp will only sync the available ones in the remote repo.

When the "tags" option is not provided:

Pulp will sync the whole repo.

Related issues:

Has duplicate Docker Support - Story #3623: As a user I can sync a registry b... **CLOSED - DUPLICATE**

Associated revisions

Revision e46d3b3e - 06/13/2018 08:23 PM - werwty

Add feature to whitelist tags on sync.

closes #3450 <https://pulp.plan.io/issues/3450>

Revision e46d3b3e - 06/13/2018 08:23 PM - werwty

Add feature to whitelist tags on sync.

closes #3450 <https://pulp.plan.io/issues/3450>

History

#1 - 03/09/2018 09:30 PM - tomckay@redhat.com

This usage would also be useful with foreman as well. I can see two ways:

1. Provide a tag filter up front to limit tags pulled down
2. Just pull the metadata files for the tags and sync the actual images later (exactly like for RPM units)

Personally, I would prefer the second option.

#2 - 05/02/2018 01:55 PM - ipanova@redhat.com

- Has duplicate Story #3137: As a user, I can view docker image Arch and Size Information through the REST API added

#3 - 05/02/2018 01:56 PM - ipanova@redhat.com

- Has duplicate deleted (Story #3137: As a user, I can view docker image Arch and Size Information through the REST API)

#4 - 05/02/2018 01:56 PM - ipanova@redhat.com

- Has duplicate Story #3623: As a user I can sync a registry by whitelisting tags (Filtered Sync) added

#5 - 05/03/2018 12:56 PM - ipanova@redhat.com

an option will be provided where a list of tags to sync can be specified. for example --tags

if option is provided:

we will sync just those tags + its' manifests + correspondent layers. Sync is performed in immediate poilicy, no on_demand.

note if the list of tags was changed, and tag X is not wanted anymore, a **manual** removal of that tag X should occur in the local repo.

note if some of invalid tags where provided in the list to sync, we will sync just available ones in the remote repo.

if option is not provided:

we will sync by default whole repo.

#6 - 05/28/2018 05:32 PM - ipanova@redhat.com

- Sprint Candidate changed from No to Yes

#7 - 05/29/2018 08:08 PM - dkliban@redhat.com

- Description updated

#8 - 05/30/2018 02:37 PM - ipanova@redhat.com

- Groomed changed from No to Yes

#9 - 06/01/2018 04:14 PM - rchan

- Sprint set to Sprint 38

#10 - 06/05/2018 05:50 PM - bizhang

- Status changed from NEW to ASSIGNED

- Assignee set to bizhang

#11 - 06/12/2018 08:03 PM - bizhang

- Status changed from ASSIGNED to POST

PR: https://github.com/pulp/pulp_docker/pull/241

#12 - 06/14/2018 04:55 PM - werwty

- Status changed from POST to MODIFIED

- % Done changed from 0 to 100

Applied in changeset [e46d3b3ee31f31fe5d83958bc96dbf1e9ffef32a](#).

#13 - 06/18/2018 03:01 PM - dkliban@redhat.com

- Platform Release set to 2.16.2

#14 - 06/18/2018 06:08 PM - dkliban@redhat.com

- Platform Release deleted (2.16.2)

#15 - 06/18/2018 11:57 PM - amacdonga@redhat.com

- Platform Release set to 2.16.2

#16 - 06/19/2018 12:09 AM - amacdonga@redhat.com

- Platform Release deleted (2.16.2)

#17 - 07/09/2018 04:10 PM - ipanova@redhat.com

- Platform Release set to 2.17.0

#18 - 07/09/2018 05:32 PM - rchan

- Sprint/Milestone set to 2.17.0

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

#19 - 08/07/2018 06:31 PM - ipanova@redhat.com

- Platform Release changed from 2.17.0 to master

#20 - 08/07/2018 06:31 PM - ipanova@redhat.com

- Platform Release deleted (master)

#21 - 08/07/2018 06:36 PM - ipanova@redhat.com

- Platform Release set to 2.17.0

#22 - 08/14/2018 11:17 PM - ipanova@redhat.com

- Status changed from MODIFIED to 5

#23 - 08/31/2018 10:43 PM - ipanova@redhat.com

- Status changed from 5 to CLOSED - CURRENTRELEASE

#25 - 04/15/2019 10:12 PM - bmbouter

- Tags Pulp 2 added